

Please note that Cypress is an Infineon Technologies Company.

The document following this cover page is marked as “Cypress” document as this is the company that originally developed the product. Please note that Infineon will continue to offer the product to new and existing customers as part of the Infineon product portfolio.

Continuity of document content

The fact that Infineon offers the following product as part of the Infineon product portfolio does not lead to any changes to this document. Future revisions will occur when appropriate, and any changes will be set out on the document history page.

Continuity of ordering part numbers

Infineon continues to support existing part numbers. Please continue to use the ordering part numbers listed in the datasheet for ordering.



PSoC 64 Standard “Secure Boot MCU”-AWS: CYS0644ABZI-S2D44 Architecture Technical Reference Manual (TRM)

PSoC 64 MCU

Document No. 002-30073 Rev. **

February 15, 2021

Cypress Semiconductor
An Infineon Technologies Company
198 Champion Court
San Jose, CA 95134-1709
www.cypress.com
www.infineon.com

Copyrights

© Cypress Semiconductor Corporation, 2021. This document is the property of Cypress Semiconductor Corporation and its subsidiaries ("Cypress"). This document, including any software or firmware included or referenced in this document ("Software"), is owned by Cypress under the intellectual property laws and treaties of the United States and other countries worldwide. Cypress reserves all rights under such laws and treaties and does not, except as specifically stated in this paragraph, grant any license under its patents, copyrights, trademarks, or other intellectual property rights. If the Software is not accompanied by a license agreement and you do not otherwise have a written agreement with Cypress governing the use of the Software, then Cypress hereby grants you a personal, non-exclusive, nontransferable license (without the right to sublicense) (1) under its copyright rights in the Software (a) for Software provided in source code form, to modify and reproduce the Software solely for use with Cypress hardware products, only internally within your organization, and (b) to distribute the Software in binary code form externally to end users (either directly or indirectly through resellers and distributors), solely for use on Cypress hardware product units, and (2) under those claims of Cypress's patents that are infringed by the Software (as provided by Cypress, unmodified) to make, use, distribute, and import the Software solely for use with Cypress hardware products. Any other use, reproduction, modification, translation, or compilation of the Software is prohibited.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CYPRESS MAKES NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS DOCUMENT OR ANY SOFTWARE OR ACCOMPANYING HARDWARE, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. No computing device can be absolutely secure. Therefore, despite security measures implemented in Cypress hardware or software products, Cypress shall have no liability arising out of any security breach, such as unauthorized access to or use of a Cypress product. CYPRESS DOES NOT REPRESENT, WARRANT, OR GUARANTEE THAT CYPRESS PRODUCTS, OR SYSTEMS CREATED USING CYPRESS PRODUCTS, WILL BE FREE FROM CORRUPTION, ATTACK, VIRUSES, INTERFERENCE, HACKING, DATA LOSS OR THEFT, OR OTHER SECURITY INTRUSION (collectively, "Security Breach"). Cypress disclaims any liability relating to any Security Breach, and you shall and hereby do release Cypress from any claim, damage, or other liability arising from any Security Breach. In addition, the products described in these materials may contain design defects or errors known as errata which may cause the product to deviate from published specifications. To the extent permitted by applicable law, Cypress reserves the right to make changes to this document without further notice. Cypress does not assume any liability arising out of the application or use of any product or circuit described in this document. Any information provided in this document, including any sample design information or programming code, is provided only for reference purposes. It is the responsibility of the user of this document to properly design, program, and test the functionality and safety of any application made of this information and any resulting product. "High-Risk Device" means any device or system whose failure could cause personal injury, death, or property damage. Examples of High-Risk Devices are weapons, nuclear installations, surgical implants, and other medical devices. "Critical Component" means any component of a High-Risk Device whose failure to perform can be reasonably expected to cause, directly or indirectly, the failure of the High-Risk Device, or to affect its safety or effectiveness. Cypress is not liable, in whole or in part, and you shall and hereby do release Cypress from any claim, damage, or other liability arising from any use of a Cypress product as a Critical Component in a High-Risk Device. You shall indemnify and hold Cypress, its directors, officers, employees, agents, affiliates, distributors, and assigns harmless from and against all claims, costs, damages, and expenses, arising out of any claim, including claims for product liability, personal injury or death, or property damage arising from any use of a Cypress product as a Critical Component in a High-Risk Device. Cypress products are not intended or authorized for use as a Critical Component in any High-Risk Device except to the limited extent that (i) Cypress's published data sheet for the product explicitly states Cypress has qualified the product for use in a specific High-Risk Device, or (ii) Cypress has given you advance written authorization to use the product as a Critical Component in the specific High-Risk Device and you have signed a separate indemnification agreement.

Cypress, the Cypress logo, Spansion, the Spansion logo, and combinations thereof, WICED, PSoC, CapSense, EZ-USB, F-RAM, and Traveo are trademarks or registered trademarks of Cypress in the United States and other countries. For a more complete list of Cypress trademarks, visit cypress.com. Other names and brands may be claimed as property of their respective owners.

Content Overview



Section A: Overview	21
1. Introduction	22
1.1 Features.....	22
1.2 Architecture.....	23
1.3 CPU Subsystem (CPUSS).....	24
1.4 System Resources Subsystem (SRSS).....	24
1.5 Analog Subsystem.....	25
1.6 Programmable Digital	26
1.7 Digital Subsystem	26
2. Getting Started	27
2.1 PSoC 6 MCU Resources	27
3. Document Organization and Conventions	28
3.1 Major Sections	28
3.2 Documentation Conventions.....	28
Section B: CPU Subsystem	32
4. CPU Subsystem (CPUSS)	33
4.1 Features.....	33
4.2 Architecture.....	34
4.3 Registers.....	36
4.4 Operating Modes and Privilege Levels	38
4.5 Instruction Set.....	39
5. SRAM Controller	40
5.1 Features.....	40
5.2 Architecture.....	40
5.3 Wait States	41
6. Inter-Processor Communication	42
6.1 Features.....	42
6.2 Architecture.....	42
6.3 Implementing Locks.....	45
6.4 Message Passing	45
7. Fault Monitoring	47
7.1 Features.....	47
7.2 Architecture.....	48
7.3 Fault Sources.....	52
7.4 Register List.....	53

8. Interrupts	54
8.1 Features.....	54
8.2 Architecture.....	55
8.3 Interrupts and Exceptions - Operation	56
8.4 Exception Sources	58
8.5 Interrupt Sources	60
8.6 Interrupt/Exception Priority	67
8.7 Enabling and Disabling Interrupts.....	67
8.8 Interrupt/Exception States	68
8.9 Stack Usage for Interrupts/Exceptions	69
8.10 Interrupts and Low-Power Modes.....	69
8.11 Interrupt/Exception – Initialization/ Configuration	69
8.12 Register List.....	70
9. Protection Units	71
9.1 Architecture.....	71
9.2 PSoC 6 Protection Architecture	72
9.3 Register Architecture	74
9.4 Bus Master Protection Attributes	76
9.5 Protection Context	77
9.6 Protection Contexts 0, 1, 2, 3	78
9.7 Protection Structure	79
10. DMA Controller (DW)	93
10.1 Features.....	93
10.2 Architecture.....	94
10.3 Channels.....	94
10.4 Descriptors.....	96
10.5 DMA Controller	101
11. DMAC Controller (DMAC)	104
11.1 Features.....	104
11.2 Architecture.....	105
11.3 Channels.....	105
11.4 Descriptors.....	107
11.5 DMAC Controller.....	112
12. Cryptographic Function Block (Crypto)	113
12.1 Features.....	113
12.2 Architecture.....	113
12.3 Definitions of Terms.....	114
12.4 Crypto Block Functions.....	115
12.5 Module Configuration and Initialization.....	116
12.6 Software Design Considerations	117
13. Program and Debug Interface	118
13.1 Features.....	118
13.2 Architecture.....	118
13.3 Serial Wire Debug (SWD) Interface	121
13.4 JTAG Interface.....	123
13.5 Programming the PSoC 6 MCU.....	126
13.6 Registers.....	126

14. Nonvolatile Memory	127
14.1 Flash Memory	127
14.2 Flash Memory Programming	130
15. eFuse Memory	132
15.1 Features.....	132
15.2 Architecture.....	133
16. “Secure Boot”	134
16.1 Features.....	136
16.2 Architecture.....	137
16.3 System Calls.....	147
16.4 Amazon FreeRTOS PSA Integration	180
Section C: System Resources Subsystem (SRSS)	181
17. Power Supply and Monitoring	183
17.1 Features.....	183
17.2 Architecture.....	184
17.3 Power Supply.....	185
17.4 Voltage Monitoring.....	187
17.5 Register List	189
18. Device Power Modes	190
18.1 Features.....	190
18.2 Architecture.....	190
18.3 Power Mode Transitions	194
18.4 Summary	198
18.5 Register List.....	198
19. Backup System	200
19.1 Features.....	200
19.2 Architecture.....	201
19.3 Power Supply.....	201
19.4 Clocking	202
19.5 Reset	203
19.6 Real-Time Clock	203
19.7 Alarm Feature	204
19.8 PMIC Control	205
19.9 Backup Registers.....	206
19.10 Register List.....	206
20. Clocking System	207
20.1 Features.....	207
20.2 Architecture.....	208
20.3 Clock Sources.....	209
20.4 Clock Generation	210
20.5 Clock Trees.....	217
20.6 CLK_HF[0] Distribution	219
20.7 Peripheral Clock Dividers	219
20.8 Clock Calibration Counters	222
21. Reset System	223
21.1 Features.....	223

21.2	Architecture.....	223
21.3	Identifying Reset Sources.....	225
21.4	Register List.....	225
22.	I/O System	226
22.1	Features.....	226
22.2	Architecture.....	227
22.3	High-Speed I/O Matrix	232
22.4	I/O State on Power Up.....	234
22.5	Behavior in Low-Power Modes	234
22.6	Interrupt	234
22.7	Peripheral Connections	235
22.8	Smart I/O	236
22.9	Registers.....	247
23.	Watchdog Timer	248
23.1	Features.....	248
23.2	Architecture.....	249
23.3	Free-running WDT	249
23.4	Multi-Counter WDTs	252
23.5	Reset Cause Detection	258
23.6	Register List.....	259
24.	Trigger Multiplexer Block	260
24.1	Features.....	260
24.2	Architecture.....	260
24.3	Register List	264
25.	Profiler	266
25.1	Features.....	266
25.2	Architecture.....	267
25.3	Using the Profiler	269
Section D:	Digital Subsystem	272
26.	Secure Digital Host Controller (SDHC)	273
26.1	Features.....	273
26.2	Block Diagram	274
26.3	Clocking	275
26.4	Bus Speed Modes	276
26.5	Power Modes.....	276
26.6	Interrupts to CPU	276
26.7	I/O Interface	277
26.8	Packet Buffer SRAM.....	278
26.9	DMA Engine.....	278
26.10	Initialization Sequence.....	279
26.11	Error Detection.....	282
27.	Serial Communications Block (SCB)	284
27.1	Features.....	284
27.2	Architecture.....	285
27.3	Serial Peripheral Interface (SPI).....	286
27.4	UART	304
27.5	Inter Integrated Circuit (I2C)	318

27.6	SCB Interrupts	332
28.	Serial Memory Interface (SMIF)	341
28.1	Features.....	341
28.2	Architecture.....	341
28.3	Memory Device Signal Interface.....	347
28.4	Triggers.....	356
28.5	Interrupts.....	357
28.6	Sleep Operation.....	357
28.7	Performance	357
29.	Timer, Counter, and PWM (TCPWM)	358
29.1	Features.....	358
29.2	Architecture.....	359
29.3	Operation Modes	363
29.4	TCPWM Registers	393
30.	Inter-IC Sound Bus	394
30.1	Features.....	394
30.2	Architecture.....	395
30.3	Digital Audio Interface Formats	395
30.4	Clocking Polarity and Delay Options	399
30.5	Interfacing with Audio Codecs	400
30.6	Clocking Features.....	400
30.7	FIFO Buffer and DMA Support	402
30.8	Interrupt Support.....	404
30.9	Watchdog Timer	405
31.	PDM-PCM Converter	406
31.1	Features.....	406
31.2	Architecture.....	407
31.3	Operating Procedure	413
32.	Universal Serial Bus (USB) Device Mode	415
32.1	Features.....	415
32.2	Architecture.....	416
32.3	Operation	417
32.4	Logical Transfer Modes	421
32.5	USB Power Modes	429
32.6	USB Device Registers	429
33.	LCD Direct Drive	431
33.1	Features.....	431
33.2	Architecture.....	431
33.3	PSoC 6 MCU Segment LCD Direct Drive.....	442
33.4	Register List	443
Section E: Analog Subsystem		444
34.	Analog Reference Block	445
34.1	Features.....	445
34.2	Architecture.....	446
34.3	Registers.....	447

35. Low-Power Comparator	448
35.1 Features.....	448
35.2 Architecture.....	449
35.3 Register List	452
36. SAR ADC	453
36.1 Features.....	453
36.2 Architecture.....	454
36.3 Registers.....	466
37. Temperature Sensor	467
37.1 Features.....	467
37.2 Architecture.....	467
37.3 Temperature Sensor Configuration	469
37.4 Algorithm.....	469
37.5 Registers.....	470
38. CapSense	471

Contents



Section A: Overview	21
1. Introduction	22
1.1 Features.....	22
1.2 Architecture.....	23
1.3 CPU Subsystem (CPUSS).....	24
1.3.1 CPU	24
1.3.2 DMA Controllers.....	24
1.3.3 Flash	24
1.3.4 SRAM.....	24
1.3.5 SROM	24
1.3.6 OTP eFuse.....	24
1.3.7 Program and Debug.....	24
1.4 System Resources Subsystem (SRSS).....	24
1.4.1 Power System.....	24
1.4.2 Clocking System	24
1.4.3 GPIO	25
1.5 Analog Subsystem.....	25
1.5.1 12-bit SAR ADC	25
1.5.2 Temperature Sensor	25
1.5.3 Low-Power Comparators	25
1.5.4 CapSense	25
1.6 Programmable Digital	26
1.6.1 Smart I/O™	26
1.7 Digital Subsystem	26
1.7.1 Secure Digital Host Controller (SDHC) with eMMC	26
1.7.2 Serial Communication Blocks (SCB)	26
1.7.3 Serial Memory Interface (SMIF).....	26
1.7.4 Audio Subsystem	26
1.7.5 Timer/Counter/PWM Block.....	26
2. Getting Started	27
2.1 PSoC 6 MCU Resources	27
3. Document Organization and Conventions	28
3.1 Major Sections	28
3.2 Documentation Conventions.....	28
3.2.1 Register Conventions.....	28
3.2.2 Numeric Naming	28
3.2.3 Units of Measure.....	29
3.2.4 Acronyms and Initializations	29

Section B: CPU Subsystem	32
4. CPU Subsystem (CPUSS)	33
4.1 Features.....	33
4.2 Architecture.....	34
4.2.1 Address and Memory Maps	34
4.3 Registers.....	36
4.4 Operating Modes and Privilege Levels	38
4.5 Instruction Set.....	39
5. SRAM Controller	40
5.1 Features.....	40
5.2 Architecture.....	40
5.3 Wait States	41
6. Inter-Processor Communication	42
6.1 Features.....	42
6.2 Architecture.....	42
6.2.1 IPC Channel.....	42
6.2.2 IPC Interrupt.....	43
6.2.3 IPC Channels and Interrupts.....	44
6.3 Implementing Locks.....	45
6.4 Message Passing	45
7. Fault Monitoring	47
7.1 Features.....	47
7.2 Architecture.....	48
7.2.1 Fault Report	48
7.2.2 Signaling Interface	50
7.2.3 Monitoring	50
7.2.4 Low-power Mode Operation.....	51
7.2.5 Using a Fault Structure	51
7.2.6 CPU Exceptions Versus Fault Monitoring	51
7.3 Fault Sources.....	52
7.4 Register List.....	53
8. Interrupts	54
8.1 Features.....	54
8.2 Architecture.....	55
8.3 Interrupts and Exceptions - Operation	56
8.3.1 Interrupt/Exception Handling.....	56
8.3.2 Level and Pulse Interrupts	57
8.3.3 Exception Vector Table	57
8.4 Exception Sources.....	58
8.4.1 Reset Exception.....	58
8.4.2 Non-Maskable Interrupt Exception.....	59
8.4.3 HardFault Exception	59
8.4.4 Memory Management Fault Exception	59
8.4.5 Bus Fault Exception	60
8.4.6 Usage Fault Exception.....	60
8.4.7 Supervisor Call (SVCall) Exception	60
8.4.8 PendSupervisory (PendSV) Exception	60
8.4.9 System Tick (SysTick) Exception	60
8.5 Interrupt Sources	60

8.6	Interrupt/Exception Priority	67
8.7	Enabling and Disabling Interrupts	67
8.8	Interrupt/Exception States	68
8.8.1	Pending Interrupts/Exceptions	68
8.9	Stack Usage for Interrupts/Exceptions	69
8.10	Interrupts and Low-Power Modes	69
8.11	Interrupt/Exception – Initialization/ Configuration	69
8.12	Register List	70
9.	Protection Units	71
9.1	Architecture	71
9.2	PSoC 6 Protection Architecture	72
9.3	Register Architecture	74
9.3.1	Protection Structure and Attributes	74
9.4	Bus Master Protection Attributes	76
9.5	Protection Context	77
9.6	Protection Contexts 0, 1, 2, 3	78
9.7	Protection Structure	79
9.7.1	Protection Violation	79
9.7.2	MPU	79
9.7.3	SMPU	80
9.7.4	PPU	80
9.7.5	Protection of Protection Structures	87
9.7.6	Protection Structure Types	89
10.	DMA Controller (DW)	93
10.1	Features	93
10.2	Architecture	94
10.3	Channels	94
10.3.1	Channel Interrupts	95
10.4	Descriptors	96
10.4.1	Address Configuration	98
10.4.2	Transfer Size	99
10.4.3	Descriptor Chaining	100
10.5	DMA Controller	101
10.5.1	Trigger Selection	101
10.5.2	Pending Triggers	101
10.5.3	Output Triggers	101
10.5.4	Status registers	101
10.5.5	DMA Performance	102
11.	DMAC Controller (DMAC)	104
11.1	Features	104
11.2	Architecture	105
11.3	Channels	105
11.3.1	Channel Interrupts	106
11.4	Descriptors	107
11.4.1	Address Configuration	109
11.4.2	Transfer Size	111
11.4.3	Descriptor Chaining	111
11.5	DMAC Controller	112
11.5.1	Trigger Selection	112
11.5.2	Channel Logic	112

11.5.3	Output Triggers	112
12.	Cryptographic Function Block (Crypto)	113
12.1	Features.....	113
12.2	Architecture.....	113
12.3	Definitions of Terms.....	114
12.4	Crypto Block Functions.....	115
12.4.1	Symmetric Key Functions	115
12.4.2	Hash Functions	115
12.4.3	Message Authentication Code (MAC) Functions	116
12.4.4	Cyclic Redundancy Code (CRC)	116
12.4.5	Random Number Generator (RNG).....	116
12.5	Module Configuration and Initialization.....	116
12.6	Software Design Considerations	117
13.	Program and Debug Interface	118
13.1	Features.....	118
13.2	Architecture.....	118
13.2.1	Debug Access Port (DAP).....	120
13.2.2	ROM Tables	120
13.2.3	Trace.....	120
13.2.4	Embedded Cross Triggering	121
13.3	Serial Wire Debug (SWD) Interface.....	121
13.3.1	SWD Timing Details	122
13.3.2	ACK Details.....	122
13.3.3	Turnaround (Trn) Period Details	123
13.4	JTAG Interface.....	123
13.5	Programming the PSoC 6 MCU.....	126
13.5.1	SWD Port Acquisition.....	126
13.5.2	SWD Programming Mode Entry.....	126
13.5.3	SWD Programming Routine Executions	126
13.6	Registers.....	126
14.	Nonvolatile Memory	127
14.1	Flash Memory	127
14.1.1	Features	127
14.1.2	Configuration.....	127
14.1.3	Flash Geometry	128
14.1.4	Flash Controller.....	129
14.1.5	Read While Write (RWW) Support.....	130
14.2	Flash Memory Programming	130
14.2.1	Features	130
14.2.2	Architecture.....	130
15.	eFuse Memory	132
15.1	Features.....	132
15.2	Architecture.....	133
16.	“Secure Boot”	134
16.1	Features.....	136
16.2	Architecture.....	137
16.2.1	Life Cycle Stages and Protection States.....	137
16.2.2	Boot Sequence	142
16.2.3	Cypress “Secure Bootloader”.....	143

16.2.4	Flash Security	144
16.2.5	Hardware-Based Encryption	144
16.2.6	Dedicated Hardware	144
16.3	System Calls	147
16.3.1	Implementation	147
16.3.2	System Call APIs	148
16.3.3	System Call Status	178
16.4	Amazon FreeRTOS PSA Integration	180
Section C: System Resources Subsystem (SRSS)		181
17. Power Supply and Monitoring		183
17.1	Features	183
17.2	Architecture	184
17.3	Power Supply	185
17.3.1	Regulators Summary	185
17.3.2	Power Pins and Rails	187
17.3.3	Power Sequencing Requirements	187
17.3.4	Backup Domain	187
17.3.5	Power Supply Sources	187
17.4	Voltage Monitoring	187
17.4.1	Power-On-Reset (POR)	187
17.4.2	Brownout-Detect (BOD)	187
17.4.3	Low-Voltage-Detect (LVD)	188
17.4.4	Over-Voltage Protection (OVP)	189
17.5	Register List	189
18. Device Power Modes		190
18.1	Features	190
18.2	Architecture	190
18.2.1	CPU Power Modes	192
18.2.2	System Power Modes	192
18.2.3	System Deep Sleep Mode	192
18.2.4	System Hibernate Mode	193
18.2.5	Other Operation Modes	193
18.3	Power Mode Transitions	194
18.3.1	Power-up Transitions	195
18.3.2	Power Mode Transitions	195
18.3.3	Wakeup Transitions	197
18.4	Summary	198
18.5	Register List	198
19. Backup System		200
19.1	Features	200
19.2	Architecture	201
19.3	Power Supply	201
19.4	Clocking	202
19.4.1	WCO with External Clock/Sine Wave Input	202
19.4.2	Calibration	202
19.5	Reset	203
19.6	Real-Time Clock	203
19.6.1	Reading RTC User Registers	203
19.6.2	Writing to RTC User Registers	203
19.7	Alarm Feature	204

19.8	PMIC Control	205
19.9	Backup Registers	206
19.10	Register List	206
20.	Clocking System	207
20.1	Features	207
20.2	Architecture	208
20.3	Clock Sources	209
20.3.1	Internal Main Oscillator (IMO)	209
20.3.2	External Crystal Oscillator (ECO)	209
20.3.3	External Clock (EXTCLK)	210
20.3.4	Internal Low-speed Oscillator (ILO)	210
20.3.5	Watch Crystal Oscillator (WCO)	210
20.4	Clock Generation	210
20.4.1	Phase-Locked Loop (PLL)	210
20.4.2	Frequency Lock Loop (FLL)	211
20.5	Clock Trees	217
20.5.1	Path Clocks	217
20.5.2	High-Frequency Root Clocks	217
20.5.3	Low-Frequency Clock	218
20.5.4	Timer Clock	218
20.5.5	Group Clocks (clk_sys)	218
20.5.6	Backup Clock (clk_bak)	218
20.6	CLK_HF[0] Distribution	219
20.6.1	CLK_FAST	219
20.6.2	CLK_PERI	219
20.6.3	CLK_SLOW	219
20.7	Peripheral Clock Dividers	219
20.7.1	Fractional Clock Dividers	219
20.7.2	Peripheral Clock Divider Configuration	219
20.8	Clock Calibration Counters	222
21.	Reset System	223
21.1	Features	223
21.2	Architecture	223
21.2.1	Power-on Reset	223
21.2.2	Brownout Reset	224
21.2.3	Watchdog Timer Reset	224
21.2.4	Software Initiated Reset	224
21.2.5	External Reset	224
21.2.6	Logic Protection Fault Reset	224
21.2.7	Clock-Supervision Logic Reset	224
21.2.8	Hibernate Wakeup Reset	224
21.3	Identifying Reset Sources	225
21.4	Register List	225
22.	I/O System	226
22.1	Features	226
22.2	Architecture	227
22.2.1	I/O Cell Architecture	228
22.2.2	Digital Input Buffer	228
22.2.3	Digital Output Driver	229
22.3	High-Speed I/O Matrix	232

22.4	I/O State on Power Up.....	234
22.5	Behavior in Low-Power Modes.....	234
22.6	Interrupt	234
22.7	Peripheral Connections	235
22.7.1	Firmware-Controlled GPIO	235
22.7.2	Analog I/O	236
22.7.3	LCD Drive	236
22.7.4	CapSense	236
22.8	Smart I/O	236
22.8.1	Overview	236
22.8.2	Block Components.....	237
22.8.3	Routing.....	244
22.8.4	Operation	246
22.9	Registers.....	247
23.	Watchdog Timer	248
23.1	Features.....	248
23.2	Architecture.....	249
23.3	Free-running WDT	249
23.3.1	Overview	249
23.3.2	Watchdog Reset	251
23.3.3	Watchdog Interrupt	251
23.4	Multi-Counter WDTs	252
23.4.1	Overview	252
23.4.2	Enabling and Disabling WDT	256
23.4.3	Watchdog Cascade Options	256
23.4.4	Watchdog Reset	258
23.4.5	Watchdog Interrupt	258
23.5	Reset Cause Detection	258
23.6	Register List.....	259
24.	Trigger Multiplexer Block	260
24.1	Features.....	260
24.2	Architecture.....	260
24.2.1	Trigger Multiplexer Group	261
24.2.2	One-to-one Trigger.....	261
24.2.3	Trigger Multiplexer Block	261
24.2.4	Software Triggers.....	263
24.3	Register List	264
25.	Profiler	266
25.1	Features.....	266
25.2	Architecture.....	267
25.2.1	Profiler Design	267
25.2.2	Available Monitoring Sources	268
25.2.3	Reference Clocks.....	268
25.3	Using the Profiler	269
25.3.1	Enable or Disable the Profiler	269
25.3.2	Configure and Enable a Counter	270
25.3.3	Start and Stop Profiling	270
25.3.4	Handle Counter Overflow.....	270
25.3.5	Get the Results	271
25.3.6	Exit Gracefully.....	271

Section D: Digital Subsystem	272
26. Secure Digital Host Controller (SDHC)	273
26.1 Features.....	273
26.1.1 Features Not Supported.....	274
26.2 Block Diagram	274
26.3 Clocking	275
26.3.1 Clock Gating	275
26.3.2 Base Clock (CLK_HF[i]) Configuration	275
26.3.3 Card Clock (SDCLK) Configuration	275
26.3.4 Timeout (TOUT) Configuration.....	275
26.4 Bus Speed Modes	276
26.5 Power Modes.....	276
26.5.1 Standby Mode.....	276
26.6 Interrupts to CPU	276
26.6.1 SDIO Interrupt.....	277
26.7 I/O Interface	277
26.7.1 Switching Signaling Voltage from 3.3 V to 1.8 V	278
26.8 Packet Buffer SRAM.....	278
26.8.1 Packet Buffer Full/Empty	278
26.9 DMA Engine.....	278
26.10 Initialization Sequence.....	279
26.10.1 Enabling SDHC.....	279
26.10.2 Card Detection	280
26.10.3 SDHC Initialization	282
26.10.4 Clock Setup.....	282
26.11 Error Detection.....	282
27. Serial Communications Block (SCB)	284
27.1 Features.....	284
27.2 Architecture.....	285
27.2.1 Buffer Modes.....	285
27.2.2 Clocking Modes	285
27.3 Serial Peripheral Interface (SPI).....	286
27.3.1 Features.....	286
27.3.2 General Description	287
27.3.3 SPI Modes of Operation.....	288
27.3.4 SPI Buffer Modes.....	293
27.3.5 Clocking and Oversampling	298
27.3.6 Enabling and Initializing SPI	300
27.3.7 I/O Pad Connection.....	301
27.3.8 SPI Registers	303
27.4 UART	304
27.4.1 Features.....	304
27.4.2 General Description	304
27.4.3 UART Modes of Operation.....	304
27.4.4 Clocking and Oversampling	314
27.4.5 Enabling and Initializing the UART	314
27.4.6 I/O Pad Connection.....	315
27.4.7 UART Registers	317
27.5 Inter Integrated Circuit (I2C)	318
27.5.1 Features.....	318
27.5.2 General Description	318

27.5.3	External Electrical Connections	319
27.5.4	Terms and Definitions	320
27.5.5	I2C Modes of Operation	320
27.5.6	I2C Buffer Modes	322
27.5.7	Clocking and Oversampling	326
27.5.8	Enabling and Initializing the I2C	329
27.5.9	I/O Pad Connections	330
27.5.10	I2C Registers	331
27.6	SCB Interrupts	332
27.6.1	SPI Interrupts	333
27.6.2	UART Interrupts	335
27.6.3	I2C Interrupts	339
28.	Serial Memory Interface (SMIF)	341
28.1	Features	341
28.2	Architecture	341
28.2.1	Tx and Rx FIFOs	343
28.2.2	MMIO Mode	344
28.2.3	XIP Mode	344
28.2.4	Cache	345
28.2.5	Arbitration	345
28.2.6	Deselect Delay	346
28.2.7	Cryptography	346
28.3	Memory Device Signal Interface	347
28.3.1	Specifying Memory Devices	347
28.3.2	Connecting SPI Memory Devices	348
28.3.3	SPI Data Transfer	353
28.3.4	Example of Setting up SMIF	354
28.4	Triggers	356
28.5	Interrupts	357
28.6	Sleep Operation	357
28.7	Performance	357
29.	Timer, Counter, and PWM (TCPWM)	358
29.1	Features	358
29.2	Architecture	359
29.2.1	Enabling and Disabling Counters in TCPWM Block	359
29.2.2	Clocking	359
29.2.3	Trigger Inputs	360
29.2.4	Trigger Outputs	361
29.2.5	Interrupts	361
29.2.6	PWM Outputs	362
29.2.7	Power Modes	362
29.3	Operation Modes	363
29.3.1	Timer Mode	364
29.3.2	Capture Mode	370
29.3.3	Quadrature Decoder Mode	373
29.3.4	Pulse Width Modulation Mode	377
29.3.5	Pulse Width Modulation with Dead Time Mode	387
29.3.6	Pulse Width Modulation Pseudo-Random Mode (PWM_PR)	390
29.4	TCPWM Registers	393

30. Inter-IC Sound Bus	394
30.1 Features.....	394
30.2 Architecture.....	395
30.3 Digital Audio Interface Formats	395
30.3.1 Standard I2S Format.....	395
30.3.2 Left Justified (LJ) Format	398
30.3.3 Time Division Multiplexed (TDM) Format.....	398
30.4 Clocking Polarity and Delay Options	399
30.5 Interfacing with Audio Codecs	400
30.6 Clocking Features.....	400
30.7 FIFO Buffer and DMA Support	402
30.8 Interrupt Support.....	404
30.9 Watchdog Timer	405
31. PDM-PCM Converter	406
31.1 Features.....	406
31.2 Architecture.....	407
31.2.1 Enable/Disable Converter	407
31.2.2 Clocking Features	407
31.2.3 Over-Sampling Ratio.....	408
31.2.4 Mono/Stereo Microphone Support	408
31.2.5 Hardware FIFO Buffers and DMA Controller Support.....	410
31.2.6 Interrupt Support.....	411
31.2.7 Digital Volume Gain	412
31.2.8 Smooth Gain Transition	412
31.2.9 Soft Mute.....	412
31.2.10 Word Length and Sign Bit Extension	412
31.2.11 High-Pass Filter	412
31.2.12 Enable/Disable Streaming	413
31.2.13 Power Modes	413
31.3 Operating Procedure	413
31.3.1 Initial Configuration	413
31.3.2 Interrupt Service Routine (ISR) Configuration	413
31.3.3 Enabling / Disabling Streaming.....	414
32. Universal Serial Bus (USB) Device Mode	415
32.1 Features.....	415
32.2 Architecture.....	416
32.2.1 USB Physical Layer (USB PHY)	416
32.2.2 Serial Interface Engine (SIE)	416
32.2.3 Arbiter	416
32.3 Operation	417
32.3.1 USB Clocking Scheme.....	417
32.3.2 USB PHY	417
32.3.3 Endpoints	418
32.3.4 Transfer Types	418
32.3.5 Interrupt Sources	418
32.3.6 DMA Support.....	420
32.4 Logical Transfer Modes	421
32.4.1 Manual Memory Management with No DMA Access	423
32.4.2 Manual Memory Management with DMA Access.....	423
32.4.3 Automatic DMA Mode	425
32.4.4 Control Endpoint Logical Transfer.....	427

32.5	USB Power Modes	429
32.6	USB Device Registers	429
33.	LCD Direct Drive	431
33.1	Features.....	431
33.2	Architecture.....	431
33.2.1	LCD Segment Drive Overview	431
33.2.2	Drive Modes.....	432
33.2.3	Recommended Usage of Drive Modes	441
33.2.4	Digital Contrast Control.....	441
33.3	PSoC 6 MCU Segment LCD Direct Drive	442
33.3.1	High-Speed and Low-Speed Master Generators	442
33.3.2	Multiplexer and LCD Pin Logic.....	443
33.3.3	Display Data Registers	443
33.4	Register List	443
Section E:	Analog Subsystem	444
34.	Analog Reference Block	445
34.1	Features.....	445
34.2	Architecture.....	446
34.2.1	Bandgap Reference Block	447
34.2.2	Zero Dependency To Absolute Temperature Current Generator (IZTAT).....	447
34.2.3	Reference Selection Multiplexers	447
34.2.4	Startup Modes.....	447
34.2.5	Low-Power Modes	447
34.3	Registers.....	447
35.	Low-Power Comparator	448
35.1	Features.....	448
35.2	Architecture.....	449
35.2.1	Input Configuration.....	449
35.2.2	Output and Interrupt Configuration	449
35.2.3	Power Mode and Speed Configuration	450
35.2.4	Hysteresis	451
35.2.5	Wakeup from Low-Power Modes	451
35.2.6	Comparator Clock	452
35.3	Register List	452
36.	SAR ADC	453
36.1	Features.....	453
36.2	Architecture.....	454
36.2.1	SAR ADC Core	455
36.2.2	SARMUX.....	460
36.2.3	SARREF	461
36.2.4	SARSEQ.....	462
36.2.5	SAR Interrupts	463
36.2.6	Trigger.....	465
36.2.7	SAR ADC Status	465
36.3	Registers.....	466
37.	Temperature Sensor	467
37.1	Features.....	467
37.2	Architecture.....	467

37.3	Temperature Sensor Configuration	469
37.4	Algorithm.....	469
37.5	Registers.....	470
38.	CapSense	471

Section A: Overview



This section encompasses the following chapters:

- [Introduction chapter on page 22](#)
- [Getting Started chapter on page 27](#)
- [Document Organization and Conventions chapter on page 28](#)

Document Revision History

Revision	Issue Date	Description of Change
**	02/15/2021	Initial version of PSoC 64 TRM

1. Introduction



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation
- [Application notes](#)
- [Code examples](#)

The PSoC[®] MCU is a scalable and reconfigurable platform architecture that supports a family of programmable embedded system controllers with Arm[®] Cortex[®] CPUs (single and multi-core). A Cortex-M0+ is used for secure operations and by default requires no code to be provided by the developer. For custom security configurations the developer may choose to add code or make changes to the default operation. The CYS0644ABZI-S2D44 product family is a combination of a dual-core microcontroller with built-in programmable peripherals. It incorporates integrated low-power flash technology, digital programmable logic, high-performance analog-to-digital and digital-to-analog conversion, Low-Power comparators, touch sensing, serial memory interface with encryption, secure digital host controller (SDHC), and standard communication and timing peripherals. The PSoC 64 family features out-of-box SecureBoot functionality with a hardware-based Root of Trust (RoT) to accelerate secure designs.

1.1 Features

The CYS0644ABZI-S2D44 product family has these characteristics:

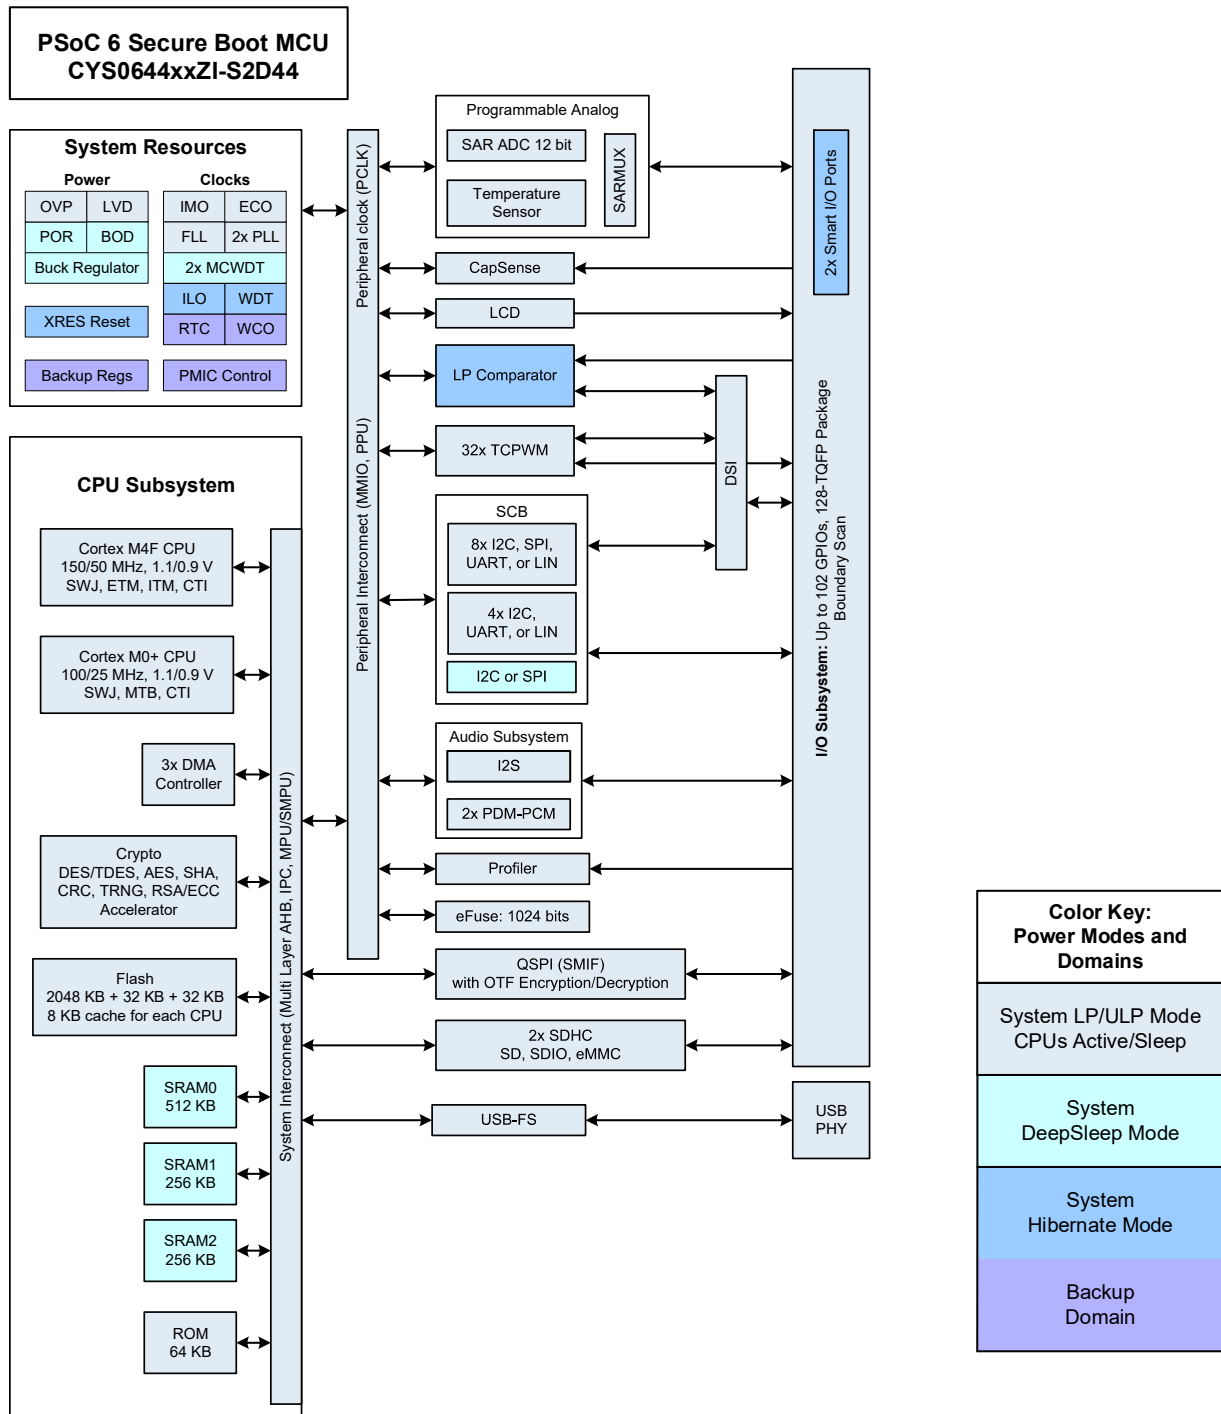
- 32-bit dual core (Arm Cortex-M4F and Arm Cortex M0+) CPU subsystem
- Hardware accelerated Crypto functions
- Secure boot system
- Hardware-based Root of Trust
- Integrated (on-chip) flash memory
- Audio subsystem with I²S interface and two PDM channels
- Serial memory interface with on-the-fly encryption and decryption
- SDHC block that supports SD, SDIO, and eMMC interfaces
- Low-power modes
- Configurable digital peripherals
- Programmable digital logic
- High-performance analog system
- Flexible and programmable interconnect
- Capacitive touch sensing (CapSense[®])
- Programmable GPIOs

This document describes each function block of the CYS0644ABZI-S2D44 in detail. In this document, PSoC 6 MCU refers to CYS0644ABZI-S2D44 unless explicitly mentioned otherwise.

1.2 Architecture

Figure 1-1 shows the major components of the PSoC 6 MCU architecture.

Figure 1-1. PSoC 6 MCU Architecture Block Diagram



The block diagram shows the device subsystems and gives a simplified view of their interconnections. The color-code shows the lowest power mode where the particular block is still functional (for example, LP comparator is functional in Deep Sleep and Hibernate modes).

1.3 CPU Subsystem (CPUSS)

1.3.1 CPU

The CPU subsystem in PSoC 6 MCUs consists of two Arm Cortex cores and their associated buses and memories: M4 with floating-point unit (FPU) and memory protection unit (MPU), and M0+ with an MPU. The Cortex M0+ provides a secure, uninterruptible boot function. This guarantees that post-boot, system integrity is checked and privileges enforced. Shared resources can be accessed through the normal Arm multi-layer bus arbitration. Exclusive accesses are supported by an inter-processor communication (IPC) scheme, which implements hardware semaphores and protection.

1.3.2 DMA Controllers

PSoC 6 MCUs have DMA controllers that support independent access to peripherals using the multilayer advanced high-performance bus (AHB).

1.3.3 Flash

PSoC 6 Cortex-M4 have a flash module with one block that can be used for EEPROM emulation for longer retention. It also has a block of flash that can be securely locked and is accessible only via a key lock that cannot be changed (one-time programmable). The flash block supports Read-While-Write (RWW) operation so that flash updates may be performed while the CPU is active.

1.3.4 SRAM

PSoC 6 MCUs have an SRAM module that can be retained in Deep Sleep mode either fully or in increments of user-designated blocks.

1.3.5 SROM

PSoC 6 MCUs have a supervisory ROM that contains boot and configuration routines. This ROM guarantees secure boot if authentication of user flash is required.

1.3.6 OTP eFuse

The OTP memory is a nonvolatile memory. When a bit is changed from the default "0" state to "1", it cannot be changed back to "0". The eFuse is not user-accessible and is used to store system-level variables, modes, and memory hashes used to validate the secure boot sequence.

1.3.7 Program and Debug

PSoC 6 MCUs have extensive support for programming, testing, debugging, and tracing both hardware and firmware. Complete debug-on-chip functionality enables full device debugging in the final system using the standard production

device. It does not require special interfaces, debugging pods, simulators, or emulators. Only the standard programming connections are required to fully support debug. The ModusToolbox integrated development environment (IDE) provides fully-integrated programming and debug support for PSoC 6 MCUs. The SWJ (SWD and JTAG) interface is fully compatible with industry-standard third-party probes. With the ability to disable debug features, with robust flash protection, and by allowing customer-proprietary functionality to be implemented in on-chip programmable blocks, the PSoC 6 MCU family provides a high level of security. Additionally, all device interfaces can be permanently disabled (device security) for applications concerned about phishing attacks due to a maliciously reprogrammed device or attempts to defeat security by starting and interrupting flash programming sequences. All programming, debug, and test interfaces are disabled when maximum device security is enabled. By default, a fixed bootloader (CySecureBootloader) is programmed during the provisioning phase. This can be replaced during provisioning by the developer, if required. Debugging is disabled by default for the Cortex-M0+ CPU for security reasons, but the developer may enable it, if required.

1.4 System Resources Subsystem (SRSS)

1.4.1 Power System

The power system confirms that voltage levels meet the requirements for the respective mode and will either delay mode entry (on power-on reset, for example) until voltage levels meet requirements or generate resets (brownout detect) when the power supply drops below specified levels. The design guarantees safe chip operation between power supply voltage dropping below specified levels and the reset. The VDD power supply feeds an on-chip LDO, which produces the core logic supply. In addition, the device includes an on-chip buck regulator that can be used to power the core.

1.4.2 Clocking System

The PSoC 6 MCU clock system provides clocks to subsystems that require clocks and switches between different clock sources without glitches. In addition, the clock system ensures that no metastable conditions occur. The clock system for PSoC 6 MCU consists of the internal main oscillator (IMO), the internal low-speed oscillator (ILO), the precision internal low-speed oscillator (PILO), the external crystal oscillator, and the watch crystal oscillator (WCO). Two phase-locked loops (PLL) and one frequency-locked loop (FLL) are used to generate high-speed clocks from either the IMO or the crystal oscillator or from an external clock supplied from a pin. The PLLs and FLL enable independent clock frequencies for peripherals.

1.4.2.1 IMO Clock Source

The IMO is the primary source of internal clocking in the PSoC 6 MCU. It is trimmed during testing to achieve the specified accuracy.

1.4.2.2 ILO Clock Source

The ILO is a very low-power oscillator, which may be used to generate clocks for peripheral operation in Deep Sleep mode. ILO-driven counters can be calibrated to the IMO to improve accuracy. Cypress provides a software component, which does the calibration.

1.4.2.3 Watchdog Timer

A watchdog timer is implemented in the clock block running from the ILO. This allows watchdog operation during Deep Sleep and Hibernate modes, and generates a watchdog reset if not serviced before the timeout occurs. The watchdog reset is recorded in the Reset Cause register.

1.4.2.4 Clock Dividers

Integer and fractional clock dividers are provided for peripheral use and timing purposes. The clock dividers are 16 and 24 bits in length to allow very fine clock control.

1.4.2.5 Reset

PSoC 6 MCUs can be reset from a variety of sources including a software reset. Reset events are asynchronous and guarantee reversion to a known state. The reset cause is recorded in a register, which allows the software to determine the cause of the reset. An XRES pin is reserved for external reset to avoid complications with configuration and multiple pin functions during power-on or reconfiguration.

1.4.3 GPIO

The GPIO pins are organized in logical entities called ports, which are eight bits in width. During power-on and reset, the blocks are forced to the disable state so as not to crowbar any inputs and/or cause excess turn-on current. A multiplexing network known as a high-speed I/O matrix (HSIOM) is used to multiplex between various signals that may connect to an I/O pin. Data output and pin state registers store, respectively, the values to be driven on the pins and the states of the pins themselves.

Every I/O pin can generate an interrupt if so enabled and each I/O port has an interrupt request (IRQ) and interrupt service routine (ISR) vector associated with it. Four GPIO pins are capable of overvoltage tolerant (OVT) operation where the input voltage may be higher than VDD (these may be used for I²C functionality to allow powering the chip off while maintaining physical connection to an operating I²C bus without affecting its functionality).

1.5 Analog Subsystem

1.5.1 12-bit SAR ADC

PSoC 6 MCUs have a 12-bit SAR ADC. The SAR is connected to a fixed set of pins through an eight-input sequencer. The sequencer cycles through the selected channels autonomously (sequencer scan) and does so with zero switching overhead (that is, the aggregate sampling bandwidth remains the same whether it is for a single channel or distributed over several channels). The sequencer switching is effected through a state machine or through firmware-driven switching. The sequencer supports the buffering of each channel to reduce CPU interrupt-service requirements. To accommodate signals with varying source impedances and frequencies, different sample times can be programmed for each channel. Also, the signal range specification through a pair of range registers (low- and high-range values) is implemented with a corresponding out-of-range interrupt if the digitized value exceeds the programmed range; this allows fast detection of out-of-range values without having to wait for a sequencer scan to be completed and the CPU to read the values and check for out-of-range values in software.

The SAR is able to digitize the output of the on-chip temperature sensor for calibration and other temperature-dependent functions. The SAR is not available in Deep Sleep and Hibernate modes because it requires a high-speed clock.

1.5.2 Temperature Sensor

The PSoC 6 MCU has an on-chip temperature sensor. This consists of a diode, which is biased by a current source that can be disabled to save power. The temperature sensor is connected to the ADC, which digitizes the reading and produces a temperature value by using a Cypress-supplied software that includes calibration and linearization.

1.5.3 Low-Power Comparators

PSoC 6 MCUs have a pair of Low-Power comparators, which can operate in Deep Sleep and Hibernate modes. This allows the analog system blocks to be disabled while retaining the ability to monitor external voltage levels in Deep Sleep and Hibernate modes. The comparator outputs are normally synchronized to avoid metastability unless operating in an asynchronous power mode (Hibernate) where the system wake-up circuit is activated by a comparator-switch event.

1.5.4 CapSense

The CapSense system, used primarily for touch sensing, can measure the self-capacitance of an electrode or the mutual capacitance between a pair of electrodes. CapSense provides industry's best-in-class signal-to-noise ratio (SNR),

high touch sensitivity, low-power operation, and superior EMI performance. CapSense touch sensing also supports liquid-tolerant operation using a driven shield signal. Any analog-capable GPIO can be used as a sensor or shield electrode.

In addition to capacitive sensing, the CapSense system can function as an ADC to measure voltage on any GPIO pin that supports the CapSense functionality. Moreover, If the CapSense block is not used for touch sensing or ADC functionality, a CapSense comparator and the two 8-bit IDACs can be used as general-purpose analog blocks. See the [PSoC 4 and PSoC 6 MCU CapSense Design Guide](#) for more details.

1.6 Programmable Digital

1.6.1 Smart I/O™

The PSoC 6 MCU has two Smart I/O blocks, which allow Boolean operations on signals going to the GPIO pins from the device subsystems or on signals coming into the device. Operation can be synchronous or asynchronous and the blocks operate in low-power modes, such as Deep Sleep and Hibernate. This allows, for example, detection of logic conditions that can indicate that the CPU should wake up instead of waking up on general I/O interrupts, which consume more power and can generate spurious wakeups.

1.7 Digital Subsystem

1.7.1 Secure Digital Host Controller (SDHC) with eMMC

The SDHC controller provides connectivity to IoT Wireless devices and external storage. The SDHC controller conforms to SDIO and eMMC specifications and have ADMA3 engines built in.

1.7.2 Serial Communication Blocks (SCB)

PSoC 6 MCU SCBs can implement communication interfaces such as I²C, UART, or SPI.

1.7.2.1 I²C Mode

The hardware I²C block implements a full multimaster and slave interface (it is capable of multimaster arbitration). This block has flexible buffering options to reduce the interrupt overhead and latency for the CPU. It also supports EzI²C, which creates a mailbox address range in the PSoC 6 MCU memory and effectively reduces the I²C communication to reading from and writing to an array in the memory. In addition, the block supports a 256-byte FIFO for receive and transmit, which, by increasing the time given for the CPU to read the data, reduces the need for clock stretching caused by the CPU not having read the data on time. The FIFO

mode is available in all channels and is useful in the absence of DMA.

The I²C peripheral is compatible with I²C Standard-mode, Fast-mode, and Fast-mode Plus devices as defined in the NXP I²C-bus specification and user manual (UM10204). The I²C bus I/O is implemented with GPIO in open-drain modes.

1.7.2.2 UART Mode

This is a full-feature UART that supports automotive single-wire interface (LIN), infrared interface (IrDA), and SmartCard (ISO7816) protocols, all of which are minor variants of the basic UART protocol. In addition, it supports the multiprocessor mode that allows the addressing of peripherals connected over common Rx and Tx lines. Common UART functions such as parity error, break detect, and frame error are supported. A 256-byte FIFO tolerates much greater CPU service latencies.

1.7.2.3 SPI Mode

The SPI mode supports full Motorola SPI, TI Secure Simple Pairing (SSP) (essentially adds a start pulse that is used to synchronize SPI codecs), and National Microwire (half-duplex form of SPI). The SPI block can use the FIFO and supports an EzSPI mode in which the data interchange is reduced to reading and writing an array in memory.

1.7.3 Serial Memory Interface (SMIF)

A serial memory interface has selectable 1-, 2-, or 4-bit widths. This block also supports on-the-fly encryption and decryption to support Execute-In-Place operation.

1.7.4 Audio Subsystem

This subsystem consists of I²S and PDM blocks. The PDM channels interface to a PDM microphone's bit-stream output.

1.7.5 Timer/Counter/PWM Block

The timer/counter/PWM block consists of counters with user-programmable period length. It has a capture register, which records the count value of an event (such as an I/O event), a period register, which is used to either stop or auto-reload the counter when its count is equal to the period register, and compare registers to generate compare value signals, which are used as PWM duty cycle outputs. The block also provides true and complementary outputs with programmable offset between them to allow the use as deadband programmable complementary PWM outputs. It also has a kill input to force outputs to a predetermined state; for example, this is used in motor-drive systems when an overcurrent state is indicated and the PWMs driving the FETs must be shut off immediately with no time for software intervention.

2. Getting Started



2.1 PSoC 6 MCU Resources

This chapter provides the complete list of PSoC 6 MCU resources that helps you get started with the device and design your applications with them. If you are new to PSoC, Cypress provides a wealth of data at www.cypress.com to help you to select the right PSoC device and quickly and effectively integrate it into your design.

The following is an abbreviated list of PSoC 6 MCU resources:

- Overview: [PSoC Portfolio](#), [PSoC Roadmap](#), [PSoC 6 MCU webpage](#)
- Product Selectors: See the [PSoC 6 MCU Product Selector Guide](#) to choose a part that suits your application. In addition, [ModusToolbox](#) includes a similar device selection tool to select devices for ModusToolbox projects.
- [Datasheets](#) describe and provide electrical specifications for each device family.
- [Application Notes](#) and [Code Examples](#) cover a broad range of topics, from basic to advanced level. Many of the application notes include code examples, which can be opened from ModusToolbox.
- [Technical Reference Manuals \(TRMs\)](#) provide detailed descriptions of the architecture and registers in each device family.
- [CapSense Design Guide](#): Learn how to design capacitive touch-sensing applications with PSoC devices.
- Development Tools
 - [ModusToolbox](#) enables cross platform code development with a robust suite of tools and software libraries.
 - [CY8CPROTO-062-4343W Prototyping Kit](#) provides PSoC 6 MCU hardware support with Wi-Fi and Bluetooth connectivity. Additional kits enable development for other PSoC 6 MCU devices with a variety of connectivity options.
 - PSoC 6 CAD libraries provide footprint and schematic support for common tools.
- Additional Resources: Visit the [PSoC 6 MCU webpage](#) for additional resources such as IBIS, BSDL models, CAD Library Files, and Programming Specifications.
- Training Videos: Visit www.cypress.com/training for a wide variety of video training resources on PSoC devices and ModusToolbox, such as:
 - [Welcome to Modus Toolbox](#)
 - [PSoC 6 BLE Pioneer Kit](#)
 - [Introduction to BLE](#)
- Technical Support
 - [Forum](#): See if your question is already answered by fellow developers of the [PSoC 6 community](#).
 - Cypress support: Visit our [support](#) page or contact a [local sales representative](#).

3. Document Organization and Conventions



This document includes the following sections:

- [Section B: CPU Subsystem on page 32](#)
- [Section C: System Resources Subsystem \(SRSS\) on page 181](#)
- [Section D: Digital Subsystem on page 272](#)
- [Section E: Analog Subsystem on page 444](#)

3.1 Major Sections

For ease of use, information is organized into sections and chapters that are divided according to device functionality.

- Section – Presents the top-level architecture, how to get started, and conventions and overview information of the product.
- Chapter – Presents the chapters specific to an individual aspect of the section topic. These are the detailed implementation and use information for some aspect of the integrated circuit.
- Glossary – Defines the specialized terminology used in this technical reference manual (TRM). Glossary terms are presented in bold, italic font throughout.
- Registers Technical Reference Manual – Supplies all device register details summarized in the technical reference manual. This is an additional document.

3.2 Documentation Conventions

This document uses only four distinguishing font types, besides those found in the headings.

- The first is the use of *italics* when referencing a document title or file name.
- The second is the use of ***bold italics*** when referencing a term described in the Glossary of this document.
- The third is the use of Times New Roman font, distinguishing equation examples.
- The fourth is the use of `Courier New` font, distinguishing code examples.

3.2.1 Register Conventions

Register conventions are detailed in the [registers TRM](#).

3.2.2 Numeric Naming

Hexadecimal numbers are represented with all letters in uppercase with an appended lowercase 'h' (for example, '14h' or '3Ah') and hexadecimal numbers may also be represented by a '0x' prefix, the C coding convention. Binary numbers have an appended lowercase 'b' (for example, '01010100b' or '01000011b'). Numbers not indicated by an 'h' or 'b' are decimal.

3.2.3 Units of Measure

This table lists the units of measure used in this document.

Table 3-1. Units of Measure

Abbreviation	Unit of Measure
bps	bits per second
°C	degrees Celsius
dB	decibels
dBm	decibels-milliwatts
fF	femtofarads
G	Giga
Hz	Hertz
k	kilo, 1000
K	kilo, 2 ¹⁰
KB	1024 bytes, or approximately one thousand bytes
Kbit	1024 bits
kHz	kilohertz (32.000)
kΩ	kilohms
MHz	megahertz
MΩ	megaohms
μA	microamperes
μF	microfarads
μs	microseconds
μV	microvolts
μVrms	microvolts root-mean-square
mA	milliamperes
ms	milliseconds
mV	millivolts
nA	nanoamperes
ns	nanoseconds
nV	nanovolts
Ω	ohms
pF	picofarads
pp	peak-to-peak
ppm	parts per million
SPS	samples per second
σ	sigma: one standard deviation
V	volts

3.2.4 Acronyms and Initializations

This table lists the acronyms and initializations used in this document

Table 3-2. Acronyms and Initializations

Acronym	Definition
ABUS	analog output bus
AC	alternating current
ADC	analog-to-digital converter
ADV	advertising
AES	Advanced Encryption Standard
AHB	AMBA (advanced microcontroller bus architecture) high-performance bus, an Arm data transfer bus
API	application programming interface
APOR	analog power-on reset
BC	broadcast clock
BCD	binary coded decimal
BESL	best effort service latency
BOD	brownout detect
BOM	bill of materials
BR	bit rate
BRA	bus request acknowledge
BRQ	bus request
CAN	controller area network
CI	carry in
CIC	cascaded integrator comb
CMAC	cipher-based message authentication code
CMP	compare
CO	carry out
COM	LCD common signal
CPHA	clock phase
CPOL	clock polarity
CPU	central processing unit
CPUSS	CPU subsystem
CRC	cyclic redundancy check
CSD	CapSense sigma delta
CSX	CapSense cross-point
CT	cipher text
CTI	cross triggering interface
CTM	cross triggering matrix
ESR	equivalent series resistance
DAC	digital-to-analog converter
DAP	debug access port
DC	direct current
DES	Data Encryption Standard
DFF	D flip-flop

Table 3-2. Acronyms and Initializations (*continued*)

Acronym	Definition
DI	digital or data input
DL	drive level
DMA	direct memory access
DMIPS	Dhrystone million instructions per second
DNL	differential nonlinearity
DO	digital or data output
DSI	digital system interconnect
DSP	digital signal processing
DSM	Deep Sleep mode
DU	data unit
DW	data wire
ECC	Elliptic-curve cryptography
ECDH	Elliptic-curve Diffie-Hellman
ECDSA	Elliptic-curve Digital Signature Algorithm
ECO	external crystal oscillator
EEPROM	electrically erasable programmable read only memory
EMIF	external memory interface
ETM	embedded trace macrocell
FB	feedback
FIFO	first in first out
FPU	floating point unit
FSR	full scale range
GAP	generic access profile
GATT	generic attribute profile
GFSK	Gaussian frequency-shift keying
GPIO	general-purpose I/O
HCI	host-controller interface
HFCLK	high-frequency clock
HMAC	hashed message authentication code
HPF	high-pass filter
HSIOM	high-speed I/O matrix
I ² C	inter-integrated circuit
I ² S	inter-IC sound
IDE	integrated development environment
ILO	internal low-speed oscillator
ITO	indium tin oxide
IMO	internal main oscillator
INL	integral nonlinearity
I/O	input/output
IOR	I/O read
IOW	I/O write
IPC	inter-processor communication

 Table 3-2. Acronyms and Initializations (*continued*)

Acronym	Definition
IPTAT	proportional to absolute temperature
IRES	initial power on reset
IRA	interrupt request acknowledge
IRK	identity resolution key
IRQ	interrupt request
ISA	instruction set architecture
ISR	interrupt service routine
ITM	instrumentation trace macrocell
IVR	interrupt vector read
IZTAT	zero dependency to absolute temperature
JSON	JavaScript object notation
JTAG	Joint Test Action Group
JWT	JSON web token
L2CAP	logical link control and adaptation protocol
LCD	liquid crystal display
LFCLK	low-frequency clock
LFSR	linear feedback shift register
LIN	local interconnect network
LJ	left justified
LL	link layer
LNA	low-noise amplifier
LP	system low-power mode
LPCOMP	Low-Power comparator
LPM	link power management
LR	link register
LRb	last received bit
LRB	last received byte
LSb	least significant bit
LSB	least significant byte
LUT	lookup table
MAC	message authentication code
MISO	master-in-slave-out
MMIO	memory mapped input/output
MOSI	master-out-slave-in
MPU	memory protection unit
MSb	most significant bit
MSB	most significant byte
MSP	main stack pointer
MTB	micro trace buffer
NI	next instant
NMI	non-maskable interrupt
NVIC	nested vectored interrupt controller
OE	output enable

Table 3-2. Acronyms and Initializations (*continued*)

Acronym	Definition
OSR	over-sampling ratio
OVP	over-voltage protection
PA	power amplifier
PC	program counter
PCB	printed circuit board
PCH	program counter high
PCL	program counter low
PD	power down
PDU	protocol data unit
PGA	programmable gain amplifier
PHY	physical layer
PLD	programmable logic device
PM	power management
PMA	PSoC memory arbiter
POR	power-on reset
PPOR	precision power-on reset
PPU	peripheral protection units
PRNG	pseudo random number generator
PRS	pseudo random sequence
PSA	Platform Security Architecture
PSoC	Programmable System-on-Chip
PSP	process stack pointer
PSR	program status register
PSRR	power supply rejection ratio
PSSDC	power system sleep duty cycle
PWM	pulse width modulator
RAM	random-access memory
RETI	return from interrupt
RF	radio frequency
RNG	random number generator
ROM	read only memory
ROT	root of trust
RPA	resolvable private address
RMS	root mean square
RW	read/write
SAR	successive approximation register
SARSEQ	SAR sequencer
SEG	LCD segment signal
SE0	single-ended zero
SC	switched capacitor
SCB	serial communication block
SHA-256	Secure Hash Algorithm
SIE	serial interface engine

 Table 3-2. Acronyms and Initializations (*continued*)

Acronym	Definition
SIMO	single input multiple output
SIO	special I/O
SNR	signal-to-noise ratio
SMPU	shared memory protection units
SOF	start of frame
SOI	start of instruction
SP	stack pointer
SPD	sequential phase detector
SPI	serial peripheral interconnect
SPIM	serial peripheral interconnect master
SPIS	serial peripheral interconnect slave
SRAM	static random-access memory
SROM	supervisory read only memory
SRSS	system resources subsystem
SSADC	single slope ADC
SSC	supervisory system call
SVCall	supervisor call
SYSClk	system clock
SWD	single wire debug
SWV	serial wire viewer
TAR	turn-around time
TC	terminal count
TCPWM	timer, counter, PWM
TD	transaction descriptors
TDM	time division multiplexed
TFF	toggle flip-flop
TIA	trans-impedance amplifier
TPIU	trace port interface unit
TRM	technical reference manual
TRNG	True random number generator
UART	universal asynchronous receiver/transmitter
ULB	system ultra low-power mode
USB	universal serial bus
USBIO	USB I/O
VTOR	vector table offset register
WCO	watch crystal oscillator
WDT	watchdog timer
WDR	watchdog reset
WIC	wakeup interrupt controller
XRES	external reset
XRES_N	external reset, active low

Section B: CPU Subsystem




This section encompasses the following chapters:

- CPU Subsystem (CPUSS) chapter on page 33
- SRAM Controller chapter on page 40
- Inter-Processor Communication chapter on page 42
- Fault Monitoring chapter on page 47
- Interrupts chapter on page 54
- Protection Units chapter on page 71
- DMA Controller (DW) chapter on page 93
- DMAC Controller (DMAC) chapter on page 104
- Cryptographic Function Block (Crypto) chapter on page 113
- Program and Debug Interface chapter on page 118
- Nonvolatile Memory chapter on page 127
- eFuse Memory chapter on page 132
- “Secure Boot” chapter on page 134

Top Level Architecture

Figure B-1. CPU System Block Diagram



4. CPU Subsystem (CPUSS)



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation
- Application notes
- Code examples

The CPU subsystem is based on dual 32-bit Arm Cortex CPUs, as [Figure 4-1](#) shows. The Cortex-M4 is the main CPU. It is designed for short interrupt response time, high code density, and high 32-bit throughput while maintaining a strict cost and power consumption budget. A secondary Cortex-M0+ CPU implements security, safety, and protection features.

The designer will implement the application using only the Cortex-M4 CPU. The Cortex-M0+ CPU implements all secure functions including initial boot sequence, setting up protection policy, and writing to nonvolatile memories, eFuse, and Flash. A secure bootloader is part of the default boot sequence. It allows device firmware to be securely updated by whatever interface the designer provides. In most cases, this does not include the device debug ports, which are usually disabled to secure the device.

This section provides only an overview of the Arm Cortex CPUs in PSoC 6 MCUs. For details, see the Arm documentation sets for [Cortex-M4](#) and [Cortex-M0+](#). Some PSoC 6 MCU parts have only one CPU. See the [device datasheet](#) for details.

4.1 Features

The PSoC 6 MCU Arm Cortex CPUs have the following features:


- Cortex-M4 has a [floating-point unit](#) (FPU) that supports single-cycle digital signal processing (DSP) instructions, and a [memory protection unit](#) (MPU). Cortex-M0+ has an [MPU](#).
- Both CPUs have 8-KB instruction caches with four-way set associativity.¹
- Maximum clock frequency of 150 MHz for the Cortex-M4 and 100 MHz for the Cortex-M0+.
- The Cortex-M4 implements a version of the Thumb instruction set based on Thumb-2 technology (defined in the *Armv7-M Architecture Reference Manual*). The Cortex-M0+ supports the Armv6-M Thumb instruction set (defined in the *Armv6-M Architecture Reference Manual*). See “[Instruction Set](#)” on [page 39](#).
- Both CPUs have [nested vectored interrupt controllers](#) (NVIC) for rapid and deterministic interrupt response. For details, see the [Interrupts chapter on page 54](#).
- Both CPUs have extensive debug support. For details, see the [Program and Debug Interface chapter on page 118](#).
 - SWJ: combined serial wire debug (SWD) and Joint Test Action Group (JTAG) ports
 - Serial wire viewer (SWV): provides real-time trace information through the serial wire output (SWO) interface
 - Breakpoints
 - Watchpoints

1. PSoC 6 does not support cache coherency. As a result when a particular row of flash that executes instructions is written/updated, the updated information will not be reflected in the cache. The cache should be cleared in the firmware during such instances. This is applicable for both CM4 and CM0+ cache – the appropriate (CM0+ and/or CM4) cache should be invalidated.

- Trace: Cortex-M4: embedded trace macrocell (ETM).
Cortex-M0+: 4-KB micro trace buffer (MTB)
- Inter-processor communication (IPC) hardware – see the [Inter-Processor Communication chapter on page 42](#).

4.2 Architecture

Figure 4-1. CPU Subsystem Block Diagram



Each CPU is a 32-bit processor with its own 32-bit datapath and a 32-bit memory interface. Each CPU has its own set of 32-bit registers. They support a wide variety of instructions in the Thumb instruction set. They support two operating modes (see [“Operating Modes and Privilege Levels” on page 38](#)).

The Cortex-M4 instruction set includes:

- Signed and unsigned, $32 \times 32 \rightarrow 32$ -bit and $32 \times 32 \rightarrow 64$ -bit, multiply and multiply-accumulate, all single-cycle
- Signed and unsigned 32-bit divides that take two to 12 cycles
- DSP instructions
- Complex memory-load and store access
- Complex bit manipulation; see the bitfield instructions in [Table 4-6](#)

The Cortex-M4 FPU has its own set of registers and instructions. It is compliant with the ANSI/IEEE Std 754-2008, IEEE Standard for Binary Floating-Point Arithmetic.

The Cortex-M0+ has a single cycle $32 \times 32 \rightarrow 32$ -bit signed multiplication instruction.

4.2.1 Address and Memory Maps

Both CPUs have a fixed address map, with shared access to memory and peripherals. The 32-bit (4 GB) address space is divided into the regions shown in [Table 4-1](#). Note that code can be executed from the code and SRAM regions.

Table 4-1. Address Map for Cortex-M4 and Cortex-M0+

Address Range	Name	Use
0x00000000 – 0x1FFFFFFF	Code	Program code region. You can also put data here. It includes the exception vector table, which starts at address 0.
0x20000000 – 0x3FFFFFFF	SRAM	Data region. You can also execute code from this region. Note that the Cortex-M4 bit-band in this region is not supported in PSoC 6.
0x40000000 – 0x5FFFFFFF	Peripheral	All peripheral registers. Code cannot be executed from this region. Note that the Cortex-M4 bit-band in this region is not supported in PSoC 6.
0x60000000 – 0x9FFFFFFF	External RAM	Not used
0xA0000000 – 0xDFFFFFFF	External device	Not used
0xE0000000 – 0xE00FFFFF	Private peripheral bus (PPB)	Provides access to peripheral registers within the CPU core.
0xE0100000 – 0xFFFFFFFF	Device	Device-specific system registers.

The device memory map shown in [Table 4-2](#) applies to both CPUs. That is, the CPUs share access to all PSoC 6 MCU memory and peripheral registers.

Table 4-2. PSoC 6 Memory Map

Address Range	Size	Name	Comments
0x0000_0000 - 0x0000_FFFF	64K	SROM	Initial ROM boot code
0x0800_0000 - 0x0800_17FF	6K	System SRAM	Bootloader SRAM
0x0800_1800 - 0x080E_BFFF	938K	User SRAM	Available for user CM4 application
0x080E_C000 - 0x080F_FFFF	80K	System SRAM	Used for secure boot, secure bootloader, and system calls (CM0+)
0x1000_0000 - 0x101C_FFFF	1856K	User Application Flash	Available for user CM4 application
0x101D_0000 - 0x101D_FFFF	64K	Bootloader	Secure bootloader
0x101E_0000 - 0x101F_FFFF	128K	Secure Code Flash	Used for secure boot, secure bootloader, and system calls (CM0+)
0x1400_0000 - 0x1400_7FFF	32K	Auxiliary Flash	EEPROM emulation (CM4)
0x1600_0000 - 0x1600_7FFF	32K	Supervisory Flash	Secure access (CM0+)

SRAM is located in the code region for both CPUs (see [Table 4-1](#)). This facilitates executing code out of SRAM. There is no physical memory located in the CPUs' SRAM region.

Note: The PSoC 6 registers CPUSS_CM0_VECTOR_TABLE_BASE and CPUSS_CM4_VECTOR_TABLE_BASE determine the location of the vector table for each CPU. A number of LS bits in each register are set to 0. As a result, there are restrictions on the location of vector tables – they must be on a 256-byte boundary for CM0+ and a 1024-byte boundary for CM4.

4.2.1.1 Wait State Lookup Tables

The wait state lookup tables show the wait states for Flash, SRAM, and ROM based on the Clk_HF0 frequency and the current power mode. SRAM and ROM have two domains for the wait states – fast clock domain (Clk_Fast) and slow clock domain (Clk_Slow); both domains are based off Clk_HF0. The following tables show the wait states for the slow clock domain. All wait states for the fast clock domain are zero. For more information on clocking see the [Clocking System chapter on page 207](#).

	Ultra-Low Power Mode	Clk_HF0 (MHz)		
		Clk_HF0 ≤ 25	25 < Clk_HF0 ≤ 100	100 < Clk_HF0
ROM/SRAM	True	0	1	1
	False	0	0	1

	Ultra-Low Power Mode	Clk_HF0 (MHz)				
		Clk_HF0 ≤ 16		16 < Clk_HF0 ≤ 33		33 < Clk_HF0
Flash	True	0		1		2
		Clk_HF0 ≤ 29		29 < Clk_HF0 ≤ 58		58 < Clk_HF0 ≤ 87
	False	0	1	2	3	4

4.3 Registers

Both CPUs have sixteen 32-bit registers, as [Table 4-3](#) shows. See the Arm documentation for details.

- R0 to R12 – General-purpose registers. R0 to R7 can be accessed by all instructions; the other registers can be accessed by a subset of the instructions.
- R13 – Stack pointer (SP). There are two stack pointers, with only one available at a time. In thread mode, the CONTROL register indicates the stack pointer to use – Main Stack Pointer (MSP) or Process Stack Pointer (PSP).
- R14 – Link register. Stores the return program counter during function calls.
- R15 – Program counter. This register can be written to control program flow.

Table 4-3. Cortex-M4 and Cortex-M0+ Registers

Name	Type ^a	Reset Value	Description
R0 – R12	RW	Undefined	R0–R12 are 32-bit general-purpose registers for data operations.
MSP (R13)	RW	[0x00000000]	The stack pointer (SP) is register R13. In thread mode, bit[1] of the CONTROL register indicates the stack pointer to use: 0 = Main stack pointer (MSP). This is the reset value. 1 = Process stack pointer (PSP). On reset, the processor loads the MSP with the value from the vector address.
PSP (R13)			
LR (R14)	RW	See note ^b	The link register (LR) is register R14. It stores the return information for subroutines, function calls, and exceptions.
PC (R15)	RW	[0x00000004]	The program counter (PC) is register R15. It contains the current program address. On reset, the processor loads the PC with the value from the vector address plus 0x00000004. Bit[0] of the value is loaded into the EPSR T-bit (see Table 4-4) at reset; it must always be 1.
PSR	RW	Undefined	The program status register (PSR) combines: Application Program Status Register (APSR). Execution Program Status Register (EPSR). Interrupt Program Status Register (IPSR).
APSR	RW	Undefined	The APSR contains the current state of the condition flags from previous instruction executions.
EPSR	RO	0x01000000	On reset, the EPSR Thumb state bit is loaded with the value bit[0] of the register [0x00000004]. It must always be 1. In Cortex-M4, other bits in this register control the state of interrupt-continuable instructions and the if-then (IT) instruction.
IPSR	RO	0	The IPSR contains the current exception number.
PRIMASK	RW	0	The PRIMASK register prevents activation of all exceptions with configurable priority.
CONTROL	RW	0	The CONTROL register controls: - The privilege level in Thread mode; see 4.4 Operating Modes and Privilege Levels . - The currently active stack pointer, MSP or PSP. - Cortex-M4 only: whether to preserve the floating-point state when processing an exception.
FAULTMASK	RW	0	Cortex-M4 only. Bit 0 = 1 prevents the activation of all exceptions except NMI.
BASEPRI	RW	0	Cortex-M4 only. When set to a nonzero value, prevents processing any exception with a priority greater than or equal to the value.

a. Describes access type during program execution in thread mode and handler mode. Debug access can differ.

b. LR reset value is 0xFFFFFFFF in Cortex-M4, undefined in Cortex-M0+.

The Cortex-M4 floating-point unit (FPU) also has the following registers:

- Thirty-two 32-bit single-precision registers, S0 to S31. These registers can also be addressed as sixteen 64-bit double-precision registers, D0 to D15.
- Five FPU control and status registers:
 - CPACR – Coprocessor Access Control Register
 - FPCCR – Floating-point Context Control Register
 - FPCAR – Floating-point Context Address Register
 - FPSCR – Floating-point Status Control Register
 - FPDSCR – Floating-point Default Status Control Register

For more information on how these registers are used, see the Arm Cortex-M4 documentation.

Use the MSR and MRS instructions to access the PSR, PRIMASK, CONTROL, FAULTMASK, and BASEPRI registers. [Table 4-4](#) and [Table 4-5](#) show how the PSR bits are assigned.

Table 4-4. Cortex-M4 PSR Bit Assignments

Bit	PSR Register	Name	Usage
31	APSR	N	Negative flag
30	APSR	Z	Zero flag
29	APSR	C	Carry or borrow flag
28	APSR	V	Overflow flag
27	APSR	Q	DSP overflow and saturation flag
26 – 25	EPSR	IC/IT	Control interrupt-continuable and IT instructions
24	EPSR	T	Thumb state bit. Must always be 1. Attempting to execute instructions when the T bit is 0 results in a HardFault exception.
23 – 20	–	–	Reserved
19 – 16	APSR	GE	Greater than or equal flags, for the SEL instruction
15 – 10	EPSR	IC/IT	Control interrupt-continuable and IT instructions
9	–	–	Reserved
8 – 0	IPSR	ISR_NUMBER	Exception number of current ISR: 0 = thread mode 1 = reserved 2 = NMI 3 = HardFault 4 = MemManage 5 = BusFault 6 = UsageFault 7 – 10 = reserved 11 = SVCall 12 = reserved for debug 13 = reserved 14 = PendSV 15 = SysTick (see “System Tick (SysTick) Exception” on page 60) 16 = IRQ0 ... 255 = IRQ240

Table 4-5. Cortex-M0+ PSR Bit Assignments

Bit	PSR Register	Name	Usage
31	APSR	N	Negative flag
30	APSR	Z	Zero flag
29	APSR	C	Carry or borrow flag
28	APSR	V	Overflow flag
27 – 25	–	–	Reserved
24	EPSR	T	Thumb state bit. Must always be 1. Attempting to execute instructions when the T bit is 0 results in a HardFault exception.
23 – 6	–	–	Reserved
5 – 0	IPSR	Exception Number	Exception number of current ISR: 0 = thread mode 1 = reserved 2 = NMI 3 = HardFault 4 – 10 = reserved 11 = SVCall 12, 13 = reserved 14 = PendSV 15 = SysTick 16 = IRQ0 ... 47 = IRQ31

4.4 Operating Modes and Privilege Levels

Both CPUs support two operating modes and two privilege levels:

■ Operating Modes:

- ❑ Thread Mode – used to execute application software. The processor enters Thread mode when it comes out of reset.
- ❑ Handler Mode – used to handle exceptions. The processor returns to Thread mode when it has finished all exception processing.

■ Privilege Levels:

- ❑ Unprivileged – the software has limited access to the MSR and MRS instructions, and cannot use the CPSID and CPSIE instructions. It cannot access the system timer, NVIC, or system control block. It may have restricted access to memory or peripherals.
- ❑ Privileged – the software can use all the instructions and has access to all resources.

In Thread mode, the CONTROL register controls whether software execution is privileged or unprivileged. In Handler mode, software execution is always privileged.

Only privileged software can write to the CONTROL register to change the privilege level. Unprivileged software can use the SVC instruction to transfer control to privileged software.

In Handler mode, the MSP is always used. The exception entry and return mechanisms automatically update the CONTROL register, which may change whether MSP/PSP is used.

In Thread mode, use the MSR instruction to set the stack pointer bit in the CONTROL register. When changing the stack pointer, use an ISB instruction immediately after the MSR instruction. This ensures that instructions after the ISB execute using the new stack pointer.

4.5 Instruction Set

Both CPUs implement subsets of the Thumb instruction set, as [Table 4-6](#) shows. The table does not show the large number of variants and conditions of the instructions. For details, see one of the Arm Cortex Generic User Guides or Technical Reference Manuals.

An instruction operand can be a register, a constant, or another instruction-specific parameter. Instructions act on the operands and often store the result in a destination register. Many instructions have restrictions on using the PC or SP for the operands or destination register. See the Arm documentation for details.

Table 4-6. Instruction Set Summary – Cortex-M4 and Cortex-M0+

Functional Group	Cortex-M4	Cortex-M0+	Brief List of Instruction Mnemonics
Memory access	✓	✓	LDR, STR, ADR, PUSH, POP
General data processing	✓	✓	Cortex-M0+: ADD, ADC, AND, ASR, BICS, CMN, CMP, EOR, LSL, LSR, MOV, MVNS, ORR, REV, ROR, RSB, SBC, SUB, SXT, UXT, TST Cortex-M4 has all of the above plus: CLZ, ORN, RRX, SADD, SAS, SSA, SSUB, TEQ, UADD, UAS, USA, USUB
Multiply and divide	✓	MUL only	MLA, MLS, MUL, SDIV, SMLA, SMLS, SMMLA, SMMLS, SMUA, SMUL, SMUS, UDIV, UMAAL, UMLAL, UMULL
Saturating	✓	–	SSAT, USAT, QADD, QSUB, QASX, QSAX, QDADD, QDSUB, UQADD, UQASX, UQSAX, UQSUB
Packing and unpacking	✓	–	PKH, SXT, SXTA, UXT, UXTA
Bitfield	✓	–	BFC, BFI, SBFX, UBFX
Branch and control	✓	✓	Cortex-M0+: B{cc}, BL, BLX, BX Cortex-M4 has all of the above plus: CBNZ, CBZ, IT, TB
Miscellaneous	✓	✓	CPSID, CPSIE, DMB, DSB, ISB, MRS, MSR, NOP, SEV, SVC, WFE, WFI Cortex-M4 has all of the above plus BKPT
Floating-point	✓	–	VABS, VADD, VCMP, VCVT, VDIV, VFMA, VFNMA, VFMS, VFNMS, VLD, VLMA, VLMS, VMOV, VMRS, VMSR, VMUL, VNEG, VNMLA, VNMLS, VNMUL, VPOP, V PUSH, VSQRT, VST, VSUB

5. SRAM Controller



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation
- Application notes
- Code examples

This chapter explains the PSoC 6 MCU SRAM Controller, its features, architecture, and wait states. The SRAM controller enables the CPU to read and write parts of the PSoC 6 SRAM.

5.1 Features

The CPUSS has up to three identical SRAM controllers; see the [device datasheet](#) for details.

The SRAM controller has the following features:

- Consists of two AHB-Lite interfaces:
 - An AHB-Lite bus interface on `clk_fast` that connects to the fast bus infrastructure
 - An AHB-Lite bus interface on `clk_slow` that connects to the slow bus infrastructure
- Supports programmable number of `clk_hf` wait states
- Supports 8-, 16-, and 32-bit accesses

5.2 Architecture

The design has two AHB-Lite interfaces that connect to the AHB-Lite infrastructure. Each AHB-Lite interface is connected to a synchronization component that translates between the interface clock (either `clk_fast` or `clk_slow`) and the high-frequency clock (`clk_hf`).

Arbitration is performed on the AHB-Lite transfers from the two ports (AHB-Lite interface). Arbitration uses device-wide bus master specific arbitration priorities. Therefore, although two AHB-Lite interfaces are provided, only one AHB-Lite transfer is accepted by the port arbitration component.

The AHB-Lite transfers are the origin for all SRAM accesses; that is, the write buffer and SRAM repair requests result from AHB-Lite transfers. The SRAM controller differentiates between the following three types of AHB-Lite transfers:

- AHB-Lite read transfers
- 32-bit AHB-Lite write transfers
- 8-bit and 16-bit AHB-Lite write transfers (also referred to as partial AHB-Lite write transfers)

Each type is described in more detail here.

AHB-Lite read transfers. An AHB-Lite read transfer is translated into an SRAM read access. If the read address matches in the write buffer, the SRAM has stale data and the write data provides the requested read data (this functionality is provided by the read merge component).

32-bit AHB-Lite write transfers. A 32-bit AHB-Lite write transfer is translated into an SRAM write access. If the write address matches in the write buffer, the matching write buffer entries have stale data and these entries are invalidated.

Partial AHB-Lite write transfers. A partial AHB-Lite write transfer is translated into an SRAM read access and an SRAM write access. The SRAM read access is the direct result of the partial write transfer and the SRAM write access is the result of a write buffer request. A partial write transfer requires an SRAM read access to retrieve the “missing” data bytes from the SRAM. If the read address matches in the write buffer, the SRAM has stale data and the write data provides the requested read data (this functionality is provided by the read merge component). The requested read data is merged with the partial write data to provide a complete 32-bit data word (this functionality is provided by the write merge component). The address and the merged write data are written to the write buffer. A future write buffer request results in an SRAM write access with the merged write data.

Only the partial AHB-Lite write transfers use the write buffer.

Write buffer. The write buffer is a temporary holding station for future SRAM write accesses.

The buffer allows SRAM write accesses to be postponed. This allows for more performance critical AHB-Lite requests to “overtake” write buffer requests. Memory consistency is guaranteed by matching the SRAM access address with the write buffer entries' addresses: a “matching” SRAM read access uses the read merge component and a matching SRAM write access invalidates the matching write buffer entries.

When the write buffer is full, an entry needs to be freed to accommodate future partial AHB-Lite write transfers. Therefore, a full write buffer raises the priority of the write buffer request path.

The write buffer is constructed as a FIFO with four entries (the order in which entries are written is the same as the order in which entries are read). Each entry consists of:

- A valid field
- An invalidated field
- A word address
- A 32-bit data word

Note that the merged write data written to the write buffer is always a 32-bit data word. Therefore, no byte mask is required.

When the write buffer is written (an entry is added): the entry valid field is set to '1' and the invalidated field is set to '0'.

When the write buffer is read (an entry is removed): the entry valid field is set to '0'. If the entry invalidate field is '1', the write buffer request path is selected for an SRAM write

access. If the entry valid field is '0', no SRAM access is performed.

On an SRAM read access, a matching entry provides write buffer merge data for the read merge component.

On an SRAM write access resulting from a 32-bit AHB-Lite write transfer, a matching entry invalidated field is set to '1'.

The state of the write buffer is reflected by `RAMi_STATUS.WB_EMPTY`. The write buffer is not retained in Deep Sleep power mode. Therefore, when transitioning to system Deep Sleep power mode, the write buffer should be empty. Note that this requirement is typically met, because a transition to Deep Sleep power mode also requires that there are no outstanding AHB-Lite transfers. If there are no outstanding AHB-Lite transfers, the write buffer gets SRAM access.

5.3 Wait States

The programmable wait states represent the number of `clk_hf` cycles for a read path through the SRAM memory to flipflops in either the fast domain (CM4 CPU) or slow domain (such as CM0+ CPU, DataWire, and DMA controller).

As the wait states are represented in `clk_hf` cycles, the wait states do not have to be reprogrammed when the fast clock domain frequency (`clk_fast`) or slow clock domain frequency (`clk_slow`) is changed. However, it may be necessary to reprogram the wait states when the high-frequency clock domain (`clk_hf`) is changed. This means the required number of wait states is a function of the `clk_hf` frequency.

The fast clock domain is timing closed at a higher frequency than the slow clock domain. Therefore, the read path through the SRAM memory to flipflops in the fast domain is faster than the read path through the SRAM memory to flipflops in the slow domain. In other words, the required number of “fast” wait states (`RAMi_CTL.FAST_WS`) should be less than or equal to the required number of “slow” wait states (`RAMi_CTL.SLOW_WS`).

The SRAM controller also has internal SRAM read paths. These paths are to flipflops in the SRAM controller in the high-frequency clock domain (`clk_hf`). For these SRAM accesses (for example, an SRAM read access to support a partial AHB-Lite write transfer), the fast wait states are used. This is because the maximum fast domain frequency (`clk_fast`) equals the high-frequency domain frequency (`clk_hf`).

6. Inter-Processor Communication



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - IPC
- Application notes
- Code examples

Inter-processor communication (IPC) provides the functionality for multiple processors to communicate and synchronize their activities. The IPC interface can also be used to communicate between two separate threads or tasks on a single core. IPC hardware is implemented using two register structures.

- IPC Channel: Communication and synchronization between processors is achieved using this structure.
- IPC Interrupt: Each interrupt structure configures an interrupt line, which can be triggered by a 'notify' or 'release' event of any IPC channel.

6.1 Features

The features of IPC are as follows:

- Implements locks for mutual exclusion between processors
- Allows sending messages between processors
- Supports up to 16 channels for communication
- Supports up to 16 interrupts, which can be triggered using notify or release events from the channels
- Allows a method to send messages between two or more tasks with just one or multiple processors

6.2 Architecture

6.2.1 IPC Channel

An IPC channel is implemented as six hardware registers, as shown in [Figure 6-1](#). The IPC channel registers are accessible to all the processors in the system.

- **IPC_ACQUIRE**: This register determines the lock feature of the IPC. The IPC channel is acquired by reading this register. If the SUCCESS field returns a '1', the read acquired the lock.
If the SUCCESS field returns a '0', the read did not acquire the lock.

Note that a single read access performs two functions:

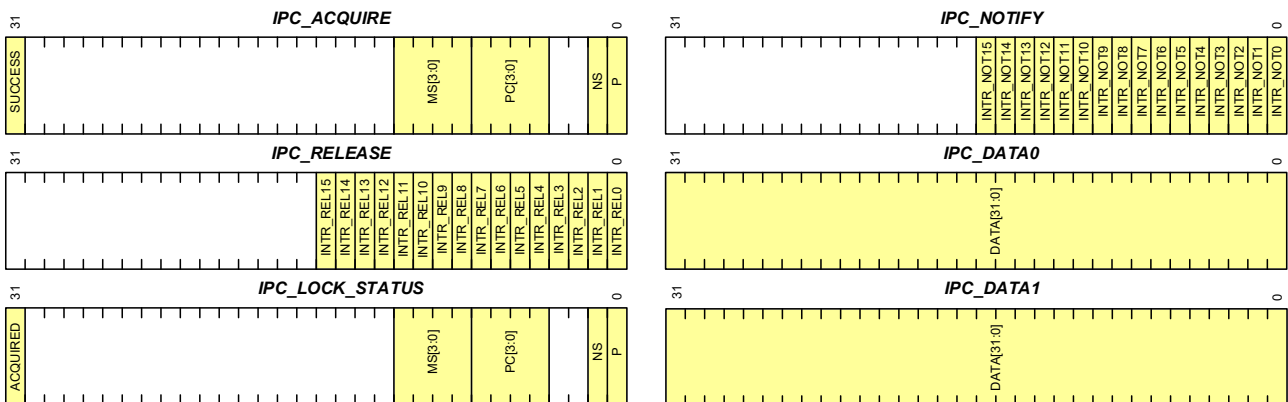
- ☐ The attempt to acquire a lock.
- ☐ Return the result of the acquisition attempt (SUCCESS field).

The atomicity of these two functions is essential in a CPU with multiple tasks that can preempt each other.

The register also has bitfields that provide information about the processor that acquired it. When acquired, this register is released by writing any value into the IPC_RELEASE register. If the register was already in an acquired state another attempt to read the register will not be able to acquire it.

- **IPC_NOTIFY:** This register is used to generate an IPC notify event. Each bit in this register corresponds to an IPC interrupt structure. The notify event generated from an IPC channel can trigger any or multiple interrupt structures.
- **IPC_RELEASE:** Any write to this register will release the IPC channel. This register also has a bit that corresponds to each IPC interrupt structure. The release event generated from an IPC channel can trigger any or multiple interrupt structures. To only release the IPC channel and not generate an interrupt, you can write a zero into the IPC release register.
- **IPC_DATA0 and IPC_DATA1:** These are two 32-bit registers with a combined size of 64 bits that are meant to hold data. These registers can be considered as the shared data memory for the channel. Typically, these registers will hold messages that need to be communicated between processors. If the messages are larger than the combined 64-bit size, place pointers in one or both of these registers.
- **IPC_LOCK_STATUS:** This register provides the instantaneous lock status for the IPC channel. The register provides details if the channel is acquired. If acquired, it provides the processor's ID, protection context, and other details. The reading of lock status provides only an instantaneous status, which can be changed in the next cycle based on the activity of other processors on the channel.

Figure 6-1. IPC Channel Structure



6.2.2 IPC Interrupt

Each IPC interrupt line in the system has a corresponding IPC interrupt structure. An IPC interrupt can be triggered by a notify or a release event from any of the IPC channels in the system. You can choose to mask any of the sources of these events using the IPC interrupt registers. Figure 6-2 shows the registers in an IPC Interrupt structure.


IPC_INTR: This register provides the instantaneous status of the interrupt sources. Note that there are 16 notify and 16 release event bits in this register. These are the notify and release events corresponding to the 16 IPC channels. When a notify event is triggered in the IPC channel 0, the corresponding Notify0 bit is activated in the interrupt registers. A write of '1' to a bit will clear the interrupt.

IPC_INTR_MASK: The bit in this register masks the interrupt sources. Only the interrupt sources with their masks enabled can trigger the interrupt.

IPC_INTR_SET: A write of '1' into this register will set the interrupt.

IPC_INTR_MASKED: This register provides the instantaneous value of the interrupts after they are masked. The value in this register is (IPC_INTR AND IPC_INTR_MASK).

Figure 6-2. IPC Interrupt Structure

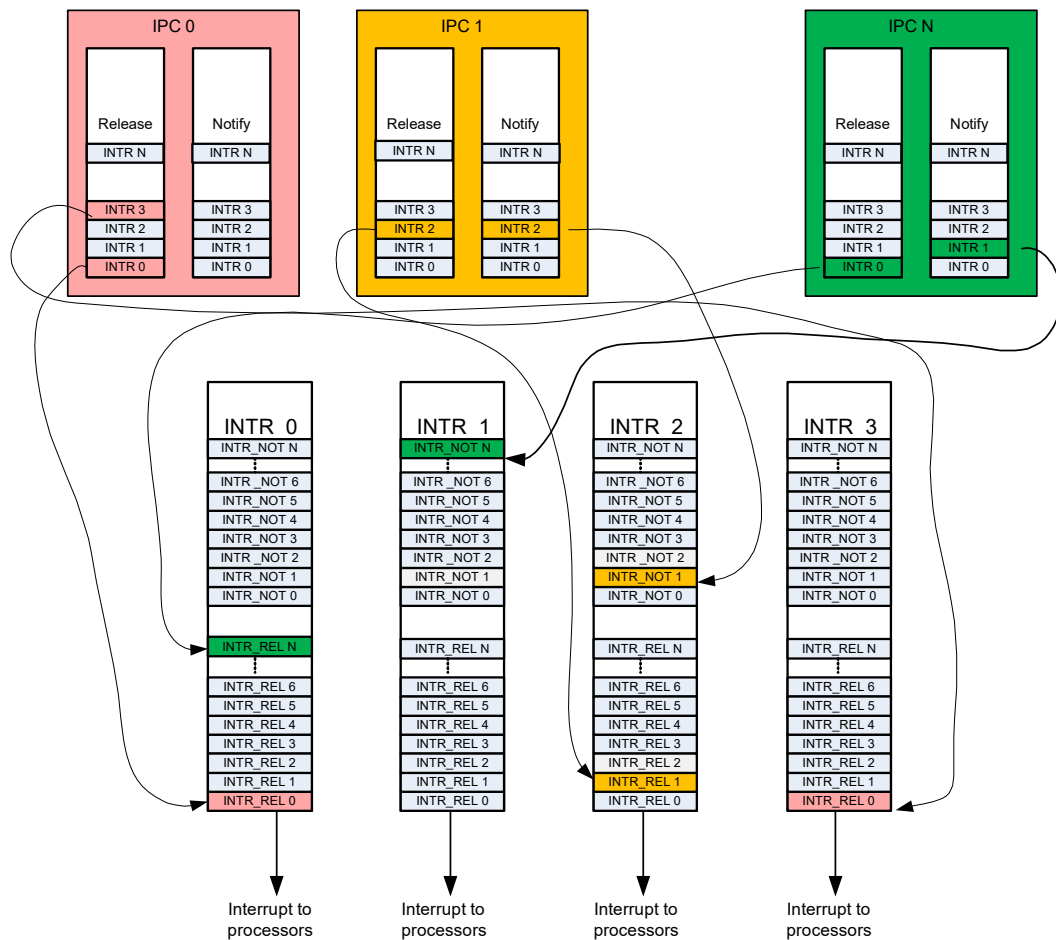


6.2.3 IPC Channels and Interrupts

The IPC block has a set of IPC interrupts associated with it. Each IPC interrupt register structure corresponds to an IPC interrupt line. This interrupt can trigger an interrupt on any of the processors in the system. The interrupt routing for processors are dependent on the device architecture.

Each IPC channel has a release and notify register, which can drive events on any of the IPC interrupts. An illustration of this relation between the IPC channels and the IPC interrupt structure is shown in Figure 6-3.

Figure 6-3. IPC Channels and Interrupts



6.3 Implementing Locks

The IPC channels can be used to implement locks. Locks are typically used in multi-core systems to implement some form of mutually exclusive access to a shared resource. When multiple processors share a resource, the processors are capable of acquiring and releasing the IPC channel. So the processor can assume an IPC channel as a lock. The semantics of this code is that access to the shared resource is gated by the processor's ownership of the channel. So the processors will need to acquire the IPC channel before they access the shared resource.

A failure to acquire the IPC channel signifies a lock on the shared resource because another processor has control of it. Note that the IPC channel will not enforce which processor acquires or releases the channel. All processors can acquire or release the IPC channel and the semantics of the code must make sure that the processor that acquires the channel is the one that releases it.

6.4 Message Passing


IPC channels can be used to communicate messages between processors. In this use case, the channel is used in conjunction with the interrupt structures. The IPC channel is used to lock the access to the data registers. The IPC channel is acquired by the sender and used to populate the message. The receiver reads the message and then releases the channel. Thus, between the sender putting data into the channel and receiver reading it, the channel is locked for all other task access. The sender uses a notify event on the receiver's IPC interrupt to denote a send operation. The receiver acts on this interrupt and reads the data from the data registers. After the reception is complete, the receiver releases the channel and can also generate a release event to the sender's IPC interrupt. Note that the action of locking the channel does not, in hardware, restrict access to the data registers. This is a semantic that should be enforced by software.

Figure 6-4 portrays an example of a sender (Processor A) sending data to a receiver (Processor B). IPC interrupt A is configured to interrupt Processor A. IPC interrupt B is configured to interrupt Processor B.

1. The sender will attempt to acquire the IPC channel by reading the IPC_ACQUIRE register. If the channel was acquired, the sender has ownership of the channel for data transmission. If the channel was not acquired, the processor should wait until the channel is free for acquisition. This can be done by polling the IPC channel's IPC_LOCK_STATUS register.
2. After the IPC channel is acquired, the sender has control of the channel for communication and places the 64-bit message data in the IPC_DATA0 and IPC_DATA1 registers.
3. Now that the message is placed in the IPC channel, the sender generates a notify event on the receiver's interrupt line. It does this by setting the corresponding bit in the IPC channel's IPC_NOTIFY register. This event creates a notify event at IPC interrupt B. If the IPC channel's notify event was enabled by setting the mask bit in the IPC interrupt B, this will generate an interrupt in the receiver.
4. When it receives IPC interrupt B, the receiver can poll the IPC_INTR_MASKED register to understand which IPC channel had triggered the notify event. Based on this, the receiver identifies the channel to read and reads from the IPC channel's IPC_DATA0 and IPC_DATA1 registers. The receiver has now received the data sent by the sender. It needs to release the channel so that other processors/processes can use it.
5. The receiver releases the channel. It also optionally generates a release event on the sender's IPC interrupt A. This will generate a release event interrupt on the sender if the corresponding channel release event was masked.

On receiving the release interrupt, the sender can act on the event based on the application requirement. It can either try to reacquire the channel for further transmission or go on to other tasks because the transmission is complete.

Figure 6-4. Sending Messages using IPC



In the previous example, the size of the data being transmitted was just 64 bits. Larger messages can be sent as pointers. The sender can allocate a larger message structure in memory and pass the pointers in the data registers. Figure 6-5 shows the usage. Note that the user code must implement the synchronization of the message read process.

- The implementation can stall the channel until the receiver has used up all the data in the message packet and the message packet can be rewritten. This is wasteful because it will stall other inter-process communications as the number of IPC channels is limited.
- The receiver can release the channel as soon as it receives the pointer to the message packet. It implements the synchronization logic in the message packet as a flag, which the sender sets on write complete and receiver clears on a read complete.

Figure 6-5. Communicating Larger Messages




Table 6-1. IPC Channels Used

Channel	Usage
0	CM0+ System Call
1	CM4 System Call
2	DAP System Call (programming)
3	CM4 -> CM0+ RPC
4	DAP -> CM0+ RPC
5	CM0+ -> CM4 RPC
6	DAP -> CM4 RPC
7-15	Unused

7. Fault Monitoring



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SysLib
- Application notes
- Code examples

Fault monitoring allows you to monitor various faults generated within the device and take actions based on the fault reported. The fault structures present in the PSoC 6 MCU monitor access violation faults at protection units (MPU, SMPU, or PPU) and flash controller bus error/fault. In addition to reporting faults, the fault structures in PSoC 6 MCUs provide a mechanism to log data from the fault sources and optionally perform soft reset.

The PSoC 6 MCU family supports two centralized fault report/monitoring structures that monitor faults generated within the device. Each fault report structure can monitor and report faults from up to 96 sources.


7.1 Features

Each PSoC 6 MCU fault report structure supports:

- Monitoring protection unit access violation faults and flash controller bus errors
- Four 32-bit data registers to record fault information
- Soft reset on fault detection while retaining the fault information
- Interrupt on fault detection
- Trigger output to DMA for fault data transfer
- Fault detected output to a pin for external fault handling

7.2 Architecture

Figure 7-1. Fault Report Structure



The PSoC 6 MCU family uses centralized fault report structures. This centralized nature allows for a system wide handling of faults simplifying firmware development. Only a single fault interrupt handler is required to monitor multiple faults. The fault report structure provides the fault source and additional fault specific information through a single set of registers; no iterative search for the fault source and fault information is required.

The fault structure can be configured to capture one or more faults as listed in Table 7-2. When a fault structure is configured to capture a specific fault, an occurrence of that fault will be recorded as a pending fault. If the fault structure has finished processing all other faults or if there are no other pending faults, the fault data will be captured into the fault structure registers. In addition, a successful capture can trigger an interrupt and be processed by either Cortex-M4 or Cortex-M0+ depending on the application requirement.

It should be noted that each fault structure is capable of capturing only one fault at a time and as long as that fault is not serviced, subsequent faults will not be captured by the fault structure. In addition to capturing faults, the fault structure can optionally perform a soft reset while retaining the fault information. This reset results in RESET_ACT_FAULT reset cause in the SRSS_RES_CAUSE register.

7.2.1 Fault Report

The PSoC 6 MCU family supports two fault report structures. Each fault report structure has a dedicated set of control and status registers. Each fault report structure captures a single fault. The captured fault information includes:

- Fault validity bit that indicates a fault is captured (VALID bit [31] of the FAULT_STRUCTx_STATUS register). This bit is set whenever a fault is captured. The bit should be cleared after processing the fault information. New faults are captured only when this bit is '0'.
- Fault index, as shown in Table 7-2, identifies the fault source (IDX bits [6:0] of FAULT_STRUCTx_STATUS)
- Additional fault information describing fault specifics (FAULT_STRUCTx_DATA0 through FAULT_STRUCTx_DATA3 registers). This additional information is fault source specific. For example, an MPU protection violation provides

information on the violating bus address, the bus master identifier, and bus access control information in only two FAULT_DATA registers. The details of the fault information for various faults is explained in [Table 7-1](#).

Table 7-1. Fault Information

Fault Source	Fault Information
MPU/SMPU violation	DATA0[31:0]: Violating address. DATA1[0]: User read. DATA1[1]: User write. DATA1[2]: User execute. DATA1[3]: Privileged read. DATA1[4]: Privileged write. DATA1[5]: Privileged execute. DATA1[6]: Non-secure. DATA1[11:8]: Master identifier. DATA1[15:12]: Protection context identifier. DATA1[31]: '0' MPU violation; '1': SMPU violation.
Master interface PPU violation	DATA0[31:0]: Violating address. DATA1[0]: User read. DATA1[1]: User write. DATA1[2]: User execute. DATA1[3]: Privileged read. DATA1[4]: Privileged write. DATA1[5]: Privileged execute. DATA1[6]: Non-secure. DATA1[11:8]: Master identifier. DATA1[15:12]: Protection context identifier. DATA1[31]: '0': PPU violation, '1': peripheral bus error.
Peripheral group PPU violation	DATA0[31:0]: Violating address. DATA1[0]: User read. DATA1[1]: User write. DATA1[2]: User execute. DATA1[3]: Privileged read. DATA1[4]: Privileged write. DATA1[5]: Privileged execute. DATA1[6]: Non-secure. DATA1[11:8]: Master identifier. DATA1[15:12]: Protection context identifier. DATA1[31:30]: '0': PPU violation, '1': timeout detected, '2': peripheral bus error.
Flash controller bus error	FAULT_DATA0[31:0]: Violating address. FAULT_DATA1[31]: '0': FLASH macro interface bus error; '1': memory hole. FAULT_DATA1[15:12]: Protection context identifier. FAULT_DATA1[11:8]: Master identifier.

7.2.2 Signaling Interface

In addition to captured fault information, each fault report structure supports a signaling interface to notify the system about the captured fault. The interface of fault report structure 'x' supports the following:


- A fault interrupt (interrupt_fault[x]). Use the `FAULT_STRUCTx_INTR`, `FAULT_STRUCTx_INTR_SET`, `FAULT_STRUCTx_INTR_MASK` and `FAULT_STRUCTx_INTR_MASKED` registers to monitor, set, and mask the `FAULT_STRUCTURE[x]`'s interrupt. Only a single interrupt cause is available, which indicates that a fault is detected. The fault report registers can be read in the interrupt handler to deduce the fault. The `FAULT` bit [0] of the `FAULT_STRUCTx_INTR_MASK` register provides a mask/enable for the interrupt. The `FAULT` bit [0] of the `FAULT_STRUCTx_INTR` register is set to '1' when a fault is captured. Setting this bit in firmware clears the interrupt.
- A DMA trigger (tr_fault[x]). The fault structure generates a DMA trigger when `VALID` bit [31] of the `FAULT_STRUCTx_STATUS` register is set. To enable the trigger, set the `TR_EN` bit [0] of the `FAULT_STRUCTx_CTL` register. The trigger can be connected to a DMA controller, which can transfer captured fault information from the fault report structure to memory and can clear the `VALID` bit [31] of the

`FAULT_STRUCTx_STATUS` register. See the [Trigger Multiplexer Block chapter on page 260](#) for more details.

- A chip output signal (fault_out[x]). The fault structure generates an output signal, which is set when `VALID` bit [31] of the `FAULT_STRUCTx_STATUS` register is set. This signal can be routed out of the device through the HSIOM (refer to the [device datasheet](#)). The output signal is enabled by setting the `OUT_EN` bit [1] of the `FAULT_STRUCTx_CTL` register. The output signal can be used to communicate non-recoverable faults to off-chip components (possibly resulting in a reset of the chip).
- A fault reset request signal (fault_reset_req[x]). The fault structure generates a soft reset when `VALID` bit [31] of the `FAULT_STRUCTx_STATUS` register is set. The reset capability is enabled by setting `RESET_REQ_EN` bit [2] of `FAULT_STRUCTx_CTL`. The reset request performs a soft reset. This reset is captured as `RESET_ACT_FAULT` in the `SRSS_RES_CAUSE` register. The fault information in `FAULT_STRUCTx_STATUS` and `FAULT_STRUCTx_DATA` registers is retained through this reset.

Because the device has a single `fault_reset_req` signal, the individual `fault_reset_req[x]` signals from the fault structures are combined into a single `fault_reset_req` signal as shown in [Figure 7-2](#).

Figure 7-2. Fault Reset Request



7.2.3 Monitoring

A central structure, which is shared by all fault report structures, keeps track of all pending faults in the system. The `FAULT_STRUCTx_PENDINGx` registers reflect what fault sources are pending and provide a single pending bit for up to 96 fault sources. The registers are mirrored in all the fault report structures; that is, they read the same value in all fault structures. The bit indexing in the registers follow the fault index captured in [Table 7-2](#). For instance, bit [0] of `FAULT_STRUCTx_PENDING0` captures a CM0+ MPU/SMPU violation and bit [1] of `FAULT_STRUCTx_PENDING1` captures a peripheral group#1 PPU violation.

The pending faults are faults that are not yet captured by a fault structure. When a pending fault is captured by a fault structure, the associated pending bit is cleared to '0'.

Each fault report structure is selective in the faults it captures. The `FAULT_STRUCTx_MASK0`,

`FAULT_STRUCTx_MASK1`, `FAULT_STRUCTx_MASK2` registers of a fault structure decide the pending faults that it captures. These faults are referred to as "enabled" faults. The `FAULT_STRUCTx_MASK` registers are unique to each fault structure. This allows for the following:

- One fault report structure is used to capture recoverable faults and one fault report structure is used to capture non-recoverable faults. The former can be used to generate a fault interrupt and the latter can be used to activate a chip output signal and/or activate a reset request.
- Two fault report structures are used to capture the same faults. The first fault is captured by the structure with the lower index (for example, fault structure 0) and the second fault is captured by the structure with the higher index (for example, fault structure 1). Note that both structures cannot capture the same fault at the same time. As soon as a fault is captured, the pending bit is cleared and the other structure will not be aware of the

fault. Fault structure 0 has precedence over fault structure 1.

The fault structure captures “enabled” faults only when VALID bit [31] of FAULT_STRUCTx_STATUS register is ‘0’. When a fault is captured, hardware sets the VALID bit [31] of the FAULT_STRUCTx_STATUS register. In addition, hardware clears the associated pending bit to ‘0’. When a fault structure is processed, firmware or a DMA transfer should clear the VALID bit [31] of the FAULT_STRUCTx_STATUS register. Note that fault capturing does not consider FAULT bit [0] of FAULT_STRUCTx_INTR register and firmware should clear the bit after servicing the interrupt, if the interrupt is enabled.

7.2.4 Low-power Mode Operation

The fault report structure functionality is available in Active and Sleep (and their LP counterparts) power modes only. The interfaces between the fault sources and fault report structures are reset in the Deep Sleep power mode. Because the fault report structure is an active functionality, pending faults (in the FAULT_STRUCTx_PENDING registers) are not retained when transitioning to Deep Sleep power mode. The fault structure’s registers can be partitioned based on the reset domain and their retention capability as follows:

- Active reset domain: FAULT_STRUCTx_PENDING, FAULT_STRUCTx_INTR, FAULT_STRUCTx_INTR_SET, and FAULT_STRUCTx_INTR_MASKED registers. These registers are not retained in Deep Sleep power mode.
- Deep Sleep reset domain: FAULT_STRUCTx_CTL, FAULT_STRUCTx_MASK, and FAULT_STRUCTx_INTR_MASK registers. These registers are retained in Deep Sleep power mode but any system reset will reset these registers to the default state.
- Hard reset domain: FAULT_STRUCTx_STATUS and FAULT_STRUCTx_DATA registers. These registers are retained through soft resets (detectable in SRSS_RES_CAUSE registers). However, hard resets such as XRES/POR/BOD will reset the registers.

7.2.5 Using a Fault Structure

Follow these steps to configure and use a fault structure:

1. Identify the faults from [Table 7-2](#) to be monitored in the system.
2. For firmware fault handling through interrupts
 - a. Set the FAULT bit [0] of the FAULT_STRUCTx_INTR_MASK register.
 - b. Set the FAULT bit [0] of the FAULT_STRUCTx_INTR register to clear any pending interrupt.

- c. Enable the FAULTx interrupt to the CPU by configuring the appropriate ISER register. Refer to the [Interrupts chapter on page 54](#).
3. For fault handling through DMA
 - a. Set the TR_EN bit [0] of the FAULT_STRUCTx_CTL register.
 - b. Route the tr_fault[x] signal to the trigger input DMA controller. Refer to the [Trigger Multiplexer Block chapter on page 260](#).
 - c. Configure and enable the DMA controller to transfer FAULT_STRUCTx_STATUS and FAULT_STRUCTx_DATA registers to memory and write back ‘0’ to FAULT_STRUCTx_STATUS register after the transfer is complete. Refer to the [DMA Controller \(DW\) chapter on page 93](#).
4. For fault handling outside the device
 - a. Set the OUT_EN bit [1] of FAULT_STRUCTx_CTL register.
 - b. Route the fault_out[x] signal to a pin through HSIOM. Refer to the [device datasheet](#).
 - c. Use the signal externally for processing the fault – generate external reset, power cycle, or log fault information.
5. Set the RESET_REQ_EN bit [2] of the FAULT_STRUCTx_CTL register, if a soft reset is required on any fault detection in the structure.
6. Clear VALID bit [31] of the FAULT_STRUCTx_STATUS register to clear any fault captured.
7. Set the fault index bits in the FAULT_STRUCTx_MASK registers for faults that need to be captured by the fault structure as explained in [7.2.3 Monitoring](#).

7.2.6 CPU Exceptions Versus Fault Monitoring

Some faults captured in [Table 7-2](#) also result in bus errors or CPU exceptions (Cortex-M4 Bus/Usage/Memory/Hard faults). The faults can be communicated in two ways:

- As a bus error to the master of the faulting bus transfer. This will result in Bus, Usage, Memory, or Hard fault exceptions in the CPU.
- As a fault in a fault report structure. This fault can be communicated as a fault interrupt to any processor in the system. This allows fault handling on a processor that is not the master of the faulting bus transfer. It is useful for faults that cause the master of the faulting transfer to become unresponsive or unreliable.

7.3 Fault Sources

The fault sources can vary between device families. [Table 7-2](#) provides the list of fault sources available in PSoC 6 MCUs.

Table 7-2. Fault Sources

Fault Index	Source	Description
0	cpuss.mpu_vio[0]	CM0+ MPU/SMPU violation
1	cpuss.mpu_vio[1]	CRYPTO MPU/SMPU violation
2	cpuss.mpu_vio[2]	DW0 MPU/SMPU violation
3	cpuss.mpu_vio[3]	DW1 MPU/SMPU violation
4	cpuss.mpu_vio[4]	DMAC MPU/SMPU violation
5	cpuss.mpu_vio[5]	SDHC0 MPU/SMPU violation
6		
7 to 15	Reserved	
16	cpuss.mpu_vio[16]	CM4 MPU/SMPU violation (System bus)
17	cpuss.mpu_vio[17]	CM4 code error
18	cpuss.mpu_vio[18]	CM4 code flash controller errors
19 to 27	Reserved	
28	peri.ms_vio[0]	CM0+ Peripheral Master Interface PPU violation
29	peri.ms_vio[1]	CM4 Peripheral Master Interface PPU violation
30	peri.ms_vio[2]	DW0 Peripheral Master Interface PPU violation
31	peri.ms_vio[3]	DW1 Peripheral Master Interface PPU violation
32	peri.group_vio[0]	Peripheral group #0 (Peripheral Clock dividers, Trigger Mux etc) PPU violation Register address range: 0x40000000 to 0x40100000
33	peri.group_vio[1]	Peripheral group #1 (Crypto block) PPU violation Register address range: 0x40100000 to 0x40200000
34	peri.group_vio[2]	Peripheral group #2 (CPUSS, SRSS, EFUSE) PPU violation Register address range: 0x40200000 to 0x40300000
35	peri.group_vio[3]	Peripheral group #3 (IOSS, LPCOMP, CSD, TCPWM, LCD, BLE) PPU violation Register address range: 0x40300000 to 0x40400000
36	peri.group_vio[4]	Peripheral group #4 (SMIF) PPU violation Register address range: 0x40400000 to 0x40500000
37		
38	peri.group_vio[6]	Peripheral group #6 (SCB) PPU violation Register address range: 0x40600000 to 0x40700000
39	Reserved	
40	Reserved	
41	peri.group_vio[9]	Peripheral group #9 (PASS) PPU violation Register address range: 0x41000000 to 0x42000000
42		
43 to 47	Reserved	
48	cpuss.flashc_main_bus_err	Flash controller bus error
49 to 95	Reserved	

7.4 Register List

Name	Description
FAULT_STRUCTx_CTL	Fault control register for enabling DMA trigger, fault output, and fault reset signals
FAULT_STRUCTx_STATUS	Fault status register that stores the validity and fault index of the currently captured fault
FAULT_STRUCTx_DATA0	Fault data register 0 that stores fault information associated with the currently captured fault
FAULT_STRUCTx_DATA1	Fault data register 1 that stores fault information associated with the currently captured fault
FAULT_STRUCTx_DATA2	Fault data register 2 that stores fault information associated with the currently captured fault
FAULT_STRUCTx_DATA3	Fault data register 3 that stores fault information associated with the currently captured fault
FAULT_STRUCTx_PENDING0	Fault pending register 0 that stores pending (not captured) faults with fault index from 0 to 31
FAULT_STRUCTx_PENDING1	Fault pending register 1 that stores pending (not captured) faults with fault index from 32 to 63
FAULT_STRUCTx_PENDING2	Fault pending register 2 that stores pending (not captured) faults with fault index from 64 to 95
FAULT_STRUCTx_MASK0	Fault mask register 0 that enables the capture of pending faults with fault index from 0 to 31 by the fault structure
FAULT_STRUCTx_MASK1	Fault mask register 1 that enables the capture of pending faults with fault index from 32 to 63 by the fault structure
FAULT_STRUCTx_MASK2	Fault mask register 2 that enables the capture of pending faults with fault index from 64 to 95 by the fault structure
FAULT_STRUCTx_INTR	Fault interrupt register that stores the unmasked status of the fault structure's interrupt
FAULT_STRUCTx_INTR_SET	Fault interrupt set register used to set the fault structure's interrupt through firmware
FAULT_STRUCTx_INTR_MASK	Fault interrupt mask register that masks fault interrupt
FAULT_STRUCTx_INTR_MASKED	Fault interrupt register that stores the masked status of the fault structure's interrupt

8. Interrupts



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SysInt
- Application notes
- Code examples

The PSoC 6 MCU family supports interrupts and CPU exceptions on both Cortex-M4 and Cortex-M0+ cores. Any condition that halts normal execution of instructions is treated as an exception by the CPU. Thus an interrupt request is treated as an exception. However, in the context of this chapter, interrupts refer to those events generated by peripherals external to the CPU such as timers, serial communication block, and port pin signals; exceptions refer to those events that are generated by the CPU such as memory access faults and internal system timer events. Both interrupts and exceptions result in the current program flow being stopped and the exception handler or interrupt service routine (ISR) being executed by the CPU. Both Cortex-M4 and Cortex-M0+ cores provide their own unified exception vector table for both interrupt handlers/ISR and exception handlers.

Note: Although this section describes how both the Cortex-M4 and the Cortex-M0+ operate, only the Cortex-M4 is accessible by the developer.

8.1 Features

The PSoC 6 MCU supports the following interrupt features:

- Supports 168 system interrupts
 - 168 Arm Cortex-M4 interrupts
 - Eight Arm Cortex-M0+ external interrupts and eight Arm Cortex-M0+ internal (software only) interrupts. The CPU supports up to 32 interrupts, but only 16 interrupts are used by PSoC 6 interrupt infrastructure. The eight external CPU interrupts support deep sleep (WIC) functionality
 - Four system interrupts can be mapped to each of the CPU non-maskable interrupt (NMI)
 - Up to 39 interrupt sources capable of waking the device from Deep Sleep power mode
- Nested vectored interrupt controller (NVIC) integrated with each CPU core, yielding low interrupt latency
- Wakeup interrupt controller (WIC) enabling interrupt detection (CPU wakeup) in Deep Sleep power mode
- Vector table may be placed in either flash or SRAM
- Configurable priority levels (eight levels for Cortex-M4 and four levels for Cortex-M0+) for each interrupt
- Level-triggered and pulse-triggered interrupt signals

8.2 Architecture

Figure 8-1. PSoC 6 MCU Interrupts Block Diagram

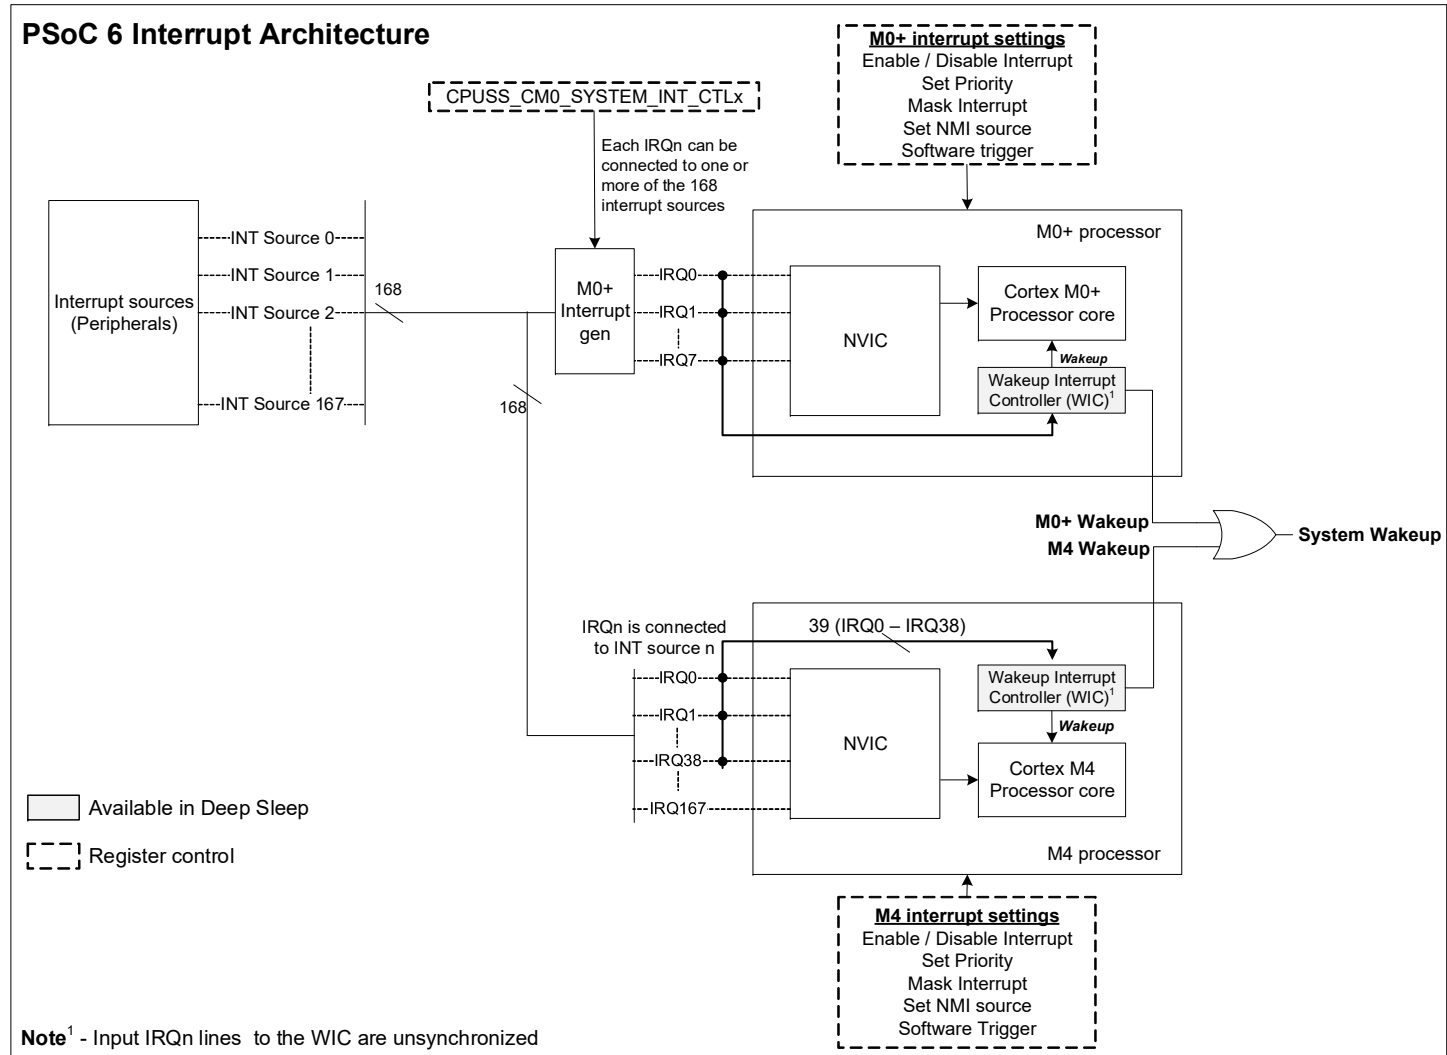


Figure 8-1 shows the PSoC 6 MCU interrupt architecture. The PSoC 6 MCU has 168 system interrupts that are generated by various peripherals. These interrupt signals are processed by the NVIC of the individual core. In the Cortex-M4 core, the system interrupt source 'n' is directly connected to IRQn. The Cortex-M0+ interrupt architecture uses eight CPU interrupts IRQ[7:0] out of the 32 available IRQn lines of the core. In the Cortex-M0+ core, the system interrupt source connection to a particular IRQn of the core is configurable and any of the 168 system interrupts can be mapped to any of the IRQ[7:0]. This ensures that all the system interrupts can be mapped onto any CPU interrupt simultaneously. Refer to [Interrupt Sources on page 60](#) for more details about the system interrupt to CPU interrupt mapping. The NVIC takes care of enabling/disabling individual interrupt IRQs, priority resolution, and communication with the CPU core. The other exceptions such as NMI and hard faults are not shown in Figure 8-1 because they are part of CPU core generated events, unlike interrupts, which are generated by peripherals external to the CPU.

In addition to the NVIC, the PSoC 6 MCU supports wakeup interrupt controllers (WIC) for each CPU and a shared interrupt synchronization block that synchronizes the interrupts to CLK_HF domain (adds two CLK_HF cycles delay for synchronization). The WIC provides detection of Deep Sleep interrupts in the Deep Sleep CPU power mode. Each CPU can individually be in Deep Sleep mode; the device is said to be in Deep Sleep mode only when both the CPUs are in Deep Sleep mode. Refer to the [Device Power Modes chapter on page 190](#) for details. The Cortex-M4 WIC block supports up to 39 interrupts that can wake up the CPU from Deep Sleep power mode. The Cortex-M0+ WIC block supports all eight interrupts. The device exits Deep Sleep mode (System Wakeup signal in Figure 8-1) as soon as one CPU wakes up. The synchronization blocks synchronize the interrupts to the CPU clock frequency as the peripheral interrupts can be asynchronous to the CPU clock frequency.

8.3 Interrupts and Exceptions - Operation

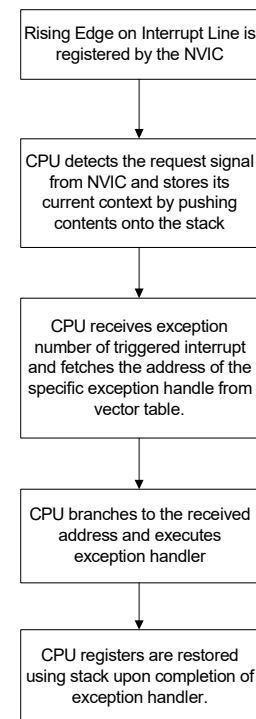
8.3.1 Interrupt/Exception Handling

The following sequence of events occurs when an interrupt or exception event is triggered:

1. Assuming that all the interrupt and exception signals are initially low (idle or inactive state) and the processor is executing the main code, a rising edge on any one of the signals is registered by the NVIC, if the interrupt or exception is enabled to be serviced by the CPU. The signal is now in a pending state waiting to be serviced by the CPU.

2. On detecting the signal from the NVIC, the CPU stores its current context by pushing the contents of the CPU registers onto the stack.
3. The CPU also receives the exception number of the triggered interrupt from the NVIC. All interrupts and exceptions have a unique exception number, as given in [Table 8-1](#). By using this exception number, the CPU fetches the address of the specific exception handler from the vector table.
4. The CPU then branches to this address and executes the exception handler that follows.
5. Upon completion of the exception handler, the CPU registers are restored to their original state using stack pop operations; the CPU resumes the main code execution.

Figure 8-2. Interrupt Handling When Triggered



When the NVIC receives an interrupt request while another interrupt is being serviced or receives multiple interrupt requests at the same time, it evaluates the priority of all these interrupts, sending the exception number of the highest priority interrupt to the CPU. Thus, a higher priority interrupt can block the execution of a lower priority ISR at any time.

Exceptions are handled in the same way that interrupts are handled. Each exception event has a unique exception number, which is used by the CPU to execute the appropriate exception handler.

8.3.2 Level and Pulse Interrupts

Both CM0+ and CM4 NVICs support level and pulse signals on the interrupt lines (IRQn). The classification of an interrupt as level or pulse is based on the interrupt source.

Figure 8-3. Level Interrupts




Figure 8-4. Pulse Interrupts




Figure 8-3 and Figure 8-4 show the working of level and pulse interrupts, respectively. Assuming the interrupt signal

is initially inactive (logic low), the following sequence of events explains the handling of level and pulse interrupts:

1. On a rising edge event of the interrupt signal, the NVIC registers the interrupt request. The interrupt is now in the pending state, which means the interrupt requests have not yet been serviced by the CPU.
2. The NVIC then sends the exception number along with the interrupt request signal to the CPU. When the CPU starts executing the ISR, the pending state of the interrupt is cleared.
3. For pulse interrupts, when the ISR is being executed by the CPU, one or more rising edges of the interrupt signal are logged as a single pending request. The pending interrupt is serviced again after the current ISR execution is complete (see Figure 8-4 for pulse interrupts).
4. For level interrupts, if the interrupt signal is still high after completing the ISR, it will be pending and the ISR is executed again. Figure 8-3 illustrates this for level triggered interrupts, where the ISR is executed as long as the interrupt signal is high.

8.3.3 Exception Vector Table

The exception vector tables (Table 8-1 and Table 8-2) store the entry point addresses for all exception handlers in Cortex-M0+ and Cortex-M4 cores. The CPU fetches the appropriate address based on the exception number.

Table 8-1. M0+ Exception Vector Table

Exception Number	Exception	Exception Priority	Vector Address
–	Initial Stack Pointer Value	Not applicable (NA)	Start_Address = 0x0000 or CM0P_SCS_VTOR ^a
1	Reset	–3, the highest priority	Start_Address + 0x04
2	Non Maskable Interrupt (NMI)	–2	Start_Address + 0x08
3	HardFault	–1	Start_Address + 0x0C
4-10	Reserved	NA	Start_Address + 0x10 to Start_Address + 0x28
11	Supervisory Call (SVCall)	Configurable (0 – 3)	Start_Address + 0x2C
12-13	Reserved	NA	Start_Address + 0x30 to Start_Address + 0x34
14	PendSupervisory (PendSV)	Configurable (0 – 3)	Start_Address + 0x38
15	System Timer (SysTick)	Configurable (0 – 3)	Start_Address + 0x3C
16	External Interrupt (IRQ0)	Configurable (0 – 3)	Start_Address + 0x40
...	...	Configurable (0 – 3)	...
31	External Interrupt (IRQ15)	Configurable (0 – 3)	Start_Address + 0x7C

a. Start Address = 0x0000 on reset and is later modified by firmware by updating the CM0P_SCS_VTOR register.

Table 8-2. Cortex-M4 Exception Vector Table

Exception Number	Exception	Exception Priority	Vector Address
–	Initial stack pointer value	–	Start_Address = 0x0000 or CM4_SCS_VTOR ^a
1	Reset	–3, highest priority	Start_Address + 0x0004
2	Non Maskable Interrupt (NMI)	–2	Start_Address + 0x0008
3	Hard fault	–1	Start_Address + 0x000C
4	Memory management fault	Configurable (0 – 7)	Start_Address + 0x0010

Table 8-2. Cortex-M4 Exception Vector Table (continued)

Exception Number	Exception	Exception Priority	Vector Address
5	Bus fault	Configurable (0 – 7)	Start _Address + 0x0014
6	Usage fault	Configurable (0 – 7)	Start _Address + 0x0018
7–10	Reserved	–	–
11	Supervisory call (SVCall)	Configurable (0 – 7)	Start _Address + 0x002C
12–13	Reserved	–	–
14	Pend Supervisory (PendSV)	Configurable (0 – 7)	Start _Address + 0x0038
15	System Tick timer (SysTick)	Configurable (0 – 7)	Start _Address + 0x003C
16	External interrupt (IRQ0)	Configurable (0 – 7)	Start _Address + 0x0040
....
182	External interrupt (IRQ166)	Configurable (0 – 7)	Start _Address + 0x02D8
183	External interrupt (IRQ167)	Configurable (0 – 7)	Start _Address + 0x02DC

a. Start Address = 0x0000 on reset and is later modified by firmware by updating CM4_SCS_VTOR register.

In [Table 8-1](#) and [Table 8-2](#), the first word (4 bytes) is not marked as exception number zero. This is because the first word in the exception table is used to initialize the main stack pointer (MSP) value on device reset; it is not considered as an exception. In the PSoC 6 MCU, both the vector tables can be configured to be located either in flash memory or SRAM. The vector table offset register (VTOR) present as part of Cortex-M0+ and Cortex-M4 system control space registers configures the vector table offset from the base address (0x0000). The CM0P_SCS_VTOR register sets the vector offset address for the CM0+ core and CM4_SCS_VTOR sets the offset for the M4 core. The VTOR value determines whether the vector table is in flash memory (0x10000000 to 0x10100000) or SRAM (0x08000000 to 0x08048000). Note that the VTOR registers can be updated only in privilege CPU mode. The advantage of moving the vector table to SRAM is that the exception handler addresses can be dynamically changed by modifying the SRAM vector table contents. However, the nonvolatile flash memory vector table must be modified by a flash memory write. Note that the exception table must be 256 byte-aligned for Cortex-M0+ and 1024 byte-aligned for Cortex-M4.

The exception sources (exception numbers 1 to 15) are explained in [8.4 Exception Sources](#). The exceptions marked as Reserved in [Table 8-1](#) are not used, although they have addresses reserved for them in the vector table. The interrupt sources (exception numbers 16 to 183) are explained in [8.5 Interrupt Sources](#).

8.4 Exception Sources

This section explains the different exception sources listed in [Table 8-1](#) and [Table 8-2](#) (exception numbers 1 to 15).

8.4.1 Reset Exception

Device reset is treated as an exception in PSoC 6 MCUs. Reset exception is always enabled with a fixed priority of –3, the highest priority exception in both the cores. When the device boots up, only the Cortex-M0+ core is available. The CM0+ core executes the ROM boot code and can enable Cortex-M4 core from the application code. The reset exception of the CM0+ core is tied to the device reset or startup. When the Cortex-M0+ core releases the Cortex-M4 reset, the M4 reset exception is executed. A device reset can occur due to multiple reasons, such as power-on-reset (POR), external reset signal on XRES pin, or watchdog reset. When the device is reset, the initial boot code for configuring the device is executed by the Cortex-M0+ out of supervisory read-only memory (SROM). The boot code and other data in SROM memory are programmed by Cypress, and are not read/write accessible to external users. After completing the SROM boot sequence, the Cortex-M0+ code execution jumps to flash memory. Flash memory address 0x10000004 (Exception#1 in [Table 8-1](#)) stores the location of the startup code in flash memory. The CPU starts executing code out of this address. Note that the reset exception address in the SRAM vector table will never be used because the device comes out of reset with the flash vector table selected. The register configuration to select the SRAM vector table can be done only as part of the startup code in flash after the reset is de-asserted. Note that the reset exception flow for Cortex-M4 is the same as Cortex-M0+. However, Cortex-M4 execution begins only after CM0+ core de-asserts the M4 reset.

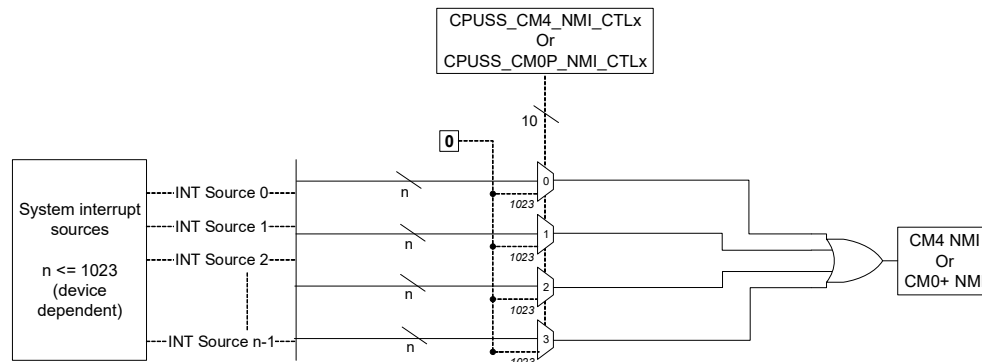
8.4.2 Non-Maskable Interrupt Exception

Non-maskable interrupt (NMI) is the highest priority exception next to reset. It is always enabled with a fixed priority of -2. Both the cores have their own NMI exception. There are three ways to trigger an NMI exception in a CPU core:

- **NMI exception from a system interrupt:** Both Cortex-M0+ and Cortex-M4 provide an option to trigger an NMI exception from up to four of the 168 system interrupts. The NMI exception triggered due to the interrupt will

execute the NMI handler pointed to by the active exception vector table. The CPUSS_CM4_NMI_CTLx and CPUSS_CM0P_NMI_CTLx registers select the interrupt source that triggers the NMI from hardware. The registers have a default value of 1023; that is, if the register is set to 1023, then that particular register does not map any interrupt source to the NMI. There are four such registers and each can map one interrupt vector to the NMI. NMI is triggered when any of the four interrupts are triggered; that is, the interrupts are logically ORed. See [Figure 8-5](#).

Figure 8-5. NMI Trigger



- **NMI exception by setting NMIPENDSET bit (user NMI exception):** An NMI exception can be triggered in software by setting the NMIPENDSET bit in the interrupt control state registers (CM0P_SCS_ICSR and CM4_SCS_ICSR). Setting this bit will execute the NMI handler pointed to by the active vector table in the respective CPU cores.
- **System Call NMI exception:** This exception is used for nonvolatile programming and other system call operations such as flash write operation and flash checksum operation. Inter processor communication (IPC) mechanism is used to implement a system call in PSoC 6 MCUs. A dedicated IPC mailbox is associated with each core (M0+ and M4) and the debug access port (DAP) to trigger a system call. The CPU or DAP acquires this dedicated mailbox, writes the system call opcode and argument to the mailbox, and notifies a dedicated IPC structure. Typically, the argument is a pointer to a structure in SRAM. This results in an NMI interrupt in the CM0+ core. Note that all the system calls are serviced by Cortex-M0+ core. A Cortex-M0+ NMI exception triggered by this method executes the NMI exception handler code that resides in SROM. Note that the NMI exception handler address is automatically initialized to the system call API located in SROM (at 0x0000000D) by the boot code. The value should be retained during vector table relocations; otherwise, no system call will be executed. The NMI handler code in SROM is not read/write accessible because it contains nonvolatile programming routines that cannot be modified by the user. The result of the system call is passed through the same IPC mechanism. For details,

refer to the [Inter-Processor Communication](#) chapter on page 42.

8.4.3 HardFault Exception

Both CM0+ and CM4 cores support HardFault exception. HardFault is an always-enabled exception that occurs because of an error during normal or exception processing. HardFault has a fixed priority of -1, meaning it has higher priority than any exception with configurable priority. A HardFault exception is a catch-all exception for different types of fault conditions, which include executing an undefined instruction and accessing an invalid memory addresses. The CPU does not provide fault status information to the HardFault exception handler, but it does permit the handler to perform an exception return and continue execution in cases where software has the ability to recover from the fault situation.

8.4.4 Memory Management Fault Exception

A memory management fault is an exception that occurs because of a memory protection-related fault. The fixed memory protection constraints determine this fault, for both instruction and data memory transactions. This fault is always used to abort instruction accesses to Execute Never (XN) memory regions. The memory management fault is only supported by the M4 core. The priority of the exception is configurable from 0 (highest) to 7 (lowest).

8.4.5 Bus Fault Exception

A Bus Fault is an exception that occurs because of a memory-related fault for an instruction or data memory transaction. This might be from an error detected on a bus in the memory system. The bus fault is supported only by the M4 core. The priority of the exception is configurable from 0 (highest) to 7 (lowest).

8.4.6 Usage Fault Exception

A Usage Fault is an exception that occurs because of a fault related to instruction execution. This includes:

- an undefined instruction
- an illegal unaligned access
- invalid state on instruction execution
- an error on exception return

The following can cause a usage fault when the core is configured to report them:

- an unaligned address on word and halfword memory access
- division by zero

The usage fault is supported only by the M4 core. The priority of the exception is configurable from 0 (highest) to 7 (lowest).

8.4.7 Supervisor Call (SVCall) Exception

Both CM0+ and CM4 cores support SVCall exception. Supervisor Call (SVCall) is an always-enabled exception caused when the CPU executes the SVC instruction as part of the application code. Application software uses the SVC instruction to make a call to an underlying operating system and provide a service. This is known as a supervisor call. The SVC instruction enables the application to issue an SVCall that requires privileged access to the system.

The priority of an SVCall exception can be configured to a value between 0 and 3 for CM0+ and 0 to 7 for CM4 core by writing to the bitfields PRI_11 of the System Handler Priority Register 2 (CM0P_SCS_SHPR2 and CM4_SCS_SHPR2). When the SVC instruction is executed, the SVCall exception enters the pending state and waits to be serviced by the CPU. The SVCALLPENDED bit in the System Handler Control and State Register (CM0P_SCS_SHCSR and CM4_SCS_SHCSR) can be used to check or modify the pending status of the SVCall exception.

8.4.8 PendSupervisory (PendSV) Exception

Both CM0+ and CM4 cores support PendSV exception. PendSV is another supervisor call related exception similar to SVCall, normally being software-generated. PendSV is always enabled and its priority is configurable similar to SVCall. The PendSV exception is triggered by setting the

PENDSVSET bit in the Interrupt Control State Register (CM0P_SCS_ICSR and CM4_SCS_ICSR). On setting this bit, the PendSV exception enters the pending state, and waits to be serviced by the CPU. The pending state of a PendSV exception can be cleared by setting the PENDSVCLR bit in the Interrupt Control State Register. The priority of a PendSV exception can be configured to a value between 0 and 3 for CM0+ and 0 to 7 for M4 by writing to the bitfields PRI_14 of the System Handler Priority Register 3. See the [Armv6-M Architecture Reference Manual](#) for more details.

8.4.9 System Tick (SysTick) Exception

Both CM0+ and CM4 cores in PSoC 6 MCUs support a system timer, referred to as SysTick, as part of their internal architecture. SysTick provides a simple, 24-bit decrementing counter for various timekeeping purposes such as an RTOS tick timer, high-speed alarm timer, or simple counter. The SysTick timer can be configured to generate an interrupt when its count value reaches zero, which is referred to as a SysTick exception. The exception is enabled by setting the TICKINT bit in the SysTick Control and Status Register (CM0P_SCS_SYST_CSR and CM4_SCS_SYST_CSR). The priority of a SysTick exception can be configured to a value between 0 and 3 for CM0+ and 0 to 7 for M4 by writing to the bitfields PRI_15 of the System Handler Priority Register 3 (SHPR3). The SysTick exception can always be generated in software at any instant by writing a one to the PENDSTSET bit in the Interrupt Control State Register. Similarly, the pending state of the SysTick exception can be cleared by writing a one to the PENDSTCLR bit in the Interrupt Control State Register.

8.5 Interrupt Sources

The PSoC 6 MCU supports 168 interrupts from peripherals. The source of each interrupt is listed in [Table 8-3](#). These system interrupts are mapped directly to Cortex-M4 core (IRQ0 to IRQ167 or exception 16 to 183). For Cortex-M0+ core, any of the 168 interrupts can be routed to any of the available eight interrupts (IRQ0 to IRQ7 or exception 16 to 23). The CPUSS_CM0_SYSTEM_INT_CTLx registers are used to make this interrupt selection in CM0+.

There are 168 CPUSS_CM0_SYSTEM_INT_CTLx registers. The CPU_INT_IDX bits [2:0] of the register selects which IRQn line of CM0+, the xth system interrupt is mapped to. The CPU_INT_VALID bit [31] of the register decides whether the xth interrupt is enabled for the core. For Cortex-M0+, multiple system interrupts can be mapped on the same CPU interrupt. Therefore, an active CPU interrupt may indicate one or multiple active system interrupts. For instance to connect Port 0 (IDX = 0) and Port 2 (IDX = 2) GPIO interrupts to IRQ0 of CM0+, CPU_INT_IDX bits should be set to '0' and CPU_INT_VALID bit should be set to '1' in CPUSS_CM0_SYSTEM_INT_CTL0 and CPUSS_CM0_SYSTEM_INT_CTL2 registers. Note that this

only connects and enables the interrupts; configuration and masking of the interrupt to the NVIC should be done as part of source peripheral configuration.

As a result of the reduction functionality in the Cortex-M0+ core, multiple system interrupts share a CPU interrupt handler as provided by the CPU's vector table. Each CPU interrupt has an associated CPUSS_CM0_INTx_STATUS register:

- SYSTEM_INT_VALID bit[31] of the register specifies if any system interrupt is active for the CPU interrupt. This bit remains set until all the system interrupts connected to the CPU interrupts are cleared at the source.
- SYSTEM_INT_IDX bits[9:0] of the register specifies the index (a number in the range [0, 1022]) of the lowest active system interrupt mapped to the corresponding CPU interrupt.

For instance, say CM0+ IRQ0 is connected to Port 0 and Port 2 GPIO interrupts. When an interrupt in both Port 0 and Port 2 are triggered simultaneously, IRQ0 in CM0+ will be triggered with CPUSS_CM0_INT0_STATUS reading 0 (Port 0 interrupt) in SYSTEM_INT_IDX bit and SYSTEM_INT_VALID bit set. After Port 0 interrupt is serviced and cleared, IRQ0 will be triggered again with SYSTEM_INT_IDX bit set to 2 (Port 2 interrupt) in the CPUSS_CM0_INT0_STATUS register. Only after servicing Port 2 interrupt and clearing, the SYSTEM_INT_VALID bit will be cleared.

The CPU interrupt handler should use the SYSTEM_INT_IDX field to index a system interrupt lookup table and jump to the system interrupt handler. The lookup table is typically located in one of the system memories. The following code illustrates the approach:

```
typedef void (* SystemIntr_Handler)(void);
void CM0_CpuIntr0_Handler (void)
{
  uint32_t system_int_idx;
  SystemIntr_Handler handler;
  if(CPUSS_CM0_INT0_STATUS.SYSTEM_INT_VALID)
  {
    system_int_idx = CPUSS_CM0_INT0_STATUS.SYSTEM_INT_IDX;
    handler = SystemIntr_Table[system_int_idx];
    handler(); // jump to system interrupt handler
  }
  else
  {
    // Triggered by SW or due to SW clear error (SW cleared a peripheral
    // interrupt flag but didn't clear the Pending flag at NVIC)
  }
}
...
void CM0_CpuIntr7_Handler (void)
{
  uint32_t system_int_idx;
  SystemIntr_Handler handler;
  if(CPUSS_CM0_INT7_STATUS.SYSTEM_INT_VALID)
  {
    system_int_idx = CPUSS_CM0_INT7_STATUS.SYSTEM_INT_IDX;
    handler = SystemIntr_Table[system_int_idx];
    handler(); // jump to system interrupt handler
  }
  else
  {
    // Triggered by SW or due to SW clear error (SW cleared a peripheral
    // interrupt flag but didn't clear the Pending flag at NVIC)
  }
}
void CM0_SystemIntr0_Handler (void)
{
  // Clear the peripheral interrupt request flag by register write
  // Read back the register, to ensure the completion of register write access
  // Handle system interrupt 0.
}
```

```
...
void CM0_SystemIntr167_Handler (void)
{
// Clear the peripheral interrupt request flag by register write
// Read back the register, to ensure the completion of register write access
// Handle system interrupt 167.
}
```

The interrupts include standard interrupts from the on-chip peripherals such as TCPWM, serial communication block, CSD block, watchdog, ADC, and so on. The interrupt generated is usually the logical OR of the different peripheral states. The peripheral interrupt status register should be read in the ISR to detect which condition generated the interrupt. These interrupts are usually level interrupts. The appropriate interrupt registers should be cleared in the ISR to deassert the interrupt. Usually a write '1' is required to clear the registers. If the interrupt register is not cleared in the ISR, the interrupt will remain asserted and the ISR will be executed continuously. See the [I/O System chapter on page 226](#) for details on GPIO interrupts.

As seen from [Table 8-3](#), 39 interrupts (IRQ0 to IRQ38) are capable of waking up the device from Deep Sleep power mode. For Cortex-M4, IRQ0 to IRQ38 directly map to these sources. However, in the Cortex-M0+, all the eight IRQ lines support Deep Sleep wakeup if they are connected to a deep sleep capable interrupt. This means the 39 Deep Sleep wakeup-capable interrupts can be connected to any of the eight IRQ lines of Cortex-M0+, if such a wakeup is desired. Therefore, reserve and use the first eight IRQ lines of Cortex-M0+ for Deep Sleep wakeup-capable sources.

Table 8-3. List of PSoC 6 Interrupts

System Interrupt	Cortex M4 Exception Number	Power Mode	Interrupt Source
NMI	2	Active	Any of the below 168 IRQ source
IRQ0	16	DeepSleep	GPIO Interrupt - Port 0
IRQ1	17	DeepSleep	GPIO Interrupt - Port 1
IRQ2	18	DeepSleep	GPIO Interrupt - Port 2
IRQ3	19	DeepSleep	GPIO Interrupt - Port 3
IRQ4	20	DeepSleep	GPIO Interrupt - Port 4
IRQ5	21	DeepSleep	GPIO Interrupt - Port 5
IRQ6	22	DeepSleep	GPIO Interrupt - Port 6
IRQ7	23	DeepSleep	GPIO Interrupt - Port 7
IRQ8	24	DeepSleep	GPIO Interrupt - Port 8
IRQ9	25	DeepSleep	GPIO Interrupt - Port 9
IRQ10	26	DeepSleep	GPIO Interrupt - Port 10
IRQ11	27	DeepSleep	GPIO Interrupt - Port 11
IRQ12	28	DeepSleep	GPIO Interrupt - Port 12
IRQ13	29	DeepSleep	GPIO Interrupt - Port 13
IRQ14	30	DeepSleep	GPIO Interrupt - Port 14
IRQ15	31	DeepSleep	GPIO Interrupt - All Ports
IRQ16	32	DeepSleep	GPIO Supply Detect Interrupt
IRQ17	33	DeepSleep	Low Power Comparator Interrupt
IRQ18	34	DeepSleep	Serial Communication Block #8 (DeepSleep capable)
IRQ19	35	DeepSleep	Multi Counter Watchdog Timer interrupt
IRQ20	36	DeepSleep	Multi Counter Watchdog Timer interrupt
IRQ21	37	DeepSleep	Backup domain interrupt
IRQ22	38	DeepSleep	Other combined Interrupts for SRSS (LVD, WDT, CLKCAL)
IRQ23	39	DeepSleep	CPUSS Inter Process Communication Interrupt #0
IRQ24	40	DeepSleep	CPUSS Inter Process Communication Interrupt #1

Table 8-3. List of PSoC 6 Interrupts

System Interrupt	Cortex M4 Exception Number	Power Mode	Interrupt Source
IRQ25	41	DeepSleep	CPUSS Inter Process Communication Interrupt #2
IRQ26	42	DeepSleep	CPUSS Inter Process Communication Interrupt #3
IRQ27	43	DeepSleep	CPUSS Inter Process Communication Interrupt #4
IRQ28	44	DeepSleep	CPUSS Inter Process Communication Interrupt #5
IRQ29	45	DeepSleep	CPUSS Inter Process Communication Interrupt #6
IRQ30	46	DeepSleep	CPUSS Inter Process Communication Interrupt #7
IRQ31	47	DeepSleep	CPUSS Inter Process Communication Interrupt #8
IRQ32	48	DeepSleep	CPUSS Inter Process Communication Interrupt #9
IRQ33	49	DeepSleep	CPUSS Inter Process Communication Interrupt #10
IRQ34	50	DeepSleep	CPUSS Inter Process Communication Interrupt #11
IRQ35	51	DeepSleep	CPUSS Inter Process Communication Interrupt #12
IRQ36	52	DeepSleep	CPUSS Inter Process Communication Interrupt #13
IRQ37	53	DeepSleep	CPUSS Inter Process Communication Interrupt #14
IRQ38	54	DeepSleep	CPUSS Inter Process Communication Interrupt #15
IRQ39	55	Active	Serial Communication Block #0
IRQ40	56	Active	Serial Communication Block #1
IRQ41	57	Active	Serial Communication Block #2
IRQ42	58	Active	Serial Communication Block #3
IRQ43	59	Active	Serial Communication Block #4
IRQ44	60	Active	Serial Communication Block #5
IRQ45	61	Active	Serial Communication Block #6
IRQ46	62	Active	Serial Communication Block #7
IRQ47	63	Active	Serial Communication Block #9
IRQ48	64	Active	Serial Communication Block #10
IRQ49	65	Active	Serial Communication Block #11
IRQ50	66	Active	Serial Communication Block #12
IRQ51	67	Active	CSD (Capsense) interrupt
IRQ52	68	Active	CPUSS DMAC, Channel #0
IRQ53	69	Active	CPUSS DMAC, Channel #1
IRQ54	70	Active	CPUSS DMAC, Channel #2
IRQ55	71	Active	CPUSS DMAC, Channel #3
IRQ56	72	Active	CPUSS DataWire #0, Channel #0
IRQ57	73	Active	CPUSS DataWire #0, Channel #1
IRQ58	74	Active	CPUSS DataWire #0, Channel #2
IRQ59	75	Active	CPUSS DataWire #0, Channel #3
IRQ60	76	Active	CPUSS DataWire #0, Channel #4
IRQ61	77	Active	CPUSS DataWire #0, Channel #5
IRQ62	78	Active	CPUSS DataWire #0, Channel #6
IRQ63	79	Active	CPUSS DataWire #0, Channel #7
IRQ64	80	Active	CPUSS DataWire #0, Channel #8
IRQ65	81	Active	CPUSS DataWire #0, Channel #9
IRQ66	82	Active	CPUSS DataWire #0, Channel #10
IRQ67	83	Active	CPUSS DataWire #0, Channel #11

Table 8-3. List of PSoC 6 Interrupts

System Interrupt	Cortex M4 Exception Number	Power Mode	Interrupt Source
IRQ68	84	Active	CPUSS DataWire #0, Channel #12
IRQ69	85	Active	CPUSS DataWire #0, Channel #13
IRQ70	86	Active	CPUSS DataWire #0, Channel #14
IRQ71	87	Active	CPUSS DataWire #0, Channel #15
IRQ72	88	Active	CPUSS DataWire #0, Channel #16
IRQ73	89	Active	CPUSS DataWire #0, Channel #17
IRQ74	90	Active	CPUSS DataWire #0, Channel #18
IRQ75	91	Active	CPUSS DataWire #0, Channel #19
IRQ76	92	Active	CPUSS DataWire #0, Channel #20
IRQ77	93	Active	CPUSS DataWire #0, Channel #21
IRQ78	94	Active	CPUSS DataWire #0, Channel #22
IRQ79	95	Active	CPUSS DataWire #0, Channel #23
IRQ80	96	Active	CPUSS DataWire #0, Channel #24
IRQ81	97	Active	CPUSS DataWire #0, Channel #25
IRQ82	98	Active	CPUSS DataWire #0, Channel #26
IRQ83	99	Active	CPUSS DataWire #0, Channel #27
IRQ84	100	Active	CPUSS DataWire #0, Channel #28
IRQ85	101	Active	CPUSS DataWire #1, Channel #0
IRQ86	102	Active	CPUSS DataWire #1, Channel #1
IRQ87	103	Active	CPUSS DataWire #1, Channel #2
IRQ88	104	Active	CPUSS DataWire #1, Channel #3
IRQ89	105	Active	CPUSS DataWire #1, Channel #4
IRQ90	106	Active	CPUSS DataWire #1, Channel #5
IRQ91	107	Active	CPUSS DataWire #1, Channel #6
IRQ92	108	Active	CPUSS DataWire #1, Channel #7
IRQ93	109	Active	CPUSS DataWire #1, Channel #8
IRQ94	110	Active	CPUSS DataWire #1, Channel #9
IRQ95	111	Active	CPUSS DataWire #1, Channel #10
IRQ96	112	Active	CPUSS DataWire #1, Channel #11
IRQ97	113	Active	CPUSS DataWire #1, Channel #12
IRQ98	114	Active	CPUSS DataWire #1, Channel #13
IRQ99	115	Active	CPUSS DataWire #1, Channel #14
IRQ100	116	Active	CPUSS DataWire #1, Channel #15
IRQ101	117	Active	CPUSS DataWire #1, Channel #16
IRQ102	118	Active	CPUSS DataWire #1, Channel #17
IRQ103	119	Active	CPUSS DataWire #1, Channel #18
IRQ104	120	Active	CPUSS DataWire #1, Channel #19
IRQ105	121	Active	CPUSS DataWire #1, Channel #20
IRQ106	122	Active	CPUSS DataWire #1, Channel #21
IRQ107	123	Active	CPUSS DataWire #1, Channel #22
IRQ108	124	Active	CPUSS DataWire #1, Channel #23
IRQ109	125	Active	CPUSS DataWire #1, Channel #24
IRQ110	126	Active	CPUSS DataWire #1, Channel #25

Table 8-3. List of PSoC 6 Interrupts

System Interrupt	Cortex M4 Exception Number	Power Mode	Interrupt Source
IRQ111	127	Active	CPUSS DataWire #1, Channel #26
IRQ112	128	Active	CPUSS DataWire #1, Channel #27
IRQ113	129	Active	CPUSS DataWire #1, Channel #28
IRQ114	130	Active	CPUSS Fault Structure Interrupt #0
IRQ115	131	Active	CPUSS Fault Structure Interrupt #1
IRQ116	132	Active	CRYPTO Accelerator Interrupt
IRQ117	133	Active	FLASH Macro Interrupt
IRQ118	134	Active	Floating Point operation fault
IRQ119	135	Active	CM0+ CTI #0
IRQ120	136	Active	CM0+ CTI #1
IRQ121	137	Active	CM4 CTI #0
IRQ122	138	Active	CM4 CTI #1
IRQ123	139	Active	TCPWM #0, Counter #0
IRQ124	140	Active	TCPWM #0, Counter #1
IRQ125	141	Active	TCPWM #0, Counter #2
IRQ126	142	Active	TCPWM #0, Counter #3
IRQ127	143	Active	TCPWM #0, Counter #4
IRQ128	144	Active	TCPWM #0, Counter #5
IRQ129	145	Active	TCPWM #0, Counter #6
IRQ130	146	Active	TCPWM #0, Counter #7
IRQ131	147	Active	TCPWM #1, Counter #0
IRQ132	148	Active	TCPWM #1, Counter #1
IRQ133	149	Active	TCPWM #1, Counter #2
IRQ134	150	Active	TCPWM #1, Counter #3
IRQ135	151	Active	TCPWM #1, Counter #4
IRQ136	152	Active	TCPWM #1, Counter #5
IRQ137	153	Active	TCPWM #1, Counter #6
IRQ138	154	Active	TCPWM #1, Counter #7
IRQ139	155	Active	TCPWM #1, Counter #8
IRQ140	156	Active	TCPWM #1, Counter #9
IRQ141	157	Active	TCPWM #1, Counter #10
IRQ142	158	Active	TCPWM #1, Counter #11
IRQ143	159	Active	TCPWM #1, Counter #12
IRQ144	160	Active	TCPWM #1, Counter #13
IRQ145	161	Active	TCPWM #1, Counter #14
IRQ146	162	Active	TCPWM #1, Counter #15
IRQ147	163	Active	TCPWM #1, Counter #16
IRQ148	164	Active	TCPWM #1, Counter #17
IRQ149	165	Active	TCPWM #1, Counter #18
IRQ150	166	Active	TCPWM #1, Counter #19
IRQ151	167	Active	TCPWM #1, Counter #20
IRQ152	168	Active	TCPWM #1, Counter #21
IRQ153	169	Active	TCPWM #1, Counter #22

Table 8-3. List of PSoC 6 Interrupts

System Interrupt	Cortex M4 Exception Number	Power Mode	Interrupt Source
IRQ154	170	Active	TCPWM #1, Counter #23
IRQ155	171	Active	SAR ADC interrupt
IRQ156	172	Active	I2S0 Audio interrupt
IRQ157	173	Active	PDM0/PCM0 Audio interrupt
IRQ158	174	Active	I2S1 Audio interrupt
IRQ159	175	Active	Profiler interrupt
IRQ160	176	Active	Serial Memory Interface interrupt
IRQ161	177	Active	USB Interrupt
IRQ162	178	Active	USB Interrupt
IRQ163	179	Active	USB Interrupt
IRQ164	180	Active	SDIO wakeup interrupt for mxsdhc
IRQ165	181	Active	Consolidated interrupt for mxsdhc for everything else
IRQ166	182	Active	EEMC wakeup interrupt for mxsdhc, not used
IRQ167	183	Active	Consolidated interrupt for mxsdhc for everything else

8.6 Interrupt/Exception Priority

Exception priority is useful for exception arbitration when there are multiple exceptions that need to be serviced by the CPU. Both M4 and M0+ cores in PSoC 6 MCUs provide flexibility in choosing priority values for different exceptions. All exceptions other than Reset, NMI, and HardFault can be assigned a configurable priority level. The Reset, NMI, and HardFault exceptions have a fixed priority of -3, -2, and -1, respectively. In PSoC 6 MCUs, lower priority numbers represent higher priorities. This means that the Reset, NMI, and HardFault exceptions have the highest priorities. The other exceptions can be assigned a configurable priority level between 0 and 3 for Cortex-M0+ and 0 to 7 for Cortex-M4.

Both M0+ and M4 support nested exceptions in which a higher priority exception can obstruct (interrupt) the currently active exception handler. This pre-emption does not happen if the incoming exception priority is the same as or lower than the active exception. The CPU resumes execution of the lower priority exception handler after servicing the higher priority exception. The CM0+ core in the PSoC 6 MCU allows nesting of up to four exceptions; the CM4 core allows up to eight exceptions. When the CPU receives two or more exceptions requests of the same priority, the lowest exception number is serviced first.

The registers to configure the priority of exception numbers 1 to 15 are explained in [Exception Sources on page 58](#).

The priority of the 32 CM0+ and 168 CM4 interrupts can be configured by writing to the respective Interrupt Priority registers (CM0P_SCS_IPR and CM4_SCS_IPR). This is a group of eight (CM0+) and 60 (CM4) 32-bit registers with each register storing the priority values of four interrupts, as given in [Table 8-4](#) and [Table 8-5](#). For CM0+, the first 16 bits of the ISER and ICER registers are valid. Refer to [Interrupt Sources on page 60](#) for details on how to map a CPU interrupt to a system interrupt in CM0+.

Table 8-4. Interrupt Priority Register Bit Definitions for Cortex-M0+ (CM0P_SCS_IPR)

Bits	Name	Description
7:6	PRI_N0	Priority of interrupt number N.
15:14	PRI_N1	Priority of interrupt number N+1.
23:22	PRI_N2	Priority of interrupt number N+2.
31:30	PRI_N3	Priority of interrupt number N+3.

Table 8-5. Interrupt Priority Register Bit definitions for Cortex-M4 (CM4_SCS_IPR)

Bits	Name	Description
7:5	PRI_N0	Priority of interrupt number N
15:13	PRI_N1	Priority of interrupt number N+1
23:21	PRI_N2	Priority of interrupt number N+2
31:29	PRI_N3	Priority of interrupt number N+3

8.7 Enabling and Disabling Interrupts

The NVICs of both CM0+ and CM4 core provide registers to individually enable and disable the interrupts in software. If an interrupt is not enabled, the NVIC will not process the interrupt requests on that interrupt line. The Interrupt Set-Enable Register (CM0P_SCS_ISER and CM4_SCS_ISER) and the Interrupt Clear-Enable Register (CM0P_SCS_ICER and CM4_SCS_ICER) are used to enable and disable the interrupts respectively. These registers are 32-bit wide and each bit corresponds to the same numbered interrupt in CM0+. For CM4 core, there are eight ISER/ICER registers. These registers can also be read in software to get the enable status of the interrupts. [Table 8-6](#) shows the register access properties for these two registers. Note that writing zero to these registers has no effect.

For CM0+, the first 16 bits of the ISER and ICER registers are valid. Refer to [Interrupt Sources on page 60](#) for details on how to map a CPU interrupt to a system interrupt in CM0+.

Table 8-6. Interrupt Enable/Disable Registers

Register	Operation	Bit Value	Comment
Interrupt Set Enable Register	Write	1	To enable the interrupt
		0	No effect
	Read	1	Interrupt is enabled
		0	Interrupt is disabled
Interrupt Clear Enable Register	Write	1	To disable the interrupt
		0	No effect
	Read	1	Interrupt is enabled
		0	Interrupt is disabled

The ISER and ICER registers are applicable only for the interrupts. These registers cannot be used to enable or disable the exception numbers 1 to 15. The 15 exceptions have their own support for enabling and disabling, as explained in [Exception Sources on page 58](#).

The PRIMASK register in the CPUs (both CM0+ and CM4) can be used as a global exception enable register to mask all the configurable priority exceptions irrespective of whether they are enabled. Configurable priority exceptions include all the exceptions except Reset, NMI, and HardFault listed in [Table 8-1](#). When the PM bit (bit 0) in the PRIMASK register is set, none of the configurable priority exceptions can be serviced by the CPU, though they can be in the pending state waiting to be serviced by the CPU after the PM bit is cleared.

8.8 Interrupt/Exception States

Each exception can be in one of the following states.

Table 8-7. Exception States

Exception State	Meaning
Inactive	The exception is not active and not pending. Either the exception is disabled or the enabled exception has not been triggered.
Pending	The exception request has been received by the CPU/NVIC and the exception is waiting to be serviced by the CPU.
Active	An exception that is being serviced by the CPU but whose exception handler execution is not yet complete. A high-priority exception can interrupt the execution of lower priority exception. In this case, both the exceptions are in the active state.
Active and Pending	The exception is being serviced by the processor and there is a pending request from the same source during its exception handler execution.

The Interrupt Control State Register (CM0P_SCS_ICSR and CM4_SCS_ICSR) contains status bits describing the various exceptions states.

- The VECTACTIVE bits ([8:0]) in the ICSR store the exception number for the current executing exception. This value is zero if the CPU does not execute any exception handler (CPU is in thread mode). Note that the value in VECTACTIVE bitfields is the same as the value in bits [8:0] of the Interrupt Program Status Register (IPSR), which is also used to store the active exception number.
- The VECTPENDING bits ([20:12]) in the ICSR store the exception number of the highest priority pending exception. This value is zero if there are no pending exceptions.
- The ISRPENDING bit (bit 22) in the ICSR indicates if a NVIC generated interrupt is in a pending state.

8.8.1 Pending Interrupts/Exceptions

When a peripheral generates an interrupt request signal to the NVIC or an exception event occurs, the corresponding exception enters the pending state. When the CPU starts executing the corresponding exception handler routine, the exception is changed from the pending state to the active state. The NVIC allows software pending of the 32 (CM0+) or 168 (CM4) interrupt lines by providing separate register bits for setting and clearing the pending states of the interrupts. The Interrupt Set-Pending register (CM0P_SCS_ISPR and CM4_SCS_ISPR) and the Interrupt Clear-Pending register (CM0P_SCS_ICPR and CM4_SCS_ICPR) are used to set and clear the pending status of the interrupt lines. These registers are 32 bits wide, and each bit corresponds to the same numbered interrupt. In the case of CM4, there are eight sets of such registers to accommodate all 168 interrupts. [Table 8-8](#) shows the register access properties for these two registers. Note that writing zero to these registers has no effect.

Table 8-8. Interrupt Set Pending/Clear Pending Registers

Register	Operation	Bit Value	Comment
Interrupt Set-Pending Register (ISPR)	Write	1	To put an interrupt to pending state
		0	No effect
	Read	1	Interrupt is pending
		0	Interrupt is not pending
Interrupt Clear-Pending Register (ICPR)	Write	1	To clear a pending interrupt
		0	No effect
	Read	1	Interrupt is pending
		0	Interrupt is not pending

Setting the pending bit when the same bit is already set results in only one execution of the ISR. The pending bit can be updated regardless of whether the corresponding interrupt is enabled. If the interrupt is not enabled, the interrupt line will not move to the pending state until it is enabled by writing to the ISER register.

Note that the ISPR and ICPR registers are used only for the peripheral interrupts. These registers cannot be used for pending the exception numbers 1 to 15. These 15 exceptions have their own support for pending, as explained in [Exception Sources on page 58](#).

8.9 Stack Usage for Interrupts/ Exceptions

When the CPU executes the main code (in thread mode) and an exception request occurs, the CPU stores the state of its general-purpose registers in the stack. It then starts executing the corresponding exception handler (in handler mode). The CPU pushes the contents of the eight 32-bit internal registers into the stack. These registers are the Program and Status Register (PSR), ReturnAddress, Link Register (LR or R14), R12, R3, R2, R1, and R0. Both Cortex-M4 and Cortex-M0+ have two stack pointers - MSP and PSP. Only one of the stack pointers can be active at a time. When in thread mode, the Active Stack Pointer bit in the Control register is used to define the current active stack pointer. When in handler mode, the MSP is always used as the stack pointer. The stack pointer always grows downwards and points to the address that has the last pushed data.

When the CPU is in thread mode and an exception request comes, the CPU uses the stack pointer defined in the control register to store the general-purpose register contents. After the stack push operations, the CPU enters handler mode to execute the exception handler. When another higher priority exception occurs while executing the current exception, the MSP is used for stack push/pop operations, because the CPU is already in handler mode. See the [CPU Subsystem \(CPUSS\) chapter on page 33](#) for details.

8.10 Interrupts and Low-Power Modes

The PSoC 6 MCU family allows device (CPU) wakeup from low-power modes when certain peripheral interrupt requests are generated. The Wakeup Interrupt Controller (WIC) block generates a wakeup signal that causes the CPU to enter Active mode when one or more wakeup sources generate an interrupt signal. After entering Active mode, the ISR of the peripheral interrupt is executed.

The Wait For Interrupt (WFI) or Wait For Event (WFE) instruction, executed by the CPU, triggers the transition into Sleep, and Deep Sleep modes. Both the WFI and WFE instructions are capable of waking up on interrupts. However, the WFE requires the interrupts to be unmasked in the CPU's Priority Mask register. Refer to the PRIMASK register definition on the [Arm website](#). In addition, the WFE instruction puts the CPU to sleep based on the status of an event bit and wakes up from an event signal, typically sent by the other CPU. WFI does not require PRIMASK unmasking and can wake up the CPU from any pending interrupt masked to the NVIC or WIC. However, WFI cannot wake up the CPU from event signals from other CPUs. The sequence of entering the different low-power modes is detailed in the [Device Power Modes chapter on page 190](#).

Chip low-power modes have two categories of interrupt sources:

- Interrupt sources that are available in the Active, Sleep, and Deep Sleep modes (watchdog timer interrupt, RTC, GPIO interrupts, and Low-Power comparators)
- Interrupt sources that are available only in the Active and Sleep modes

When using the WFE instruction in CM4, make sure to call the WFE instruction twice to properly enter and exit Sleep/ Deep Sleep modes. This behavior comes from the event register implementation in Arm v7 architecture used in Cortex-M4. According to the ARM V7 architecture reference manual (Section B1.5.18 Wait For Event and Send Event):

- A reset clears the event register.
- Any WFE wakeup event, or the execution of an exception return instruction, sets the event register.
- A WFE instruction clears the event register.
- Software cannot read or write the value of the event register directly.

Therefore, the first WFE instruction puts CM4 to sleep and second WFE clears the event register after a WFE wakeup, which sets the event register. So the next WFE will put the core to sleep.

Note that this behavior is not present in Arm v6 architecture used in Cortex-M0+. Therefore, in CM0+ only one WFE instruction is sufficient to successfully enter or exit Sleep and Deep Sleep modes.

8.11 Interrupt/Exception – Initialization/ Configuration

This section covers the different steps involved in initializing and configuring exceptions in the PSoC 6 MCU.

1. Configuring the Exception Vector Table Location: The first step in using exceptions is to configure the vector table location as required - either in flash memory or SRAM. This configuration is done as described in [Exception Vector Table on page 57](#).

The vector table should be available in SRAM if the application must change the vector addresses dynamically. If the table is located in flash, then a flash write operation is required to modify the vector table contents. The ModusToolbox IDE uses the vector table in SRAM by default.

2. Configuring Individual Exceptions: The next step is to configure individual exceptions required in an application, as explained in earlier sections.
 - a. Configure the exception or interrupt source; this includes setting up the interrupt generation conditions. The register configuration depends on the specific exception required. Refer to the respective peripheral chapter to know more about the interrupt configuration supported by them.

- b. Define the exception handler function and write the address of the function to the exception vector table. [Table 8-1](#) gives the exception vector table format; the exception handler address should be written to the appropriate exception number entry in the table.
- c. For Cortex-M0+, define and enable the additional system interrupt handler table and functions as explained in [Interrupt Sources on page 60](#).
- d. Set up the exception priority, as explained in [Interrupt/Exception Priority on page 67](#).
- e. Enable the exception, as explained in [Enabling and Disabling Interrupts on page 67](#).

8.12 Register List

Table 8-9. Register List

Register Name	Description
CPUSS_CM0_NMI_CTLx	Cortex-M0+ NMI control registers (4 registers)
CPUSS_CM0_SYSTEM_INT_CTL0x	Cortex-M0+ interrupt control 0 registers for x th system interrupt (168 registers)
CPUSS_CM0_SYSTEM_INT_CTL1x_STATUS	Cortex-M0+ interrupt control 1 status registers for x th CPU interrupt (8 registers)
CPUSS_CM0_INT_CTL2	Cortex-M0+ interrupt control 2
CPUSS_CM0_INT_CTL3	Cortex-M0+ interrupt control 3
CPUSS_CM0_INT_CTL4	Cortex-M0+ interrupt control 4
CPUSS_CM0_INT_CTL5	Cortex-M0+ interrupt control 5
CPUSS_CM0_INT_CTL6	Cortex-M0+ interrupt control 6
CPUSS_CM0_INT_CTL7	Cortex-M0+ interrupt control 7
CPUSS_CM4_NMI_CTLx	Cortex-M4 NMI control registers (4 registers)
SYSTEM_CM0P_SCS_ISER	Cortex-M0+ interrupt set-enable register
SYSTEM_CM0P_SCS_ICER	Cortex-M0+ interrupt clear enable register
SYSTEM_CM0P_SCS_ISPR	Cortex-M0+ interrupt set-pending register
SYSTEM_CM0P_SCS_ICPR	Cortex-M0+ interrupt clear-pending register
SYSTEM_CM0P_SCS_IPR	Cortex-M0+ interrupt priority register
SYSTEM_CM0P_SCS_ICSR	Cortex-M0+ interrupt control state register
SYSTEM_CM0P_SCS_VTOR	Cortex-M0+ vector table offset register
SYSTEM_CM0P_SCS_AIRCR	Cortex-M0+ application interrupt and reset control register
SYSTEM_CM0P_SCS_SHPR2	Cortex-M0+ system handler priority register 2
SYSTEM_CM0P_SCS_SHPR3	Cortex-M0+ system handler priority register 3
SYSTEM_CM0P_SCS_SHCSR	Cortex-M0+ system handler control and state register
SYSTEM_CM4_SCS_ISER	Cortex-M4 interrupt set-enable register
SYSTEM_CM4_SCS_ICER	Cortex-M4 interrupt clear enable register
SYSTEM_CM4_SCS_ISPR	Cortex-M4 interrupt set-pending register
SYSTEM_CM4_SCS_ICPR	Cortex-M4 interrupt clear-pending register
SYSTEM_CM4_SCS_IPR	Cortex-M4 interrupt priority registers
SYSTEM_CM4_SCS_ICSR	Cortex-M4 interrupt control state register
SYSTEM_CM4_SCS_VTOR	Cortex-M4 vector table offset register
SYSTEM_CM4_SCS_AIRCR	Cortex-M4 application interrupt and reset control register
SYSTEM_CM4_SCS_SHPR2	Cortex-M4 system handler priority register 2
SYSTEM_CM4_SCS_SHPR3	Cortex-M4 system handler priority register 3
SYSTEM_CM4_SCS_SHCSR	Cortex-M4 system handler control and state register

9. Protection Units



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

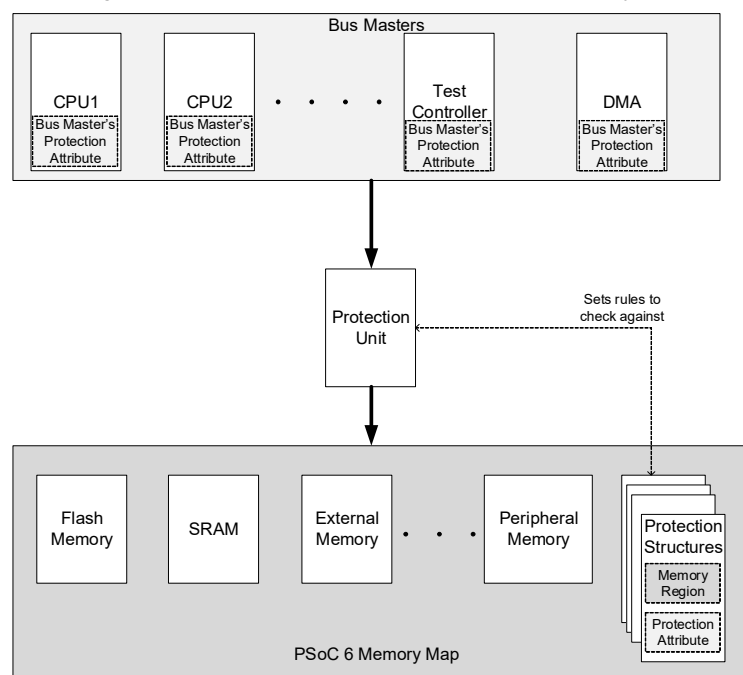
- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - Protection Unit
- Application notes
- Code examples

Protection units are implemented in the PSoC 6 MCU to enforce security based on different operations. A protection unit allows or restricts bus transfers. The rules are enforced based on specific properties of a transfer. The rules that determine protection are implemented in protection structures (a register structure). A protection structure defines the protected address space and the protection attributes. The hardware that evaluates these protection structures, to restrict or permit access, is the protection unit. The PSoC device has different types of protection units such as MPU, S MPU, and PPU. Each have a distinct set of protection structures, which helps define different protection regions and their attributes.

9.1 Architecture

Figure 9-1 shows a conceptual view of implementation of the PSoC protection system.

Figure 9-1. Conceptual View of PSoC Protection System



The functioning of a secure system is based on the following:

- **Bus masters:** This term refers to the bus masters in the architecture. In a PSoC 6 device, an example of a bus master is a Cortex-M core, DMA, or a test controller.
- **Protection units:** Protection units are the hardware engines that enforce the protection defined by protection structures. There are three types of protection units, acting at different levels of memory access with different precedence and priority of protection – MPU, SMPU, and PPU.
- **Protection structure:** A protection structure is a register structure in memory that sets up the rules based on which each protection unit will evaluate a transfer. Each protection unit associates itself to multiple protection structures. The protection structure associated with a protection unit are evaluated in the order starting with the protection structure with the largest index. For example, if there are 16 protection structures associated with a protection unit, then the evaluation of a transfer starts from protection structure 15 and counts down. Physically a protection structure is a register structure in the memory map that defines a protection rule. Each protection structure constitutes the following:
 - Defines a memory region on which the rule is applied. It designates what the bus transfer needs to be evaluated against this protection structure.
 - Base address
 - Size of memory block
 - A set of protection attributes
 - R/W/X
 - User/privilege
 - Secure/non Secure
 - Protection context
- **Protection attributes:** These are properties based on which a transfer is evaluated. There are multiple protection attributes. The set of protection attributes available for a protection structure depends on the protection unit it is associated with. Protection attributes appear in two places:
 - **Protection structures:** Protection attributes associated with a protection structure set the rules for access based on these attributes.
 - **Bus master's protection attribute:** Each bus master has its own access attributes, which define the bus master's access privileges. Some of these attributes, such as secure/non-secure, are set for a master. Other attributes such as protection context and user/privilege attribute are dynamic attributes, which change based on bus master's context and state.

In summary, a PSoC 6 device has protection units that act as a gate for any access to the PSoC memory map. The rules for protection are set by the protection structures.

Each bus master is qualified by its own protection attribute. For every bus transfer, the protection unit compares the bus master's protection attribute and accessed address against the rules set in the protection structures and decides on providing or denying access.

9.2 PSoC 6 Protection Architecture

When there is a memory (SRAM/flash/peripheral) access by a bus master, the access is evaluated by a protection unit against the protection attributes set in protection structures for the memory location being accessed. If the bus master's protection attributes satisfy the protection attributes set in the protection structures, then access is allowed by the protection unit. If there is an access restriction, a fault condition is triggered and a bus error occurs. Thus protection units secure bus transfer address range either in memory locations (SRAM/flash) or peripheral registers. From an architectural perspective, there is no difference between memory protection and peripheral protection. However, from an implementation perspective, separate memory and peripheral protection is provided.

Two types of protection units, memory protection units (MPU) and shared memory protection units (SMPU), are provided in the CPU subsystem (CPUSS) to protect memory locations. A separate protection unit type is provided for peripheral protection (PPU) in the PERI:

- A bus master may have a dedicated MPU. In a CPU bus master, the MPU is typically implemented as part of the CPU and is under control of the OS/kernel. In a non-CPU bus master, the MPU is typically implemented as part of the bus infrastructure and under control of the OS/kernel of the CPU that "owns or uses" the bus master. If a CPU switches tasks or if a non-CPU switches ownership, the MPU settings are typically updated by OS/kernel software. The different MPU types are:
 - An MPU that is implemented as part of the CPU. This type is found in the Arm CM0+ and CM4 CPUs.
 - An MPU that is implemented as part of the bus infrastructure. This type is found in bus masters such as crypto and test controller. The definition of this MPU type follows the Arm MPU definition (in terms of memory region and access attribute definition) to ensure a consistent software interface.
- SMPUs are intended for implementing protection in a situation with multiple bus masters. These protection units implement a concept called Protection Context. A protection context is a pseudo state of a bus master, which can be used to determine access attributes across multiple masters. The protection context is a protection attribute not specific to a bus master. The SMPUs can distinguish between different protection contexts; they can also distinguish secure from non-secure accesses.

This allows for an effective protection in a multi-core scenario.

- PPU are protection units provided in the PERI register space for peripheral protection. The PPU attributes are similar to the SMPU, except that they are intended for protecting the peripheral space. Refer to the [registers TRM](#) for details. The PPUs are intended to distinguish between different protection contexts and to distinguish secure from non-secure accesses and user mode accesses from privileged mode accesses. There are two types of PPU structures.
 - Fixed PPUs implement protection for fixed address regions that typically correspond to a specific peripheral
 - Programmable PPUs allows the user to program the address region to be protected

The platform's DMA controller does not have an MPU. Instead, a DMA controller channel inherits the access control attributes of the bus transfer that programmed the channel.

The definition of SMPU and PPU follows the MPU definition and adds the capability to distinguish accesses from different protection contexts (the MPU does not include support for a protection context). If security is required, the SMPU and possibly PPUs MMIO registers must be controlled by a secure CPU that enforces system-wide protection.

[Figure 9-2](#) gives an overview of the location of MPUs, SMPUs, and PPUs in the system. Note that a peripheral group PPU needs to provide access control only to the peripherals within a peripheral group (group of peripherals with a shared bus infrastructure).

As mentioned, the MPU, SMPU, and PPU protection functionality follows the Arm MPU definition:

- Multiple protection structures are supported.
- Each structure specifies an address range in the unified memory architecture and access attributes. An address range can be as small as 32 bytes.

A protection violation is caused by a mismatch between a bus master's access attributes and the protection structure and access attributes for the memory region configured in the protection structure.

A bus transfer that violates a protection structure results in a bus error.

For AXI transfers, the complete address range is matched. If a transfer references multiple 32-byte regions (the smallest protection structure address range is 32 bytes), multiple cycles are required for matching – one cycle per 32-byte region.

Protection violations are captured in the fault report structure to allow for failure analysis. The fault report structures can generate an interrupt to indicate the

occurrence of a fault. This is useful if the violating bus master cannot resolve the bus error by itself, but requires another CPU bus master to resolve the bus error on its behalf.

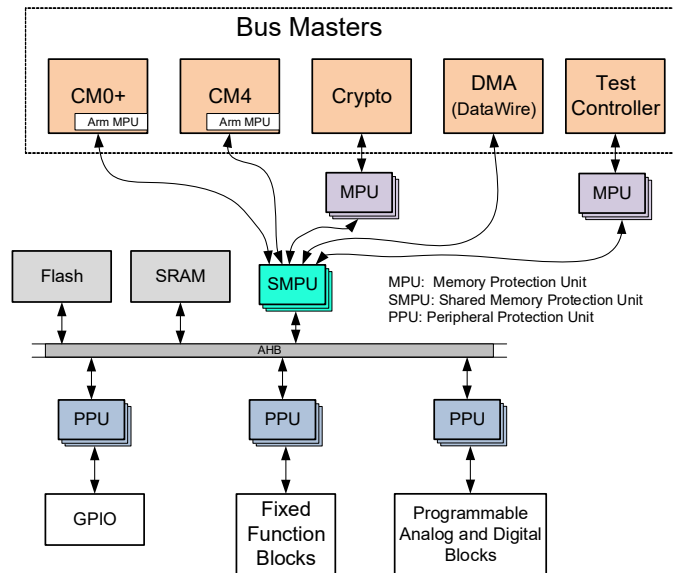
A protection violation results in a bus error and the bus transfer will not reach its target. An MPU or SMPU violation that targets a peripheral will not reach the associated protection evaluation (PPU). In other words, MPU and SMPU have a higher priority over PPU.

Protection unit addresses the following:

- Security requirements. This includes prevention of malicious attacks to access secure memory or peripherals. For example, a non-secure master should not be able to access key information in a secure memory region.
- Safety requirements. This includes detection of accidental (non-malicious) software errors and random hardware errors. Enabling failure analysis is important so the root cause of a safety violation can be investigated. For example, analyzing a flash memory failure on a device that is returned from the field should be possible.

To address security requirements, the Cortex M0+ is used as a 'secure CPU'. This CPU is considered a trusted entity.

Figure 9-2. PSoC 6 Protection Architecture



The different types of protection units cater to different use cases for protection.

Table 9-1. Protection Use Cases

Protection Unit Type	Use	Number of Instances
Arm MPU	Used to protect memory between tasks within a single Arm core. A task in one of the Arm cores can protect its memory from access by another task in the same core.	
MPU	Same as the Arm MPU, but for other bus masters such as the test controller or crypto, which do not have a built-in MPU in their block IP.	
SMPU	Used to protect memory addresses that are shared between multiple bus masters.	
Fixed PPU protection structures	These protect specific peripheral memory space. The protection structures have a preprogrammed memory region and can be used only to protect the peripheral it was intended for.	
Programmable PPU protection structures	These protect the peripheral space but the memory region is not fixed. So users can easily program it to protect any space in the peripheral memory region.	

9.3 Register Architecture

The protection architecture has different conceptual pieces and different sets of registers correspond to each of these concepts.

9.3.1 Protection Structure and Attributes

The MPU, SMPU, and PPU protection structure definition follows the Arm definition. Each protection structure is defined by:

- An address region
- Access control attributes

A protection structure is always aligned on a 32-byte boundary in the memory space. Two registers define a protection structure: ADDR (address register) and ATT (attribute register). This structure alignment and organization allow straightforward protection of the protection structures by the protection scheme. This is discussed later in this chapter.

Address region: The address region is defined by:

- The base address of a region as specified by ADDR.ADDR.
- The size of a region as specified by ATT.REGION_SIZE.

- Individual disables for eight subregions within the region, as specified by ADDR.SUBREGION_DISABLE.

The REGION_SIZE field specifies the size of a region. The region size is a power of 2 in the range of [256 B, 4 GB]. The base address ADDR specifies the start of the region, which must be aligned to the region size. A region is partitioned into eight equally sized sub-regions. The SUBREGION_DISABLE field specifies individual enables for the sub-regions within a region. For example, a REGION_SIZE of "0x08" specifies a region size of 512 bytes. If the start address is 0x1000:5400 (512-byte aligned), the region ranges from 0x1000:5400 to 0x1000:55ff. This region is partitioned into the following eight 64-byte subregions:

subregion 0 from 0x1000:5400 to 0x1000:543f

subregion 1 from 0x1000:5440 to 0x1000:547f

...

subregion 7 from 0x1000:55c0 to 0x1000:55ff.

If the SUBREGION_DISABLE is 0x82 (bitfields 1 and 7 are '1'), subregions 1 and 7 are disabled; subregions 0, 2, 3, 4, 5, and 6 are enabled.

In addition, an ATT.ENABLED field specifies whether the region is enabled. Only enabled regions participate in the protection "matching" process. Matching identifies if a bus transfer address is contained within an enabled subregion (SUBREGION_DISABLE) of an enabled region (ENABLED).

Protection attributes: The protection attributes specify access control to the region (shared by all subregions within the region). Access control is performed by comparing against a bus master's protection attributes of the bus master performing the transfer. The following access control fields are supported:

- Control for read accesses in user mode (ATT.UR field).
- Control for write accesses in user mode (ATT.UW field).
- Control for execute accesses in user mode (ATT.UX field).
- Control for read accesses in privileged mode (ATT.PR field).
- Control for write accesses in privileged mode (ATT.PW field).
- Control for execute accesses in privileged mode (ATT.PX field).
- Control for secure access (ATT.NS field).
- Control for individual protection contexts (ATT.PC_MASK[15:0], with MASK[0] always constant at 1). This protection context control field is present only for the SMPU and PPU.

The execute and read access control attributes are orthogonal. Execute transfers are typically read transfers. To allow execute and read transfers in user mode, both ATT.UR and ATT.UX must be set to '1'. To allow data and read transfers in user mode, only ATT.UR must be set to '1'. In addition, the ATT.PC_MATCH control field is supported, which controls "matching" and "access evaluation" processes. This control field is present only for the SMPU and PPU protection structures.

For example, only protection context 2 can access a specific address range. These accesses are restricted to read and write secure accesses in privileged mode. The access control fields are programmed as follows:

- ATT.UR is 0: read accesses in user mode not allowed.
- ATT.UW is 0: write accesses in user mode not allowed.
- ATT.UX is 0: execute accesses in user mode not allowed.
- ATT.PR is 1: read accesses in privileged mode allowed.
- ATT.PW is 1: write accesses in privileged mode allowed.
- ATT.PX is 0: execute accesses in privileged mode not allowed.
- ATT.NS is 0: secure access required.
- ATT.PC_MASK is 0x0005: protection context 1 and 3 accesses enabled (all other protection contexts are disabled).
- ATT.PC_MATCH is 0: the PC_MASK field is used for access evaluation. Three separate access evaluation subprocesses are distinguished:
 - A subprocess that evaluates the access based on read/write, execute, and user/privileged access attributes.
 - A subprocess that evaluates the access based on the secure/non-secure attribute.

- A subprocess that evaluates the access based on the protection context index (used only by the SMPU and PPU when ATT.PC_MATCH is 0).

If all access evaluations are successful, access is allowed. If any process evaluation is unsuccessful, access is not allowed. Matching the bus transfer address and access evaluation of the bus transfer (based on access attributes) are two independent processes:

- Matching process. For each protection structure, the process identifies whether a transfer address is contained within the address range. This identifies the “matching” regions.
- Access evaluation process. For each protection structure, the process evaluates the bus transfer access attributes against the access control attributes.

A protection unit typically has multiple protection structures and evaluates the protection structures in decreasing order. The first matching structure provides the access control attributes for the evaluation of the transfer's access attributes. In other words, higher-indexed structures take precedence over lower-indexed structures.

The following pseudo code illustrates the process.

```
match = 0;
for (i = n-1; i >= 0; i--) // n: number of protection regions
    if (Match ("transfer address", "protection context"
              "MMIO registers ADDR and ATT of protection structure i")) {
        match = 1; break;
    }
}

if (match)
    AccessEvaluate ("transfer access attributes", "protection context"
                  "MMIO register ATT of protection structure i");
else
    "access allowed"
```

Notes:

- If no protection structure provides a match, the access is allowed.
- If multiple protection structures provide a match, the access control attributes for the access evaluation are provided by the protection structure with the highest index.

An example of using the PC_MATCH feature is as follows. Two SMPU structures are configured to protect the same address range:

- Case 1: SMPU#2: PC = 3, PC_MATCH = 0 SMPU#1: PC = 2, PC_MATCH = 0 To access the master of protection context 2, SMPU#2 has highest index and the address match, but attributes do not match; therefore, access is restricted. The SMPU#1 is not evaluated because the PC_MATCH is 0.
- Case 2: SMPU#2: PC = 3, PC_MATCH = 1 SMPU#1: PC = 2, PC_MATCH = 0 The SMPU#2 address matches but PC does not match, and is skipped because PC_MATCH is 1. SMPU#1 is evaluated and the address and attributes match; therefore, access is allowed.

As mentioned, the protection unit evaluates the protection structures in decreasing order. From a security requirements perspective, this is of importance: a non-secure protection context must not be able to add protection structures that have a higher index than the protection structures that provide secure access. The protection structure with a higher index can be programmed to allow non-secure accesses. Therefore, in a secure system, the higher programmable protection structures are protected to only allow restricted accesses

9.4 Bus Master Protection Attributes

The protection structures set up the rules for different memory regions and their access attributes. The bus master's own protection attributes are used by the protection units to regulate access, based on rules set by the protection structures. Not all bus masters provide all protection attributes that are associated with a bus transfer. Some examples are:

- None of the bus masters has a native protection context attribute. This must be set dynamically based on the task being executed by the bus master.

- The Arm Cortex M4 and Arm Cortex M0+ CPUs provide a user/privilege attribute, but do not provide a secure/non-secure attribute natively. This must be set at a system level.

To ensure system-wide restricted access, missing attributes are provided by register fields. These fields may be set during the boot process or by the secure CPU.

- The SMPU MS_CTL.PC_MASK[] and MPU MS_CTL.PC[] register fields provide protection context functionality.
- The SMPU MS_CTL.P register field provides the user/privileged attribute for those masters that do not provide their own attribute.
- The SMPU MS_CTL.NS register provides the secure/non-secure attribute for those masters that do not provide their own attribute.
- Masters that do not provide an execute attribute have the execute attribute set to '0'.

The DMA controller channels inherit the access control attributes of the bus transfers that configured the DMA channel.

- All the bus masters in the system have SMPU and MPU MS_CTL registers associated with them.
- The MPU MS_CTL.PC_SAVED field (and associated protection context 0 functionality, which is discussed later in the chapter) is only present for the CM0+ master.
- The SMPU MS_CTL.P, MS_CTL.NS, and MS_CTL.PC_MASK fields are not present for the DMA. The bus transfer attributes are provided through "inheritance": the bus transfer attributes are from the master that owns the DMA channel that initiated the bus transfer.
- The MPU MS_CTL register is not present for the DMA masters. The protection context (PC) bus transfer attribute is provided through inheritance.

9.5 Protection Context

Each bus master has a MPU MS_CTL.PC[3:0] protection context field. This protection context is used as the protection context attribute for all bus transfers that are initiated by the master. The SMPUs and PPU allow or restrict bus transfers based on the protection context attribute.

Multiple masters can share a protection context. For example, a CPU and a crypto controlled by the CPU may share a protection context (the CPU and crypto PC[] fields are the same). Therefore, the CPU and crypto share the SMPU and PPU access restrictions.

In the PSoC 64 family, the protection context levels have been pre-defined. See [Table 9-2](#).

Table 9-2. Protection Context Assignment

PC	Secure	Description	CPU	Privileged	User
0	Yes	Hardware Root-of-Trust, SysCall access only	CM0+	SysCall	<not used>
1	Yes	Cypress Trusted Environment and Functions	CM0+	BootLoader/SPM/Drivers	Trusted Functions
2	Yes	Third-party Trusted Functions (optional)	CM0+	BootLoader/SPM/Drivers	Trusted Functions
3	Yes	Third-party Trusted Functions (optional)	CM0+	BootLoader/SPM/Drivers	Trusted Functions
4	Yes	Third-party Trusted Functions (optional)	CM0+	BootLoader/SPM/Drivers	Trusted Functions
5	No	<reserved>	—	—	—
6	No	Application	CM4	OperatingSystem/Drivers	Application
7	No	DAP SYSAP	DAP	<not used>	Debugger/Tester

A bus master protection context is changed by reprogramming the master's PC[] field. Changing a protection context is required for CPU bus masters that may transition between multiple tasks, each catering to different protection contexts. As the protection context allows or restricts bus transfers, changes to the protection context should be controlled and should not compromise security. Furthermore, changes to the protection context should incur limited CPU overhead to allow for frequent protection context changes. Consider a case in which a CPU executes two software stacks with different protection contexts. To this end, each bus master has an SMPU MS_CTL.PC_MASK[15:0] protection context mask field that identifies what protection contexts can be programmed for the bus master:

- The protection context field MS_CTL.PC[3:0]. This register is controlled by the associated bus master and has the same access restrictions as the bus master's MPU registers.

- The protection context mask field MS_CTL.PC_MASK[15:0]. This register is controlled by the secure CPU and has the same access restrictions as the S MPU registers.

The PC_MASK[] field is a “hot-one” field that specifies whether the PC[] field can be programmed with a specific protection context. Consider an attempt to program PC[] to ‘3’:

- If PC_MASK[3] is ‘1’, PC[] is set to “3”.
- If PC_MASK[3] is ‘0’, PC[] is not changed.

9.6 Protection Contexts 0, 1, 2, 3

The PSoC 6 MCU supports protection contexts to isolate software execution for security and safety purposes. Protection contexts are used to restrict access to memory and peripheral resources.

Out of a maximum of 16 protections contexts (PCs), four PCs are treated special: the entry to special PCs 0, 1, 2, and 3 is hardware-controlled. For each PC *i*, a programmable exception handler address is provided (CM0_PCi_HANDLER.ADDR[31:0]). A CPU exception handler fetch, which returns a handler address that matches the programmed CM0_PCi_HANDLER.ADDR[31:0] address value, causes the CM0+ PC to be changed to PC *i* by hardware. However, if the current PC is already 0, 1, 2, or 3, the current PC is not changed (an attempt to change the PC results in an AHB-Lite bus error). This ensures that CPU execution in PC 0, 1, 2, or 3 cannot be interrupted/preempted by CPU execution in another PC 0, 1, 2, or 3. In other words, CPU execution in PC 0, 1, 2, or 3 requires cooperative multi-tasking between the different PCs. This means that hand-overs between different PCs are software-scheduled/controlled. A security implementation requires PC software to clean/zeroize information that it wants to keep confidential from other PC software.

The following pseudo code gives a description of the hardware control over the CM0+ current PC “pc”:

```

if      (CM0_PC_CTL.VALID[0] & (addr == CM0_PC0_HANDLER.ADDR)) {match_new, pc_new} = {1, 0};
else if (CM0_PC_CTL.VALID[1] & (addr == CM0_PC1_HANDLER.ADDR)) {match_new, pc_new} = {1, 1};
else if (CM0_PC_CTL.VALID[2] & (addr == CM0_PC2_HANDLER.ADDR)) {match_new, pc_new} = {1, 2};
else if (CM0_PC_CTL.VALID[3] & (addr == CM0_PC3_HANDLER.ADDR)) {match_new, pc_new} = {1, 3};
else
    match_new = 0;

if ("exception handler fetch")
begin
    if (match_new)
    begin
        if      (pc == pc_new)      ;
        else if (CM0_PC_CTL.VALID[pc]) "AHB-Lite bus error";
        else
            {pc, pc_saved} = {pc_new, pc};    // "push"
        end
    else
    begin
        {pc, pc_saved} = {pc_saved, pc_saved};    // "pop"
    end
end
end
  
```

Note that each of the protection special PCs 0, 1, 2, and 3 have a dedicated CM0_PC_CTL.VALID[*i*] field to specify that the PC's exception handler address is provided through CM0_PCi_HANDLER.ADDR[31:0]. If a PC's exception handler address is not provided, the PC is treated as an ordinary PC (PCs 4, 5, ..., 15). Note that the current PC “pc” and a saved PC “pc_saved” implement a two-entry stack. Hardware pushes the current PC to the stack upon entry of a special exception handler. Hardware then pushes the saved PC from the stack upon entry of an ordinary exception handler. An attempt to enter a special exception handler from a special exception handler with a different PC results in an AHB-Lite bus error (which causes the CPU to enter the bus fault exception handler). This scenario should not occur in a carefully designed cooperative multi-tasking software implementation.

Of the four special PCs, PC 0 is treated differently:

- It is the default PC value after a Deep Sleep reset.
- It has unrestricted access; that is, it is not affected by the CPUSS and PERI protection schemes. Therefore, the Cypress boot code software always starts execution in PC 0. The boot code software initializes the protection structures and initializes the CM0_PCi_HANDLER registers. The boot code is considered “trusted” software and its unrestricted access

in PC 0 is not considered a protection concern. After initialization of the protection information, the access to the protection information itself is typically restricted for all other PCs (the boot code software deploys the restrictions) and the protection information provides specific restricted access to the other special PCs and ordinary PCs (1, 2, ..., 15). As part of the boot process configuration, the following things happen:

- ❑ The boot process sets MPUn.MS_CTL.PC_SAVED to a context that is not a special context. This is required to prevent accidental/malicious incorrect entry into a special context.
- ❑ If the system configuration does not use PC=0 after boot is complete, the boot process disables all future use of PC=0 by making ineffective CPUSS.CM0_PC0_HANDLER by clearing CPUSS.CM0_PC_CTL.VALID[0]. Furthermore, the protection region CPUSS.BOOT is assigned to be accessible from PC=0 only, which ensures these settings can no longer be changed after boot.
- ❑ The boot process then changes the current PC by changing MPUn.MS_CTL.PC to a protection context that is not a special context. This is done for all currently active masters (which typically is only the CM0+).
- ❑ When the above is complete, the system executes a regular (not special) context and can enter into special contexts only through the use of interrupts using the special handler addresses.

9.7 Protection Structure

9.7.1 Protection Violation

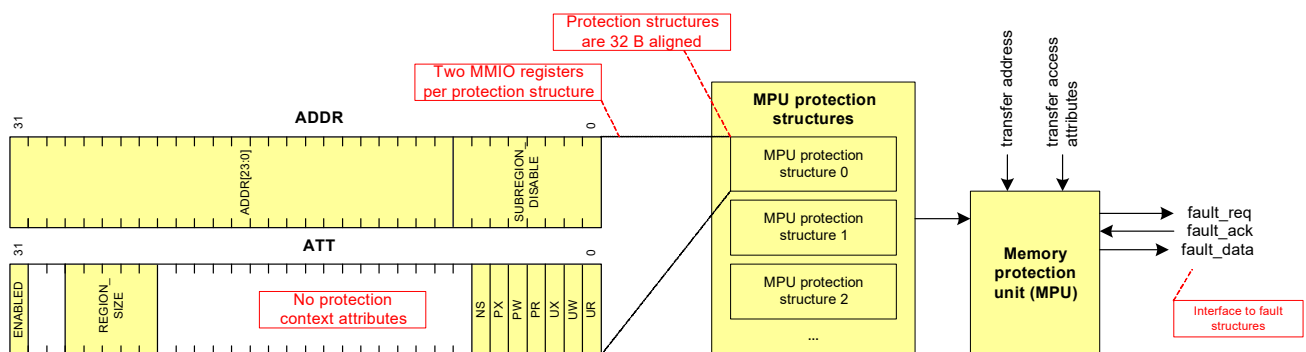
If an MPU, S MPU, or PPU detects a not-allowed transfer, the bus transfer results in a bus error. The bus transfer does not reach its target memory location or peripheral register. In addition, information on the violating bus transfer is communicated to the fault report structure.

9.7.2 MPU

The MPUs are situated in the CPUSS and are associated to a single master. An MPU distinguishes user and privileged accesses from a single bus master. However, the capability exists to perform access control on the secure/non-secure attribute.

As an MPU is associated to a single master, the MPU protection structures do not provide protection context control attributes.

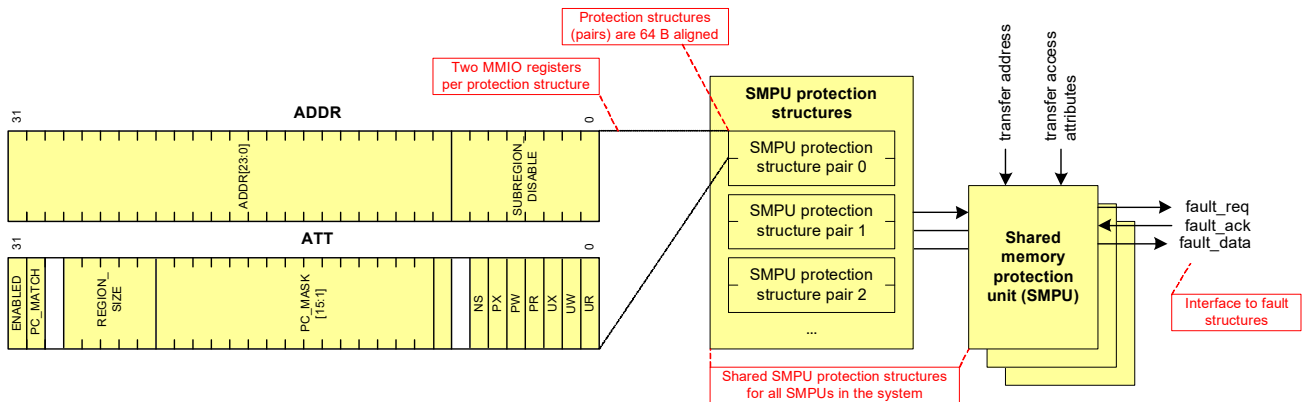
Figure 9-3. MPU Functionality



9.7.3 SMPU

The SMPU is situated in the CPUSS and is shared by all bus masters. The SMPU distinguishes between different protection contexts and distinguishes secure from non-secure accesses. However, the capability exists to perform access control on the user/privileged mode attribute. Although all SMPUs are predefined in the PSoC 64 family, the definition of how they word and are programmed are provided to provide a better understanding of the system security.

Figure 9-4. SMPU Functionality




Note that a single set of SMPU region structures provides the same protection information to all SMPUs in the systems.

9.7.4 PPU

The PPU are situated in the PERI block and are associated with a peripheral group (a group of peripherals with a shared AHB-Lite bus infrastructure). A PPU is shared by all bus masters. The PPU distinguishes between different protection contexts; it also distinguishes secure from non-secure accesses and user mode from privileged mode accesses. Some PPUs have been preallocated to protect secure sections of the device. The PPUs pre-allocated include:

- Access to Flash programming control
- Access to eFuse programming
- Access to debug port configuration

Figure 9-5. PPU Functionality



There are two types of PPU structures: fixed and programmable.

- The fixed PPU structures protect fixed areas of memory and hence a specific predetermined peripheral region. In other words, the ADDR, SUBREGION_DISABLE, and REGION_SIZE fields are fixed for a specific device. Refer to the [registers TRM](#) for a definition of fixed PPUs and the address regions they protect. Their protection attributes can be configured for each protection context. The following table summarizes the different fixed PPU structures in CY8C62xx devices and details the regions they protect.

Table 9-3. Fixed PPU Structures

Protection Structure	Base Address of Protected Region	Size of Protected Region	Protected Peripheral/Block(s)
PERI_MS_PPU_FX0	0x40000000	8 KB	PERI
PERI_MS_PPU_FX1	0x40004010	4 B	PERI_GR0
PERI_MS_PPU_FX2	0x40004030	4 B	PERI_GR1
PERI_MS_PPU_FX3	0x40004050	4 B	PERI_GR2
PERI_MS_PPU_FX4	0x40004060	32 B	PERI_GR3
PERI_MS_PPU_FX5	0x40004080	32 B	PERI_GR4
PERI_MS_PPU_FX6	0x400040C0	32 B	PERI_GR6
PERI_MS_PPU_FX7	0x40004120	32 B	PERI_GR9
PERI_MS_PPU_FX8	0x40004140	32 B	PERI_GR10
PERI_MS_PPU_FX9	0x40008000	32 KB	PERI_TR
PERI_MS_PPU_FX10	0x40100000	1 KB	CRYPTO
PERI_MS_PPU_FX11	0x40101000	2 KB	CRYPTO
PERI_MS_PPU_FX12	0x40102000	256 B	CRYPTO
PERI_MS_PPU_FX13	0x40102100	4 B	CRYPTO
PERI_MS_PPU_FX14	0x40102120	4 B	CRYPTO

Table 9-3. Fixed PPU Structures

Protection Structure	Base Address of Protected Region	Size of Protected Region	Protected Peripheral/Block(s)
PERI_MS_PPU_FX15	0x40108000	4 KB	CRYPTO
PERI_MS_PPU_FX16	0x40200000	1 KB	CPUSS
PERI_MS_PPU_FX17	0x40201000	4 KB	CPUSS
PERI_MS_PPU_FX18	0x40202000	512 B	CPUSS
PERI_MS_PPU_FX19	0x40208000	1 KB	CPUSS_CM0_SYSTEM_INT
PERI_MS_PPU_FX20	0x4020A000	1 KB	CPUSS_CM4_SYSTEM_INT
PERI_MS_PPU_FX21	0x40210000	256 B	FAULT_STRUCT0
PERI_MS_PPU_FX22	0x40210100	256 B	FAULT_STRUCT1
PERI_MS_PPU_FX23	0x40220000	32 B	IPC_STRUCT0
PERI_MS_PPU_FX24	0x40220020	32 B	IPC_STRUCT1
PERI_MS_PPU_FX25	0x40220040	32 B	IPC_STRUCT2
PERI_MS_PPU_FX26	0x40220060	32 B	IPC_STRUCT3
PERI_MS_PPU_FX27	0x40220080	32 B	IPC_STRUCT4
PERI_MS_PPU_FX28	0x402200A0	32 B	IPC_STRUCT5
PERI_MS_PPU_FX29	0x402200C0	32 B	IPC_STRUCT6
PERI_MS_PPU_FX30	0x402200E0	32 B	IPC_STRUCT7
PERI_MS_PPU_FX31	0x40220100	32 B	IPC_STRUCT8
PERI_MS_PPU_FX32	0x40220120	32 B	IPC_STRUCT9
PERI_MS_PPU_FX33	0x40220140	32 B	IPC_STRUCT10
PERI_MS_PPU_FX34	0x40220160	32 B	IPC_STRUCT11
PERI_MS_PPU_FX35	0x40220180	32 B	IPC_STRUCT12
PERI_MS_PPU_FX36	0x402201A0	32 B	IPC_STRUCT13
PERI_MS_PPU_FX37	0x402201C0	32 B	IPC_STRUCT14
PERI_MS_PPU_FX38	0x402201E0	32 B	IPC_STRUCT15
PERI_MS_PPU_FX39	0x40221000	16 B	IPC_INTR_STRUCT0
PERI_MS_PPU_FX40	0x40221020	16 B	IPC_INTR_STRUCT1
PERI_MS_PPU_FX41	0x40221040	16 B	IPC_INTR_STRUCT2
PERI_MS_PPU_FX42	0x40221060	16 B	IPC_INTR_STRUCT3
PERI_MS_PPU_FX43	0x40221080	16 B	IPC_INTR_STRUCT4
PERI_MS_PPU_FX44	0x402210A0	16 B	IPC_INTR_STRUCT5
PERI_MS_PPU_FX45	0x402210C0	16 B	IPC_INTR_STRUCT6
PERI_MS_PPU_FX46	0x402210E0	16 B	IPC_INTR_STRUCT7
PERI_MS_PPU_FX47	0x40221100	16 B	IPC_INTR_STRUCT8
PERI_MS_PPU_FX48	0x40221120	16 B	IPC_INTR_STRUCT9
PERI_MS_PPU_FX49	0x40221140	16 B	IPC_INTR_STRUCT10
PERI_MS_PPU_FX50	0x40221160	16 B	IPC_INTR_STRUCT11
PERI_MS_PPU_FX51	0x40221180	16 B	IPC_INTR_STRUCT12
PERI_MS_PPU_FX52	0x402211A0	16 B	IPC_INTR_STRUCT13
PERI_MS_PPU_FX53	0x402211C0	16 B	IPC_INTR_STRUCT14
PERI_MS_PPU_FX54	0x402211E0	16 B	IPC_INTR_STRUCT15
PERI_MS_PPU_FX55	0x40230000	64 B	PROT_SMPU
PERI_MS_PPU_FX56	0x40234000	4 B	PROT_MPU0
PERI_MS_PPU_FX57	0x40235400	1 KB	PROT_MPU5
PERI_MS_PPU_FX58	0x40235800	1 KB	PROT_MPU6
PERI_MS_PPU_FX59	0x40237800	4 B	PROT_MPU14

Table 9-3. Fixed PPU Structures

Protection Structure	Base Address of Protected Region	Size of Protected Region	Protected Peripheral/Block(s)
PERI_MS_PPU_FX60	0x40237C00	1 KB	PROT_MPU15
PERI_MS_PPU_FX61	0x40240000	8 B	FLASHC_FLASH
PERI_MS_PPU_FX62	0x40240008	4 B	FLASHC_FLASH
PERI_MS_PPU_FX63	0x40240200	256 B	FLASHC
PERI_MS_PPU_FX64	0x40240400	128 B	FLASHC_CM0
PERI_MS_PPU_FX65	0x40240480	128 B	FLASHC_CM4
PERI_MS_PPU_FX66	0x40240500	4 B	FLASHC_CRYPT0
PERI_MS_PPU_FX67	0x40240580	4 B	FLASHC_DW0
PERI_MS_PPU_FX68	0x40240600	4 B	FLASHC_DW1
PERI_MS_PPU_FX69	0x40240680	4 B	FLASHC_DMAC
PERI_MS_PPU_FX70	0x40240700	4 B	FLASHC_EXT_MS0
PERI_MS_PPU_FX71	0x40240780	4 B	FLASHC_EXT_MS1
PERI_MS_PPU_FX72	0x4024F000	4 KB	FLASHC_FM
PERI_MS_PPU_FX73	0x40260000	256 B	PWR
PERI_MS_PPU_FX74	0x40260100	16 B	
PERI_MS_PPU_FX75	0x40260180	16 B	WDT
PERI_MS_PPU_FX76	0x40260200	128 B	MCWDT
PERI_MS_PPU_FX77	0x40260300	256 B	CLK
PERI_MS_PPU_FX78	0x40260400	1 KB	CLK
PERI_MS_PPU_FX79	0x40260800	8 B	RES
PERI_MS_PPU_FX80	0x40267000	4 KB	PWR_TRIM, CLK_TRIM
PERI_MS_PPU_FX81	0x4026FF00	128 B	PWR_TRIM, CLK_TRIM
PERI_MS_PPU_FX82	0x40270000	64 KB	BACKUP
PERI_MS_PPU_FX83	0x40280000	128 B	DW00
PERI_MS_PPU_FX84	0x40290000	128 B	DW10
PERI_MS_PPU_FX85	0x40280100	128 B	DW0_CRC
PERI_MS_PPU_FX86	0x40290100	128 B	DW1_CRC
PERI_MS_PPU_FX87	0x40288000	64 B	DW0_CH_STRUCT0
PERI_MS_PPU_FX88	0x40288040	64 B	DW0_CH_STRUCT1
PERI_MS_PPU_FX89	0x40288080	64 B	DW0_CH_STRUCT2
PERI_MS_PPU_FX90	0x402880C0	64 B	DW0_CH_STRUCT3
PERI_MS_PPU_FX91	0x40288100	64 B	DW0_CH_STRUCT4
PERI_MS_PPU_FX92	0x40288140	64 B	DW0_CH_STRUCT5
PERI_MS_PPU_FX93	0x40288180	64 B	DW0_CH_STRUCT6
PERI_MS_PPU_FX94	0x402881C0	64 B	DW0_CH_STRUCT7
PERI_MS_PPU_FX95	0x40288200	64 B	DW0_CH_STRUCT8
PERI_MS_PPU_FX96	0x40288240	64 B	DW0_CH_STRUCT9
PERI_MS_PPU_FX97	0x40288280	64 B	DW0_CH_STRUCT10
PERI_MS_PPU_FX98	0x402882C0	64 B	DW0_CH_STRUCT11
PERI_MS_PPU_FX99	0x40288300	64 B	DW0_CH_STRUCT12
PERI_MS_PPU_FX100	0x40288340	64 B	DW0_CH_STRUCT13
PERI_MS_PPU_FX101	0x40288380	64 B	DW0_CH_STRUCT14
PERI_MS_PPU_FX102	0x402883C0	64 B	DW0_CH_STRUCT15
PERI_MS_PPU_FX103	0x40288400	64 B	DW0_CH_STRUCT16
PERI_MS_PPU_FX104	0x40288440	64 B	DW0_CH_STRUCT17

Table 9-3. Fixed PPU Structures

Protection Structure	Base Address of Protected Region	Size of Protected Region	Protected Peripheral/Block(s)
PERI_MS_PPU_FX105	0x40288480	64 B	DW0_CH_STRUCT18
PERI_MS_PPU_FX106	0x402884C0	64 B	DW0_CH_STRUCT19
PERI_MS_PPU_FX107	0x40288500	64 B	DW0_CH_STRUCT20
PERI_MS_PPU_FX108	0x40288540	64 B	DW0_CH_STRUCT21
PERI_MS_PPU_FX109	0x40288580	64 B	DW0_CH_STRUCT22
PERI_MS_PPU_FX110	0x402885C0	64 B	DW0_CH_STRUCT23
PERI_MS_PPU_FX111	0x40288600	64 B	DW0_CH_STRUCT24
PERI_MS_PPU_FX112	0x40288640	64 B	DW0_CH_STRUCT25
PERI_MS_PPU_FX113	0x40288680	64 B	DW0_CH_STRUCT26
PERI_MS_PPU_FX114	0x402886C0	64 B	DW0_CH_STRUCT27
PERI_MS_PPU_FX115	0x40288700	64 B	DW0_CH_STRUCT28
PERI_MS_PPU_FX116	0x40298000	64 B	DW1_CH_STRUCT0
PERI_MS_PPU_FX117	0x40298040	64 B	DW1_CH_STRUCT1
PERI_MS_PPU_FX118	0x40298080	64 B	DW1_CH_STRUCT2
PERI_MS_PPU_FX119	0x402980C0	64 B	DW1_CH_STRUCT3
PERI_MS_PPU_FX120	0x40298100	64 B	DW1_CH_STRUCT4
PERI_MS_PPU_FX121	0x40298140	64 B	DW1_CH_STRUCT5
PERI_MS_PPU_FX122	0x40298180	64 B	DW1_CH_STRUCT6
PERI_MS_PPU_FX123	0x402981C0	64 B	DW1_CH_STRUCT7
PERI_MS_PPU_FX124	0x40298200	64 B	DW1_CH_STRUCT8
PERI_MS_PPU_FX125	0x40298240	64 B	DW1_CH_STRUCT9
PERI_MS_PPU_FX126	0x40298280	64 B	DW1_CH_STRUCT10
PERI_MS_PPU_FX127	0x402982C0	64 B	DW1_CH_STRUCT11
PERI_MS_PPU_FX128	0x40298300	64 B	DW1_CH_STRUCT12
PERI_MS_PPU_FX129	0x40298340	64 B	DW1_CH_STRUCT13
PERI_MS_PPU_FX130	0x40298380	64 B	DW1_CH_STRUCT14
PERI_MS_PPU_FX131	0x402983C0	64 B	DW1_CH_STRUCT15
PERI_MS_PPU_FX132	0x40298400	64 B	DW1_CH_STRUCT16
PERI_MS_PPU_FX133	0x40298440	64 B	DW1_CH_STRUCT17
PERI_MS_PPU_FX134	0x40298480	64 B	DW1_CH_STRUCT18
PERI_MS_PPU_FX135	0x402984C0	64 B	DW1_CH_STRUCT19
PERI_MS_PPU_FX136	0x40298500	64 B	DW1_CH_STRUCT20
PERI_MS_PPU_FX137	0x40298540	64 B	DW1_CH_STRUCT21
PERI_MS_PPU_FX138	0x40298580	64 B	DW1_CH_STRUCT22
PERI_MS_PPU_FX139	0x402985C0	64 B	DW1_CH_STRUCT23
PERI_MS_PPU_FX140	0x40298600	64 B	DW1_CH_STRUCT24
PERI_MS_PPU_FX141	0x40298640	64 B	DW1_CH_STRUCT25
PERI_MS_PPU_FX142	0x40298680	64 B	DW1_CH_STRUCT26
PERI_MS_PPU_FX143	0x402986C0	64 B	DW1_CH_STRUCT27
PERI_MS_PPU_FX144	0x40298700	64 B	DW1_CH_STRUCT28
PERI_MS_PPU_FX145	0x402A0000	16 B	DMAC
PERI_MS_PPU_FX146	0x402A1000	256 B	DMAC_CH0
PERI_MS_PPU_FX147	0x402A1100	256 B	DMAC_CH1
PERI_MS_PPU_FX148	0x402A1200	256 B	DMAC_CH2
PERI_MS_PPU_FX149	0x402A1300	256 B	DMAC_CH3

Table 9-3. Fixed PPU Structures

Protection Structure	Base Address of Protected Region	Size of Protected Region	Protected Peripheral/Block(s)
PERI_MS_PPU_FX150	0x402C0000	128 B	EFUSE
PERI_MS_PPU_FX151	0x402C0800	512 B	EFUSE
PERI_MS_PPU_FX152	0x402D0000	64 KB	PROFILE
PERI_MS_PPU_FX153	0x40300000	8 B	HSIOM_PRT0_PORT
PERI_MS_PPU_FX154	0x40300010	8 B	HSIOM_PRT1_PORT
PERI_MS_PPU_FX155	0x40300020	8 B	HSIOM_PRT2_PORT
PERI_MS_PPU_FX156	0x40300030	8 B	HSIOM_PRT3_PORT
PERI_MS_PPU_FX157	0x40300040	8 B	HSIOM_PRT4_PORT
PERI_MS_PPU_FX158	0x40300050	8 B	HSIOM_PRT5_PORT
PERI_MS_PPU_FX159	0x40300060	8 B	HSIOM_PRT6_PORT
PERI_MS_PPU_FX160	0x40300070	8 B	HSIOM_PRT7_PORT
PERI_MS_PPU_FX161	0x40300080	8 B	HSIOM_PRT8_PORT
PERI_MS_PPU_FX162	0x40300090	8 B	HSIOM_PRT9_PORT
PERI_MS_PPU_FX163	0x403000A0	8 B	HSIOM_PRT10_PORT
PERI_MS_PPU_FX164	0x403000B0	8 B	HSIOM_PRT11_PORT
PERI_MS_PPU_FX165	0x403000C0	8 B	HSIOM_PRT12_PORT
PERI_MS_PPU_FX166	0x403000D0	8 B	HSIOM_PRT13_PORT
PERI_MS_PPU_FX167	0x403000E0	8 B	HSIOM_PRT14_PORT
PERI_MS_PPU_FX168	0x40302000	32 B	HSIOM_AMUX_SPLIT
PERI_MS_PPU_FX169	0x40302200	16 B	
PERI_MS_PPU_FX170	0x40310000	64 B	GPIO_PRT0
PERI_MS_PPU_FX171	0x40310080	64 B	GPIO_PRT1
PERI_MS_PPU_FX172	0x40310100	64 B	GPIO_PRT2
PERI_MS_PPU_FX173	0x40310180	64 B	GPIO_PRT3
PERI_MS_PPU_FX174	0x40310200	64 B	GPIO_PRT4
PERI_MS_PPU_FX175	0x40310280	64 B	GPIO_PRT5
PERI_MS_PPU_FX176	0x40310300	64 B	GPIO_PRT6
PERI_MS_PPU_FX177	0x40310380	64 B	GPIO_PRT7
PERI_MS_PPU_FX178	0x40310400	64 B	GPIO_PRT8
PERI_MS_PPU_FX179	0x40310480	64 B	GPIO_PRT9
PERI_MS_PPU_FX180	0x40310500	64 B	GPIO_PRT10
PERI_MS_PPU_FX181	0x40310580	64 B	GPIO_PRT11
PERI_MS_PPU_FX182	0x40310600	64 B	GPIO_PRT12
PERI_MS_PPU_FX183	0x40310680	64 B	GPIO_PRT13
PERI_MS_PPU_FX184	0x40310700	64 B	GPIO_PRT14
PERI_MS_PPU_FX185	0x40310040	16 B	GPIO_PRT0
PERI_MS_PPU_FX186	0x403100C0	16 B	GPIO_PRT1
PERI_MS_PPU_FX187	0x40310140	16 B	GPIO_PRT2
PERI_MS_PPU_FX188	0x403101C0	16 B	GPIO_PRT3
PERI_MS_PPU_FX189	0x40310240	16 B	GPIO_PRT4
PERI_MS_PPU_FX190	0x403102C0	16 B	GPIO_PRT5
PERI_MS_PPU_FX191	0x40310340	16 B	GPIO_PRT6
PERI_MS_PPU_FX192	0x403103C0	16 B	GPIO_PRT7
PERI_MS_PPU_FX193	0x40310440	16 B	GPIO_PRT8
PERI_MS_PPU_FX194	0x403104C0	16 B	GPIO_PRT9

Table 9-3. Fixed PPU Structures

Protection Structure	Base Address of Protected Region	Size of Protected Region	Protected Peripheral/Block(s)
PERI_MS_PPU_FX195	0x40310540	16 B	GPIO_PRT10
PERI_MS_PPU_FX196	0x403105C0	16 B	GPIO_PRT11
PERI_MS_PPU_FX197	0x40310640	16 B	GPIO_PRT12
PERI_MS_PPU_FX198	0x403106C0	16 B	GPIO_PRT13
PERI_MS_PPU_FX199	0x40310740	8 B	GPIO_PRT14
PERI_MS_PPU_FX200	0x40314000	64 B	GPIO_INTR
PERI_MS_PPU_FX201	0x40315000	8 B	
PERI_MS_PPU_FX202	0x40320800	256 B	SMARTIO_PRT8
PERI_MS_PPU_FX203	0x40320900	256 B	SMARTIO_PRT9
PERI_MS_PPU_FX204	0x40350000	64 KB	LPCOMP
PERI_MS_PPU_FX205	0x40360000	4 KB	CSD
PERI_MS_PPU_FX206	0x40380000	64 KB	TCPWM0
PERI_MS_PPU_FX207	0x40390000	64 KB	TCPWM1
PERI_MS_PPU_FX208	0x403B0000	64 KB	LCD0
PERI_MS_PPU_FX209	0x403F0000	64 KB	USBFS0
PERI_MS_PPU_FX210	0x40420000	64 KB	SMIF0
PERI_MS_PPU_FX211	0x40460000	64 KB	SDHC0
PERI_MS_PPU_FX212	0x40470000	64 KB	SDHC1
PERI_MS_PPU_FX213	0x40600000	64 KB	SCB0
PERI_MS_PPU_FX214	0x40610000	64 KB	SCB1
PERI_MS_PPU_FX215	0x40620000	64 KB	SCB2
PERI_MS_PPU_FX216	0x40630000	64 KB	SCB3
PERI_MS_PPU_FX217	0x40640000	64 KB	SCB4
PERI_MS_PPU_FX218	0x40650000	64 KB	SCB5
PERI_MS_PPU_FX219	0x40660000	64 KB	SCB6
PERI_MS_PPU_FX220	0x40670000	64 KB	SCB7
PERI_MS_PPU_FX221	0x40680000	64 KB	SCB8
PERI_MS_PPU_FX222	0x40690000	64 KB	SCB9
PERI_MS_PPU_FX223	0x406A0000	64 KB	SCB10
PERI_MS_PPU_FX224	0x406B0000	64 KB	SCB11
PERI_MS_PPU_FX225	0x406C0000	64 KB	SCB12
PERI_MS_PPU_FX226	0x40A00000	4 KB	PDM0
PERI_MS_PPU_FX227	0x40A10000	4 KB	I2S0
PERI_MS_PPU_FX228	0x40A11000	4 KB	I2S1

- The programmable PPU structures can have configurable address regions. Similar to fixed-PPU structures, the protection attributes of programmable PPU structures can be configured for each protection context. Programmable PPU structures are similar to SMPU structures but are intended to be used with the peripheral register space. These protection structures are typically used to protect registers in a specific block, which are not covered by the resolution of fixed PPU structures.

9.7.5 Protection of Protection Structures

The MPU, SMPU, and PPU-based protection architecture is consistent and provides the flexibility to implement different system-wide protection schemes. Protection structures can be set once at boot time or can be changed dynamically during device execution. For example, a CPU RTOS can change the CPU's MPU settings; a secure CPU can change the SMPU and PPU's settings. But such a system will be left insecure if there is no way to protect the protection structures themselves. There must be a way to restrict access to the protection structures.

The protection of protection structures is achieved using another protection structure. For this reason, protection structures are defined in pairs of master and slave. We refer to the slave and master protection structures as a protection pair. Note that the address range of the master protection structure is known, that is, it is constant.

The first (slave) protection structure protects the resource and the second (master) protection structure protects the protection (address range of the second protection structure includes both the master and slave protection structures).

The protection architecture is flexible enough to allow for variations:

- Exclusive peripheral ownership can be shared by more than two protection contexts.
- The ability to change ownership is under control of a single protection context, and exclusive or non-exclusive peripheral ownership is shared by multiple protection contexts.

Note that in secure systems, typically a single secure CPU is used. In these systems, the ability to change ownership is assigned to the secure CPU at boot time and not dynamically changed. Therefore, you must assign the secure CPU its own, dedicated protection context.

Both PPU and SMPU is intended to distinguish between different protection contexts and to distinguish secure from non-secure accesses. Therefore, both PPU and SMPU protection use protection structure pairs. In the SMPU, the slave protection structure provides SMPU protection information and the master protection structure provides PPU protection information (the master and slave protection structures are registers).

Table 9-4. Protection Units Reserved by the System

PPU Type	PPU Define	Name	Description	Comment
Programmable	PERI_PPU_PR4	PERI.PERI_GROUP_STRUCT[]	Peripheral master enables	
Fixed	PERI_GR_PPU_SL_PERI_GR1	PERI.PERI_GROUP[1]	Peripheral interconnect	Not used
Fixed	PPU Defi+A1:D1	CPUSS-M4.CRYPTO	Cryptography component	
Fixed	PERI_GR_PPU_SL_PERI_GR2	PERI.PERI_GROUP[2]	Peripheral interconnect	Not used
Fixed	PERI_GR_PPU_SL_CPUSS	CPUSS-M4.CPUSS	CPU subsystem (CPUSS)	
Programmable	PERI_PPU_PR0	CPUSS-M4.CPUSS.*	CPU subsystem (CPUSS)	
Programmable	PERI_PPU_PR1	CPUSS-M4.CPUSS.TRIM*	CPU subsystem (CPUSS)	
Fixed	PERI_GR_PPU_SL_IPC	CPUSS-M4.IPC	IPC	
Region	PERI_GR_PPU_RG_IPC_STRUCT0	IPC_STRUCT0	IPC structure protection	CM0+ SysCall
Region	PERI_GR_PPU_RG_IPC_STRUCT1	IPC_STRUCT1	IPC structure protection	CM4 SysCall
Region	PERI_GR_PPU_RG_IPC_STRUCT2	IPC_STRUCT2	IPC structure protection	DAP SysCall
Region	PERI_GR_PPU_RG_IPC_STRUCT3	IPC_STRUCT3	IPC structure protection	CM4 -> CM0+ RPC
Region	PERI_GR_PPU_RG_IPC_STRUCT4	IPC_STRUCT4	IPC structure protection	DAP -> CM0+ RPC
Region	PERI_GR_PPU_RG_IPC_STRUCT5	IPC_STRUCT5	IPC structure protection	CM0+ -> CM4 RPC
Region	PERI_GR_PPU_RG_IPC_STRUCT6	IPC_STRUCT6	IPC structure protection	DAP -> CM4 RPC
Region	PERI_GR_PPU_RG_IPC_INTR_STRUCT0	IPC_INTR_STRUCT0	IPC interrupt structure protection	CM0+ NMI
Region	PERI_GR_PPU_RG_IPC_INTR_STRUCT1	IPC_INTR_STRUCT1	IPC structure protection	CM0+ RPC IRQ
Region	PERI_GR_PPU_RG_IPC_INTR_STRUCT2	IPC_INTR_STRUCT2	IPC structure protection	CM4 RPC IRQ
Fixed	PERI_GR_PPU_SL_PROT	CPUSS-M4.PROT	Protection	
Region	PERI_GR_PPU_RG_SMPU	CPUSS-M4.PROT.SMPU	SMPU protection	
Region	PERI_GR_PPU_RG_MPU_CM0P	CPUSS-M4.PROT.MPU_CM0P	MPU protection for CM0+	
Region	PERI_GR_PPU_RG_MPU_CRYPTO	CPUSS-M4.PROT.MPU_CRYPTO	MPU protection for Crypto	
Region	PERI_GR_PPU_RG_MPU_TC	CPUSS-M4.PROT.MPU_TC	MPU protection for test controller	
Fixed	PERI_GR_PPU_SL_FLASHC	CPUSS-M4.FLASHC	Flash controller	
Fixed	PERI_GR_PPU_SL_SRSS	SRSS.SRSS	SRSS core registers	

Table 9-4. Protection Units Reserved by the System

PPU Type	PPU Define	Name	Description	Comment
Fixed	PERI_GR_PPU_SL_BACKUP	SRSS.BACKUP	SRSS backup domain	
Fixed	PERI_GR_PPU_SL_DW1	CPUSS-M4.DW[1]	Datawire controller	
Region	PERI_GR_PPU_RG_DW1_DW_CH_STRU CT0	DW1_DW_CH_STRUCT0	DataWire #1 channel protection (channels 0..3)	Can be allocated to a single PC in user mode (1/2/3/4)
Region	PERI_GR_PPU_RG_DW1_DW_CH_STRU CT1	DW1_DW_CH_STRUCT1	DataWire #1 channel protection (channels 0..3)	Can be allocated to a single PC in user mode (1/2/3/4)
Region	PERI_GR_PPU_RG_DW1_DW_CH_STRU CT2	DW1_DW_CH_STRUCT2	DataWire #1 channel protection (channels 0..3)	Can be allocated to a single PC in user mode (1/2/3/4)
Region	PERI_GR_PPU_RG_DW1_DW_CH_STRU CT3	DW1_DW_CH_STRUCT3	DataWire #1 channel protection (channels 0..3)	Can be allocated to a single PC in user mode (1/2/3/4)
Fixed	PERI_GR_PPU_SL_EFUSE	EFUSE	eFuse block (only applies to eFuse data public region (00-40)	Only applies to eFuse public data region
Programmable	PERI_PPU_PR2	EFUSE.EFUSE_DATA[]	eFuse control registers	
Programmable	PERI_PPU_PR3	EFUSE.EFUSE_DATA[]	eFuse data region private	
Fixed	PERI_GR_PPU_SL_PERI_GR3	PERI.PERI_GROUP[3]	Peripheral interconnect	
Fixed	PERI_GR_PPU_SL_PERI_GR4	PERI.PERI_GROUP[4]	Peripheral interconnect	
Fixed	PERI_GR_PPU_SL_SMIF	SMIF	Serial memory interface	
Fixed	PERI_GR_PPU_SL_PERI_GR6	PERI.PERI_GROUP[6]	Peripheral interconnect	
Fixed	PERI_GR_PPU_SL_PERI_GR9	PERI.PERI_GROUP[9]	Peripheral interconnect	
Fixed	PERI_GR_PPU_SL_PERI_GR10	PERI.PERI_GROUP[10]	Peripheral interconnect	

9.7.6 Protection Structure Types

Different protection structure types are used because some resources, such as peripheral registers, have a fixed address range. Protection of protection structures requires pairs of neighboring protection structures.

Three types of protection structures with a consistent register interface are described here:

- Programmable protection structures. These are 32-byte protection structures with a programmable address range. These structures are used by the MPUs.
- Fixed protection structure pairs. These are 64-byte master/slave protection structure pairs, consisting of two 32-byte protection structures. These structures are used by the PPU. Both structures have a fixed, constant address region. The master structure has the UX and PX attributes as constant '0' (execution is never allowed) and the UR and PR attributes as constant '1' (reading is always allowed). The slave structure has the UX and PX attributes as constant '1'.
- Programmable protection structure pairs. These are 64-byte master/slave protection structure pairs, consisting of two 32-byte protection structures. These structures are used by the PPU and SMPU. The master structure has a fixed, constant address region. The slave structure

has a programmable address region. The master structure has the UX and PX attributes as constant '0' (execution is never allowed) and the UR and PR attributes as constant '1' (reading is always allowed). The PPU slave structure has the UX and PX attributes as constant '1'. The SMPU slave structure has programmable UX and PX attributes.

Note that the master protection structure in a protection structure pair is required only to address security requirements. The distinction between the three protection structure types is an implementation optimization. From an architectural perspective, all PPU protection structures are the same, with the exception that for some protection structures the address range is fixed and not programmable.

As mentioned earlier, a protection unit evaluates the protection regions in decreasing protection structure index order. The protection structures are evaluated in the following order:

- Fixed protection structures for specific peripherals or peripheral register address ranges.
- Programmable protection structures.

In other words, fixed structures take precedence over programmable structures.

Figure 9-6. PPU and SMPU Pairs

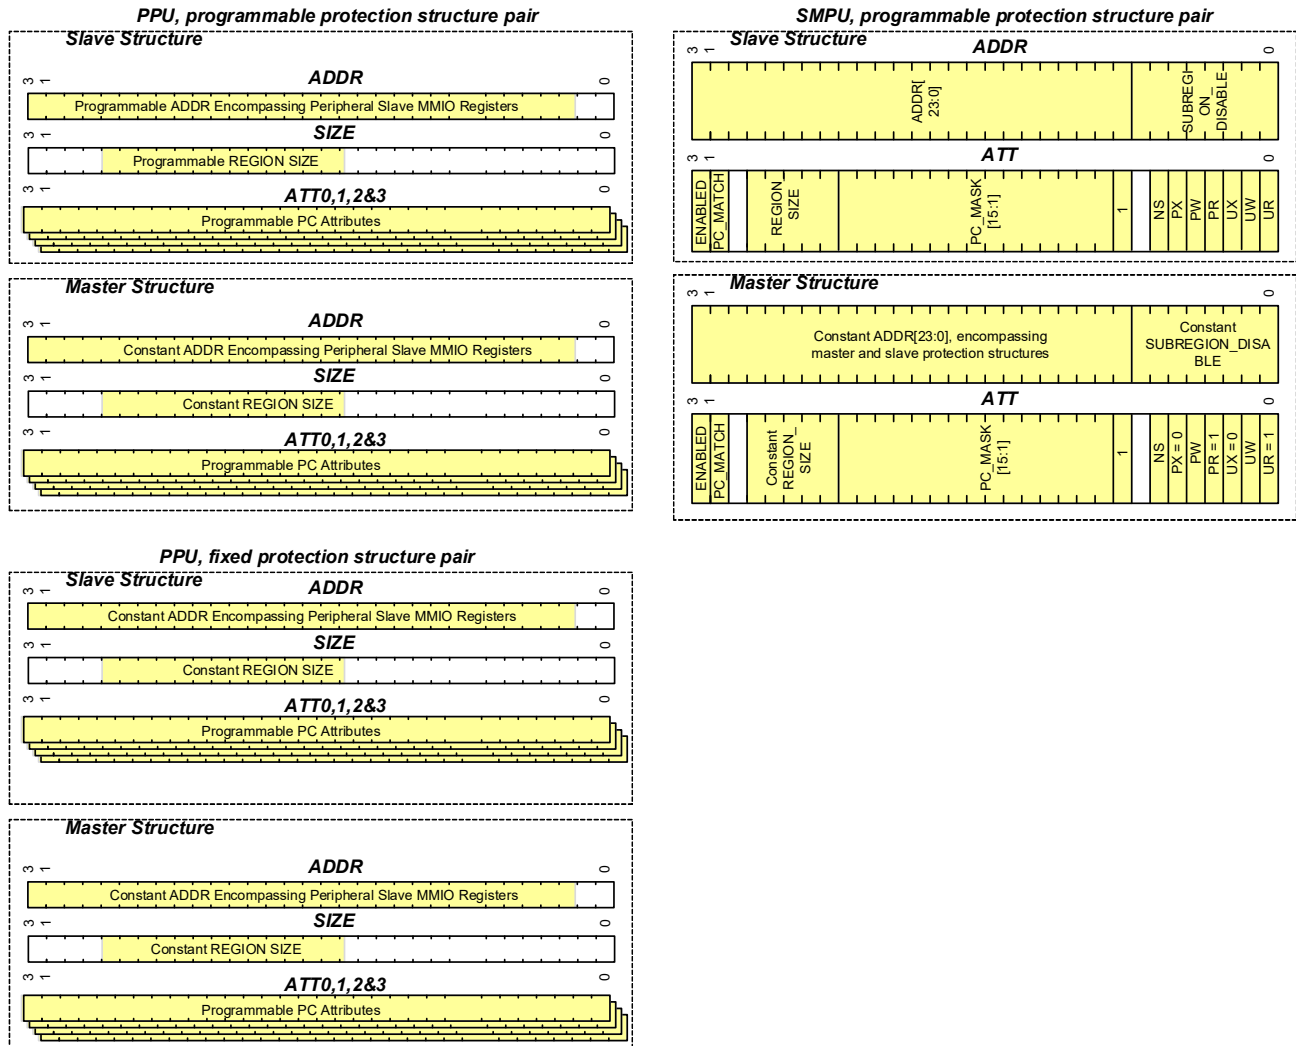


Figure 9-7. Fixed Protection Structure Pair

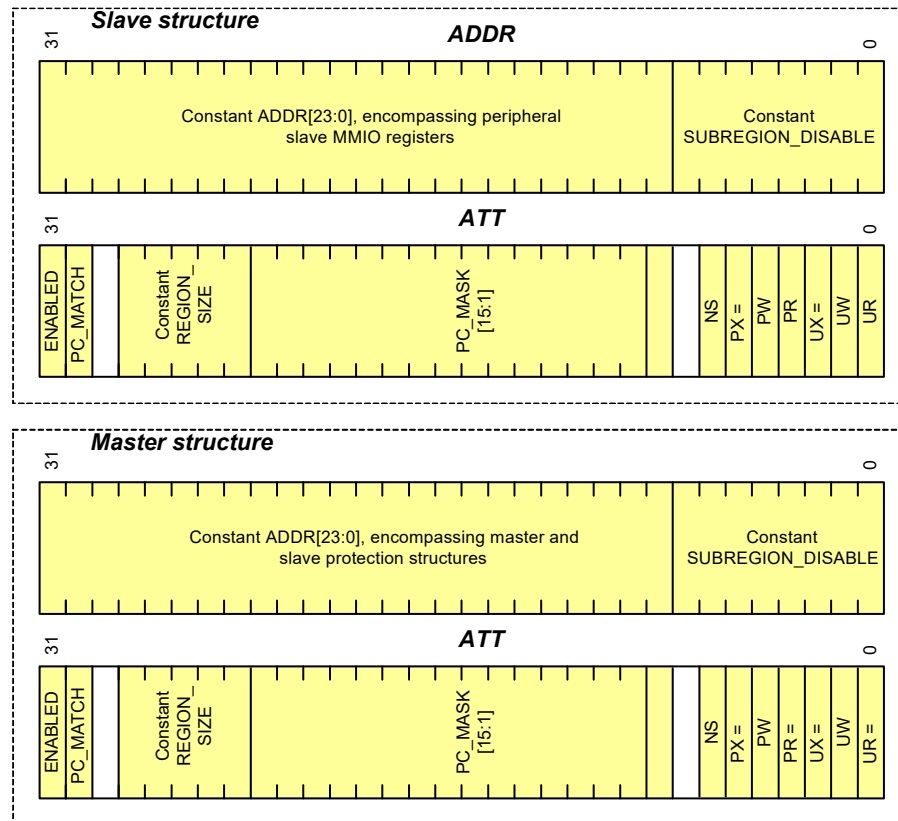
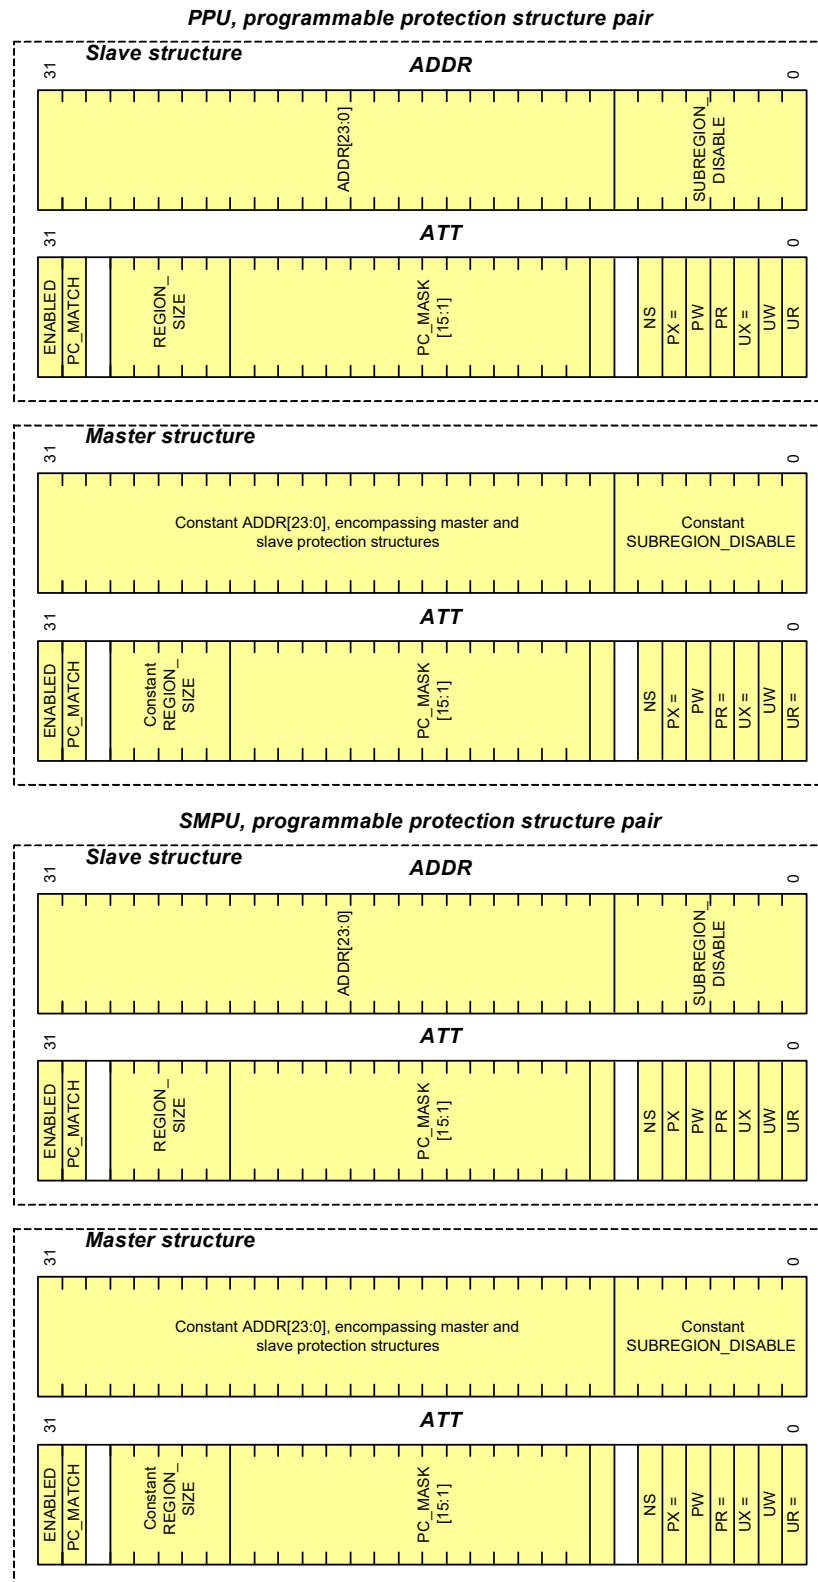
PPU, fixed protection structure pair


Figure 9-8. Programmable Protection Structure Pair



10. DMA Controller (DW)



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - DMA
- Application notes
- Code examples

The DMA transfers data to and from memory, peripherals, and registers. These transfers occur independent from the CPU. The DMA can be configured to perform multiple independent data transfers. All data transfers are managed by a channel. There can be up to 32 channels in the DMA. The number of channels in the DMA controller can vary with devices. Refer to the [device datasheet](#) for the number of channels supported in the device. A channel has an associated priority; channels are arbitrated according to their priority.


10.1 Features

The DMA controller has the following features:

- Supports up to 29 channels per DMA controller; see the [device datasheet](#) for details
- Supports multiple DMA controller instances in a device
- Four levels of priority for each channel
- Descriptors are defined in memory and referenced to the respective channels
- Supports single, 1D, or 2D transfer modes for a descriptor
- Supports transfer of up to 65536 data elements per descriptor
- Configurable source and destination address increments
- Supports 8-bit, 16-bit, and 32-bit data widths at both source and destination
- Configurable input trigger behavior for each descriptor
- Configurable interrupt generation in each descriptor
- Configurable output trigger generation for each descriptor
- Descriptors can be chained to other descriptors in memory

10.2 Architecture

Figure 10-1. DMA Controller



A data transfer is initiated by an input trigger. This trigger may originate from the source peripheral of the transfer, the destination peripheral of the transfer, CPU software, or from another peripheral. Triggers provide Active/Sleep functionality and are not available in Deep Sleep and Hibernate power modes.

The data transfer details are specified by a descriptor. Among other things, this descriptor specifies:

- The source and destination address locations and the size of the transfer.
- The actions of a channel; for example, generation of output triggers and interrupts. See the [Interrupts chapter on page 54](#) for more details.
- Data transfer types can be single, 1D, or 2D as defined in the descriptor structure. These types define the address sequences generated for source and destination. 1D and 2D transfers are used for “scatter gather” and other useful transfer operations.

connects to a specific system trigger through a trigger multiplexer that is outside the DMA controller.

Channel priority: A channel is assigned a priority (CHI_CTL.PRIO) between 0 and 3, with 0 being the highest priority and 3 being the lowest priority. Channels with the same priority constitute a priority group. Priority decoding determines the highest priority pending channel, which is determined as follows.

- The highest priority group with pending channels is identified first.
- Within this priority group, round-robin arbitration is applied.

Channel state: At any given time, one channel actively performs a data transfer. This channel is called the active channel. A channel can be in one of four channel states. The active channel in a DW controller can be determined by reading the DWx_STATUS[ACTIVE] and DWx.STATUS[CH_IDX].

Pending state of a channel is determined by reading the DW_CH_STRUCT_CH_STATUS[PENDING] associated with that channel. If a channel is enabled and is not in the Pending or Active state, then it is considered blocked.

10.3 Channels

The DMA controller supports multiple independent data transfers that are managed by a channel. Each channel

Table 10-1. Channel States

Channel State	Description
Disabled	The channel is disabled by setting CHI_CTL.ENABLED to '0'. The channel trigger is ignored in this state.
Blocked	The channel is enabled and is waiting for a trigger to initiate a data transfer.
Pending	The channel is enabled and has received an active trigger. In this state, the channel is ready to initiate a data transfer but waiting for it to be scheduled.
Active	The channel is enabled, has received an active trigger and has been scheduled. It is actively performing data transfers. If there are multiple channels pending, the highest priority pending channel is scheduled.

The data transfer associated with a trigger is made up of one or more 'atomic transfers' or 'single transfers'; see [Table 10-2](#) for a better understanding. A single trigger could be configured to transfer multiple "single transfers".

A channel can be marked preemptable (CHi_CTL.PREEMPTABLE). If preemptable, and there is a higher priority pending channel, then that higher priority channel can preempt the current channel between single transfers. If a channel is preempted, the existing single transfer is completed; the current channel goes to pending state and the higher priority channel is serviced. On completion of the higher priority channel's transfer, the pending channel is resumed. Note that preemption has an impact on the data transfer rates of the channel being preempted. Refer to ["DMA Performance" on page 102](#) for these performance implications.

A channel has two access control attributes that are used by the shared memory protection units (SMPUs) and peripheral protection units (PPUs) for access control. These fields are typically inherited from the master that modified the channel's control register.

- The Privileged Mode (CHi_CTL.P) attribute can be set to privileged or user.
- The Non-secure (CHi_CTL.NS) attribute can be set to secure or non-secure.

A descriptor associated with each channel describes the data transfer. The descriptor is stored in memory and CHi_CURR_PTR provides the descriptor address associated with channel "i" and Chi_IDX provides the current X and Y indices into the descriptor.

A channel's descriptor state is encoded as part of the channel's register state. The following registers provide a channel's descriptor state:

- CH_CTL. This register provides generic channel control information.
- CH_CURR_PTR. This register provides the address of the memory location where the current descriptor is located. The user firmware must initialize this register. If the descriptors are chained, the DMA hardware automatically sets this register to the next descriptor pointer.
- CH_IDX. This register provides the current X and Y indices of the channel into the current descriptor. User firmware must initialize this register. DMA hardware sets the X and Y indices to 0, when advancing from the current descriptor to the next descriptor in a descriptor list.

Note that channel state is retained in Deep Sleep power mode.

10.3.1 Channel Interrupts

Every DMA channel has an interrupt line associated with it. The INTR_TYPE parameter in the descriptor determines the

event that will trigger the interrupt for the channel. In addition each DMA channel has INTR, INTR_SET, INTR_MASK, and INTR_MASKED registers to control their respective interrupt lines. INTR_MASK can be used to mask the interrupt from the DMA channel. The INTR and INTR_SET can be used to clear and set the interrupt, respectively, for debug purposes.

The DW_CH_STRUCT_CH_STATUS[INTR_CAUSE] field provides the user a means to determine the cause of the interrupt being generated. The following are different values for this register:

- 0: No interrupt generated
- 1: Interrupt based on transfer completion configured based on INTR_TYPE field in the descriptor
- 2: Source bus error
- 3: Destination bus error
- 4: Misaligned source address
- 5: Misaligned destination address
- 6: Current descriptor pointer is null
- 7: Active channel is in disabled state
- 8: Descriptor bus error
- 9-15: Not used.




For error related interrupt causes (INTR_CAUSE is 2, 3,..., 8), the channel is disabled (hardware sets CH_CTL.ENABLED to '0').

The bus errors are typically caused by incompatible accesses to the addresses in question. This may be due to those addresses being protected or having read or write restrictions. Source and destination bus errors can also occur due to mismatch in data sizes (see ["Transfer Size" on page 99](#)).

10.4 Descriptors

The data transfer between a source and destination in a channel is configured using a descriptor. Descriptors are stored in memory. The descriptor pointer is specified in the DMA channel registers. The DMA controller does not modify the descriptor and treats it as read only. A descriptor is a set of up to six 32-bit registers that contain the configuration for the transfer in the associated channel. There are three types of descriptors.

Table 10-2. Descriptor Types

Descriptor Type	Description
Single transfer	Transfers a single data element 
1D transfer	Performs a one-dimensional “for loop”. This transfer is made up of X number of single transfers 
2D transfer	Performs a two-dimensional “for loop”. This transfer is made up of Y number of 1D transfers 
CRC transfer	This performs a one-dimensional “for loop” similar to the 1D transfer. However, the source data is not transferred to a destination. Instead, a CRC is calculated over the source data. The CRC configuration is provided through a set of registers that is shared by all DMA channels. The assumption is that the DMA channels use the CRC functionality mutually exclusively in time. These registers are: CRC_CTL, CRC_DATA_CTL, CRC_POL_CTL, CRC_LFSR_CTL, CRC_REM_CTL, and CRC_REM_RESULT. Note that the CRC configuration is the same as the Crypto CRC configuration.

Single Transfer:

The following pseudo code illustrates a single transfer.

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
DST_ADDR[0] = (t_DATA_SIZE) SRC_ADDR[0];
```

1D Transfer:

The following pseudo code illustrates a 1D transfer. Note that the 1D transfer is represented by a loop with each iteration executing a single transfer.

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
for (X_IDX = 0; X_IDX <= X_COUNT; X_IDX++) {
  DST_ADDR[X_IDX * DST_X_INCR] =
    (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR];
}
```

2D Transfer:

The following pseudo code illustrates a 2D transfer. Note that the 2D transfer is represented by a loop with each iteration executing an inner loop, which is the 1D transfer.

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
for (Y_IDX = 0; Y_IDX <= Y_COUNT; Y_IDX++) {
  for (X_IDX = 0; X_IDX <= X_COUNT; X_IDX++) {
    DST_ADDR[X_IDX * DST_X_INCR + Y_IDX * DST_Y_INCR] =
      (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR + Y_IDX * SRC_Y_INCR];
  }
}
```

CRC Transfer:

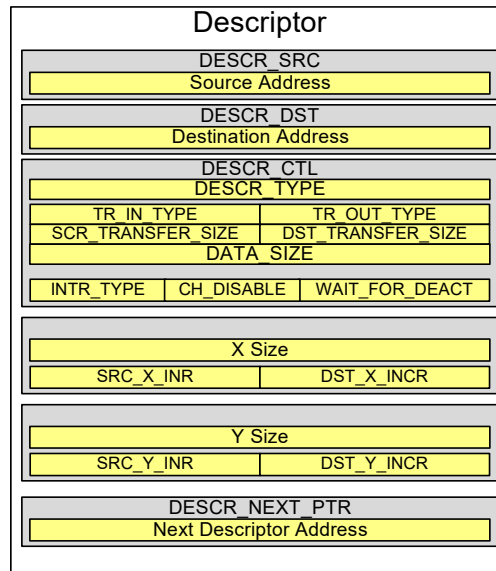
The following psuedo code illustrates CRC transfer.

```
// DST_ADDR is a pointer to an address location where the calculated CRC is stored.
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
CRC_STATE = CRC_LFSR_CTL;
for (X_IDX = 0; X_IDX <= X_COUNT; X_IDX++) {
  Update_CRC (CRC_STATE, (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR]);
}
DST_ADDR = CRC_STATE;
```

The parameters in the descriptor help configure the different aspects of the transfers explained.

Figure 10-2 shows the structure of a descriptor.

Figure 10-2. Descriptor Structure



10.4.1 Address Configuration

Source and Destination Address: The source and destination addresses are set in the respective registers in the descriptor. These set the base addresses for the source and destination location for the transfer. In case the descriptor is configured to transfer a single element, this field holds the source/destination address of the data element. If the descriptor is configured to transfer multiple elements with source address or destination address or both in an incremental mode, this field will hold the address of the first element that is transferred.

DESCR_TYPE: This field configures whether the descriptor has a single, 1D, or 2D type.

Trigger input type, TR_IN_TYPE: This field determines how the DMA engine responds to input trigger signal. This field can be configured for one of the following modes:

- Type 0: A trigger results in execution of a single transfer. Regardless of the DESCR_TYPE setting, a trigger input will trigger only a single element transfer. For example, in a 1D transfer, the DMA will transfer only one data element in every trigger.
- Type 1: A trigger results in the execution of a single 1D transfer. If the DESCR_TYPE was set to single transfer, the trigger signal will trigger the single transfer specified by the descriptor. For a DESCR_TYPE set to 1D transfer, the trigger signal will trigger the entire 1D transfer configured in the descriptor. For a 2D transfer, the trigger signal will trigger only a single iteration of the Y loop transfer.
- Type 2: A trigger results in execution of the current descriptor. Regardless of DESCR_TYPE, the trigger will execute the entire descriptor. If there was a next descriptor configured for the current descriptor, this

trigger setting will not automatically trigger the next descriptor.

- Type 3: A trigger results in execution of the current descriptor and also triggering the next descriptor. The execution of the next descriptor from this point will be determined by the TR_IN_TYPE setting of the next descriptor.

Trigger out type, TR_OUT_TYPE: This field determines what completion event will generate the output trigger signal. This field can be configured to one of the following modes:

- Type 0: Generates a trigger output for completion of every single element transfer.
- Type 1: Generates a trigger output for completion of a 1D transfer
- Type 2: Generates a trigger output for completion of the current descriptor. This trigger output is generated independent of the state of the DESCR_NEXT_PTR.
- Type 3: Generates a trigger output on completion of the current descriptor, when the current descriptor is the last descriptor in the descriptor chain. This means a trigger is generated when the descriptor execution is complete and the DESCR_NEXT_PTR is '0'.

Interrupt Type, INTR_TYPE: This field determines which completion event will generate the output interrupt signal. This field can be configured to one of the following modes:

- Type 0: Generates an interrupt output for completion of every single element transfer.
- Type 1: Generates an interrupt output for completion of a 1-D transfer.

- Type 2: Generates an interrupt output for completion of the current descriptor. This interrupt output is generated independent of the state of the `DESCR_NEXT_PTR`.
- Type 3: Generates an interrupt output on completion of the current descriptor, when the current descriptor is the last descriptor in the descriptor chain. This means an interrupt is generated when the descriptor execution is complete and the `DESCR_NEXT_PTR` is '0'.

WAIT_FOR_DEACT: When the DMA transfer based on the `TR_IN_TYPE` is completed, the data transfer engine checks the state of trigger deactivation. The data transfer on the second trigger is initiated only after deactivation of the first. The `WAIT_FOR_DEACT` parameter will determine when the trigger signal is considered deactivated. The first DMA transfer is activated when the trigger is activated, but the transfer is not considered complete until the trigger is deactivated. This field is used to synchronize the controller's data transfers with the agent that generated the trigger. This field has four settings:

- 0 – Pulse Trigger: Do not wait for deactivation. When a trigger is detected, the transfer is initiated. After completing the transfer, if the trigger is still active then it is considered as another trigger and the subsequent transfer is initiated immediately.
- 1 – Level-sensitive waits four slow clock cycles after the transfer to consider as a deactivation. When a trigger is detected, the transfer is initiated. After completing the transfer, if the trigger is still active then it is considered as another trigger after waiting for four cycles. Then, a subsequent transfer is initiated. The transfer corresponding to the trigger is considered complete only at the end of the four additional cycles. Even trigger output events will be affected based on this delay. This parameter adds a four-cycle delay in each trigger transaction and hence affects throughput.
- 2 – Level-sensitive waits 16 slow clock cycles after the transfer to consider as a deactivation. When a trigger is detected, the transfer is initiated. After completing the transfer, if the trigger is still active then it is considered as another trigger after waiting for 16 cycles. Then, a subsequent transfer is initiated. The transfer corresponding to the trigger is considered complete only at the end of the 16 additional cycles. Even trigger output events will be affected based on this delay. This parameter adds a 16-cycle delay in each trigger transaction and hence affects throughput.
- 3 – Pulse trigger waits indefinitely for deactivation. The DMA transfer is initiated after the trigger signal deactivates. The next transfer is initiated only if the trigger goes low and then high again. A trigger signal that remains active or does not transition to zero between two transaction will simply stall the DMA channel.

The `WAIT_FOR_DEACT` field is used in a system to cater to delayed response of other parts of the system to actions of the DMA. Consider an example of a TX FIFO

that has a trigger going to the DMA when its not full. Free space in FIFO will trigger a DMA transfer to the FIFO, which in turn will deactivate the trigger. However, there can be a delay in this deactivation by the agent, which may cause the DMA to have initiated another transfer that can cause a FIFO overflow. This can be avoided by using the four or 16 clock cycle delays.

X Count: This field determines the number of single element transfers present in the X loop (inner loop). This field is valid when the `DESCR_TYPE` is set to 1D or 2D transfer.

Source Address Increment (X loop) (`SCR_X_INCR`): This field configures the index by which the source address is to be incremented for every iteration in an X loop. The field is expressed in multiples of `SRC_TRANSFER_SIZE`. This field is a signed number and hence may be decrementing or incrementing. If the source address does not need to be incremented, you can set this parameter to zero.

Destination Address Increment (X loop) (`DST_X_INCR`): This field configures the index by which the destination address is to be incremented, for every iteration in an X loop. The field is expressed in multiples of `DST_TRANSFER_SIZE`. This field is a signed number and hence may be decrementing or incrementing. If the destination address does not need to be incremented, you can set this parameter to zero.

Y Count: This field determines the number of 1-D transfers present in the Y loop (outer loop). This field is valid when the `DESCR_TYPE` is set to 2-D transfer.

Source Address Increment (Y loop) (`SCR_Y_INCR`): This field configures the index by which the source address is to be incremented, for every iteration in a Y loop. The field is expressed in multiples of `SRC_TRANSFER_SIZE`. This field is a signed number and hence may be decrementing or incrementing. If the source address does not need to be incremented, you can set this parameter to zero.

Destination Address Increment (X loop) (`DST_Y_INCR`): This field configures the index by which the destination address is to be incremented, for every iteration in a Y loop. The field is expressed in multiples of `DST_TRANSFER_SIZE`. This field is a signed number and hence may be decrementing or incrementing. If the destination address does not need to be incremented, you can set this parameter to zero.

Channel Disable (`CH_DISABLE`): This field specifies whether the channel is disabled or not after completion of the current descriptor (independent of the value of the `DESCR_NEXT_PTR`). A disabled channel will ignore its input triggers.

10.4.2 Transfer Size

The word width for a transfer can be configured using the transfer/data size parameter in the descriptor. The settings

are diversified into source transfer size, destination transfer size, and data size. The data size parameter (DATA_SIZE) sets the width of the bus for the transfer. The source and destination transfer sizes set by SCR_TRANSFER_SIZE and DST_TRANSFER_SIZE can have a value either the DATA_SIZE or 32 bit. DATA_SIZE can have a 32-bit, 16-bit, or 8-bit setting.

The source and destination transfer size for the DMA must match the addressable width of the source and destination, regardless of the width of data that must be moved. The DATA_SIZE parameter will correspond to the width of the actual data. For example, if a 16-bit PWM is used as a destination for DMA data, the DST_TRANSFER_SIZE must

be set to 32 bit to match the width of the PWM register, because the peripheral register width for the TCPWM block (and most PSoC 6 MCU peripherals) is always 32-bit wide. However, in this example the DATA_SIZE for the destination may still be set to 16 bit because the 16-bit PWM only uses two bytes of data. SRAM and Flash are 8-bit, 16-bit, or 32-bit addressable and can use any source and destination transfer sizes to match the needs of the application.

Table 10-3 summarizes the possible combinations of the transfer size settings and its description.

Table 10-3. Transfer Size Settings

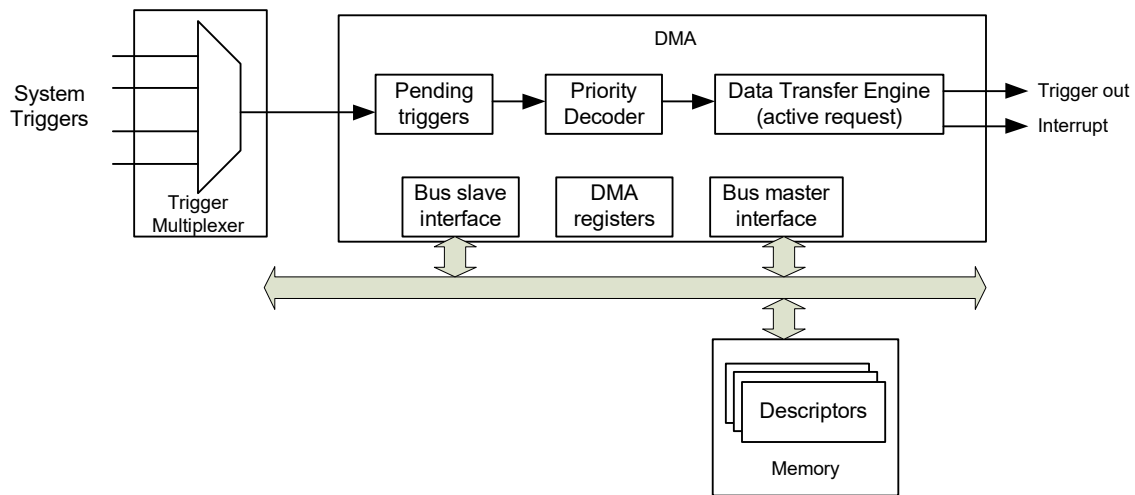
DATA_SIZE	SCR_TRANSFER_SIZE	DST_TRANSFER_SIZE	Typical Usage	Description
8-bit	8-bit	8-bit	Memory to Memory	No data manipulation
8-bit	32-bit	8-bit	Peripheral to Memory	Higher 24 bits from the source dropped
8-bit	8-bit	32-bit	Memory to Peripheral	Higher 24 bits zero padded at destination
8-bit	32-bit	32-bit	Peripheral to Peripheral	Higher 24 bits from the source dropped and higher 24 bits zero padded at destination
16-bit	16-bit	16-bit	Memory to Memory	No data manipulation
16-bit	32-bit	16-bit	Peripheral to Memory	Higher 16 bits from the source dropped
16-bit	16-bit	32-bit	Memory to Peripheral	Higher 16 bits zero padded at destination
16-bit	32-bit	32-bit	Peripheral to Peripheral	Higher 16 bits from the source dropped and higher 16-bit zero padded at destination
32-bit	32-bit	32-bit	Peripheral to Peripheral	No data manipulation

10.4.3 Descriptor Chaining

Descriptors can be chained together. The DESCR_NEXT_PTR field contains a pointer to the next descriptor in the chain. A channel executes the next descriptor in the chain when it completes executing the current descriptor. The last descriptor in the chain has DESCR_NEXT_PTR set to '0' (NULL pointer). A descriptor chain is also referred to as a descriptor list. It is possible to have a circular list; in a circular list, the execution continues indefinitely until there is an error or the channel or the controller is disabled by user code.

10.5 DMA Controller

Figure 10-3. DMA Controller Overview



10.5.1 Trigger Selection

Trigger signals can be generated from different sections of the chips. A trigger multiplexer block helps route these trigger signals to the destination. The DMA is one such destination of triggers. The trigger multiplexer block is outside the DMA block and is discussed in the [Trigger Multiplexer Block chapter on page 260](#).

10.5.2 Pending Triggers

Pending triggers keep track of activated triggers by locally storing them in pending bits. This is essential because multiple channel triggers may be activated simultaneously, whereas only one channel can be served by the data transfer engine at a time. This component enables the use of both level-sensitive (high/'1') and pulse-sensitive (two high/'1' clk_slow cycles) triggers.

- Level-sensitive triggers are associated with a certain state, for example, a FIFO being full. These triggers remain active as long as the state is maintained. It is not required to track pending level-sensitive triggers in the DMA controller because the triggers are maintained outside the controller.
- Pulse-sensitive triggers are associated with a certain event, for example, an ADC sample has become available. It is essential to track these triggers in the DMA controller because the trigger pulse may disappear before it is served by the data transfer engine. Pulse triggers should be high/'1' for two clk_slow cycles.

The **priority decoder** determines the highest priority pending channel.

The **data transfer engine** is responsible for the data transfer from a source location to a destination location. When idle, the data transfer engine is ready to accept the

highest priority activated channel. It is also responsible for reading the channel descriptor from memory.

Master I/F is an AHB-Lite bus master that allows the DMA controller to initiate AHB-Lite data transfers to the source and destination locations as well as to read the descriptor from memory.

Slave I/F is an AHB-Lite bus slave that allows the main CPU to access DMA controller control/status registers.

10.5.3 Output Triggers

Each channel has an output trigger. This trigger is high for two slow clock cycles. The trigger is generated on the completion of a data transfer. At the system level, these output triggers can be connected to the trigger multiplexer component. This connection allows a DMA controller output trigger to be connected to a DMA controller input trigger. In other words, the completion of a transfer in one channel can activate another channel or even reactivate the same channel.

DMA output triggers also connect to digital system interconnects (DSI) and some DSI signals connect to the trigger multiplexer inputs. Trigger outputs routing to other DMA channels or other peripheral trigger inputs is achieved using the trigger multiplexer. Refer to the [Trigger Multiplexer Block chapter on page 260](#).

10.5.4 Status registers

The controller status register (DWx_STATUS) contains the following information.

- ACTIVE – Active channel present, yes/no.
- P – Active channel's access control user/privileged
- NS – Active channel's access control secure/non-secure

- CH_IDX – Active channel index if there is an active channel
- PRIO – Active channel priority
- PREEMPTABLE – Active channel pre-emptable or not
- STATE – State of the DW controller state machine. The following states are specified:
 - Default/inactive state
 - Loading descriptor: This state is when the controller has recognized the channel that was triggered and become active and is now loading its respective descriptor.
 - Loading data element: Reading data from source address
 - Storing data element: Writing data into the destination address
 - Update of active channel control information
 - Waiting for input trigger deactivation

10.5.5 DMA Performance

The DMA block works on the clk_slow domain and hence all clocks described in this section are in clk_slow units.

Every time a DMA channel is triggered the DMA hardware goes through the following steps:

- Trigger synchronization
- Detection, priority decoding, and making channel pending
- Start state machine and load channel configuration
- Load DMA descriptor
- Load next DMA descriptor pointer
- Moving first element of data from source to destination.

Each of these steps involve multiple cycles for completion. [Table 10-4](#) shows the number of cycles needed for each step.

Table 10-4. DMA Steps and Performance

Operation	Cycles
Trigger Synchronization	2
Detection, priority decoding and making channel pending	1
Start state machine and load channel config	3
Load descriptors	4 for single transfer 5 for 1-D transfer 6 for 2-D transfer
Load next pointer	1
Moving data from source to destination	3
Total	14 for single transfer

For subsequent transfers on a preloaded descriptor, cycles are needed only to move the data from source to destination. Therefore, transfers such as 1-D and 2-D, which are not preempted, incurs all the cycles only for the first transfer; subsequent transfers will cost three cycles.

Based on the configuration of TRIG_IN_TYPE, the trigger synchronization cycles may be incurred for each single element transfer or for each 1-D transfer.

The descriptor is four words long for a single transfer type, five words for 1-D transfer, and six words for a 2-D transfer. Hence, the number of cycles needed to fetch a descriptor will vary based on its type.

Another factor to note is the latency in data or descriptor fetch due to wait states or bus latency.

The DMA performance for different types of transfers can be summarized as follows.

- Single transfer
 - 14 cycles per transfer + latency due to wait states or bus latency
- 1D transfer
 - To transfer n data elements
$$\text{Number of cycles} = 12 + n * 3 + m$$

m is the total number of wait states seen by DMA while loading or storing descriptors or data. An additional cycle is required for the first transfer, to load the X-Loop configuration register.
- 2D transfer
 - If the 2 D transfer is transferring n elements then

Number of cycles = $13 + n * 3 + m$

m is total number of wait states seen by DMA while loading or storing descriptors or data. Two additional cycles are required for the first transfer, to load the X-loop and Y-Loop configuration register.

Note: Descriptors in memory and memory wait states will also affect the descriptor load delay.

- Wait states: Memory accesses can have a wait state associated with them. These wait states need to be accounted into the calculation of throughput.
- Channel arbitration: Some time channels are not immediately made active after reception of trigger. This is due to other active channels in the system. This can lead to multiple cycles being lost before the channel is even made active.
- Preemption: The choice of making a DMA channel preemptable impacts its performance. This is because every time a channel is preempted:
 - The channel is in a pending state for as long as the higher priority channel is running
 - On resumption, the channel descriptor needs to be fetched again. This is additional cycles for every resume. So if there are a large number of high-priority channels, making a low-priority channel preemptable can have adverse effects on its throughput. On the other hand, if there is a low-priority channel that is transferring a large amount of data, then not making it preemptable can starve other high-priority channels for too long.

Sometimes, users can also distribute channels across multiple DW blocks to avoid conditions of preemption and deal with the contention at the bus arbitration level.

- Bus arbitration: Several bus masters access the bus, including the CPU cores and multiple DMA (DW) and DMAC. This makes any access to data movement over the bus subject to arbitration with other masters. Actions such as fetching the descriptor or data can be stalled by arbitration. The arbitration of the bus is based on the arbitration scheme configured in PROT_SMPU_MSx_CTL[PRIQ]
- Transfer width: The width of the transfer configured by the Data_size parameter in the descriptor is important in the transfer throughput calculation. 32-bit transfers are four times faster than 8-bit transfers.

11. DMAC Controller (DMAC)



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - DMAC
- Application notes
- Code examples

The DMAC transfers data to and from memory, peripherals, and registers. These transfers occur independent from the CPU. The DMAC can be configured to perform multiple independent data transfers. All data transfers are managed by a channel. There can be up to 32 channels in the DMAC. The number of channels in the DMAC controller can vary with devices. Refer to the [device datasheet](#) for the number of channels supported in the device. A channel has an associated priority; channels are arbitrated according to their priority. The main difference between the DMA (DW) and DMAC relate to their usage. The DMA (DW) is meant as a small data size, transactional DMA, which would typically be used to transfer bytes between peripherals such as, from ADC to RAM. Using the DMA (DW) for large transaction is expensive on a system due to its relatively low performance. The DMAC is a transactional DMA. It is more efficient than the DMA (DW) and should be used to transfer large amounts of data. The DMAC has dedicated channel logic for all channels. Furthermore, the DMAC also includes a 12-byte FIFO for temporary data storage. This results in increased memory bandwidth for the DMAC. The DMAC also supports an additional transfer mode called memory copy and scatter transfer.

11.1 Features

The DMAC controller has the following features:

- Supports up to channels per DMAC controller; see the [device datasheet](#) for details
- Four levels of priority for each channel
- Descriptors are defined in memory and referenced to the respective channels
- Supports single, 1D, 2D, Memory-copy, or Scatter-transfer modes for a descriptor
- Supports transfer of up to 2^{32} data elements per descriptor
- Configurable source and destination address increments
- Supports 8-bit, 16-bit, and 32-bit data widths at both source and destination
- Configurable input trigger behavior for each descriptor
- Configurable interrupt generation in each descriptor
- Configurable output trigger generation for each descriptor
- Descriptors can be chained to other descriptors in memory

11.2 Architecture

A data transfer is initiated by an input trigger. This trigger may originate from the source peripheral of the transfer, the destination peripheral of the transfer, CPU software, or from another peripheral. Triggers provide Active/Sleep functionality and are not available in Deep Sleep and Hibernate power modes.

The data transfer details are specified by a descriptor. Among other things, this descriptor specifies:

- The source and destination address locations and the size of the transfer.
- The actions of a channel; for example, generation of output triggers and interrupts. See the [Interrupts chapter on page 54](#) for more details.
- Data transfer types can be single, 1D, or 2D as defined in the descriptor structure. These types define the address sequences generated for source and destination. 1D and 2D transfers are used for “scatter gather” and other useful transfer operations.

11.3 Channels

The DMAC controller supports multiple independent data transfers that are managed by a channel. Each channel connects to a specific system trigger through a trigger multiplexer that is outside the DMAC controller.

Channel priority: A channel is assigned a priority (CHi_CTL.PRIO) between 0 and 3, with 0 being the highest priority and 3 being the lowest priority. Channels with the same priority constitute a priority group. Priority decoding determines the highest priority pending channel, which is determined as follows.

- The highest priority group with pending channels is identified first.
- Within this priority group, round-robin arbitration is applied. Round-robin arbitration within a priority group gives the highest priority to the lower channel indices in the priority group.

The data transfer associated with a trigger is made up of one or more ‘atomic transfers’ or ‘single transfers’; see [Table 11-1](#) for a better understanding. A single trigger could be configured to transfer multiple “single transfers”.

Channel Registers:

- CH_CTL. This register provides generic channel control information.
- CH_CURR_PTR. This register provides the memory location address where the current descriptor is located. Software needs to initialize this register. Hardware sets this register to the current descriptor’s next descriptor pointer, when advancing from the current descriptor to the next descriptor in a chained descriptor list. When this field is “0”, there is no valid descriptor.

- CH_IDX. This register provides the current X and Y indices of the channel into the current descriptor. Software needs to initialize this register. Hardware sets the X and Y indices to 0, when advancing from the current descriptor to the next descriptor in a descriptor list.
- CH_SRC. This register provides the current address of the source location.
- CH_DST. This register provides the current address of the destination location.
- CH_DESCR_STATUS. This register provides the validity of other CH_DESCR registers.
- CH_DESCR_CTL. This register contains a copy of DESCR_CTL of the currently active descriptor.
- CH_DESCR_SRC. This register contains a copy of DESCR_SRC of the currently active descriptor.
- CH_DESCR_DST. This register contains a copy of DESCR_DST of the currently active descriptor.
- CH_DESCR_X_INCR. This register contains a copy of DESCR_X_INCR of the currently active descriptor.
- CH_DESCR_X_SIZE. This register contains a copy of DESCR_X_SIZE of the currently active descriptor.
- CH_DESCR_Y_INCR. This register contains a copy of DESCR_Y_INCR of the currently active descriptor.
- CH_DESCR_Y_SIZE. This register contains a copy of DESCR_Y_SIZE of the currently active descriptor.
- CH_DESCR_NEXT_PTR. This register contains a copy of DESCR_NEXT_PTR of the currently active descriptor.
- INTR. This register contains the interrupts that are currently activated for this channel
- INTR_SET. Writing ‘1’ to the appropriate bit in this register sets the corresponding INTR field to 1.
- INTR_MASK. Mask for corresponding field in the INTR register.
- INTR_MASKED. Logical AND of corresponding INTR and INTR_MASK fields.

Note that channel state is retained in Deep Sleep power mode.

A channel has three access control attributes that are used by the shared memory protection units (SMPUs) and peripheral protection units (PPUs) for access control. These fields are typically inherited from the master that modified the channel’s control register.

- The Privileged Mode (CHi_CTL.P) attribute can be set to privileged or user.
- The Non-secure (CHi_CTL.NS) attribute can be set to secure or non-secure.
- The Protection context (CHi_CTL.PC) attribute can be set to one of the protection contexts.

11.3.1 Channel Interrupts

Every DMAC channel has an interrupt line associated with it. The INTR_TYPE parameter in the descriptor determines the event that will trigger the interrupt for the channel. In addition, each DMAC channel has INTR, INTR_SET, INTR_MASK, and INTR_MASKED registers to control their respective interrupt lines. INTR_MASK can be used to mask the interrupt from the DMA channel. The INTR and INTR_SET can be used to clear and set the interrupt, respectively, for debug purposes.

The DMAC_CHx_INTR registers allow the user to set different interrupt causes. The following causes are available to configure:




- bit 0: Interrupt based on transfer completion configured based on INTR_TYPE field in the descriptor
- bit 1: Source bus error
- bit 2: Destination bus error
- bit 3: Misaligned source address
- bit 4: Misaligned destination address
- bit 5: Current descriptor pointer is null
- bit 6: Active channel is in disabled state
- bit 7: Descriptor bus error

For error-related interrupt causes (every state other than Completion), the channel is disabled (hardware sets CH_CTL.ENABLED to '0').

11.4 Descriptors

The data transfer between a source and destination in a channel is configured using a descriptor. Descriptors are stored in memory. The descriptor pointer is specified in the DMAC channel registers. The DMAC controller does not modify the descriptor and treats it as read only. A descriptor is a set of up to six 32-bit registers that contain the configuration for the transfer in the associated channel. There are three types of descriptors.

Table 11-1. Descriptor Types

Descriptor Type	Description
Single transfer	Transfers a single data element 
1D transfer	Performs a one-dimensional "for loop". This transfer is made up of X number of single transfers 
2D transfer	Performs a two-dimensional "for loop". This transfer is made up of Y number of 1D transfers 
Scatter transfer	This descriptor type is intended to write a set of 32-bit data elements, whose addresses are "scattered" around the address space. The size of the descriptor is four 32-bit words. DESCR_CTL, DESCR_SRC, DESCR_X_SIZE, and DESCR_NEXT_PTR.

Single Transfer:

The following pseudo code illustrates a single transfer.

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
DST_ADDR[0] = (t_DATA_SIZE) SRC_ADDR[0];
```

1D Transfer:

The following pseudo code illustrates a 1D transfer. Note that the 1D transfer is represented by a loop with each iteration executing a single transfer.

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
for (X_IDX = 0; X_IDX <= X_COUNT; X_IDX++) {
  DST_ADDR[X_IDX * DST_X_INCR] =
    (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR];
}
```

2D Transfer:

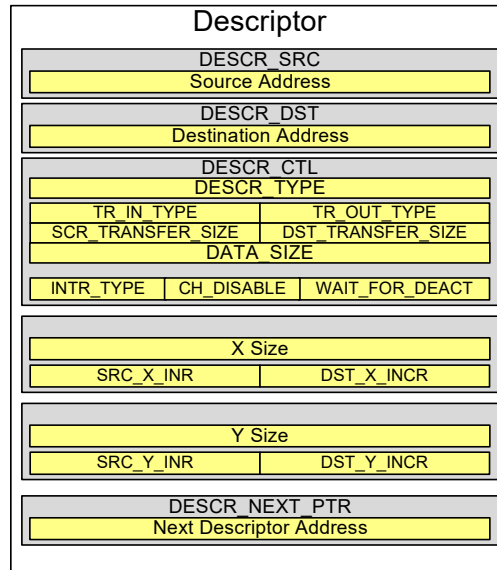
The following pseudo code illustrates a 2D transfer. Note that the 2D transfer is represented by a loop with each iteration executing an inner loop, which is the 1D transfer.

```
// DST_ADDR is a pointer to an object of type defined by DST_TRANSFER_SIZE
// SRC_ADDR is a pointer to an object of type defined by SRC_TRANSFER_SIZE
// t_DATA_SIZE is the type associated with the DATA_SIZE
for (Y_IDX = 0; Y_IDX <= Y_COUNT; Y_IDX++) {
  for (X_IDX = 0; X_IDX <= X_COUNT; X_IDX++) {
    DST_ADDR[X_IDX * DST_X_INCR + Y_IDX * DST_Y_INCR] =
      (t_DATA_SIZE) SRC_ADDR[X_IDX * SRC_X_INCR + Y_IDX * SRC_Y_INCR];
  }
}
```

The parameters in the descriptor help configure the different aspects of the transfers explained.

Figure 11-2 shows the structure of a descriptor.

Figure 11-1. Descriptor Structure



11.4.1 Address Configuration

Source and Destination Address: The source and destination addresses are set in the respective registers in the descriptor. These set the base addresses for the source and destination location for the transfer. In case the descriptor is configured to transfer a single element, this field holds the source/destination address of the data element. If the descriptor is configured to transfer multiple elements with source address or destination address or both in an incremental mode, this field will hold the address of the first element that is transferred.

DESCR_TYPE: This field configures whether the descriptor has a single, 1D, 2D type, memory copy, or scatter transfer.

Source data prefetch, DATA_PREFETCH: When enabled, source data transfers are initiated as soon as the channel is enabled; the current descriptor pointer is not "0" and there is space available in the channel's data FIFO. When the input trigger is activated, the trigger can initiate destination data transfers with data that is already in the channel's data FIFO. This effectively shortens the initial delay of the data transfer.

Note: Data prefetch should be used with care to ensure that data coherency is guaranteed and that prefetches do not cause undesired side effects.

Trigger input type, TR_IN_TYPE: This field determines how the DMAC engine responds to input trigger signal. This field can be configured for one of the following modes:

- **Type 0:** A trigger results in execution of a single transfer. Regardless of the DESCR_TYPE setting, a trigger input

will trigger only a single element transfer. For example, in a 1D transfer, the DMAC will transfer only one data element in every trigger.

- **Type 1:** A trigger results in the execution of a single 1D transfer. If the DESCR_TYPE was set to single transfer, the trigger signal will trigger the single transfer specified by the descriptor. For a DESCR_TYPE set to 1D transfer, the trigger signal will trigger the entire 1D transfer configured in the descriptor. For a 2D transfer, the trigger signal will trigger only a single iteration of the Y loop transfer. If the descriptor type is "memory copy", the trigger results in the execution of a memory copy transfer.
- **Type 2:** A trigger results in execution of the current descriptor. Regardless of DESCR_TYPE, the trigger will execute the entire descriptor. If there was a next descriptor configured for the current descriptor, this trigger setting will not automatically trigger the next descriptor.
- **Type 3:** A trigger results in execution of the current descriptor and also triggering the next descriptor. The execution of the next descriptor from this point will be determined by the TR_IN_TYPE setting of the next descriptor.

Trigger out type, TR_OUT_TYPE: This field determines what completion event will generate the output trigger signal. This field can be configured to one of the following modes:

- Type 0: Generates a trigger output for completion of every single element transfer.
- Type 1: Generates a trigger output for completion of a 1D transfer. If the descriptor type is “memory copy”, the output trigger is generated after the execution of a memory copy transfer. If the descriptor type is “scatter”, the output trigger is generated after the execution of a scatter transfer.
- Type 2: Generates a trigger output for completion of the current descriptor. This trigger output is generated independent of the state of the DESCR_NEXT_PTR.
- Type 3: Generates a trigger output on completion of the current descriptor, when the current descriptor is the last descriptor in the descriptor chain. This means a trigger is generated when the descriptor execution is complete and the DESCR_NEXT_PTR is ‘0’.

Interrupt Type, INTR_TYPE: This field determines which completion event will generate the output interrupt signal. This field can be configured to one of the following modes:

- Type 0: Generates an interrupt output for completion of every single element transfer.
- Type 1: Generates an interrupt output for completion of a 1-D transfer. If the descriptor type is “memory copy”, the interrupt is generated after the execution of a memory copy transfer. If the descriptor type is “scatter” the interrupt is generated after the execution of a scatter transfer.
- Type 2: Generates an interrupt output for completion of the current descriptor. This interrupt output is generated independent of the state of the DESCR_NEXT_PTR.
- Type 3: Generates an interrupt output on completion of the current descriptor, when the current descriptor is the last descriptor in the descriptor chain. This means an interrupt is generated when the descriptor execution is complete and the DESCR_NEXT_PTR is ‘0’.

WAIT_FOR_DEACT: When the DMAC transfer based on the TR_IN_TYPE is completed, the data transfer engine checks the state of trigger deactivation. The data transfer on the second trigger is initiated only after deactivation of the first. The WAIT_FOR_DEACT parameter will determine when the trigger signal is considered deactivated. The first DMAC transfer is activated when the trigger is activated, but the transfer is not considered complete until the trigger is deactivated. This field is used to synchronize the controller's data transfers with the agent that generated the trigger. This field has four settings:

- 0 – Pulse Trigger: Do not wait for deactivation.
- 1 – Level-sensitive waits four SYSCLK cycles: The DMAC trigger is deactivated after the level trigger signal is detected for four cycles.

- 2 – Level-sensitive waits 16 SYSCLK cycles: The DMAC transfer is initiated if the input trigger is seen to be activated after 16 clock cycles.
- 3 – Pulse trigger waits indefinitely for deactivation. The DMAC transfer is initiated after the trigger signal deactivates. The next transfer is initiated only if the trigger goes low and then high again.

X Size: This field determines the number of single element transfers present in the X loop (inner loop). This field is valid when the DESCR_TYPE is set to 1D or 2D transfer. For the “memory copy” descriptor type, (X_COUNT + 1) is the number of transferred Bytes. For the “scatter” descriptor type, ceiling (X_COUNT/2) is the number of (address, write data) initialization pairs processed.

Source Address Increment (X loop) (SCR_X_INCR): This field configures the index by which the source address is to be incremented for every iteration in an X loop. The field is expressed in multiples of SRC_TRANSFER_SIZE. This field is a signed number and hence may be decrementing or incrementing. If the source address does not need to be incremented, you can set this parameter to zero.

Destination Address Increment (X loop) (DST_X_INCR): This field configures the index by which the destination address is to be incremented, for every iteration in an X loop. The field is expressed in multiples of DST_TRANSFER_SIZE. This field is a signed number and hence may be decrementing or incrementing. If the destination address does not need to be incremented, you can set this parameter to zero.

Y Size: This field determines the number of 1-D transfers present in the Y loop (outer loop). This field is valid when the DESCR_TYPE is set to 2-D transfer.

Source Address Increment (Y loop) (SCR_Y_INCR): This field configures the index by which the source address is to be incremented, for every iteration in a Y loop. The field is expressed in multiples of SRC_TRANSFER_SIZE. This field is a signed number and hence may be decrementing or incrementing. If the source address does not need to be incremented, you can set this parameter to zero.

Destination Address Increment (X loop) (DST_Y_INCR): This field configures the index by which the destination address is to be incremented, for every iteration in a Y loop. The field is expressed in multiples of DST_TRANSFER_SIZE. This field is a signed number and hence may be decrementing or incrementing. If the destination address does not need to be incremented, you can set this parameter to zero.

Channel Disable (CH_DISABLE): This field specifies whether the channel is disabled or not after completion of the current descriptor (independent of the value of the DESCR_NEXT_PTR). A disabled channel will ignore its input triggers.

11.4.2 Transfer Size

The word width for a transfer can be configured using the transfer/data size parameter in the descriptor. The settings are diversified into source transfer size, destination transfer size, and data size. The data size parameter (DATA_SIZE) sets the width of the bus for the transfer. The source and destination transfer sizes set by SCR_TRANSFER_SIZE and DST_TRANSFER_SIZE can have a value either the DATA_SIZE or 32 bit. DATA_SIZE can have a 32-bit, 16-bit, or 8-bit setting.

The source and destination transfer size for the DMAC must match the addressable width of the source and destination, regardless of the width of data that must be moved. The DATA_SIZE parameter will correspond to the width of the

actual data. For example, if a 16-bit PWM is used as a destination for DMAC data, the DST_TRANSFER_SIZE must be set to 32 bit to match the width of the PWM register, because the peripheral register width for the TCPWM block (and most PSoC 6 MCU peripherals) is always 32-bit wide. However, in this example the DATA_SIZE for the destination may still be set to 16 bit because the 16-bit PWM only uses two bytes of data. SRAM and Flash are 8-bit, 16-bit, or 32-bit addressable and can use any source and destination transfer sizes to match the needs of the application.

Table 11-2 summarizes the possible combinations of the transfer size settings and its description.

Table 11-2. Transfer Size Settings


DATA_SIZE	SCR_TRANSFER_SIZE	DST_TRANSFER_SIZE	Description
8-bit	8-bit	8-bit	No data manipulation
8-bit	32-bit	8-bit	Higher 24 bits from the source dropped
8-bit	8-bit	32-bit	Higher 24 bits zero padded at destination
8-bit	32-bit	32-bit	Higher 24 bits from the source dropped and higher 24 bits zero padded at destination
16-bit	16-bit	16-bit	No data manipulation
16-bit	32-bit	16-bit	Higher 16 bits from the source dropped
16-bit	16-bit	32-bit	Higher 16 bits zero padded at destination
16-bit	32-bit	32-bit	Higher 16 bits from the source dropped and higher 16-bit zero padded at destination
32-bit	32-bit	32-bit	No data manipulation

11.4.3 Descriptor Chaining

Descriptors can be chained together. The DESCR_NEXT_PTR field contains a pointer to the next descriptor in the chain. A channel executes the next descriptor in the chain when it completes executing the current descriptor. The last descriptor in the chain has DESCR_NEXT_PTR set to '0' (NULL pointer). A descriptor chain is also referred to as a descriptor list. It is possible to have a circular list; in a circular list, the execution continues indefinitely until there is an error or the channel or the controller is disabled by user code.

11.5 DMAC Controller

Figure 11-2. DMAC Controller Overview



11.5.1 Trigger Selection

Trigger signals can be generated from different sections of the chips. A trigger multiplexer block helps route these trigger signals to the destination. The DMAC is one such destination of triggers. The trigger multiplexer block is outside the DMAC block and is discussed in the [Trigger Multiplexer Block chapter on page 260](#).

11.5.2 Channel Logic

The channel logic keeps track of pending triggers for each channel and initiates the transfer corresponding to the active descriptor, based on availability of the bus and arbitration by priority decoder.

The **priority decoder** determines the highest priority pending channel.

Master I/F is an AHB-Lite bus master that allows the DMAC controller to initiate AHB-Lite data transfers to the source and destination locations as well as to read the descriptor from memory.

Slave I/F is an AHB-Lite bus slave that allows the main CPU to access DMAC controller control/status registers.

11.5.3 Output Triggers

Each channel has an output trigger. This trigger is high for two slow clock cycles. The trigger is generated on the completion of a data transfer. At the system level, these output triggers can be connected to the trigger multiplexer component. This connection allows a DMAC controller output trigger to be connected to a DMAC controller input trigger. In other words, the completion of a transfer in one channel can activate another channel or even reactivate the same channel.

DMAC output triggers routing to other DMAC channels or other peripheral trigger inputs is achieved using the trigger multiplexer. Refer to the [Trigger Multiplexer Block chapter on page 260](#).

12. Cryptographic Function Block (Crypto)



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - Crypto
- Application notes
- Code examples

The Cryptographic block (Crypto) provides hardware implementation and acceleration of cryptographic functions. Implementation in hardware takes less time and energy than the equivalent firmware implementation. In addition, the block provides True Random Number generation functionality in silicon, which is not available in firmware.

12.1 Features

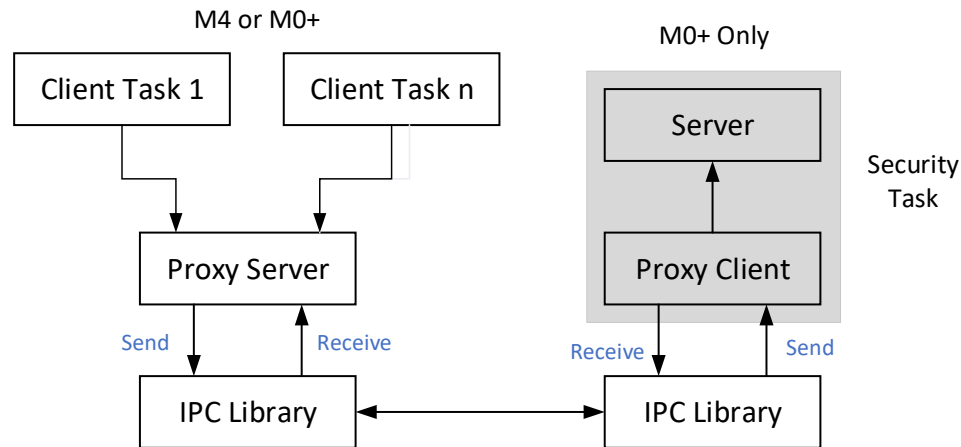
- Symmetric key encryption and decryption
- Hashing
- Message authentication
- Random number generation (pseudo and true)
- Cyclic redundancy checking
- Utility functions such as enable/disable, interrupt settings, and flags

12.2 Architecture

The Crypto block is implemented as a secure block and access to it is through inter-processor calls (IPC); direct access at a register level is not permitted for security reasons. A client-server model is used to access Crypto functions. The client will call a server running on the M0+. User tasks (clients) may run on either the M4 or M0+; the function calls will effectively call a proxy server, which will call the security server (providing Crypto functions) running on the M0+.

The IPC (see the [Inter-Processor Communication chapter on page 42](#)) mechanism is used to implement a remote procedure call (RPC) model that will be used to make calls to the security server code. Conceptually, the scheme is as follows.

Figure 12-1. Crypto Block Architecture



This model provides security for Crypto tasks that may otherwise expose or compromise key material. An API, as documented in the Peripheral Driver Library (PDL) documentation, will provide access to Crypto functions for general user code.

12.3 Definitions of Terms

The following terms are used in the description of this block:

- **Plaintext:** An unencrypted message
- **Ciphertext:** An encrypted message
- **Block cipher:** An encryption function for fixed-size blocks of data. This function takes a fixed-size key and a block of plaintext data from the message and encrypts it to generate ciphertext. Block ciphers are reversible, the function performed on a block of encrypted data will decrypt it.
- **Block cipher mode:** A mode of encrypting a message using block ciphers for messages of arbitrary length. The message is padded so that its length is an integer multiple of the block size. Electronic Code Book (ECB), Cipher Block Chaining (CBC), and Cipher Feedback (CFB) are all modes of using block ciphers to create an encrypted message of arbitrary length.
- **Data Encryption Standard (DES):** A block cipher that is now obsolete but supported for legacy reasons. It uses a 56-bit key and a 64-bit message block size.
- **3DES (or TDES):** Triple DES uses the DES operation and three DES encryptions in sequence—encrypt with DES with one 56-bit key, decrypt with a second 56-bit key, and then encrypt again either with the first key or a third 56-bit key. The block size is 64-bits.
- **Advanced Encryption Standard (AES):** This block cipher was designed to replace DES and 3DES. It is the block cipher standard used currently. It uses keys that can be 128, 192, or 256 bits in length and the message block size is 128 bits. This is the U.S. government standard. AES is also used for message authentication.
- **Secure Hash Algorithm (SHA):** This function takes a message of arbitrary length and reduces it to a fixed-length residue or message digest after performing a series of mathematically-defined operations, which guarantees that any change in the message will change the hash value. It is used for message authentication by transmitting a message with a hash value appended to it and recalculating the message hash value using the same algorithm at the recipient's end. If the hashes differ then the message is corrupted.
- **Message Authentication Code (MAC):** MACs are used to verify that a received message has not been altered. This is done by first computing a MAC value at the sender's end and appending it to the transmitted message. When the message is received, the MAC is computed again and checked against the MAC value transmitted with the message. If they do not match, the message has been altered. Either a hash algorithm (such as SHA) or a block cipher (such as AES) can be used to produce the MAC value. Keyed MAC schemes use a secret key along with the message, thus the key value must be known to be able to compute the MAC value.
- **Hash Message Authentication Code (HMAC):** This method uses a key along with the message to compute the MAC value using a hash algorithm.
- **Cipher-based Message Authentication Code (CMAC):** This method uses a key along with the message to compute the MAC value using the AES block cipher algorithm.
- **Pseudo Random Number Generator (PRNG):** This is based on Linear Feedback Shift Registers, which generate a sequence starting from a non-zero seed.
- **True Random Number Generator (TRNG):** A block that generates a number that is statistically random and

based on some physical random variation, cannot be duplicated by running the process again.

- **Symmetric key cryptography:** Refers to using a common known key to encrypt and decrypt messages (a shared secret between sender and receiver). An efficient method used for encrypting and decrypting messages when the authenticity of the other party has been established. DES (now obsolete), 3DES, and AES (currently used) are well-known symmetric cryptography methods.
- **Asymmetric key cryptography:** Also referred to as Public and Private key methods. If you want to receive a message, you must publish a public key (can be up to 4096 bits). Anyone wishing to send a message to the publisher of the public key encrypts the message with the public key; this message can now only be decrypted with the private key (the other prime factor held secret by the recipient). The message is now sent over any channel to the recipient who can decrypt it with the private, secret key. The same process is used to send messages to the sender of the original message. Asymmetric cryptography relies on the mathematical impracticality (usually related to the processing power available at any given time) of factoring the keys. Common, computationally intensive, asymmetric algorithms are RSA and ECC.

12.4 Crypto Block Functions

The following is a list of currently supported (at an API level) functions.

12.4.1 Symmetric Key Functions

Data Encryption Standard (DES): DES is a block cipher, an encryption function that works on fixed-size blocks of data. The DES block size is 64 bits and it uses a 56-bit key. DES is now considered obsolete, but is retained for compatibility with legacy systems. With block ciphers, encryption is considered a forward cipher; block ciphers have a complementary operation called decryption (inverse cipher), which has the same parameters as encryption and uses either the same key or a transform of that key.

DES is described in FIPS 46-3, which is archived at <http://csrc.nist.gov/publications/fips/fips46-3/fips46-3.pdf>. Note

that DES is not recommended for use anymore and support is provided only for legacy and compatibility reasons.

Triple DES: TDES (also known as Triple-DES or 3DES) is the DES algorithm invoked thrice with three different keys; the first pass encrypts with one key, the second pass decrypts with the second, and the third pass encrypts with the third key.

TDES is specified in <http://csrc.nist.gov/publications/nistpubs/800-67-Rev1/SP-800-67-Rev1.pdf>.

Advanced Encryption Standard (AES): The AES operation works on a 128-bit block size and uses keys of 128, 192, or 256 bits in length. The block cipher modes supported are ECB, CBC, CFB, and CTR.

The ECB mode is invoked by calling a function, which takes the following parameters: a pointer to the key, the function mode (forward/encryption or inverse/decryption), a pointer to the source block, and a pointer to the result block.

The CBC function is more secure than EBC; it relies on an "Exclusive-Or" of each plaintext block with the previous ciphertext block. Thus, block-sized plaintext patterns cannot be used to derive information about cipher characteristics as with EBC. The CBC function is invoked with the same parameters as EBC, with the addition of a pointer to an initial value.

The four modes supported are in the Arm mbed TLS library. AES is specified in FIPS 197: <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>. Operating modes for the block cipher are described in Special Publication 800-38A at <http://csrc.nist.gov/publications/nistpubs/800-38a/sp800-38a.pdf>.

12.4.2 Hash Functions

Secure Hash Algorithm (SHA): The SHA produces a message digest from a message of arbitrary size. It takes as arguments pointers to the message, the size of the message, a pointer to the initial digest value, and the mode. The mode specifies which SHA algorithm to use. Note that SHA does not use a key but starts with a pre-defined hash value. It supports the following message sizes and produces the stated message digest size.

Table 12-1. SHA Message Size

Algorithm	Message Size (bits)	Block Size (bits)	Word Size (bits)	Message Digest Size (bits)
SHA-1	< 2 ⁶⁴	512	32	160
SHA-224	< 2 ⁶⁴	512	32	224
SHA-256	< 2 ⁶⁴	512	32	256

Table 12-1. SHA Message Size

Algorithm	Message Size (bits)	Block Size (bits)	Word Size (bits)	Message Digest Size (bits)
SHA-384	< 2 ¹²⁸	1024	64	384
SHA-512	< 2 ¹²⁸	1024	64	512
SHA-512/224	< 2 ¹²⁸	1024	64	224
SHA-512/256	< 2 ¹²⁸	1024	64	256

SHA is specified in FIPS 180-4, Secure Hash Standard, at <http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.180-4.pdf>.

12.4.3 Message Authentication Code (MAC) Functions

Hashed Message Authentication Code (HMAC)

This function produces a message digest from a message of arbitrary size, but also uses a private key so that the hash cannot be calculated without the knowledge of the private key. HMAC is implemented using SHA described earlier; in addition to the parameters described in Table 12-1, it takes two additional parameters: a pointer to the key and the length of the key.

HMAC is specified in FIPS 198-1 at http://csrc.nist.gov/publications/fips/fips198-1/FIPS-198-1_final.pdf.

Cipher-based Message Authentication Code (CMAC)

This is a block-based message authentication algorithm that uses the AES block cipher to produce a message digest (analogous to HMAC, which uses a hash function). This function uses a key to encrypt the first block of the message after which the encrypted block produced is XOR-ed with the second message block and encrypted. The result of that operation is XOR-ed with the third message block and so on. The last message block is mixed with sub-keys derived from the encryption key, XOR-ed with the preceding encryption value, and encrypted to produce the authentication output.

The function implementing CMAC using AES is called with a pointer to the key, the key length, and pointers to the message block, a constant used for sub-key generation, and a pointer to temporary storage.

CMAC is specified at http://csrc.nist.gov/publications/nistpubs/800-38B/SP_800-38B.pdf.

12.4.4 Cyclic Redundancy Code (CRC)

The CRC block generates a CRC remainder with a programmable polynomial of up to 32 bits. The CRC function considers the following as arguments: the polynomial, the initial state of the LFSR, data byte reversal and XOR-ing with a pattern in a register, remainder reversal and XOR-ing with a pattern in a register, the data size, and a pointer to the data block. It returns the CRC remainder value.

12.4.5 Random Number Generator (RNG)

Pseudo Random Number Generator (PRNG): The PRNG generates pseudo random numbers based on three Linear Feedback Shift Registers (LFSRs). The three registers are initialized to non-zero seed values. The function takes as argument the length of the argument to be returned and returns a pseudo-random number.

True Random Number Generator (TRNG): The TRNG generates true random numbers with programmable bit size. It uses up to size polynomial generators with fixed and programmable polynomials. The function takes as argument the length of the argument to be returned and returns a true random number.

12.5 Module Configuration and Initialization

The Crypto module can be enabled or disabled using the Enable and Disable functions that take no arguments. The operational mode of the block is through a command (and parameters as required) FIFO, which is written with a command and parameters as required. Execution of the command invokes the appropriate operation such as DES, TDES, and SHA.

The interrupt model is based on four registers that contain interrupt flags resulting from various causes, bits for masking interrupts, software writable interrupt flags, and interrupt state bits (logical AND of mask and interrupt flag bits).

An initialization function is provided for setting interrupt flags, which must be responded to, to set interrupt flags in software, masking interrupts, and reading interrupt state.

Details of APIs provided are in the Peripheral Driver Library (PDL) documentation.

12.6 Software Design Considerations

Some system calls use the crypto block to perform the requested function. Care should be taken not to access the crypto block from the CM4 when it is in use by the secure CM0+. System calls are executed by the secure CM0+. [Table 16-4](#) lists all the system calls and indicates which use the crypto block. It is good practice to use a mutex or semaphore so that an application with multiple tasks can prevent two or more tasks attempting to access the crypto block at the same time whether by direct access or a system call.

13. Program and Debug Interface



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation
- Application notes
- Code examples

The PSoC 6 MCU Program and Debug interface provides a communication gateway for an external device to perform programming or debugging. The external device can be a Cypress-supplied programmer and debugger, or a third-party device that supports programming and debugging. The serial wire debug (SWD) or the JTAG interface can be used as the communication protocol between the external device and PSoC 6 MCUs. During the provisioning process, the developer can set the level of access to the Cortex-M4, M0+, memory, and MMIO registers to the level required. During the development phase, the debug ports may be open to all resources; however, when the product is ready for shipment, the debug ports may be totally disabled.

13.1 Features

- Supports programming and debugging through the JTAG or SWD interface.
- CM4 supports 4-bit ETM tracing, serial wire viewer (SWV), and printf() style debugging through the single-wire output (SWO) pin. CM0+ supports Micro Trace Buffer (MTB) with 4 KB dedicated RAM.
- Supports Cross Triggering Interface (CTI) and Cross Triggering Matrix (CTM).
- CM0+ supports four hardware breakpoints and two watchpoints. CM4 supports six hardware breakpoints and four watchpoints.
- Provides read and write access to all memory and registers in the system while debugging, including the Cortex-M4 and Cortex-M0+ register banks when the core is running or halted.

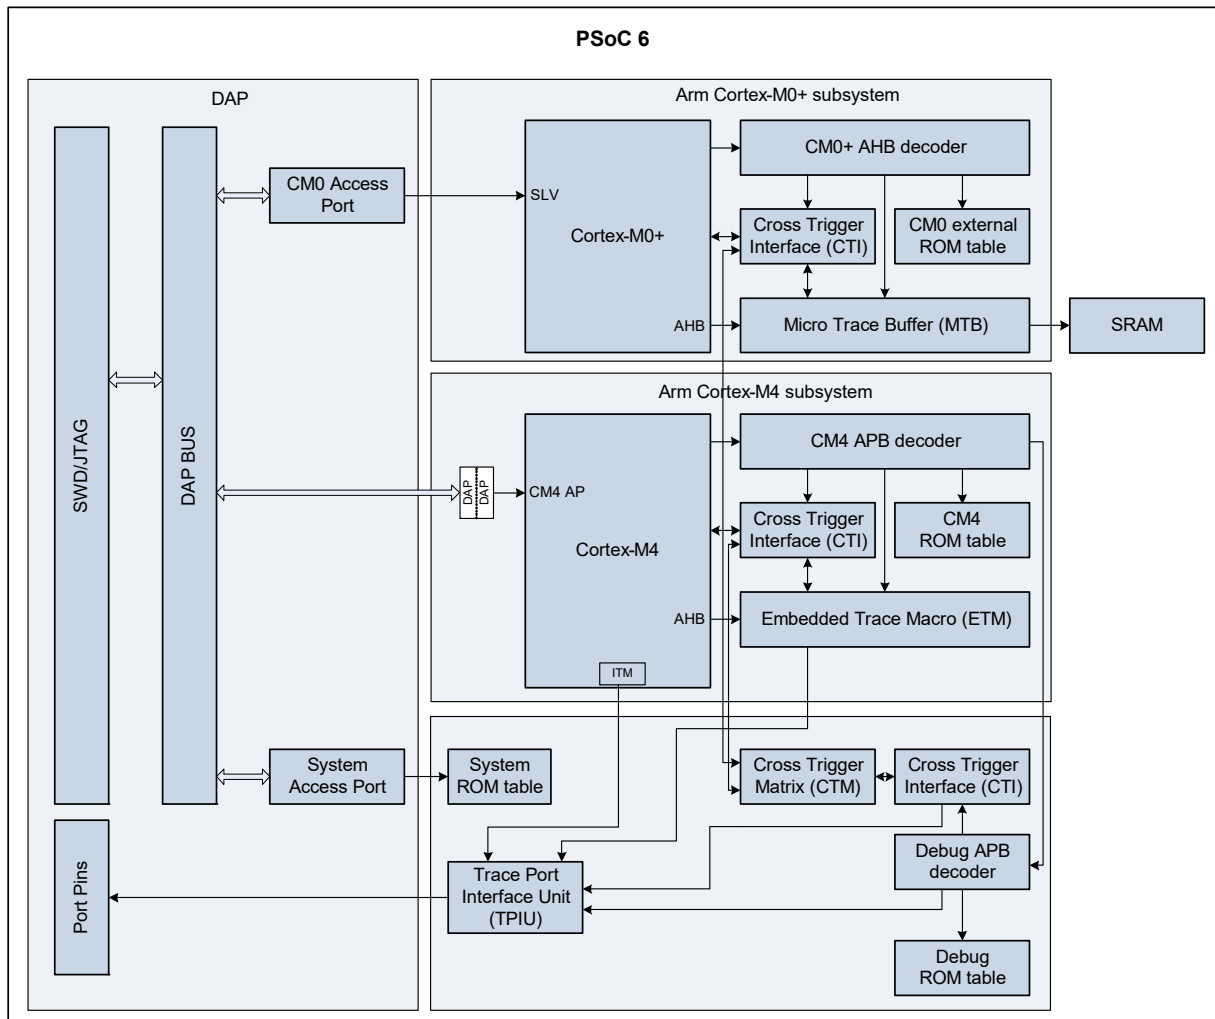
13.2 Architecture

Figure 13-1 shows the block diagram of the program and debug interface in the PSoC 6 MCU. The debug and access port (DAP) acts as the program and debug interface. The external programmer or debugger, also known as the “host”, communicates with the DAP of the PSoC 6 MCU “target” using either the SWD or JTAG interface. The debug physical port pins communicate with the DAP through the high-speed I/O matrix (HSIOM). See the [I/O System chapter on page 226](#) for details on HSIOM.

The debug infrastructure is organized in the following four groups:

- DAP (provides pin interfaces through which the debug host can connect to the chip)
- Cortex-M0+ core debug components (The Cortex-M0+ CPU debug port is disabled by default.)
- Cortex-M4 core debug components
- Other debug infrastructure (includes the CM4 tracing, the CTM, and the System ROM table)

Figure 13-1. Program and Debug Interface



The DAP communicates with the Cortex-M0+ CPU using the Arm-specified advanced high-performance bus (AHB) interface. AHB is the systems interconnect protocol used inside the device, which facilitates memory and peripheral register access by the AHB master. The PSoc 6 MCU has six AHB masters – Arm CM4 CPU core, Arm CM0 CPU core, Datawire0, Datawire1, Crypto, and DAP. The external host can effectively take control of the entire device through the DAP to perform programming and debugging operations.

The following are the various debug and trace components:

- Debug components
 - JTAG and SWD for debug control and access
- Trace source components
 - Micro trace buffer (MTB-M0+) for tracing Cortex-M0+ program execution
 - Embedded trace macrocell (ETM-M4) for tracing Cortex-M4 program execution
- Trace sink components
 - Trace port interface unit (TPIU) to drive the trace information out of the chip to an external trace port analyzer
- Cross-triggering components
 - Cross-trigger interface (CTI)
 - Cross-trigger matrix (CTM)
- ROM tables

13.2.1 Debug Access Port (DAP)

The DAP consists of a combined SWD/JTAG interface (SWJ) that also includes the SWD listener. The SWD listener decides whether the JTAG interface (default) or SWD interface is active. Note that JTAG and SWD are mutually exclusive because they share pins.

The debug port (DP) connects to the DAP bus, which in turn connects to one of three Access Ports (AP), namely:

- The CM0-AP, which connects directly to the AHB debug slave port (SLV) of the CM0+ and gives access to the CM0+ internal debug components. This also allows access to the rest of the system through the CM0+ AHB master interface. This provides the debug host the same view as an application running on the CM0+. This includes access to the MMIO of other debug components of the Cortex M0+ subsystem. These debug components can also be accessed by the CM0+ CPU, but cannot be reached through the other APs or by the CM4 core. The CM0-AP is disabled with the default provisioning script, but may be changed by the developer.
- The CM4-AP located inside the CM4 gives access to the CM4 internal debug components. The CM4-AP also allows access to the rest of the system through the CM4 AHB master interfaces. This provides the debug host the same view as an application running on the CM4 core. Additionally, the CM4-AP provides access to the debug components in the CM4 core through the External Peripheral Bus (EPB). These debug components can also be accessed by the CM4 CPU, but cannot be reached through the other APs or by the CM0+ core.
- The System-AP, which through an AHB mux gives access to the rest of the system. This allows access to the System ROM table, which cannot be reached any other way. The System ROM table provides the chip ID but is otherwise empty.

13.2.1.1 DAP Security

Although there are three debug access ports (CM4, CM0+, and SYS) only two of them (CM4 and SYS) are available to the user in the Secure Boot devices, by default. The CM0+ access port may be enabled if required for development purposes, but needs to be closed when the product is in production. The CM4 access port is open by default to allow users to develop their application. After the code development cycle is complete and the product is moved to manufacturing, both the CM4 and SYS access ports may be totally disabled by configuring the proper Debug Policy. See the [“Secure Boot” chapter on page 134](#) for more information on security and security policies.

13.2.1.2 DAP Power Domain

Almost all the debug components are part of the Active power domain. The only exception is the SWD/JTAG-DP,

which is part of the Deep Sleep power domain. This allows the debug host to connect during Deep Sleep, while the application is ‘running’ or powered down. This enables in-field debugging for low-power applications in which the chip is mostly in Deep Sleep.

After the debugger is connected to the chip, it must bring the chip to the Active state before any operation. For this, the SWD/JTAG-DP has a register (DP_CTL_STAT) with two power request bits. The two bits are CDBGPWRUPREQ and CSYSPWRUPREQ, which request for debug power and system power, respectively. These bits must remain set for the duration of the debug session.

Note that only the two SWD pins (SWCLKTCK and SWDIOTMS) are operational during the Deep Sleep mode – the JTAG pins are operational only in Active mode. The JTAG debug and JTAG boundary scan are not available when the system is in Deep Sleep mode. JTAG functionality is available only after a chip power-on-reset.

13.2.2 ROM Tables

The ROM tables are organized in a tree hierarchy. Each AP has a register that contains a 32-bit address pointer to the base of the root ROM table for that AP. For PSoc 6 MCUs, there are three such root ROM tables.

Each ROM table contains 32-bit entries with an address pointer that either points to the base of the next level ROM table. Each ROM table also contains a set of ID registers that hold JEDEC compliant identifiers to identify the manufacturer, part number, and major and minor revision numbers. For all ROM tables in PSoc 6 MCUs, these IDs are the same. Each ROM table and CoreSight compliant component also contains component identification registers.

13.2.3 Trace

The micro trace buffer (MTB-M0+) component captures the program execution flow from Cortex-M0+ CPU and stores it in a local SRAM memory. This information can be read by an external debug tool through JTAG/SWD interface to construct the program execution flow.

The embedded trace macro (ETM) component connected to Cortex-M4 captures the program execution flow from Cortex-M4 CPU and generates trace output on its advanced trace bus (ATB) interface. The instrumentation trace macrocell (ITM), which is inside Cortex-M4, also generates trace output on its ATB interface. These two ATB interfaces (from ETM-M4 and ITM) are connected a trace port interface unit (TPIU).

The TPIU drives the external pins of a trace port (through IOSS interface), so that the trace can be captured by an external trace port analyzer (TPA). For more details, refer to the *Arm Debug Interface Architecture Specification ADIV5.0 to ADIV5.2*.

13.2.4 Embedded Cross Triggering

The Arm CoreSight includes Embedded Cross Triggering (ECT) to communicate events between debug components. These events are particularly useful with tracing and multicore platforms. For example trigger events can be used to:

- Start or stop both CPUs at (almost) the same time
- Start or stop instruction tracing based on trace buffer being full or not or based on other events

CoreSight uses two components to support ECT, namely a CTI and a CTM, both of which are used in PSoC 6 MCUs.

The CTI component interfaces with other debug components, sending triggers back and forth and synchronizing them as needed. The CTM connects several CTIs, thus allowing events to be communicated from one CTI to another.

The PSoC 6 MCU has three CTIs, one for each CPU and one for the trace components in the debug structure. These three CTIs are connected together through the CTM. The CM4 CTI is located in the fast clock domain and the other two CTIs and the CTM are all located in the same slow-frequency clock domain. For more details, refer to the Arm documentation.

13.3 Serial Wire Debug (SWD) Interface

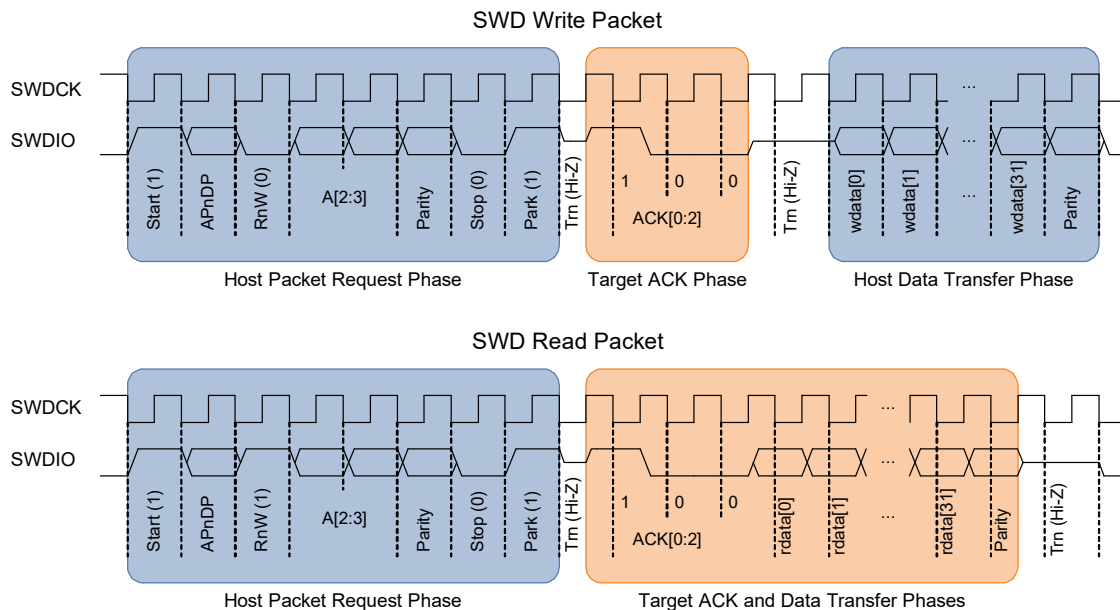
The PSoC 6 MCU supports programming and debugging through the SWD interface. The SWD protocol is a packet-based serial transaction protocol. At the pin level, it uses a single bidirectional data signal (SWDIO) and a unidirectional clock signal (SWDCK). The host programmer always drives the clock line, whereas either the host or the target drives the data line. A complete data transfer (one SWD packet) requires 46 clocks and consists of three phases:

- **Host Packet Request Phase** – The host issues a request to the PSoC 6 MCU target.
- **Target Acknowledge Response Phase** – The PSoC 6 MCU target sends an acknowledgement to the host.
- **Data Transfer Phase** – The host or target writes data to the bus, depending on the direction of the transfer.

When control of the SWDIO line passes from the host to the target, or vice versa, there is a turnaround period (T_{rn}) where neither device drives the line and it floats in a high-impedance (Hi-Z) state. This period is either one-half or one and a half clock cycles, depending on the transition.

Figure 13-2 shows the timing diagrams of read and write SWD packets.

Figure 13-2. SWD Write and Read Packet Timing Diagrams



The sequence to transmit SWD read and write packets are as follows:

1. Host Packet Request Phase: SWDIO driven by the host

- a. The start bit initiates a transfer; it is always logic 1.
- b. The AP not DP (APnDP) bit determines whether the transfer is an AP access – 1b or a DP access – 0b.

- c. The Read not Write bit (RnW) controls which direction the data transfer is in. 1b represents a 'read from' the target, or 0b for a 'write to' the target.
 - d. The Address bits (A[3:2]) are register select bits for AP or DP, depending on the APnDP bit value.
Note: Address bits are transmitted with the LSb first.
 - e. The parity bit contains the parity of APnDP, RnW, and ADDR bits. It is an even parity bit; this means, when XORed with the other bits, the result will be 0.
If the parity bit is not correct, the header is ignored by the PSoC 6 MCU; there is no ACK response (ACK = 111b). The programming operation should be aborted and retried again by following a device reset.
 - f. The stop bit is always logic 0.
 - g. The park bit is always logic 1.
2. Target Acknowledge Response Phase: SWDIO driven by the target
 - a. The ACK[2:0] bits represent the target to host response, indicating failure or success, among other results. See [Table 13-1](#) for definitions. **Note:** ACK bits are transmitted with the LSb first.
 3. Data Transfer Phase: SWDIO driven by either target or host depending on direction
 - a. The data for read or write is written to the bus, LSb first.
 - b. The data parity bit indicates the parity of the data read or written. It is an even parity; this means when XORed with the data bits, the result will be 0.
If the parity bit indicates a data error, corrective action should be taken. For a read packet, if the host detects a parity error, it must abort the programming operation and restart. For a write packet, if the target detects a parity error, it generates a FAULT ACK response in the next packet.

According to the SWD protocol, the host can generate any number of SWDCK clock cycles between two packets with SWDIO low. Three or more dummy clock cycles should be generated between two SWD packets if the clock is not free-running or to make the clock free-running in IDLE mode.

The SWD interface can be reset by clocking the SWDCK line for 50 or more cycles with SWDIO high. To return to the idle state, clock the SWDIO low once.

13.3.1 SWD Timing Details

The SWDIO line is written to and read at different times depending on the direction of communication. The host drives the SWDIO line during the Host Packet Request phase and, if the host is writing data to the target, during the Data Transfer phase as well. When the host is driving the SWDIO line, each new bit is written by the host on falling SWDCK edges, and read by the target on rising SWDCK edges. The target drives the SWDIO line during the Target Acknowledge Response phase and, if the target is reading

out data, during the Data Transfer phase as well. When the target is driving the SWDIO line, each new bit is written by the target on rising SWDCK edges, and read by the host on falling SWDCK edges.

[Table 13-1](#) and [Figure 13-2](#) illustrate the timing of SWDIO bit writes and reads.

Table 13-1. SWDIO Bit Write and Read Timing

SWD Packet Phase	SWDIO Edge	
	Falling	Rising
Host Packet Request	Host Write	Target Read
Host Data Transfer		
Target Ack Response	Host Read	Target Write
Target Data Transfer		

13.3.2 ACK Details

The acknowledge (ACK) bitfield is used to communicate the status of the previous transfer. OK ACK means that previous packet was successful. A WAIT response requires a data phase. For a FAULT status, the programming operation should be aborted immediately. [Table 13-2](#) shows the ACK bit-field decoding details.

Table 13-2. SWD Transfer ACK Response Decoding

Response	ACK[2:0]
OK	001b
WAIT	010b
FAULT	100b
NO ACK	111b

Details on WAIT and FAULT response behaviors are as follows:

- For a WAIT response, if the transaction is a read, the host should ignore the data read in the data phase. The target does not drive the line and the host must not check the parity bit as well.
- For a WAIT response, if the transaction is a write, the data phase is ignored by the PSoC 6 MCU. But, the host must still send the data to be written to complete the packet. The parity bit corresponding to the data should also be sent by the host.
- A WAIT response means that the PSoC 6 MCU is processing the previous transaction. The host can try for a maximum of four continuous WAIT responses to see whether an OK response is received. If it fails, then the programming operation should be aborted and retried again.
- For a FAULT response, the programming operation should be aborted and retried again by doing a device reset.

13.3.3 Turnaround (Trn) Period Details

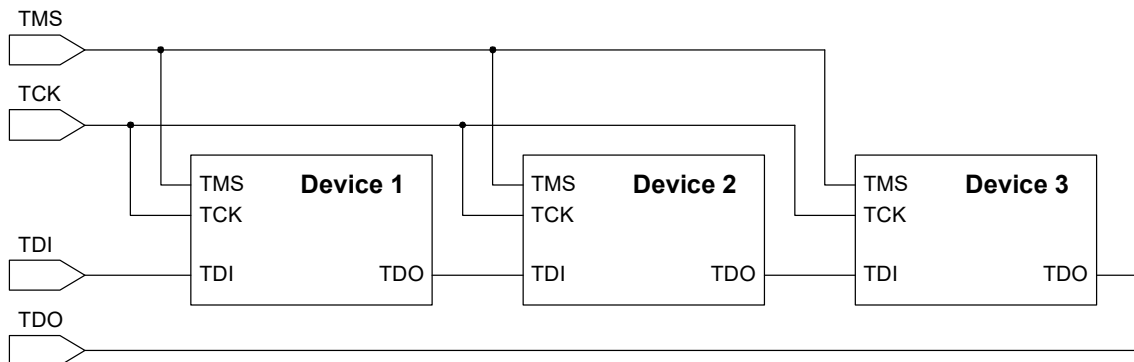
There is a turnaround period between the packet request and the ACK phases, as well as between the ACK and the data phases for host write transfers, as shown in [Figure 13-2](#). According to the SWD protocol, the Trn period is used by both the host and target to change the drive modes on their respective SWDIO lines. During the first Trn period after the packet request, the target starts driving the ACK data on the SWDIO line on the rising edge of SWDCK. This action ensures that the host can read the ACK data on the next falling edge. Thus, the first Trn period lasts only one-half cycle. The second Trn period of the SWD packet is one and a half cycles. Neither the host nor the PSoC 6 MCU should drive the SWDIO line during the Trn period.

13.4 JTAG Interface

In response to higher pin densities on ICs, the Joint Test Action Group (JTAG) proposed a method to test circuit

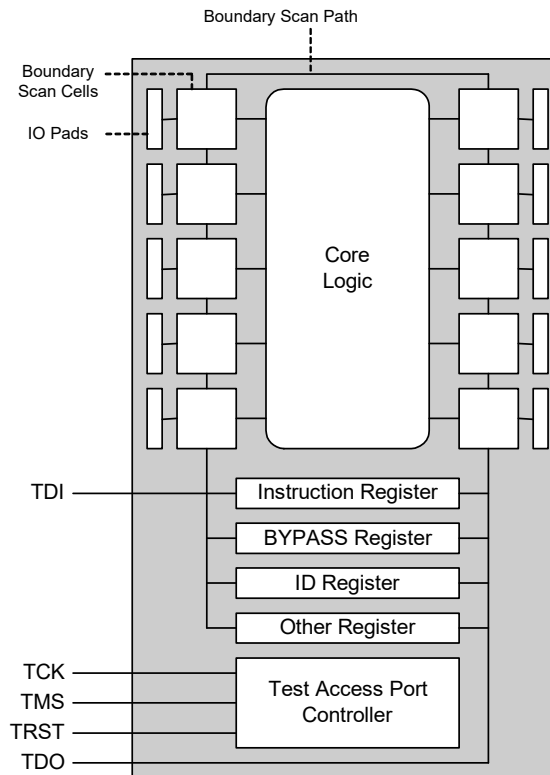
boards by controlling the pins on the ICs (and reading their values) via a separate test interface. The solution, later formalized as IEEE Standard 1149.1-2001, is based on the concept of a serial shift register routed across all of the pins of the IC – hence the name “boundary scan.” The circuitry at each pin is supplemented with a multipurpose element called a boundary scan cell. In PSoC 6 MCUs, most GPIO port pins have a boundary scan cell associated with them (see the GPIO block diagrams in the [I/O System chapter on page 226](#)). The interface used to control the values in the boundary scan cells is called the Test Access Port (TAP) and is commonly known as the JTAG interface. It consists of three signals: Test Data In (TDI), Test Data Out (TDO), and Test Mode Select (TMS). Also included is a clock signal (TCK) that clocks the other signals. TDI, TMS, and TCK are all inputs to the device and TDO is the output from the device. This interface enables testing multiple ICs on a circuit board, in a daisy-chain fashion, as shown in [Figure 13-3](#).

Figure 13-3. JTAG Interface to Multiple ICs on a Circuit Board



The JTAG interface architecture within each device is shown in [Figure 13-4](#). Data at TDI is shifted in, through one of several available registers, and out to TDO.

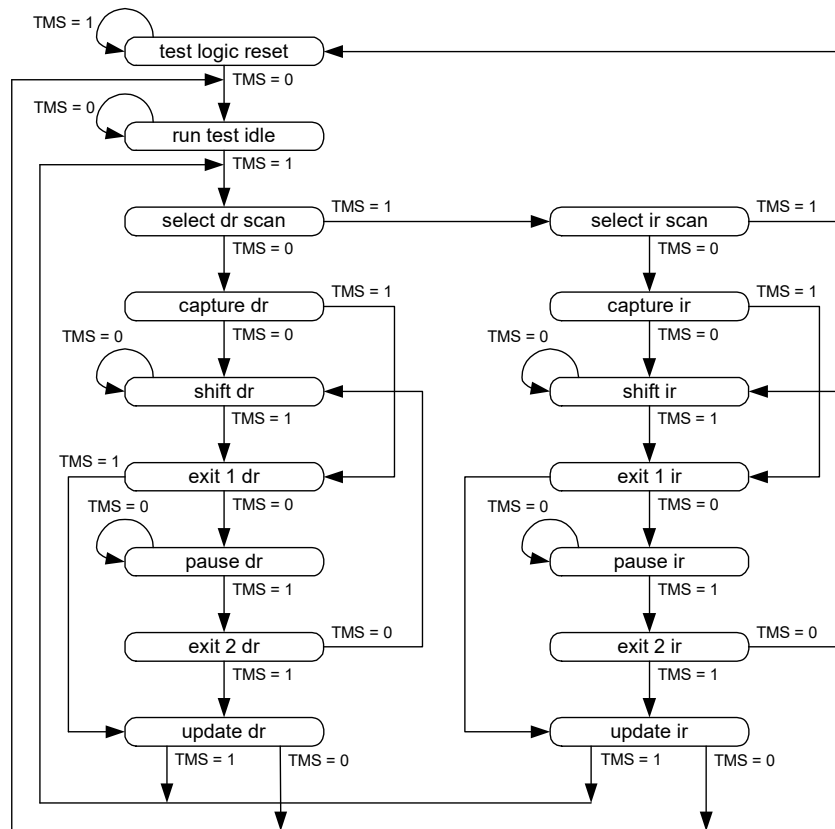
Figure 13-4. JTAG Interface Architecture



The TMS signal controls a state machine in the TAP. The state machine controls which register (including the boundary scan path) is in the TDI-to-TDO shift path, as shown in [Figure 13-5](#). The following terms apply:

- IR - the instruction register
- DR - one of the other registers (including the boundary scan path), as determined by the contents of the instruction register
- capture - transfer the contents of a DR to a shift register, to be shifted out on TDO (read the DR)
- update - transfer the contents of a shift register, shifted in from TDI, to a DR (write the DR)

Figure 13-5. TAP State Machine



The registers in the TAP are:

- Instruction – Typically two to four bits wide, holds the current instruction that defines which data register is placed in the TDI-to-TDO shift path.
- Bypass – one bit wide, directly connects TDI with TDO, causing the device to be bypassed for JTAG purposes.
- ID – 32 bits wide, used to read the JTAG manufacturer/part number ID of the device.
- Boundary Scan Path (BSR) – Width equals the number of I/O pins that have boundary scan cells, used to set or read the states of those I/O pins.

Other registers may be included in accordance with the device manufacturer specifications. The standard set of instructions (values that can be shifted into the instruction register), as specified in IEEE 1149, are:

- EXTEST – Causes TDI and TDO to be connected to the BSR. The device is changed from its normal operating mode to a test mode. Then, the device's pin states can be sampled using the capture dr JTAG state, and new values can be applied to the pins of the device using the update dr state.
- SAMPLE – Causes TDI and TDO to be connected to the BSR, but the device remains in its normal operating mode. During this instruction, the BSR can be read by the capture dr JTAG state to take a sample of the functional data entering and leaving the device.
- PRELOAD – Causes TDI and TDO to be connected to the BSR, but the device is left in its normal operating mode. The instruction is used to preload test data into the BSR before loading an EXTEST instruction.

Optional, but commonly available, instructions are:

- IDCODE – Causes TDI and TDO to be connected to an IDCODE register.
- INTEST – Causes TDI and TDO to be connected to the BSR. While the EXTEST instruction allows access to the device pins, INTEST enables similar access to the corelogic signals of a device

For more information, see the IEEE Standard, available at www.ieee.org.

13.5 Programming the PSoC 6 MCU

The PSoC 64 MCU is programmed using the industry-standard SWD port. Refer to the [PSoC 6 MCU Programming Specifications](#) for complete details on the programming algorithm, timing specifications, and hardware configuration required for programming. Direct from the factory, only the CM4 can be programmed until the device is provisioned. Refer to the [PSoC 64 Secure MCU Secure Boot SDK User Guide](#) and the “Secure Boot” chapter on [page 134](#) for more details.

13.5.1 SWD Port Acquisition

13.5.1.1 SWD Port Acquire Sequence

The first step in device programming is for the host to acquire the target's SWD port. The host first performs a device reset by asserting the external reset (XRES) pin. After removing the XRES signal, the host must send an SWD connect sequence for the device within the acquire window to connect to the SWD interface in the DAP.

The debug access port must be reset using the standard Arm command. The DAP reset command consists of more than 49 SWDCK clock cycles with SWDIO asserted high. The transaction must be completed by sending at least one SWDCK clock cycle with SWDIO asserted low. This sequence synchronizes the programmer and the chip.

Read_DAP() refers to the read of the IDCODE register in the debug port. The sequence of line reset and IDCODE read should be repeated until an OK ACK is received for the IDCODE read or a timeout (2 ms) occurs. The SWD port is said to be in the acquired state if an OK ACK is received within the time window and the IDCODE read matches with that of the Cortex-M0+ DAP.

13.5.2 SWD Programming Mode Entry

After the SWD port is acquired, the host must enter the device programming mode within a specific time window. This is done by setting the TEST_MODE bit (bit 31) in the test mode control register (MODE register). The debug port should also be configured before entering the device programming mode. Timing specifications and pseudo code for entering the programming mode are detailed in the [PSoC 6 MCU Programming Specifications](#) document.

13.5.3 SWD Programming Routine Executions

When the device is in programming mode, the external programmer can start sending the SWD packet sequence for performing programming operations such as flash erase, flash program, checksum verification, and so on. The programming routines are explained in the [Nonvolatile Memory chapter on page 127](#). The exact sequence of calling the programming routines is given in the [PSoC 6 MCU Programming Specifications](#) document.

13.6 Registers

Table 13-3. List of Registers

Register Name	Description
CM0P_DWT	Cortex M0+ Data Watchpoint and Trace (DWT) registers
CM0P_BP	Cortex M0+ BreakPoint (BP) registers
CM0P_ROM	Cortex M0+ CPU Coresight ROM table
CM0P_CTI	Cortex M0+ Cross-Trigger Interface (CTI) registers
CM0P_MTB	Cortex M0+ Micro Trace Buffer (MTB) registers
CM4_ITM	Cortex M4 Instrumentation Trace Macrocell (ITM) registers
CM4_DWT	Cortex M4 Data Watchpoint and Trace (DWT) registers
CM4_FPB	Cortex M4 Flash Patch and Breakpoint (FPB) registers
CM4_SCS	Cortex M4 System Control Space (SCS) registers
CM4_ETM	Cortex M4 Embedded Trace Macrocell (ETM) registers
CM4_CTI	Cortex M4 Cross-Trigger Interface (CTI) registers
CM4_ROM	Cortex M4 CPU Coresight ROM table
TRC_TPIU	System Trace Coresight Trace Port Interface Unit (TPIU) registers

14. Nonvolatile Memory



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation
- Application notes
- Code examples

Nonvolatile memory refers to the flash and SROM memory in the PSoC 6 MCU. This chapter explains the geometry and capability of the flash memory. It also lists the SROM API functions that are used to program the flash memory.

14.1 Flash Memory

The PSoC 6 flash offers high bulk program performance and supports ultra-low-power operation. Flash is typically used to store CPU instructions and data when the device power is off. Flash may be written, but the process is much slower and more restrictive than for SRAM.

14.1.1 Features

This section lists the features of PSoC 6 flash.

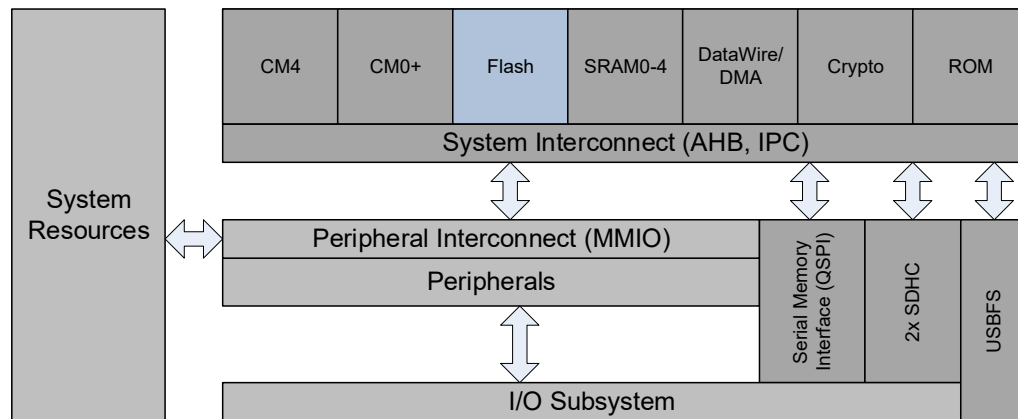
- 512-byte row size; minimum programmable unit
- Supports the Read While Write (RWW) feature with a RWW sector size of 256KB
- 10-year retention
- Endurance of 10 k program cycles

14.1.2 Configuration

14.1.2.1 Block Diagram

Flash is part of the CPU subsystem. The Cortex-M4 and Cortex M0+, as well as other bus masters, can access flash via the AHB.

Figure 14-1. Block Diagram



14.1.3 Flash Geometry

The flash is divided into three regions: the application region, the supervisory flash region (SFlash), and the Auxiliary Flash (AUXFlash) region. The flash has an Erase Disturb mechanism in which writes to rows affect the endurance of other rows in the same sector. It is recommended that the AUXFlash region is used for frequently-updated data. For data that changes infrequently or code images, the application flash region can be used.

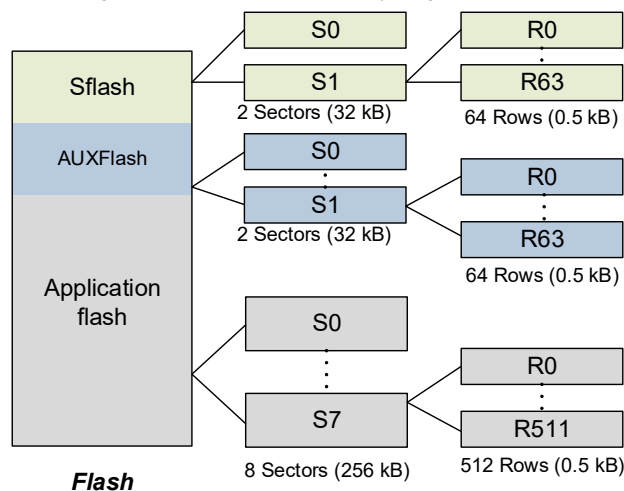
The SFlash region is used to store trim parameters, system configuration parameters, protection and security settings, boot code, and other Cypress proprietary information. Read access to this region is permitted, but program/erase access is limited. The application region is used to store code images or data. Common use of AUXFlash is to use if for EEPROM emulation.

Each region divides into sectors and pages. The sector is the largest division of the region and consists of a number of 512 byte rows.

Table 14-1. Flash Geometry

Application Flash						SFlash and AUXFlash				
Read Width	KB	Sectors	KB/Sector	Rows/Sector	KB/Row	KB	Sectors	KB/Sector	Rows/Sector	KB/Row
128	2048	8	256	512	0.5	64	2	32	64	0.5

Figure 14-2. Flash Geometry Organization



14.1.4 Flash Controller

Access to the flash memory is enabled through the flash controller. The flash controller interfaces with the AHB-Lite bus and provides flash access for the CM0+, CM4, Crypto, DataWire, and debug. The flash controller generates a bus error if:

- A flash read access to a sector that is currently being programmed/erased

- A read access to a memory hole in the flash memory region. A memory hole is defined as a location that is not occupied by the application flash, AUXFlash, or SFlash. Note that the AUXFlash space and SFlash space can be non-powers of 2 for some flash macros

The flash controller also provides the registers which support configuration of flash accesses.

Table 14-2. Flash Controller Retention Registers

Register Name	Description
FLASHC_FLASH_CTL	Flash controller control register
FLASHC_FLASH_PWR_CTL	Flash enable control
FLASHC_FLASH_CMD	Flash commands for cache/buffer invalidation
FLASHC_CM0/CM4_STATUS	CM0/4 Interface status
FLASHC_CM0/CM4_CA_CTL	CM0/4 Cache control
FLASHC_*Peripheral*_BUFF_CTL	Buffer control register where *peripheral* may be Crypto, DMA, DWx (DataWire x = 0, 1), or EXT_MSx (External master; x = 0, 1)

14.1.4.1 Wait State Count

FLASHC_FLASH_CTL supports configuration of flash wait cycles. If clk_hf is greater than the maximum operating frequency of the flash memory, it is necessary to insert wait cycles when accessing the flash memory by setting the appropriate value in the FLASHC_FLASH_CTL.MAIN_WS register. Set the wait cycles as follows.

HF Clock Frequency	FLASH_CTL.MAIN_WS[3:0]
$\text{clk_hf} \leq 33 \text{ MHz}$	0
$33 < \text{clk_hf} \leq 66 \text{ MHz}$	1
$66 < \text{clk_hf} \leq 99 \text{ MHz}$	2
$99 < \text{clk_hf} \leq 133 \text{ MHz}$	3
$133 < \text{clk_hf} \leq 150 \text{ MHz}$	4

14.1.4.2 Power Modes

FLASHC_FLASH_PWR_CTL provides enable bits for the flash memory. Software can turn off all regions of the flash memory by setting the ENABLE and ENABLE_HV fields to 0.

The wakeup time of flash memory is 10 μs . Software should wait at least 10 μs before reading from flash after it is re-enabled through the FLASHC_FLASH_PWR_CTL register.

The flash controller provides functionality down to the Deep Sleep power mode. In the Deep Sleep power mode, the following flash controller information is retained:

- The retention MMIO registers (listed in Table 14-2)
- The cache data structure

Note that buffer information (in the AHB-Lite buffer interfaces and in the synchronization logic) is not retained. Losing buffer information after Deep Sleep transition has limited performance impact.

14.1.4.3 CPU Caches

The flash controller provides 8 kB caches for both the CM0+ and CM4 CPUs. Each cache is a four-way set associative with a least recently used (LRU) replacement scheme. These caches can be enabled/disabled through the FLASHC_CM0/4_CA_CTL.CA_EN registers.

Register	Bit Field and Bit Name	Description
CM0_CA_CTL[32:0]	CA_EN[31]	Cache enable: 0: Disabled 1: Enabled
CM4_CA_CTL[32:0]	CA_EN[31]	Cache enable: 0: Disabled 1: Enabled

The cache supports faster flash memory reads when enabled. On a read transfer “miss”, however, a normal flash controller access will occur.

Cache Prefetch. The caches support pre-fetching through the CM0/4_CA_CTL.PREF_EN register.

Table 14-3. CM0/4 Cache Control Prefetch Enable Register Values and Bit Field

Register	Bit Field and Bit Name	Description
CM0_CA_CTL[32:0]	PREF_EN[30]	Prefetch enable: 0: Disabled 1: Enabled
CM4_CA_CTL[32:0]	PREF_EN[30]	Prefetch enable: 0: Disabled 1: Enabled

If prefetch is enabled, a cache miss results in a 16 B refill for the missing data and a 16 B prefetch for the next sequential data. The prefetch data is stored in a temporary buffer and is only copied to the cache when a read transfer “misses” and requires that data.

sector are never accessed or can disable prefetch during a flash write/erase operation.

14.1.5 Read While Write (RWW) Support

The PSoC 6 MCU supports read operations on one area while programming/erasing in another area. This is implemented to support firmware upgrades and parallel tasks in the dual-core system. The application flash contains eight RWW sectors, UFLASH0 to UFLASH7, each 256KB in size. The AUXFlash and SFlash are additional RWW sectors apart from the main flash.

The RWW feature is available between sectors – you can read/execute from one sector while there is an ongoing write/erase operation in another sector. However, when the code execution/read is in the last 16 bytes of a given sector (say sector 0) and the flash write/erase operation is in the next sector (sector 1), an RWW violation may occur if prefetch is enabled. This is because prefetch will fetch the next 16 bytes of data, which is part of sector 1 while a write operation is underway in the same sector. This will result in a fault and should be considered during firmware design. Firmware can be designed to place dead code in the last 16 bytes of every sector making sure the last 16 bytes of a

14.2 Flash Memory Programming

14.2.1 Features


- SROM API library for flash management through system calls such as Program Row and Erase Flash
- System calls can be performed using CM0+, CM4, or DAP

14.2.2 Architecture

Flash programming operations are implemented as system calls. System calls are executed out of SROM in Protection Context 0. System calls are executed inside CM0+ NMI. The system call interface makes use of IPC to initiate an NMI to CM0+.

System calls can be performed by CM0+, CM4, or DAP. Each of them have a reserved IPC structure (used as a mailbox) through which they can request CM0+ to perform a system call. Each one acquires the specific mailbox, writes the opcode and argument to the data field of the mailbox, and notifies a dedicated IPC interrupt structure. This results in an NMI interrupt in CM0+. The following diagram illustrates the system call interface using IPC.

Figure 14-3. System Call Interface Using IPC



The PSoC 6 MCU's IPC component carries only a single 32-bit argument. This argument is either a pointer to SRAM or a formatted opcode or argument value that cannot be a valid SRAM address. The encoding used for DAP and the CM4 or CM0+ is slightly different.

DAP. If (opcode + argument) is less than or equal to 31 bits, store them in the data field and set the LSb of the data field as '1'. Upon completion of the call, a return value is passed in the IPC data register. For calls that need more argument data, the data field is a pointer to a structure in SRAM (aligned on a word boundary) that has the opcode and the argument. So it is a pointer if and only if the LSb is 0.

CM4 or CM0+. A pointer is always used to structure SRAM. Commands that are issued as a single word by DAP can still be issued by CM0+ or CM4, but use an SRAM structure instead.

The NMI interrupt handler for system calls works as follows.

- If the ROM boot process code is not initialized in the protection state (PROTECTION is still at its default/reset value UNKNOWN), the NMI calls have no effect and the handler returns.
- A jump table is used to point to the code in ROM or flash. This jump table is located in ROM or flash (as configured in SFlash).

The IPC mechanism is used to return the result of the system call. Two factors must be considered.

- The result is to be passed in SRAM: CM0+ writes the result in SRAM and releases the IPC structure. The requester knows that the result is ready from the RELEASE interrupt.
- The result is scalar (≤ 32 bits) and there is no SRAM to pass the result: in this case, the CM0+ writes the result to the data field of the IPC structure and releases it. The requester can read the data when the IPC structure lock is released. The requester polls the IPC structure to know when it is released.

External programmers program the PSoC 6 MCU flash memory using the JTAG or SWD protocol by sending the commands to the DAP. The programming sequence for PSoC 6 MCUs with an external programmer is given in [PSoC 6 MCU Programming Specifications](#). Flash memory can also be programmed by the CM4/CM0+ CPU by accessing the IPC interface. This type of programming is typically used to update a portion of the flash memory as part of a bootloader operation, or other application requirement, such as updating a lookup table stored in the flash memory. All write operations to flash memory, whether from the DAP or from the CPU, are done through the CM0+.

15. eFuse Memory



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - eFuse
- Application notes
- Code examples

The eFuse memory consists of a set of eFuse bits. When an eFuse bit is programmed, or “blown”, its value can never be changed. Some of the eFuse bits are used to store various unchanging device parameters, including critical device factory trim settings, device life cycle stages (see the [“Secure Boot” chapter on page 134](#)), DAP security settings, and encryption keys.

The eFuse memory is not directly accessible in the PSoC 64 devices, by the CM4 application. It is used to store hash codes and other security information.

15.1 Features

The PSoC 6 MCU eFuses have the following features:

- A total of 1024 eFuse bits. All bits are reserved by the Secure Boot system and are not available to the user.
- The eFuse bits are programmed one at a time, in a manufacturing environment. The eFuse bits cannot be programmed in the field.
- Multiple eFuses can be read at the bit or byte level through a PDL API function call or an SROM call. An unblown eFuse reads as logic 0 and a blown eFuse reads as logic 1. There are no hardware connections from eFuse bits to elsewhere in the device.
- SROM system calls are available to program and read eFuses. See the [Nonvolatile Memory chapter on page 127](#). For detailed information on programming eFuses, see the [PSoC 6 MCU Programming Specifications](#).

15.2 Architecture

The PSoC 6 MCU eFuses can be programmed only when V_{DDIO0} is at 2.5 V. Since eFuses are programmed during provisioning, V_{DDIO0} must be at 2.5 V during that time as well. For more information, see the [PSoC 6 MCU Programming Specifications](#).

Table 15-1 shows the usage of the PSoC 6 MCU eFuse bytes.

Table 15-1. PSoC 6 MCU eFuse Byte Assignments

Offset	No. of Bytes	Name	Description
0	20	Reserved	Reserved for PSoC 6 MCU system usage
20	16	SECURE_CMACE	Secure objects 128-bit CMACE
36	2	FLASH_BOOT_SIZE	Flash boot image size
38	1	SECURE_CMACE_ZEROS	Number of zeros in SECURE_CMACE
39	2	DEAD_ACCESS_RESTRICT	Access restrictions in DEAD life cycle stage
41	2	SECURE_ACCESS_RESTRICT	Access restrictions in SECURE life cycle stage
43	1	LIFECYCLE_STAGE	Life cycle state
44	8	Reserved	Reserved for PSoC 6 MCU system usage
52	12	Unused	Not used
64	64	CUSTOM_DATA	Custom data
64	17	ASSET_HASH	The HASH and HASH ZEROS from objects filled in TOC1 except DIE_ID and Trim Table – Secure Flashboot image and Cypress public key.
81	13	SECURE_HASH2	12 bytes of HASH and HASH zero count from device identity – device public and private key
94	17	SECURE_HASH3	The HASH and HASH zero count from provisioning data objects filled in TOC3 – Cypress Bootloader image and provisioning JWT packet.
111	17	UDS	Unique Device Secret and Zero count.

16. “Secure Boot”



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation
- Application notes
- Code examples

The PSoC 64 family has a fully integrated security system built into the device. There is no need for the customer to write a custom secure boot or secure bootloader. The multi-stage boot process requires that each step must be signed by an authorized source. The PSoC 64 family provides a means to implement a full Root-of-Trust (RoT) secured system. This RoT is initially owned by Cypress/Infineon and is transferred to the user/OEM during the provisioning process, using a Cypress signed JSON Web Token (JWT). Provisioning uses the SWD interface to communicate the signed JWTs to the device, which include transfer of security keys, debug policies, and so on. Refer to the [PSoC 64 Secure MCU Secure Boot SDK User Guide](#) for more information on provisioning and format of JWT packets.

The PSoC 64 family devices contain two CPUs, a CM0+ and a CM4. The CM0+ is designated the secure processor and performs all security policy evaluations and system calls. The CM4 executes only the user (unsecure) application. Protection units are used to divide and isolate the system memory resources, such as Flash and SRAM between the secure and unsecure CPUs. Any register access that may compromise the security of the device is protected by one of the many PPU (Peripheral Protection Units). This includes registers that control Flash writes, system clock tree, protection unit configuration registers, power levels, etc.

Basic Definitions

Before continuing, you must understand some terms that will be used throughout this document. Many of these terms will be discussed in more detail within this section.

- **Chain of Trust (CoT):** The root of trust begins with the Cypress code residing in ROM, which cannot be altered. Chain of Trust is established by validating the blocks of software starting from the root of trust located in ROM.
- **Code Signing:** The process of calculating a hash of the code binary and encrypting the hash with a private key.
- **Debug Access Port (DAP):** Interface between an external debugger/programmer and PSoC 6 MCU for programming and debugging. This allows connection to one of three access ports (AP), CM0_AP, CM4_AP, and System_AP. The System_AP can only access SRAM, Flash, and MMIOs, not the CPUs.
- **Entrance Exam:** A test used just prior to provisioning to verify the device has not been tampered with and contains ONLY the code and data that was programmed at the device factory.
- **Digest:** A digest is computed by applying a hash function on a data set, such as the binary of an application. This digest is then encrypted with a private key to form a digital signature.
- **Digital Signature:** Encrypting of the digest (hash of a data set). For example, the encrypted hash of the user application.
- **Flash (User):** This is the flash memory that is used to store your application code.
- **Flashboot:** This is part of the boot system that performs three basic tasks:
 - ❑ Sets up the debug port based on the lifecycle stage.
 - ❑ Validates the user application before executing it.
 - ❑ Validate and execute the device policy.

- **HAL:** (Hardware Abstraction Layer) High level driver library that is independent of underlying hardware.
- **Hash:** A crypto algorithm that generates a repeatable but unique signature for a given block of data. This function is non-reversible.
- **HSM:** (Hardware Security Module) This is a physical computing or programming device that protects and manages digital keys, it can perform digital software signing and encryption as well as other cryptographic functions. In the context of the PSoC 64 Secure MCU, the HSM is a device programming engine placed in a physically secure facility.
- **IPC:** Inter-Processor Communication hardware used to facilitate communication between the two CPU cores.
- **JSON:** An open standard file format to store value pairs and array data types.
- **JWT:** JSON Web Token format used to transmit and store provisioning information to the device.
- **Lifecycle:** This is the security mode in which the device is operating.
- **Memory Protection Unit (MPU):** Used to isolate memory sections from different bus masters.
- **MMIO:** Memory-Mapped Input/Output, usually refers to registers that control the hardware I/O.
- **PC:** Protection Context, although PC most often refers to a Program Counter, in this document it refers to the Protection Context state. This parameter allows each bus master a security state level from 0 to 15. A bus master can be assigned a PC value that stays static or that is changed during application execution. PC provides a more precise way of applying memory restrictions. PC=0 is a special case which allows any bus master to have full access to the entire memory space including registers. The PC state works together with protection units.
- **PDL:** (Peripheral Driver Library) Low level driver library that supports all hardware blocks for the PSoC 64. It is provided as part of ModusToolbox.
- **Peripheral Protection Unit (PPU):** PPU's are used to restrict access to a peripheral or set of peripherals to only one or a specific set of bus masters.
- **Protection Units:** These hardware blocks are used to limit bus master access to memory (SRAM, ROM, flash) or hardware (peripheral) registers. They include MPU, PPU, and Shared Memory Protection Unit (SMPU).
- **Provisioning:** The process by which keys, policies and secrets are injected into the device. Once provisioned, the device can be accessed or modified only with the keys injected adhering to the relevant policies.
- **PSA:** (Platform Security Architecture) The PSoC 64 devices include a subset of the PSA Cryptographic API (Crypto API) that provides a portable interface to cryptographic operations. These functions are performed by the secured CM0+ CPU but are accessible by the non-secure CM4 CPU for user applications.
- **Public-Key Cryptography (PKC):** Otherwise known as asymmetrical cryptography. Public-key cryptography is an encryption technique that uses a paired public and private key (or asymmetric key) algorithm for secure data. It is used to secure a message or block of data. The private key is used to encrypt data and must be kept secure, and the public key is used to decrypt but can be disseminated widely.
 - **Public Key:** The public key can be shared, but it should be authenticated or secure so it cannot be modified.
 - **Private Key:** The private key must be kept in a secure location, so it cannot be viewed or stolen. It is used to encrypt a block of data that will be decrypted using an associated public key.
- **RMA:** Return Merchandise Authorization
- **Rollback Counter:** Special counter accessed by secure boot code that holds the value of the latest valid image and used in an anti-rollback protection mechanism. The goal of anti-rollback protection is to prevent downgrading the device to an older version of its software that has been deprecated due to security concerns.
- **SALT:** A salt is a random number or data that is added to data to be hashed so that devices with identical data, will have a different hash value.
- **Security Policy:** This is the set of rules that the designer imposes to determine what resources are protected from outside tampering or between the internal CPUs.
- **Serial Memory Interface (SMIF):** A SPI (Serial Peripheral Interface) communication interface to serial memory devices, including NOR Flash, SRAM, and non-volatile SRAM.
- **SFlash:** Supervisor flash memory. This memory partition in flash contains several areas that include system trim values, flash boot executable code, public key storage, etc. After the device transitions into a secure mode, it can no longer be changed.
- **SHA-256:** A cryptographic hash algorithm used to create a signature for a block of data or code. This hash algorithm produces a 256-bit unique signature of the data no matter the size of the data block.

- **Shared Memory Protection Unit (SMPU):** SMPUs are used to allow access to a specific memory space (flash, SRAM, or registers) to only one or a specific set of bus masters.
- **SYS-AP:** System Access Port which is part of the Serial Wire Interface for programming and debugging. This port is used for transferring the provisioning packets to the device.
- **System Calls:** Functions such as flash write functions that are executed by the Arm® Cortex® M0+ CPU (CM0+) from ROM.
- **Table of Contents 1: (TOC1)** This is an area in SFlash that is used to store pointers to the trim values, flash boot entry points, etc. It is used only by boot code in ROM and is not editable by the designer.
- **Table of Contents 2: (TOC2)** This table contains only a fixed value and the CRC of itself. It is of little value in PSoC 64 devices, but may be used for future derivations.
- **Trusted Function:** Secure functions that run on the Secure CPU. These functions may be created either by the chip vendor or the OEM.
- **TFM:** (Trusted Firmware-M) It is a Trusted Execution Environment for the CM0+ CPU which provide secure runtime services like secure storage, cryptography, and attestation.
- **UDS:** (Unique Device Secret) A code that is unique to each device. Generated by hashing data in a device that is different for each device manufactured.

16.1 Features

The PSoC 64 MCU provides the following device security features:

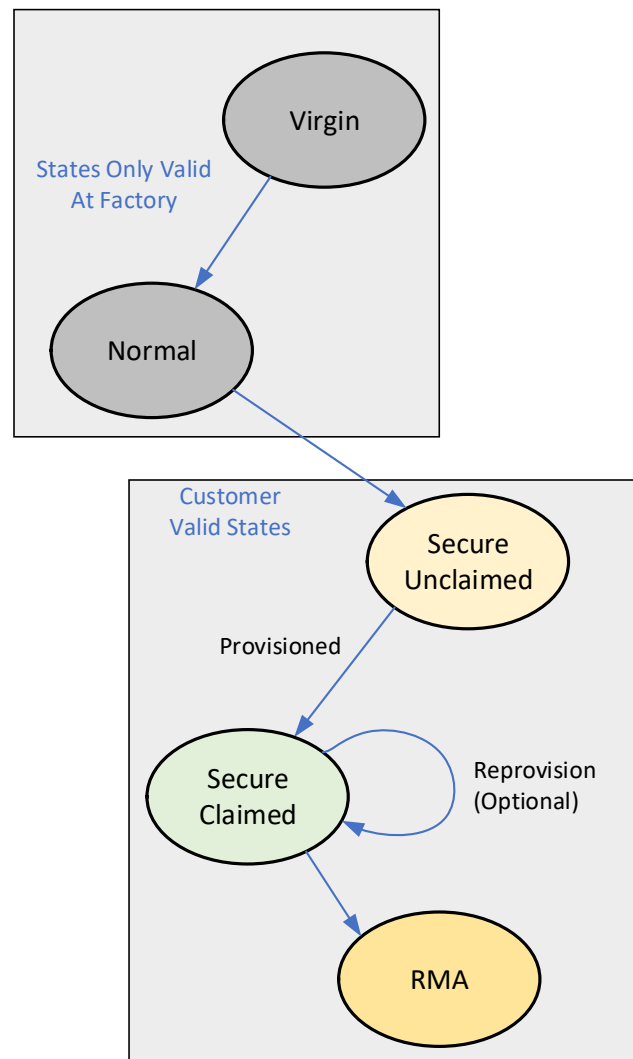
- Built-in immutable key storage area that can be loaded with required OEM keys and other user keys.
- Secure method to transfer Root-of-Trust to OEM public/private key.
- Policy interpreter that allows the user to program key usage, debug access, and hardware/firmware wounding.
- Protection units that provide protection and isolation of flash, SRAM, and registers between the secure (CM0+) and nonsecure (CM4) CPU.
- Cryptographic hardware block that provides encryption and decryption of data and code.

16.2 Architecture

16.2.1 Life Cycle Stages and Protection States

PSoC 64 MCUs have configurable, nonvolatile life cycle stages. Life cycle stages follow a strict, irreversible progression governed by writing to eFuse - a 1024-bit nonvolatile memory area with each bit being one time programmable (OTP). The eFuse is used to store Life Cycle Stage information, security configuration, and several HASH signatures to validate the secure trim values, flashboot code, secure data structures, etc. The secure boot system uses all the eFuse bits and they are not usable by the users main (non-secure) application. See the [eFuse Memory chapter on page 132](#) for more details of the eFuse.

Figure 16-1. PSoC 6 MCU Life Cycle Stage Transitions



The life cycle stages can move only from one stage to the next and cannot be reversed. For example, when the life cycle has moved from **Secure Unclaimed** to **Secure Claimed**, it cannot be moved back to **Secure Unclaimed**. It may remain in the **Secure Claimed** mode or move to **RMA** (Returned Merchandise Authorization) by the OEM. Once the device has been moved to the **RMA** state, it can never be changed to the **Secure Unclaimed** or **Secure Claimed** state which makes it unusable in a production product. The OEM would only advance to the **RMA** state if the device needed to be evaluated for failure analysis.

16.2.1.1 Usage of hashing

The PSoC 64 devices use hashing extensively to authenticate data and code stored in the device. Below is a summary of each of the hashes generated and stored in the device. The code/data that cannot change throughout the life of the device is stored directly in eFuse. Hashes for items that can change are encrypted with the device private key and stored in flash memory.

All hashes are generated with the SHA-256 algorithm, but the ones stored in eFuse, only the most significant 128 bits are stored with additional 8-bits stored that contains the number of zeros in the 128-bits of hash. Hashes stored in flash use the full 256-bits.

FACTORY_HASH: This hash covers objects from TOC1 which includes Trim values, Die ID, FlashBoot code, etc. It is created during the transition from Virgin to Normal mode and verified by ROMBoot before the transition from Normal to Secure mode. After the transition to Secure mode, it is no longer used. The FACTORY_HASH will be different for each part because it contains both the unique trim values and unique die ID.

ASSET_HASH: This hash mainly covers FlashBoot and the CY Public Key. It is generated during the transition from Virgin to Normal mode. This hash is used by the entrance exam to verify the device has not been tampered with prior to provisioning. The ASSET_HASH is the same for all parts with the same version and part number.

SECURE_HASH: This hash covers objects from TOC1 and TOC2. It is generated during the transition from Normal to Secure mode. Once the device is in Secure mode, it is used to validate FlashBoot by ROMBoot before FlashBoot is executed.

SECURE_HASH2: This hash covers the Device and Group Private/Public keys. It is created and stored at the end of the identity creation step when the device changes from the Secure_Unclaimed state to the Secure_Claimed state. This hash is used to validate the keys and certificates stored on the device.

SECURE_HASH3: This hash is for the OEM public key, Product ID, and optionally for the bootloader and provisioned packets if the "reprovisioning" policy is not enabled. It is calculated and written in eFuse at the end of the provisioning when the "complete" flag is set true.

SECURE_HASH4: This is an optional hash for the Bootloader and provisioned packets if the reprovisioning policy is set to true. It is calculated and signed by the Device Private key and stored in flash at the end of the provisioning.

16.2.1.2 VIRGIN Stage

This stage is valid only at the Cypress factory. When a device is powered up for the first time at Cypress/Infineon factory, it will be in VIRGIN mode. The following steps take place while the device is in VIRGIN mode before the transition to NORMAL mode:

1. Device Testing
2. Trim (calibration) values stored in SFlash
3. Unique ID (DIE_ID) written to eFuse
4. Flashboot written in SFlash and Main Flash
5. Cypress public key is stored in the device.
6. Generate FACTORY_HASH based on the TOC1 table and write to eFuse
7. Generate ASSET_HASH which includes FlashBoot and Cy Pub Key, then write it to eFuse
8. Transition to NORMAL mode by writing NORMAL flag in eFuse

16.2.1.3 NORMAL Stage

The Normal life cycle stage is valid only in the Cypress factory. The device hardware can be fully tested in this mode since there are no code restrictions. Before the device leaves the factory, it is moved to Secure mode. Just prior to moving a device to Secure mode, the SECURE_HASH is created based on elements in TOC1 and TOC2 and is stored in eFuse. The SECURE_HASH will be used once the device is in Secure mode during the boot up procedure. The final step to transition to Secure mode is to set the SECURE eFuse bit. After this bit is set and the device is rebooted, it can never revert back to Normal or Virgin modes.

16.2.1.4 SECURE UNCLAIMED

There are two sub-states in the SECURE state, SECURE_UNCLAIMED and SECURE_CLAIMED. The SECURE_UNCLAIMED state is entered when the device transitions from NORMAL to SECURE. The device leaves the Cypress factory in this state. When the device is in the SECURE_UNCLAIMED mode, both the SECURE_HASH2 and SECURE_HASH3 are equal to 0 and have not been stored in eFuse.

When the device leaves the Cypress factory, it is in the SECURE (Unclaimed) mode with the following configuration:

- Secure life-cycle mode with HRoT (Hardware Root of Trust) (PC=0) protection settings applied
- Contains CY Public key
- CM0+ AP disabled
- CM4 core available to run code starting at 0x1000_0000 for basic hardware (PCB) testing.
- CM4 AP enabled
- Secure Hash 2 and Secure Hash 3 values are all zeros (un programmed)

Devices directly from Cypress are "owned" by Cypress/Infineon (SECURE_UNCLAIMED) until they are provisioned by the OEM. To transfer the ownership to the OEM, a JWT (JSON Web Token) with the OEM's public key is signed by Cypress, then loaded into the device. This transfers the root of trust (RoT) from Cypress to the OEM. From this point on, only code, keys, and policies signed by the OEM can execute on the device. The policies, keys, and OEM boot code (first code executed after Flashboot) are hashed and the hash stored in eFuse. This puts the device in SECURE_CLAIMED mode.

The Serial Wire Debugger (SWD) interface is used to transfer provisioning information to the device, as well as for firmware programming and debug. JWT provisioning structures are used to transfer policies and keys to the PSoC 64 devices. The JWTs need to be signed by the proper authority (OEM) to be accepted by the device. The actual format of the JWTs is defined in the [PSoC 64 Secure MCU Secure Boot SDK User Guide](#).

The Sys-AP interface will be open, but with restricted access to any CM0+ area such as Flashboot, SROM, and stack space used by that code. The CM4 debug port defaults as open, and the CM0+ debug port will be closed by default. While in the SECURE_UNCLAIMED state, developers may download and debug code only in the CM4 for evaluation and test purposes without transferring the ownership to the OEM. Unsecure CM4 test code may be downloaded into user flash space with the CM4's vector table starting at the default location of 0x1000_0000. When the device boots in this mode, it will check for a valid reset vector and stack location only. If these vectors are valid, it will start the CM4 at the location pointed to by the reset vector. The CM4 will have access to only nonsecure memory and register space. Memory areas that contain the private key, control writing and reading Flash and eFuse, protection unit configuration, and so on are all secured with protection units, so they cannot be modified by unsecure firmware. Systems calls provide access to flash memory writing, but only in user flash area that is not protected.

16.2.1.5 SECURE CLAIMED

After the provisioning process is complete, the device is now in the SECURE_CLAIMED mode and owned by the OEM and will only run authenticated code from the OEM. It will no longer recognize application code signed by Cypress or anyone else for that matter. During provisioning the OEM has the option to enable the device to be reprovisioned, so that the bootloader, keys, and policies may be changed. When the device is in the SECURE_CLAIMED mode, both the SECURE_HASH2 and SECURE_HASH3 will be non-zero and stored in eFuse.

When the OEM is ready to ship the product, the devices will most likely be provisioned so that the debug ports are disabled for security reasons. The debug port settings are just one of the several security parameters that can be configured with the policy file. A method in the user's code should be provided to transfer code upgrades into the code update staging area, if code updates are required. The default Cypress Secure Bootloader will evaluate the code in the staging area with the code that is currently in the execution area. If the code in the staging area has been validated and is more recent, it will be copied into the execution area, where it will be authenticated and executed.


16.2.1.6 Provisioning and Reprovisioning

The provisioning process of the PSoC 64 secure parts achieves the following goals:

- Transfer RoT from Cypress to the user/OEM
- Transfer of secure assets such as keys, certificates, and secure bootloader into the device
- Set up policies that govern the debug and boot-up behavior of the device

- Device generates its own Private and Public Key pair
- This is achieved by sending a series of signed JSON Web Tokens (JWTs) to the part. The high-level flow of provisioning is shown in the following diagram.

Figure 16-2. Provisioning Scheme

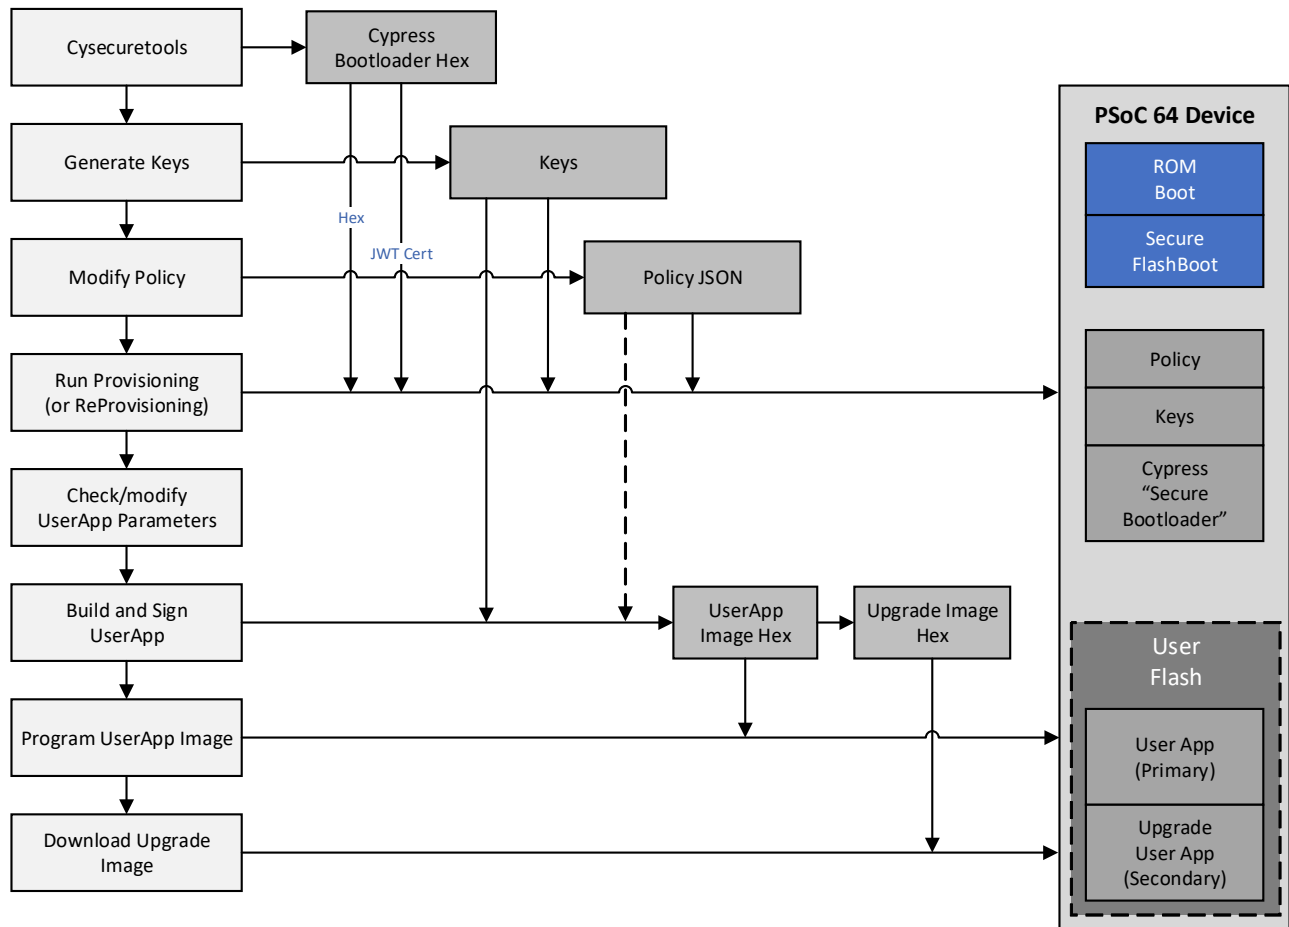


Provisioning may occur at any time in the `SECURE_UNCLAIMED` mode with the `ProvisionKeysAndPolicies` system call. Production provisioning should be performed in a secure facility. The OEM may choose to allow the device to be re-provisioned at a later time, or not allow the device provisioning properties to be altered. Once the device provisioning allows re-provisioning, it cannot be changed to disallow re-provisioning. Also, once a device is provisioned with the re-provisioning option disabled, the device can never be re-provisioned. The diagram below shows the flow for provisioning and programming of the device.

The OEM application should be programmed into the device before the transition from `SECURE_UNCLAIMED` to `SECURE_CLAIMED` if the device is provisioned for manufacturing and has its debug ports closed. If the device is provisioned for development with the debug ports open, the OEM code does not have to be loaded into the device at this time. The first code executed after Flashboot, such as Cypress "Secure Bootloader" needs to be programmed into the device before the transition to `SECURE_CLAIMED` mode, since it will be hashed and cannot be reprogrammed.

A sample of the provisioning policies in JWT format is provided as part of the PSoC 64 Secure MCU Secure Boot SDK User Guide. CySecureTools Python scripts are provided to perform the initial provisioning and as an example of how the communication is performed between the development computer and the PSoC 64 device. Cypress partners such as Data I/O, can provide the provisioning services and key management required for production, using a secure Hardware Security Module (HSM). Customers are free to implement their own HSM as well. A key will need to be generated by the customer and signed by Cypress to perform the initial RoT transfer to the OEM.

Figure 16-3. OEM Provisioning and Programming Flow



The last step of the provisioning process is dependent on whether or not reprovisioning has been enabled in the policy. Enabling reprovisioning is achieved by setting one or more of the parameters in the "reprovision" policy to true in the policy json file. The reprovisioning parameters include "keys_and_policies" and "boot_loader".

Once provisioned, specific memory objects are hashed to secure the configuration. The objects hashed to create the SECURE_HASH3 are dependent on which if any of the reprovision params are set to true or false. The following is a list of objects that are hashed for SECURE_HASH3.

- OEM Pub Key
- Product ID
- Provisioned JWT (If reprovision->keys_and_policies is false)
- Chain of Trust JWT (If reprovision->keys_and_policies is false)
- Image Cert JWT (If reprovision->boot_loader is false)
- Bootloader image object included in "image_address" and "image_size" values in the "image_cert" provisioned packet (if reprovision->boot_loader is false)

The SECURE_HASH4 is only generated if one of more of the reprovision parameters is set to true. The following is the list of objects that may be included in the SECURE_HASH4 hash.

- Provisioned JWT (If reprovision->keys_and_policies is true)
- Chain of Trust JWT (If reprovision->keys_and_policies is true)
- Image Cert JWT (If reprovision->boot_loader is true)

- Bootloader image object included in "image_address" and "image_size" values in the "image_cert" provisioned packet (if reprovision->boot_loader is true)

If the SECURE_HASH4 is generated, it is signed by the Device Private Key (Internally generated during provisioning) and stored in Flash. SECURE_HASH3 and SECURE_HASH4 are validated each time the device comes out of reset. If the calculated hash does not match the stored hash the device will not continue to the user application and will loop forever in the Dead mode.

Because SECURE_HASH4 is placed in the flash it can be re-calculated and re-signed each time valid ProcessProvisionCmd request is used to reprovision the device.

Note: "reprovision" policy configuration cannot be changed after first provisioning. When device is originally provisioned with "reprovision" enabled, it cannot be reprovisioned with "reprovision" feature disabled.

The CM0+ (secure processor) is by default set to a Protection Context (PC) equal to 0 at reset. This is effectively super user mode in which the CM0+ can read, write, or execute all code and data. When the device finishes executing the boot code, it switches the PC to "1" before jumping to the bootloader or first executable from the OEM.

16.2.1.7 RMA

Customers can transition the device to the RMA mode (from SECURE_CLAIMED) when they want Cypress to perform failure analysis on the device. Before provisioning, the RMA section of the Debug policy should be updated in the provisioning JWT file. This policy should include any areas of Flash that may contain proprietary code or sensitive data. This area will be erased automatically when the device transitions into RMA mode. See the [PSoC 64 Secure MCU Secure Boot SDK User Guide](#) for more information on the format of the JWT file.

To place a device into the RMA mode, the internal code must execute the [TransitionToRMA](#) system call. A JWT packet that contains the DIE_ID that is signed by a specified installed public key is passed as a parameter to the system call. The OEM has the option to pre-generate this packet and store it in flash, or to provide a method to generate and transfer it to the device through an OEM defined interface. The TransitionToRMA system call is only valid when the device is in the SECURE_CLAIMED mode.

After the device is moved into RMA mode, when it boots, it does not attempt to execute any code in user flash; instead, it waits for the [OpenRMA](#) system call through the debug port. The OpenRMA system call requires the same signed packet as a parameter that was used to place it into the RMA mode in the first place. When the OEM sends the part to Cypress for failure analysis, the signed JWT packet that contains the device DIE_ID would also need to be supplied. The JWT packet will only work for the part with the corresponding DIE_ID.


In RMA mode, the chip accessibility is similar to that in the SECURE_UNCLAIMED mode with the following exceptions:

- Device can no longer be provisioned
- CM4 debug port will be open
- CM0p debug port will be closed
- Only CM4 can run code
- No system calls can be made by the CM4
- CM4 code vector table will start at 0x1000_0000

16.2.2 Boot Sequence

A high-level view of the boot sequence is as follows. ROM boot does the basic trimming and initializing of hardware. It validates that Flashboot has not been altered in any way before it jumps to execute it. Flashboot validates all the policies and the bootloader by verifying SECURE_HASH2, SECURE_HASH3 and SECURE_HASH4. SECURE_HASH4 is only validated if reprovisioning was enabled. Next, the bootloader determines if the user application is valid or whether it needs to be updated. After the user application is validated, the CM4 is enabled and starts to execute the application. The CM0+ stays in a low-power mode until a system call is executed or it is reset.

Figure 16-4. Boot Sequence

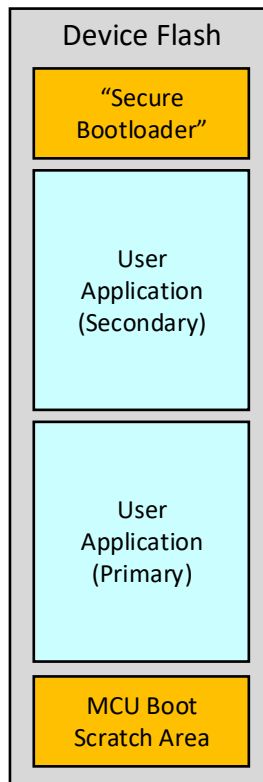


16.2.3 Cypress "Secure Bootloader"

The Cypress "Secure Bootloader" is a port of the industry-standard MCU Boot. The available Flash is divided into 4 sections, Bootloader (the actual code), Flash Scratch area, and two application areas, primary and secondary. The Primary and Secondary sections are of equal size and consume most of the Flash area. Originally the Primary and Secondary areas were referred to as Slot 0 and Slot 1 respectively. The bootloader itself is executed by the CM0+ CPU. The user (non-secure) application is only executed by the CM4.

The application code only executes in the Primary section. The Secondary section is used to hold the application firmware upgrade. The CM4 updates the Secondary section when updated firmware is received from a local host or via a wireless connection. It will not be verified until the device resets and the bootloader will execute. If the firmware is validated, it will be copied into the Primary section where it will be executed. The firmware update may be encrypted for transport and stored in the Secondary area as such. If the code in the Secondary location is encrypted, the bootloader will decrypt it prior to copying it to the Primary location. The boot upgrade policy determines if the firmware update will be encrypted and what key is used to decrypt it. The code in the Primary location is never encrypted.

Figure 16-5. Cypress Secure Bootloader



The user application (CM4) is responsible for transferring the updated firmware via whatever interface the application uses, into the Secondary memory location. Updates may come through a serial port, I2C, SD Card, WiFi, BLE, etc. The bootloader (CM0+) has nothing to do with the transfer of the upgraded firmware to the Secondary location.

16.2.4 Flash Security

PSoC 6 MCUs include a flexible flash-protection system that controls access to flash memory. This feature is designed to secure proprietary code, but it can also be used to protect against inadvertent writes to the bootloader or other protected areas of flash.

Flash protection is provided by a Shared Memory Protection Unit (SMPU). The SMPU is intended to distinguish between different protection contexts and to distinguish secure from non-secure accesses. The system function that performs flash programming first looks at the SMPU settings and will not allow to programming or erasing flash blocks protected by the SMPUs. All SMPUs have been configured or reserved during the boot process. The CM4 user code cannot modify any of these settings since they are protected with PPUs. For more details, see the [Protection Units chapter on page 71](#).

16.2.5 Hardware-Based Encryption

The PSoC 6 MCU has a cryptographic block (Crypto) that provides hardware implementation and acceleration of cryptographic functions. It implements symmetric key encryption and decryption, hashing, message authentication, random number generation (pseudo and true), and cyclic redundancy checking. It can work with internal as well as external memory. See the [Cryptographic Function Block \(Crypto\) chapter on page 113](#) and [Serial Memory Interface \(SMIF\) chapter on page 341](#) for more details.

The Crypto Block is accessible either via System calls or directly by the CM4, but not both at the same time. Care should be taken not to make a System call that requires Crypto functions while using the crypto block hardware directly from the CM4. System calls are executed by the secure processor (CM0+). They will always take priority and will block access to the crypto block from the CM4 or any other bus master until the system call is complete. If the CM4 attempts to access the crypto block when in use by the CM0+ because of a system call, or a system call takes control of the crypto block while the CM4 is using it, a hard fault will occur.

16.2.6 Dedicated Hardware

Some of the PSoC 64 hardware is reserved by the secure boot system and is not accessible from the user's (non-secure) application. Attempting to access memory or peripherals reserved by the secure boot system will cause the system to hard fault. Some of the flash and SRAM is used by the secure boot code as well. See the device datasheet for the specific memory map that documents the system memory map. The following sections document the other resources reserved by the secure boot system.

16.2.6.1 IPC Channels and IPC Interrupts

The PSoC 64 devices have 16 IPC channels. Channels 0 thru 9 are used or reserved by the system. The table below shows the usage of each of these channels.

Table 16-1. IPC Channels Usage

IPC Channels	Name	Usage
0	CY_IPC_CHAN_SYSCALL_CM0	System calls for the CM0 processor. PyOCD and OpenOCD.
1	CY_IPC_CHAN_SYSCALL_CM4	System calls for the CM4 processor. PyOCD and OpenOCD.
2	CY_IPC_CHAN_SYSCALL_DAP	System calls for the DAP (programming, OpenOCD)
3	CY_IPC_CHAN_SEMA	Reserved for CM4 -> CM0+ RPC: not used by Secure FW and Debugger.
4	PRA	Used for Protected Register Access
5	CY_IPC_CHAN_CYPPIPE_EP0	Reserved for CM0+ -> CM4 RPC: not used by Secure FW and Debugger.
6	CY_IPC_CHAN_CYPPIPE_EP1	Reserved for DAP -> CM4 RPC
7	Reserved	Reserved for future use

Table 16-1. IPC Channels Usage

IPC Channels	Name	Usage
8	TFM	Reserved PSoC 64
9	TFM	Reserved PSoC 64
10	Unused	-
11	Unused	-
12	Unused	-
13	Unused	-
14	Unused	-
15	Unused	-

Table 16-2. IPC Interrupt Usage

IPC Interrupt	Usage
0,1,2	SysCalls (#0 for CM0+, #1 for CM4, and #2 for DAP)
3,4	CM0/CM4 Pipe communications
5	PRA Driver
6	TFM
7	SysPm
8	TFM
9-15	Unused

16.2.6.2 Protection Context

Although Protection Context (PC) is not one specific block of hardware, it is used for all the protection units such as SMPU and PPUs. Multiple protection contexts are used to isolate the different security levels within the PSoC 64. The CM0+ makes use of several of them during the boot sequence, bootloading, system calls, etc. Protection context 6 is used for the user application code that runs on the CM4 CPU. The non-secure (CM4) application cannot change this value. The following table details how each of the protection contexts is used.

Table 16-3. Protection Contexts

PC	Secure	Description	CPU	Privileged	User
0	Yes	Hardware Root-of-Trust, SysCall access only	CM0+	SysCall	<not used>
1	Yes	Cypress Trusted Environment & Functions	CM0+	BootLoader/SPM/Drivers	Trusted Functions
2	Yes	3rd Party Trusted Functions (optional*)	CM0+	BootLoader/SPM/Drivers	Trusted Functions
3	Yes	3rd Party Trusted Functions (optional*)	CM0+	BootLoader/SPM/Drivers	Trusted Functions
4	Yes	3rd Party Trusted Functions (optional*)	CM0+	BootLoader/SPM/Drivers	Trusted Functions
5	No	<reserved>	-	-	-
6	No	Application	CM4	OperatingSystem/Drivers	Application
7	No	DAP SYSAP	DAP	<not used>	Debugger/Tester

***Note:** Future releases of the device firmware may allow for 3rd party trusted functions, not available at this time.

16.2.6.3 SMPUs

SMPUs are used to protect areas of memory (flash, sflash, and ROM) between multiple bus masters. The secure boot system reserves all the available 16 SMPUs.

16.2.6.4 Protected Register Access (PRA)

PSoC 64 devices use Peripheral Protection Units (PPU) to protect the system from invalid configurations that could potentially cause the system to be unstable or indirectly allow access to registers and memory that are protected. In some cases, the unsecure application running on the CM4 still must access these registers for several system functions. The Protected Register Access (PRA) driver provides a secure method in which to access these registers, but at the same time protect the system from register combinations that could compromise the security of the device.

The PRA driver implements a bridge between the CM4 and the CM0+ via IPC. Several PDL drivers that access the protected registers have been modified specifically for the PSoC 64 to allow controlled access to these registers. Some of these registers do not pose any security risk, but due to the coarseness of the fixed protection register blocks, may have protected the register. In this case the PRA library provides a virtual direct access via an IPC channel. The functions that access these registers have been updated to automatically use the PRA driver for access.

In the case where more control is required to protect a register due to specific values or sequence of operations, the entire function runs on the secure CM0+. When the non-secure user application calls one of these functions, the parameters are passed to the CM0+ through the IPC channel and the code is executed by the CM0+.

- The PRA driver is not intended to be used by developers directly.
- The supplied PDL and HAL drivers have incorporated the PRA functions so that little or no code changes are required.
- How the PRA deals with registers can be summarized in 6 cases.
 - **DIRECT:** This means that the CM4 has full direct access to this register and the PRA driver is not used to access this register. This is the case for most registers that are not accessed by the listed driver libraries.
 - **UNCONSTRAINED:** Register that fall under this case have been protected because of the coarseness of the fixed PPU that secures the block containing that register. The register itself should have full access by the CM4 CPU. The PRA driver passes any read/write data directly from/to this register without any restrictions.
 - **REGISTER_POLICY:** Data written to a register in this group will be checked for proper range and bit settings. If the data is outside this range an error will be generated and returned to the CM4.
 - **SECURE_ONLY:** Register in this group can only be accessed with a PDL function, that use the PRA driver. This PDL function will be executed on the CM0+. The CM4 function call will use the IPC interface to transfer the parameters of the function call to the CM0+ where it will be executed by the CM0+ with parameter checking and proper sequence relative to writing to other registers. Any errors will be passed back to the CM4 via the IPC interface.
 - **FUNCTION_POLICY:** Register that fall under this classification are primarily those used to set the frequency and clock source for CLK_HF0 (CPU clock). An entire configuration of CLK_HF0 is evaluated for correctness before setting any registers. Also, wait states and system voltages are evaluated to verify that the device is operating per specification.
 - **NO ACCESS:** The few registers that fall into this category are inaccessible by the CM4 in any way. These registers have been protected during the boot process and cannot be modified by the CM4 with or without the PRA driver. Any attempt to access these registers will cause a page fault.

The register TRM will tag any of the registers that are affected by the PRA driver. This will appear at the top of the register definition with a "Protected Register Access:" label. The type of PRA used for that register will be shown after the label. For example, if a register is of type "SECURE_ONLY", the label will show "Protected Register Access: SECURE_ONLY".

Most PDL drivers are not affected or use the PRA driver. Only the following PDL driver libraries are affected by this driver:

- **LVD (Low-Voltage-Detect)**
 - UNCONSTRAINED register policy is used for the (LVD).
- **GPIO (General Purpose Input/Output)**
 - The GPIO ports are only affected if that port is used for an external clock such as an ECO or WCO. When an external clock or oscillator is enabled, the remaining pins on that port that aren't used by the clock or oscillator, are affected by the PRA driver. This means that any access to these pins in firmware incur an access penalty because the PRA driver must be invoked. For example, if the WCO is used on port pin P0[0] and P0[1], access to pins P0[7:2] in firmware will be significantly slower. This is because the CM4 CPU cannot have direct access to these pins and must communicate with the CM0+ secure CPU to read and write to these pins. The PRA driver takes care of this automatically when using the standard PDL GPIO library. Any attempt to access this port directly without the use of the PDL GPIO library will cause a CPU fault. This restriction also includes the HSIOM register associated with the port. The only ports affected by this is Port 0 (External Clock and WCO) and Port 12 (External Crystal Oscillator), and only if the security policy

defines an oscillator or external clock on that port. This is done to make sure that unsecure software cannot directly affect the clock that affects the operation of the secure CPU (CM0+).

- SysLib (System Library)
 - REGISTER_POLICY is used for the Cy_SysLib_ClearFlashCacheAndBuffer() and Cy_SysLib_ClearResetReason() functions.
 - The Cy_SysLib_SetWaitStates() implementation was made empty for PSoC64. The wait states are automatically updated on any clock and power configuration changes.
- SysClk (System Clock)
 - Access to SysClk driver is via the FUNCTION_POLICY
- SysPm (System Power Management)
 - When calling the function Cy_SysPm_BuckSetVoltage1() it uses the PRA driver to enable the buck register.
- SysTick (ARM System Timer)
 - SysTick has the option of having an alternate clock source. This alternate clock source is common for both the CM4 and the CM0+. The option of selecting this clock source is up to the each of the CPUs which is not a security issue. Selecting the source of the alternate clock source may be a security issue. If the CM4 was able to set the alternate systick clock source it could affect the operation of the CM0+ code execution. All other systick functions are not affected by the PRA.
- WDT (Watch Dog Timer)
 - Access policy is UNCONSTRAINED for the WDT registers.
- Flash (Flash System Routine)
 - The registers associated with programming of Flash and eFuses are restricted. eFuses cannot be programmed directly by the CM4. To access Flash programming, the CM4 must use standard system calls.

A list of affected register and functions is provided in the cy_pra.h file located at PROJECT_DIRECTORY\libs\mtb-pdl-cat1\include. Each register name has a prefix of "CY_PRA_INDX" or "CY_PRA_INDEX". Driver functions that are executed by the Secure CM0+ instead of the CM4, such as defined by the "FUNCTION_POLICY", or "SECURE_ONLY" are listed in the same header file. These functions will have a prefix "CY_PRA_driver_FUNC" where "driver" is the driver PLD library. For example, "CY_PRA_CLK_FUNC_ECO_ENABLE" is for the "Cy_SysClk_EcoEnable" function.

The execution time of the functions that access the protected registers is increased on the PSoC 64 devices. This is because the affected library function is merely a wrapper that sends the function parameters to the CM0+ and is executed by the CM0+. The parameters are transferred to the CM0+ using an IPC channel and shared SRAM. The result is then passed back through the IPC channel to the CM4. For more complicated functions the additional overhead may not be notable, but for simpler functions such as those that just set register values the overhead will be more noticeable.

- The access to the protected register may take around by 20 times longer compared to unprotected one.
- The initial device configuration based on the device configuration depends on actual configuration, but may take up to 40 times longer.
- The transition Active to DeepSleep may take 2 times longer.

For more information about the PRA driver and how it works, see the PSoC 6 Peripheral Driver Library documentation for the PRA driver.

16.3 System Calls

16.3.1 Implementation

16.3.1.1 System Call via CM0+ or CM4

System calls can be made from the CM0+ or CM4 during code execution. CM0+ or CM4 should acquire the IPC_STRUCT reserved for them and provide arguments in either of the methods described above and notify IPC interrupt 0 to trigger a system call. Although the CM0+ is the security processor, it can perform a system call to itself. However, after the bootup sequence has completed and the CM4 is up and running, the CM0+ will not make system calls, but will only execute them.

16.3.1.2 System Call via DAP

If the device is acquired, then the boot ROM enters "busy-wait loop" and waits for commands issued by the DAP. For a detailed description on acquiring the device see the [PSoC 6 MCU Programming Specifications](#).

16.3.1.3 Exiting from a System Call

When the API operation is complete, CM0+ will release the IPC structure that initiated the system call. If an interrupt is required upon release, the corresponding mask bit should be set in IPC_INTR_STRUCT.INTR_MASK.RELEASE[i]. The exit code must also restore the CM0+ protection context (MPU.MS_PC_CTL.MS_PC) to the one that was backed up in MPU.MS_PC_CTL.MS_PC_SAVED.

16.3.1.4 SRAM Usage

2KB of SRAM [TOP_OF_SRAM – 2KB, TOP_OF_SRAM] is reserved for system calls and 1KB of SRAM [TOP_OF_SRAM – 3KB, TOP_OF_SRAM – 2KB] is used by SROM boot and should not be part of noint RAM. TOP_OF_SRAM is the last address of SRAM.

16.3.2 System Call APIs

SROM has two categories of APIs:

There are three main categories of System Calls:

- Flash management APIs (FLS) - These APIs provide the ability to program, erase, and test the flash macro.
- System management APIs (SYS) - These APIs provide the ability to perform system tasks such as checksum and blowing eFuse.
- Provisioning APIs (PROV) - These APIs are used for the process of provisioning the device.

Table 16-4 shows a summary of the APIs.

Table 16-4. List of System Calls

System Call	Opcode	Description	API Category	Access Allowed	Crypto Block Used
Cypress ID	0x00	Returns die ID, major/minor ID, and protection state	SYS	CM0+, CM4, DAP	No
Read eFuse Byte	0x03	Reads addressed eFuse byte	SYS	CM0+, DAP	No
Write Row	0x05	Pre-program, erase, and program the addressed flash row	FLS	CM0+, CM4, DAP	No
Program Row	0x06	Programs the addressed flash row	FLS	CM0+, CM4, DAP	No
Erase All	0x0A	Erases full flash memory. When part of flash sector is protected by protection settings, this patch verify permission and clears flash first sector by sector, then row by row.	FLS	CM0+, CM4, DAP	No
Checksum	0x0B	Reads either the whole flash or a row of flash, and returns the sum of each byte read	FLS	CM0+, CM4, DAP	No
FmTransitionToLpUlp	0x0C	Configures FLASH macro as per desired low power mode	SYS	CM0+, CM4, DAP	No
Compute Basic Hash	0x0D	Computes the hash value of a given flash region	SYS	CM0+, CM4, DAP	Yes
Erase Sector	0x14	Erases Sector in flash memory. If part of flash sector is protected by secure protection settings, flash is cleared row by row.	FLS	CM0+, CM4, DAP	No
Soft Reset	0x1B	Provides system reset to either or both cores.	SYS	DAP	No
Erase Row	0x1C	Erases the addressed flash row	FLS	CM0+, CM4, DAP	No
Erase Subsector	0x1D	Erases the addressed flash sector	FLS	CM0+, CM4, DAP	No
ReadUniqueID	0x1F	Returns the unique ID of the die from SFlash	SYS	CM0+, CM4, DAP	No
ReadFuseByteMargin	0x2B	Marginally reads eFuse	SYS	CM0+, CM4, DAP	No
AcquireResponse	0x32	Sets or clears response expected by the programming tools to the second word of PUB RAM.	SYS	DAP	No

Table 16-4. List of System Calls (*continued*)

System Call	Opcode	Description	API Category	Access Allowed	Crypto Block Used
PSACrypto	0x35	Wraparound function to call PSA Crypto APIs with PC0 from Bootloader, SPM or any other application. PC of CRYPTO block is changed to PC=0 during this call to have access to the Key storage for different cryptographic operations. After operation complete, PC of CRYPTO is returned to default value. Device Private Key is still protected from export by PSA Crypto interface.	SYS	CM4, DAP	Yes
RollBackCounter	0x36	This system call is used to update and read rollback counter. Updates the rollback counter to higher value only.	SYS	CM4, DAP	
GetProvDetails	0x37	Syscall returns provisioning packet (JWT), templates, or public keys strings in JSON format.	SYS	DAP	No
Provisioning System Calls: The following systems calls are intended for provisioning and are not meant to be called from the CM4.					
Region Hash	0x31	This function is used to compute the Hash of a given region.	PROV	DAP	YES
ProcessProvisionCmd	0x33	This function is used to inject various keys and policies, including the OEM RoT key. These keys and policies are certified by CY and the OEM and verified by this function. A certificate chain is used in the verification of the CY and OEM signatures. For CY, the initial certificate is signed by the HSM using a private key generated by HSM. The second certificate certifies the corresponding public key. The second certificate is signed by CY whose public key is stored in the device. A similar scheme is used for the OEM certificate. This function converts the device from SECURE UNCLAIMED to SECURE CLAIMED life cycle.	PROV	DAP	Yes
EncryptedProgramming	0x34	This function is used to request an encrypted firmware images programming during firmware programming on the production line. The data is transferred in the encrypted form and is decrypted on the device. Application developers can use this process to protect their firmware from theft in the supply chain.	PROV	DAP	Yes
Debug Port Control System Call					
DAPControl	0x3A	This function controls DAP access if it is "allowed" by the debug policy.	SYS	CM0+, CM4	No
RMA System Calls					
TransitionToRMA	0x3B	Converts parts from SECURE to RMA lifecycle stage.	SYS	DAP	Yes
OpenRMA	0x29	Enables full access to device in RMA lifecycle stage upon successful execution	SYS	DAP	Yes
Attestation					
Attestation	0x3C	Implements attestation service - calculates memory hashes and signs a certificate with a device information	SYS	CM0+, CM4, DAP	Yes

Note: The provisioning system calls are meant to be executed during factory programming and provisioning. They are not allowed to be called by the CM4 user application. They are documented in the [PSoC 64 Secure MCU Secure Boot SDK User Guide](#).

The following sections provide detailed information on each system call.

16.3.2.1 Cypress ID

This function returns a 12-bit family ID, 16-bit silicon ID, 8-bit revision ID, and the current device protection mode. These values are returned to the IPC_STRUCT.DATA register if invoked with IPC_STRUCT.DATA[0] set to '1'. Parameters are passed through the IPC_STRUCT.DATA register.

Note that only 32 bits are available to store the return value in the IPC structure. Therefore, the API takes a parameter ID type based on which it will return family ID and revision ID if the ID type is set to '0', and silicon ID and protection state if the ID type is set to '1'.

Table 16-5. Cypress ID

Cypress IDs	Memory Location	Data
Family ID [7:0]	0xF0000FE0	Part Number [7:0]
Family ID [11:8]	0xF0000FE4	Part Number [3:0]
Major Revision	0xF0000FE8	Revision [7:4]
Minor Revision	0xF0000FEC	Rev and Minor Revision Field [7:4]
Silicon ID	SFlash	Silicon ID [15:0]
Protection state	MMIO	Protection [3:0]

Parameters if DAP is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:24]	0x00	Silicon ID opcode.
Bits [15:8]	0 - returns 0. Read family ID and revision ID from SFlash 1 - returns 16-bit silicon ID and protection state	ID type.
Bits [0]	0x1	Indicates that all the arguments are passed in DATA.

Parameters if CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x00	Silicon ID opcode.
Bits [15:8]	0 - returns 12-bit family ID and revision ID 1 - returns 16-bit silicon ID and protection state	ID type.
Bits [0]		Not used (don't care).

Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	If error, Bits [27:0] contain error status Status code (see System Call Status for details).
Bits [27:24]	0xFF	Not used (don't care).

Address	Return Value	Description
Bits [23:20]	If ID Type = 0 Major Revision ID If ID Type = 1 Life-cycle stage 0: VIRGIN 1: NORMAL 2: SEC_W_DBG 3: SECURE 4: RMA	See the device datasheet for these values.
Bits [19:16]	If ID type = 0, Minor Revision ID If ID type = 1, Protection state 0: UNKNOWN 1: VIRGIN 2: NORMAL 3: SECURE 4: DEAD	
Bits [15:8]	If ID type = 0, Family ID Hi If ID type = 1, Silicon ID Hi	See the device datasheet for silicon ID values for different part numbers.
Bits [7:0]	If ID type = 0, Family ID Lo If ID type = 1, Silicon ID Lo	

Return if CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	If error, Bits [27:0] contain error status Status code (see System Call Status for details).
Bits [27:24]	0xXX	Not used (don't care).
Bits [23:20]	If ID Type = 0 Major Revision ID If ID Type = 1 Life-cycle stage 0: VIRGIN 1: NORMAL 2: SEC_W_DBG 3: SECURE 4: RMA	See the PSoC 6 MCU Programming Specifications for these values.
Bits [19:16]	If ID type = 0, Minor Revision ID If ID type = 1, Protection state 0: UNKNOWN 1: VIRGIN 2: NORMAL 3: SECURE 4: DEAD	
Bits [15:8]	If ID type = 0, Family ID Hi If ID type = 1, Silicon ID Hi	See the device datasheet for silicon ID values for different part numbers.
Bits [7:0]	If ID type = 0, Family ID Lo If ID type = 1, Silicon ID Lo	

16.3.2.2 Read eFuse Byte

This function returns the eFuse contents of the addressed byte. The read value of a blown eFuse bit is '1' and that of an unblown eFuse bit is '0'. These values are returned to the IPC_STRUCT.DATA register. Parameters are passed through the IPC_STRUCT.DATA register.

Parameters if DAP is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:24]	0x03	Read eFuse bit opcode.
Bits [23:8]	eFuse Address	Refer to the eFuse Memory chapter on page 132 for more details.
Bits [0]	0x1	Indicates all arguments are passed in DATA.

Parameters if CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x03	Read eFuse bit opcode.
Bits [23:8]	eFuse Address	Refer to the eFuse Memory chapter on page 132 for more details.
Bits [7:0]	0xXX	Not used (don't care).

Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	If error, Bits [27:0] contain error status Status code (see System Call Status for details).
Bits [27:24]	0xXX	Not used (don't care).
Bits [23:0]	eFuse byte	Byte read from eFuse if status is success; otherwise, error code.

Return if CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	If error, Bits [27:0] contain error status Status code (see System Call Status for details).
Bits [27:24]	0xXX	Not used (don't care).
Bits [23:0]	eFuse byte	Byte read from eFuse if status is success; otherwise, error code.

16.3.2.3 Write Row

This function is used to program the flash. You must provide data to be loaded and the flash address to be programmed. The WriteRow parameter performs pre-program and erase, and then programs the flash row with contents from the row latch.

The PSoC 6 MCU supports the Read While Write (RWW) feature, which allows flash to be read from a RWW sector that is not programmed/erased during a program/erase of another RWW sector. Each row is of 512 bytes in size.

The API is implemented in three phases to make it non-blocking. The first phase sets up the flash for pre-program and erase operations and returns to the user code by exiting from NMI without releasing the IPC structure that invoked the API. Now, user code and interrupts can be handled but no NMI can be invoked.

Upon completion of the erase, an interrupt is generated by the flash macro, which will invoke the second phase of the WriteRow API to complete the ongoing erase operation successfully and start the program operation. API returns from NMI to user code after it sets up for program operation.

Upon completion of the program, an interrupt is generated by the flash macro, which will invoke the third phase of the WriteRow API to complete the ongoing program operation successfully; this completes the WriteRow API. SROM API will now return the pass or fail status and releases the IPC structure.

This API can also be called in blocking mode by setting the blocking parameter as '1', in which case the API will return only after all flash operation completes.

The API returns a fail status if you do not have write access to flash according to SPMU settings. See the [CPU Subsystem \(CPUSS\) chapter on page 33](#) for more details. After flash program operation is complete, the API will optionally compare the flash row with the contents in row latch for data integrity check. The function returns 0xF0000022 if the data integrity check fails.

Note: To perform flash writes, the device must be configured such that VCCD is greater than 0.99 V. You must use the Cy_PRA_SystemConfig() function to change the system LDO voltage. More information on this function can be found in the PRA (Protected Register Access) driver section of the Peripheral Driver Library (PDL) documentation. The PRA driver is used to prevent any combination of hardware settings that could destabilize the system and make it more susceptible to attacks.

Parameters if DAP/CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x05	Write Row opcode.
Bits [23:16]	0xFF	Not used
Bits [15:8]	Blocking: 0x01 – API blocks CM0+ Other values - Non-blocking	
Bit [2]	0 - Read operation is allowed on the sector that is not being erased/ programmed 1 - Read operation is stalled until the erase/program operation is complete	
Bit [1]	0 - HV cycles are firmware controlled 1 - HV cycles are hardware controlled	
SRAM_SCRATCH + 0x04		
Bits [23:16]	Verify row: 0 - Data integrity check is not performed 1 - Data integrity check is performed	
Bits [31:24], Bits [15:0]	0xFF	Not used (don't care).
SRAM_SCRATCH + 0x08		
Bits [31:0]		Flash address to be programmed. This should be provided in 32-bit system address format. For example, to program the second half-word, provide either of the byte address 0x10000003 or 0x10000004.
SRAM_SCRATCH + 0x0C		
Bits [31:0]	Data word 0 (data provided should be proportional to the data size provided, data to be programmed into LSbs)	
SRAM_SCRATCH + 0x0C + n*0x04		
Bits [31:0]	Data word n (data to be programmed)	

Return if DAP/CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS /Program ongoing in background 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error code (if any)	See System Call Status for details.

16.3.2.4 Program Row

This function programs the addressed flash row. You must provide the data to be loaded and flash address to be programmed. The flash row should be in the erased state before calling this function.

The function is implemented in two phases to make it non-blocking. The first phase sets up the flash for program operation and returns to user code by exiting from NMI without releasing the IPC structure that invoked the API. Now user code and interrupts can be handled but no NMI can be invoked.

Upon completion of the program operation, an interrupt is generated by the flash macro, which will invoke the second phase of the ProgramRow API to complete the ongoing program operation successfully. The SROM API will return the pass or fail status and releases the IPC structure.

After flash program operation is complete, the API will optionally compare the flash row with the contents in the row latch for data integrity check. It returns STATUS_PL_ROW_COMP_FA if data integrity check fails. The values are returned to the IPC_STRUCT.DATA register. Parameters are passed through the IPC_STRUCT.DATA register.

Note: To perform flash writes, the device must be configured such that V_{CCD} is more than 0.99 V. You must use the Cy_PRA_SystemConfig() function to change the system LDO voltage. More information on this function can be found in the PRA (Protected Register Access) driver section of the Peripheral Driver Library (PDL) documentation. The PRA driver is used to prevent any combination of hardware settings that could destabilize the system and make it more susceptible to attacks.

Parameters if DAP/CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x06	Program Row opcode.
Bits [23:16]	Skip blank check: 0x01 – Skips the blank check step Other – Perform blank check	
Bits [15:8]	Blocking: 0x01 – API blocks CM0+ Other values - Non-blocking	
Bit [2]	0 - Read operation is allowed on the sector that is not being erased/ programmed 1 - Read operation is stalled until the erase/program operation is complete	
Bit [1]	0 - HV cycles are firmware controlled 1 - HV cycles are hardware controlled	
SRAM_SCRATCH + 0x04		
Bits [23:16]	Verify row: 0 - Data integrity check is not performed 1 - Data integrity check is performed	

Address	Value to be Written	Description
Bits [15:8]	Data location: 0 - row latch 1 - SRAM	
Bits [7:0]	0xXX	Not used (don't care).
SRAM_SCRATCH + 0x08		
Bits [31:0]		Flash address to be programmed. This should be provided in 32-bit system address format. For example, to program the second half-word, provide either of the byte address 0x1000003 or 0x1000004.
SRAM_SCRATCH + 0x0C		
Bits [31:0]	Data word 0 (data provided should be proportional to data size provided, data to be programmed into LSbs)	
SRAM_SCRATCH + 0x0C + n*0x04		
Bits [31:0]	Data word n (data to be programmed)	

Return if DAP/CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS /Program ongoing in background 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error code (if any)	See System Call Status for details.

16.3.2.5 Erase All

This function erases all the user Flash area, not the secured memory. The API returns fail status if user does not have write access to FLASH. API will check if all data is 0 to confirm if erase successful. API will return CHECKSUM_NON_ZERO error status if a non-zero word is encountered in the available flash.

The values are returned to the IPC_STRUCT_DATA register. Parameters are passed through the IPC_STRUCT_DATA register.

Note: To perform flash writes (or erase), the device must be configured such that V_{CCD} is greater than 0.99 V. You must use the Cy_PRA_SystemConfig() function to change the system LDO voltage. More information on this function can be found in the PRA (Protected Register Access) driver section of the Peripheral Driver Library (PDL) documentation. The PRA driver is used to prevent any combination of hardware settings that could destabilize the system and make it more susceptible to attacks.

Parameters if DAP is Master

Address	Return Value	Description
IPC_STRUCT_DATA Register		
Bits [31:24]	0x0A	Erase All opcode.
Bits [2]	0 - Read operation is allowed on the sector that is not being erased/ programmed 1 - Read operation is stalled until the erase/program operation is complete	
Bits [1]	0 - HV cycles are firmware controlled 1 - HV cycles are hardware controlled	
Bits [0]	1	Indicates that all the arguments are passed in DATA.

Parameters if CM0+/CM4 is Master

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x0A	Erase All opcode
Bits [1]	0 - HV cycles are firmware controlled 1 - HV cycles are hardware controlled	

Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	If error Bits [27:0] contain error status Status code (see System Call Status for details).
Bits [27:24]	0xFF	Not used (don't care)
Bits [23:0]	Checksum	Checksum if status is SUCCESS

Return if CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	If error Bits [27:0] contain error status Status code (see System Call Status for details).
Bits [27:24]	0xFF	Not used (don't care)
Bits [23:0]	Checksum	Checksum if status is SUCCESS

16.3.2.6 Checksum

This function reads either the entire flash or a row of flash, and returns the sum of each byte read. Bytes 1 and 2 of the parameters select whether the checksum is performed on the entire flash or on a row of flash. This function will inherit the identity of the master that called the function. Hence if a non-secure master requests for either the whole or page checksum of a secured flash, then the fault exception will be raised by the hardware.

The values are returned to the IPC_STRUCT.DATA register. Parameters are passed through the IPC_STRUCT.DATA register.

Parameters if DAP is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:24]	0x0B	Checksum opcode.
Bits [23:22]	0 - application 1 - EE emulation other - supervisory	Flash region.
Bits [21]	0 - row 1 - whole flash	Whole flash.
Bits [20:8]		Row ID.
Bits [0]	1	Indicates that all the arguments are passed in DATA.

Parameters if CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x0B	Checksum opcode.
Bits [23:22]	0 - application 1 - EE emulation 2 - supervisory	Flash region.
Bits [21]	0 - row 1 - whole flash	Whole flash.
Bits [20:8]		Row ID.
Bits [0]	0	

Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	If error, Bits [27:0] contain error status Status code (see System Call Status for details).
Bits [27:24]	0xXX	Not used (don't care)
Bits [23:0]	Checksum	Checksum if status is SUCCESS

Return if CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	If error, Bits [27:0] contain error status Status code (see System Call Status for details).
Bits [27:24]	0xXX	Not used (don't care)
Bits [23:0]	Checksum	Checksum if status is SUCCESS

16.3.2.7 FmTransitionToLpUlp

This system call configures the required flash macro from trim bits when transitioning from LP to ULP mode or vice-versa. This function takes approximately 20 µs; reads from flash are stalled during this time.

Parameters if DAP/CM0+/CM4 is Master and SRAM is used

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		

Address	Value to be Written	Description
Bits [31:24]	0x0C	Opcode for FmTransitionToLpUlp.
Bits [23:2]	0xXX	Not used (don't care)
Bits [1]	0 - Transition To LP 1 - Transition to ULP	Transition mode
Bits [0]	0xXX	Not used (don't care)

Parameters if DAP/CM0+/CM4 is Master and SRAM is not used

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:24]	0x0C	Opcode for FmTransitionToLpUlp.
Bits [23:2]	0xXX	Not used (don't care)
Bits [1]	0 - Transition To LP 1 - Transition to ULP	Transition mode
Bits [0]	0 - Parameters are in address pointed to by IPC_DATA 1 - Parameters are in IPC_DATA	Parameter location

Return if DAP/CM0+/CM4 Invoked the System Call and SRAM is used

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status (if any)	See System Call Status for details

Return if DAP/CM0+/CM4 Invoked the System Call and SRAM is not used

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status (if any)	See System Call Status for details

16.3.2.8 Compute Basic Hash

This function generates the hash of the flash region provided using the formula:

$$H(n+1) = \{H(n) * 2 + \text{Byte}\} \% 127; \text{ where } H(0) = 0$$

This function returns an invalid address status if called on an out-of-bound flash region. Note that CM0+ will inherit the protection context of the master, which invoked it before performing hash. The values are returned to the IPC_STRUCT.DATA register. Parameters are passed through the IPC_STRUCT.DATA register.

Parameters if DAP/CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x0D	Compute hash opcode.
Bits [23:16]	0xXX	Not used (don't care).
Bits [15:8]	0x01 - CRC8 SAE others - Basic hash	Type
Bits [7:0]	0xXX	Not used (don't care)
SRAM_SCRATCH + 0x04		
Bits [31:0]		Start address (32-bit system address of the first byte of the data).
SRAM_SCRATCH + 0x08		
Bits [31:0]	0 – 1 byte 1 – 2 bytes, and so on	Number of bytes.

Return if DAP/CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	If error, Bits [27:0] contain error status Status code (see System Call Status for details).
Bits [27:24]	0xXX	Not used (don't care).
Bits [23:0]	Hash of the data	Hash of data if status is SUCCESS.

16.3.2.9 Erase Sector

This function erases the specified sector. Each sector consists of eight rows. The values are returned to the IPC_STRUCT.DATA register. Parameters are passed through the IPC_STRUCT.DATA register.

Note: To perform flash writes (or erase), the device must be configured such that VCCD is more than 0.99 V. You must use the Cy_PRA_SystemConfig() function to change the system LDO voltage. More information on this function can be found in the PRA (Protected Register Access) driver section of the Peripheral Driver Library (PDL) documentation. The PRA driver is used to prevent any combination of hardware settings that could destabilize the system and make it more susceptible to attacks.

Parameters if DAP/CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x14	Erase Sector opcode.
Bits [23:16]	0x01 - Set FM interrupt mask Other - Do not set FM interrupt mask	Interrupt mask

Address	Value to be Written	Description
Bits [15:8]	Blocking: 0x01 – API blocks CM0+ Other values – Non-blocking	
Bit [1]	0 - HV cycles are firmware controlled 1 - HV cycles are hardware controlled	
SRAM_SCRATCH + 0x04		
Bits [31:0]	Flash Address	Flash address to be erased. Should be provided in 32-bit system address format. For example, to erase the second sector you need to provide the 32-bit system address of any of the bytes in the second sector.

Return if DAP/CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error code (if any)	See System Call Status for details.

16.3.2.10 Soft Reset

This function resets the system by setting CM0+ AIRCR system reset bit. This will result in a system-wide reset, except debug logic. This API can also be used for selective reset of only the CM4 core based on the 'Type' parameter. The values are returned to the IPC_STRUCT.DATA register. Parameters are passed through the IPC_STRUCT.DATA register.

Parameters if DAP is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:24]	0x1B	Soft Reset opcode.
Bits [7:1]	0 - System reset 1 - Only CM4 resets	Type of reset
Bits [0]	0x1	Indicates all arguments are passed in DATA.

Parameters if CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x1B	Soft Reset opcode.
Bits [7:1]	0 - System reset 1 - Only CM4 resets	Type of reset
Bits [0]	0xFF	Not used (don't care).

Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error code (if any)	See System Call Status for details.

Return if CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error code (if any)	See System Call Status for details.

16.3.2.11 Erase Row

This function erases the specified row. You must provide the address of the row that needs to be erased. The values are returned to the IPC_STRUCT.DATA register. Parameters are passed through the IPC_STRUCT.DATA register.

Note: To perform flash writes (or erase), the device must be configured such that VCCD is greater than 0.99 V. You must use the Cy_PRA_SystemConfig() function to change the system LDO voltage. More information on this function can be found in the PRA (Protected Register Access) driver section of the Peripheral Driver Library (PDL) documentation. The PRA driver is used to prevent any combination of hardware settings that could destabilize the system and make it more susceptible to attacks.

Parameters if DAP/CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x1C	Erase row opcode.
Bits [23:16]	0xFF	
Bits [15:8]	0x01 - API blocks CM0+ Other - Non-blocking	
Bit [2]	0 - Read operation is allowed on the sector that is not being erased/programmed 1 - Read operation is stalled until the erase/program operation is complete	
Bit [1]	0 - HV cycles are firmware controlled 1 - HV cycles are hardware controlled	
SRAM_SCRATCH + 0x04		
Bits [31:0]	Address	Flash address to be erased. This should be provided in the 32-bit system address format. For example, to erase the second row you need to provide the 32-bit system address of any of the bytes in the second row.

Return if DAP/CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error code (if any)	See System Call Status for details.

16.3.2.12 Erase Subsector

This function erases the specified subsector. The values are returned to the IPC_STRUCT.DATA register. Parameters are passed through the IPC_STRUCT.DATA register.

Note: To perform flash writes (or erase), the device must be configured such that VCCD is greater than 0.99 V. You must use the Cy_PRA_SystemConfig() function to change the system LDO voltage. More information on this function can be found in the PRA (Protected Register Access) driver section of the Peripheral Driver Library (PDL) documentation. The PRA driver is used to prevent any combination of hardware settings that could destabilize the system and make it more susceptible to attacks.

Parameters if DAP/CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x1D	Erase subsector opcode.
Bits [23:16]	0xFF	Not used (don't care).
Bits [15:8]	0x01 - API blocks CM0+ Other - Non-blocking	
Bit [2]	0 - Read operation is allowed on the sector that is not being erased/ programmed 1 - Read operation is stalled until the erase/program operation is complete	
Bit [1]	0 - HV cycles are firmware controlled 1 - HV cycles are hardware controlled	
SRAM_SCRATCH + 0x04		
Bits [31:0]	Address	Flash address to be erased. This should be provided in 32-bit system address format. For example, to erase the second subsector you need to provide the 32-bit system address of any of the bytes in the second subsector.

Return if DAP/CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [23:0]	Error code (if any)	See System Call Status for details.

16.3.2.13 ReadUniqueID

This function returns the unique ID of the die from SFlash.

Parameters if DAP/ CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x1F	Opcode for ReadUniqueID
Bits [15:0]	Not Used (Don't care)	

Return if DAP/ CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details)
Bits [27:0]	Die Lot if success Error status if failed	If error, Bits [27:0] contain error status See System Call Status for details
SRAM_SCRATCH + 0x04		
Bits [31:0]	DIE_ID0	
SRAM_SCRATCH + 0x08		
Bits [31:0]	DIE_ID1	

16.3.2.14 ReadFuseByteMargin

This API returns the eFuse contents of the addressed byte read marginally. The read value of a blown bit is '1' and of a not blown bit is '0'.

Parameters if DAP is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:24]	0x2B	Opcode for ReadFuseByteMargin
Bits [23:20]	0 - Low resistance, -50% from nominal 1 - Nominal resistance (default read condition) 2 - High resistance (+50% from nominal) Other - 100% from nominal	Margin control
Bits [19:8]	eFuse address (0, 1... 511)	
Bit [0]	1 - Indicates all arguments are passed from DATA	

Parameters if CM0+/CM4 is Master

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		

Address	Value to be Written	Description
Bits [31:24]	0x2B	Opcode for ReadFuseByteMargin
Bits [23:20]	0 - Low resistance, -50% from nominal 1 - Nominal resistance (default read condition) 2 - High resistance (+50% from nominal) Other - 100% from nominal	Margin control
Bits [19:8]	eFuse address (0, 1... 511)	
Bit [0]	1 - Indicates all arguments are passed from DATA	

Return if DAP Invoked the System Call

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Byte read from eFuse	See System Call Status for details.

Return if CM0+/CM4 Invoked the System Call

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Byte read from eFuse or Error status	If error, Bits [27:0] contain error status See System Call Status for details.

16.3.2.15 AcquireResponse

This system call sets or clears response expected by the programming tools to the second word of PUB RAM.

Parameters

Address	Value to be written	
IPC_STRUCT.DATA Register		
Bits [31:24]	0x32	Acquire Response opcode
Bits [23:16]	0x00	Not used
Bits [15:8]	Mode	Clear = 0, Set 1-255
Bits [7:0]	0x01	

Return

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status (if any)	See System Call Status for details.

16.3.2.16 PSACrypto

This syscall is implemented to expose some Mbed Crypto library functions. These functions can be used by Cypress Bootloader (or any other bootloader instance) to optimize memory footprint.

Key Slots in FB Mbed Crypto Key Storage

Provisioned Secure Flashboot initializes keys of Mbed Crypto key storage located in the internal protected RAM. Their slot numbers are available in the following table for use as a parameter for Mbed Crypto API system calls.

Slot #	Key Usage	Key Algorithm	Key Description
1	PSA_KEY_USAGE_DERIVE PSA_KEY_TYPE_ECC_KEYPAIR (PSA_ECC_CURVE_SECP256R1)	PSA_ALG_ECDH PSA_ALG_HKDF PSA_ALG_SHA_256	Device Private Key for ECDH Key agreement by <code>psa_key_agreement()</code> API.
2	PSA_KEY_USAGE_SIGN PSA_KEY_USAGE_VERIFY PSA_KEY_TYPE_ECC_KEYPAIR (PSA_ECC_CURVE_SECP256R1)	PSA_ALG_ECDSA PSA_ALG_SHA_256	Device private key for signing (attestation service). This is a copy of the key slot 1. Mbed Crypto does not allow using the same key slot for key derivation and signature generation/verification.
3	PSA_KEY_USAGE_VERIFY PSA_KEY_USAGE_EXPORT PSA_KEY_TYPE_ECC_PUBLIC_KEY (PSA_ECC_CURVE_SECP256R1)	PSA_ALG_ECDSA PSA_ALG_SHA_256	Cypress public key
4	PSA_KEY_USAGE_VERIFY PSA_KEY_USAGE_EXPORT PSA_KEY_TYPE_ECC_PUBLIC_KEY (PSA_ECC_CURVE_SECP256R1)	PSA_ALG_ECDSA PSA_ALG_SHA_256	HSM public key
5	PSA_KEY_USAGE_VERIFY PSA_KEY_USAGE_EXPORT PSA_KEY_TYPE_ECC_PUBLIC_KEY (PSA_ECC_CURVE_SECP256R1)	PSA_ALG_ECDSA PSA_ALG_SHA_256	OEM public key
6-10	PSA_KEY_USAGE_VERIFY PSA_KEY_USAGE_EXPORT PSA_KEY_TYPE_ECC_PUBLIC_KEY (PSA_ECC_CURVE_SECP256R1)	PSA_ALG_ECDSA PSA_ALG_SHA_256	Custom public keys provisioned during ProvisionKeysAndPolicies SysCall
11	PSA_KEY_USAGE_DERIVE PSA_KEY_TYPE_DERIVE	PSA_ALG_HKDF PSA_ALG_SHA_256	AES 256-bit key derived from 128-bit UDS for Key Derivation by <code>psa_key_derivation()</code> API
12	PSA_KEY_USAGE_DERIVE PSA_KEY_TYPE_ECC_KEYPAIR (PSA_ECC_CURVE_SECP256R1)	PSA_ALG_ECDH PSA_ALG_HKDF PSA_ALG_SHA_256	Device Group Key for ECDH Key agreement by <code>psa_key_agreement()</code> API.
13-16			Reserved

PsaCrypto System Call

The PsaCrypto system calls adhere to the ARM PSA PSA (Platform Security Architecture) Cryptography API. During each syscall, the Crypto MMIOs are protected and accessible only by PC = 0. Programmable PPU MS_PPU_PROG_STRUCT_NR-6 is used with ROM boot syscalls. When the syscall is complete, the Crypto hardware is cleaned and disabled.

Arguments

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	(Address of SRAM where the API parameters are stored. Must be a 32-bit aligned address)
SRAM_SCRATCH		
Bits [31:24]	0x35	PSA opcode
Bits [23:16]	0xXX	NA
Bits [15:8]	1: psa_export_public_key 2: psa_get_key_attributes 3: psa_key_derivation_input_key 4: psa_key_derivation_input_bytes 5: psa_key_derivation_abort 6: psa_key_derivation_key_agreement 7: psa_key_derivation_output_bytes 8: reserved 9: psa_import_key 10: psa_destroy_key 11: psa_cipher_decrypt_setup 12: psa_cipher_set_iv 13: psa_cipher_update 14: psa_cipher_finish 15: psa_generate_random 16: psa_hash_setup 17: psa_hash_update 18: psa_hash_finish 19: psa_asymmetric_sign 20: reserved 21: Cy_FB_Keys_CreateKey 22: Cy_FB_Keys_LoadKeyHandle 23: Cy_FB_Keys_CloseKey 24: psa_key_derivation_output_key 25: psa_generate_key 26: psa_export_key 27: Cy_FB_Get_Keys_Count 28: reserved 29: psa_key_derivation_setup 30: psa_mac_verify_setup 31: psa_mac_verify_update 32: psa_mac_verify_finish 33: Cy_FB_memcpy	PSA Crypto Function#:
Bit [7:0]	0x00	
SRAM_SCRATCH + 0x04		
Bits [31:0]	Address	SRAM_SCRATCH2_ADDR (Address of SRAM where the extended set of API parameters are stored. Must be a 32-bit aligned address)
SRAM_SCRATCH2		
Bits [31:0] * N	PSA function parameters	Parameters in order of how they are described in the PSA Crypto specification. Each parameter has a 4-byte size.

Return

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status (if any)	See System Call Status for details.

16.3.2.17 RollBackCounter

Updates the rollback counter to a higher value only. This function never erases flash, only writes '1's to empty flash bits (OTP emulation) to minimize time of operation and failure risk.

This syscall is used by Cypress Bootloader to prevent firmware reversion during firmware update.

Arguments

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	Address	SRAM_SCRATCH_ADDR (Address of SRAM where the API parameters are stored. Must be a 32-bit aligned address)
SRAM_SCRATCH		
Bits [31:24]	0x36	Rollback Counter opcode
Bits [23:8]	0-15	Counter index
Bits [15:8]	Mode: Read/Write	0 = Read 1-255 = Write
Bits [7:0]	0x00	Not used
Bit [7:0]	0x00	
SRAM_SCRATCH + 0x04		
Bits [31:0]	Address	SRAM_SCRATCH2_ADDR (Address of SRAM where the extended set of API's parameters are stored. Must be a 32-bit aligned address)
SRAM_SCRATCH2		
Bits [31:0]	New Counter Value	Counter Value (In or Out parameter depend on Mode. Decrease is not supported for write operation)

Return

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status (if any)	See System Call Status for details.
SRAM_SCRATCH + 0x04		
Bits [31:0]	Address of SRAM_SCRATCH2	Address to where counter value is located
SRAM_SCRATCH2		
Bits [31:0]	Counter Value	Result of read or updated counter

16.3.2.18 GetProvDetails

Syscall returns provisioning packet (JWT), templates, or public keys strings in JSON format.

Arguments

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH ADDR	Address of SRAM where the API parameters are stored. Must be a 32-bit aligned address
SRAM_SCRATCH		
Bits [31:24]	0x37	GetProvDetails opcode
Bits [23:8]	0XXXXXX	Not used
SRAM_SCRATCH + 0x04		
Bits [31:0]	SRAM_SCRATCH2 ADDR	Address of SRAM where the API parameters are stored. Must be a 32-bit aligned address
SRAM_SCRATCH2		
Bits [31:0]	KeyID	Supports KeyID from PSACrypto section above. Other IDs: 0x100 - FB_POLICY_JWT 0x101 - FB_POLICY_TEMPL_BOOT 0x102 - FB_POLICY_TEMPL_DEBUG 0x2xx - FB_POLICY_CERTIFICATE, where xx is a certificate index in the "chain_of_trust" array of the provisioned packet. 0x300 - FB_POLICY_IMG_CERTIFICATE

Return

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status	See System Call Status for details.
SRAM_SCRATCH + 0x04		
Bits [31:0]	SRAM_SCRATCH2_ADDR	Address of SRAM where the extended set of API response is stored
SRAM_SCRATCH2		
Bits [31:0]	Length	Length of string returned
SRAM_SCRATCH2+ 0x04		
Bits [31:0]	Address	Address of response string

16.3.2.19 Region Hash

Generates the hash of the provided region. When byte 1 of the SRAM_SCRATCH parameter is set to Compare, this API compares the provided region to the provided Compare value in byte 2 of SRAM_SCRATCH and returns CY_FB_INVALID_FLASH_OPERATION error if procedure fails.

API returns invalid address status if called on out of bound Flash region.

This API checks if the client has read access to the requested memory region by looking into DAP MPU and SMPUs. If the client does not have read access, then CY_FB_SYSCALL_ADDR_PROTECTED status is returned.

Arguments

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH ADDR	Address of SRAM where the API parameters are stored. Must be a 32-bit aligned address
SRAM_SCRATCH		
Bits [31:24]	0x31	Region Hash opcode
Bits [23:16]	Compare value	Type: 0x02 - SHA256 0xFF - Compare
Bits [15:8]	Mode	0x02 -> Compute SHA256 Hash 0xFF -> Perform compare function
SRAM_SCRATCH + 0x04		
Bits [31:0]	SRAM_SCRATCH2 ADDR	Address of SRAM where the extended set of API's parameters are stored. Must be a 32-bit aligned address
SRAM_SCRATCH2		
Bits [31:0]	Length	Number of bytes to compare or compute Hash

Return

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status (if any)	See System Call Status for details.
SRAM_SCRATCH + 0x04		
Bits [31:0]	SRAM_SCRATCH2_ADDR	(Address of SRAM where the extended set of API parameters are stored. Must be a 32-bit aligned address) (not used when Type = Compare)
SRAM_SCRATCH2		
Bits [31:0]	Size	Size in bytes of the Hash (not used when Type = Compare)
SRAM_SCRATCH2+ 0x04		
Bits [31:0]	Address	Address where Hash of the data are stored (not used when Type = Compare)

16.3.2.20 ProcessProvisionCmd

The ProcessProvisionCmd system call injects RoT credentials (when complete parameter is set to 0) and changes the device life-cycle mode from SECURE UNCLAIMED to SECURE CLAIMED (when complete parameter is set to 1). Setting the complete parameter to 0 can also be used to verify that the JWT packet is syntactically correct, without making the transition to SECURE CLAIMED. Also, with the complete parameter set to 0, the response packet, which includes the device, die, product IDs, device public key, and OEM key, will still be returned before the transition.

Note: Before calling this syscall with the complete parameter set to 1, ensure that the Cypress Bootloader image and TOC3 are programmed to FLASH.

Syscall handler receives rot_cmd.jwt or prov_cmd.jwt packet, and performs following the tasks. Refer to the [PSoC 64 Secure MCU Secure Boot SDK User Guide](#) for format and definition of all JWT packets used for provisioning.

This system call receives JWT packet, and performs following tasks:

cy_auth

- Validates input packet

- ❑ Verifies CY signature of cy_auth packet
- ❑ Verifies CY public key in cy_auth packet
- ❑ Import HSM public key
- ❑ Verifies HSM signature of full a provisioning JWT packet.
- ❑ Verifies DIE_ID and DEV_ID range in packet with device identity

create_identity = true:

- Generates UDS, Device Public and Private key (if not received in same packet)
 - ❑ Encrypts Device and received Group Key by derived from UDS key and stores to flash. AES-CBC with PKCS7 padding is used for encryption.
 - ❑ Stores UDS to eFuses, protected by PPU from read and write.
- Calculates and blows SECURE_HASH2 from encrypted package.

rot_auth

- Validates input packet
 - ❑ Import OEM public key
 - ❑ Verifies OEM signature of rot_auth packet
 - ❑ Compares HSM public key in cy_auth and rot_auth packet
 - ❑ Verifies that Expiration time of cy_auth is not less than Expiration time of the rot_auth packet.
- Store OEM public key and prod_id in flash

prod_id

- Verifies prod_id with the stored in flash during rot_auth packet process

chain_of_trust, prov_req

- Stores JWT packet to flash when OEM signature is correct.

image_cert

- Verifies OEM signature of input packet and store to flash.
- Extracts "image_address" and "image_size" values of Bootloader image
- Verifies hash for the Bootloader image with hash value from Bootloader FW Image Distribution packet
- Bootloader image must be programmed to FLASH before this token execution.

complete = true

- Verifies policies against templates (only min/max values and types).
- Calculates and blows SECURE_HASH3 from:
 - ❑ OEM Public Key
 - ❑ Product ID
 - ❑ provisioning JWT packet (prov_cmd and chain_of_trust) if "reprovision"->"keys_and_policies" = false
 - ❑ Bootloader image and image certificate if "reprovision"->"boot_loader" = false
- Calculates SECURE_HASH4, sign it by Device Private key and store to flash from:
 - ❑ provisioning JWT packet (prov_cmd and chain_of_trust) if "reprovision"->"keys_and_policies" = true
 - ❑ Bootloader image and image certificate if "reprovision"->"boot_loader" = true

Generates JWT response packet signed by Device Private Key with:

- Device ID (silicon ID, family ID)
- DIE ID (lot, wafer, xpos, ypos, day, month, year)
- Device Public Key

- Group Public Key
- OEM Public Key
- Product ID
- Complete status

Note: Secure FlashBoot in SECURE CLAIMED mode validates SECURE_HASH3 during every boot before transferring control to the Cypress Bootloader.

Note: Secure FlashBoot applies Bootloader flash protection region during every boot to avoid unexpected clearing of this image.

Note: Secure FlashBoot validates signature of SECURE_HASH4 during every boot if "reprovision"->"keys_and_policies" = true or "reprovision"->"boot_loader" = true

Note: VDDIO0 must be at 2.5 volts in order to complete this command.

Arguments

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x33	Opcode for ProcessProvisionCmd
Bits [23:0]	0XXXXXX	Not used
SRAM_SCRATCH + 0x04		
Bits [31:0]	SRAM_SCRATCH2 ADDR	Address of SRAM where the extended set of API's parameters are stored. Must be a 32-bit aligned address
SRAM_SCRATCH2		
Bits [31:0]	Length	Number of bytes (including itself)
SRAM_SCRATCH2 + 0x04		
Array [Length-4]	The JWT packet, signed by HSM Private Key, should contain: <ol style="list-style-type: none"> 1. JWT packet(cy_auth) signed by Cypress Private key with: <ol style="list-style-type: none"> 1. HSM Public Key 2. Cypress Public Key 2. Optionally JWT packet(rot_auth) signed by OEM private key with: <ol style="list-style-type: none"> 1. HSM Public Key 2. OEM Public Key When complete parameter is set to true: <ol style="list-style-type: none"> 3. JWT packet (prov_req) signed by OEM Private Key for provisioning with policies 4. JWT packet (image_cert) signed by any Private Key referenced and provisioned in prov_req->boot_upgrade->firmware->boot_auth with Bootloader image hash. 5. Optionally Certificates (chain_of_trust) 	

Return

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status (if any)	See System Call Status for details.
SRAM_SCRATCH + 0x04		

Address	Return Value	Description
Bits [31:0]	SRAM_SCRATCH2_ADDR	Address of SRAM where the short set of API parameters are stored. Must be a 32-bit aligned address
SRAM_SCRATCH2		
Bits [31:0]	Length	Number of bytes in response
SRAM_SCRATCH2+ 0x04		
Bits [31:0]	Address	Address where JWT response packet is stored.

16.3.2.21 EncryptedProgramming

The EncryptedProgramming syscall consist of three parts:

1. Init

- Validates the programming request.
- Decrypts the decryption key.
 - ECDH key agreement by using Device or Group Private/Public keys and HSM or OEM public key, shared secret is generated.
 - Generate key material from the shared secret, salt and info: AES key and IV
 - Import AES key to the key storage.
 - Init cipher decrypt setup.
 - Set Initialization Vector to be used with AES decryption Key.
 - Import decryption AES key to the key storage.

- Sets up for encrypted programming

2. Data

- Sets Initialization Vector, decrypts the data and programs it to flash

3. Finish

- Cleans the initialized data

Note: To perform flash writes, the device must be configured such that VCCD is greater than 0.99 V. You must use the Cy_PRA_SystemConfig() function to change the system LDO voltage. More information on this function can be found in the PRA (Protected Register Access) driver section of the Peripheral Driver Library (PDL) documentation. The PRA driver is used to prevent any combination of hardware settings that could destabilize the system and make it more susceptible to attacks.

Arguments

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x34	Opcode for EncryptedProgramming
SRAM_SCRATCH + 0x04		
Bits [31:0]	SRAM_SCRATCH2_ADDR	Address of SRAM where the short set of API parameters are stored. Must be a 32-bit aligned address

Address	Value to be Written	Description
SRAM_SCRATCH2		
Bits [31:0]	Length	Number of bytes (Init: the length if encryption header, Program: Data size to program)
SRAM_SCRATCH2 + 0x04		
Array [Length-4]	When Mode is set to Init: <ul style="list-style-type: none"> ■ AES key size in bytes(1 byte): 128 bits = 16 bytes, 256 bits = 32 bytes ■ Encrypted AES key (32 bytes) <ul style="list-style-type: none"> □ ECDH key agreement algorithm should be used with Device or Group Private/Public keys and HSM/OEM public key. ■ Encrypted Initialization Vector (16 bytes) ■ Padding (16 bytes) ■ Salt (16 bytes) and Label (16 bytes) to derivate key encryption key When Mode is set to Program: <ul style="list-style-type: none"> ■ Flash address to be programmed (4 bytes). ■ Encrypted data array with padding to decode and program one flash row (512 bytes+16bytes padding). 	

Return

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status (if any)	See System Call Status for details.

16.3.2.22 DAPControl

This SysCall allows users to control DAP access during run-time. It works only when a particular CPU ID Debug Access Port has permission set to "allowed" in the debug policy. The control field in the debug policy defines whether this SysCall can control the DAP with (control is set to "certificate") or without (control is set to "open") a signed certificate (JWT).

Arguments

Arguments if IPC_STRUCT.DATA[0] = 1 for a case when signed JWT is not required by debug policy:

Address	Value to be Written	Description
IPC_STRUCT.DATA Register		
Bits [31:24]	0x3A	Opcode for DAPControl
Bits [23:16]	Desired State	0 - Disable 1 - Enable
Bits [15:9]	DAP AP	0 - M0p 1 - M4 2 - System
Bits [7:0]	0x01	No SRAM needed

Arguments if IPC_STRUCT.DATA[0] = 1

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x3A	Opcode for DAPControl
Bits [23:16]	Desired State	0 - Disable 1 - Enable
Bits [15:8]	DAP AP	0 - M0p 1 - M4 2 - System
Bits [7:0]	0x01	No SRAM needed
SRAM_SCRATCH + 0x04		
Bits [31:0]	SRAM_SCRATCH2_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH2		
Bits [31:0]	Length	Number of bytes in packet below
SRAM_SCRATCH2+ 0x04		
Array [Length]	The JWT packet, signed by the Private Key mentioned in the Debug Policy/DAP AP/key with a "auth" object including valid DIE_ID range.	

Return

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status (if any)	See System Call Status for details.

16.3.2.23 TransitionToRMA

Converts parts from SECURE to RMA life cycle stage.

Before transitioning to RMA stage, syscall destroys fuse and flash contents five times as specified in the debug policy <destroy_fuses> and <destroy_flash>. When successful, it sets IPC3_DATA with a value of 0x000000EA to enable full access to device in the RMA life cycle stage.

After each reset, ROM boot code will wait on OpenRMA syscall to open full access again. DAP will only have access via SYSTEM AP to IPC MMIOs and one-sixteenth of RAM0. Only OpenRMA system call is allowed before successful OpenRMA execution.

Arguments

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x3B	Opcode for TransitionToRMA
Bits [23:0]	0XXXXXX	Not used

Address	Return Value	Description
SRAM_SCRATCH + 0x04		
Bits [31:0]	SRAM_SCRATCH2_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH2		
Bits [31:0]	Length	Length of JWT packet
SRAM_SCRATCH2+ 0x04		
Array [Length]	The JWT packet, signed by the Private Key mentioned in the Debug Policy/RMA/key, with a "auth" object including valid DIE_ID range. DIE_ID must be present in the signed JWT packet in the following format: "die_id": {"lot": "integer", "wafer": "integer", "xpos": "integer", "ypos": "integer", "sort": "integer", "minor": "integer", "day": "integer", "month": "integer", "year": "integer"}	

Return

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status	See System Call Status for details.

16.3.2.24 OpenRMA

When successful, it sets IPC3_DATA with a value of 0x000000EA to enable full access to device in the RMA life cycle stage.

The interface and parameters are similar to the TransitionToRMA syscall with the difference that SRAM_SCRATCH and SRAM_SCRATCH2 should be placed in the beginning of RAM: 0x08000000.

Arguments

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	0x08000000	The scratch SRAM must be at the beginning of the device SRAM.
SRAM_SCRATCH (0x08000000)		
Bits [31:24]	0x29	Opcode for TransitionToRMA
Bits [23:0]	0xFFFFFFFF	Not used
SRAM_SCRATCH + 0x04 (0x08000004)		
Bits [31:0]	0x08000008	
SRAM_SCRATCH2 (0x08000008)		
Bits [31:0]	Length	Length of JWT packet
SRAM_SCRATCH2+ 0x04		
Array [Length]	The JWT packet, signed by the Private Key mentioned in the Debug Policy/RMA/key, with a "auth" object including valid DIE_ID range. DIE_ID must be present in the signed JWT packet in the following format: "die_id": {"lot": "integer", "wafer": "integer", "xpos": "integer", "ypos": "integer", "sort": "integer", "minor": "integer", "day": "integer", "month": "integer", "year": "integer"}	

Return

Address	Return Value	Description
IPC_STRUCT.DATA		
Bits [31:0]	0xA = SUCCESS 0xF = ERROR	This values is written into the data register only if the command is successful, otherwise the value is not defined.
SRAM_SCRATCH (0x08000000)		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status	See System Call Status for details.

16.3.2.25 Attestation

Calculates hashes of memory regions provided in an input list. Also calculates signature of a certificate that attests the device state at the moment of signing. Signature is calculated for the following data structure:

- Server random number (uint32_t)
- Syscall random number (uint32_t)
- Device UID (SFLASH->DIE_LOT array, 11 bytes)
- Device Identity (cy_flashDeviceKeyData array, 512 bytes)
- OEM Public key and Product ID (cy_flashProvKeyData array, 512 bytes)
- Chain of trust (cy_flashChainOfTrust array, 5 kbytes)
- Image certificate (cy_flashImgCertJWT array, 1 kbytes)
- Policy package (cy_flashProvisionJWT array, 10 kbytes)
- Number of memory regions (uint32_t)
- for (each memory region)
 - Memory region address (uint32_t)
 - Memory region size (uint32_t)
 - Memory region content (uint8_t array)
 - Memory region hash (calculated with the syscall random number at the beginning) (uint8_t array)

Signature is calculated without any additional padding/aligning between different certificate fields. Memory regions can be from SRAM, flash, AUXflash, or SFlash. SMIF and peripheral address space are not supported.

AttestationSysCall Algorithm

1. Generate syscall random number
2. Init signature hash calculation
3. Update signature hash with Server/Syscall random numbers, DevUID, Dev public key, policy package
4. Update signature hash with number of regions
5. Check whether the array with memory regions has at least read access allowed for the caller (to avoid side channel attacks)
6. for (each specified memory region) {
 - Start region hash calculation with the syscall random number
 - Update signature and region hash with a region content (in parallel using the same data)
 - Update signature hash with the region hash
 - Verify writing rights of the caller and sufficient memory size
 - Output the region hash to corresponding memory
7. Sign the signature hash with a device private key

8. Clear Crypto block internal memories and used stack

A real-life usage example: Consider a server that has information about all connected devices in the field, their IDs, public keys, and available versions of firmware. The server wants to know the exact state of one of the devices – whether it is attacked, which firmware versions it has, and which configuration. The server sends a request to the device to send back a signed certificate with hashes and content of specific memory regions. Application (through SPM service) calls the SysCall, which calculates hashes of the specified regions and creates a signature of a data certificate with specific structure. SPM code creates a package with all the data required, includes their calculated hashes and signature, and returns it to the application. Application sends it back to the server.

Arguments

Address	Return Value	Description
IPC_STRUCT.DATA Register		
Bits [31:0]	SRAM_SCRATCH_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH		
Bits [31:24]	0x3C	Opcode for Attestation
Bits [23:0]	0XXXXXX	Not used
SRAM_SCRATCH + 0x04		
Bits [31:0]	SRAM_SCRATCH2_ADDR	SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH2		
Bits [31:0]	PSA signing algorithm	PSA signing algorithm (contains hash algorithm type used both for signature and memory region hashes). Only PSA_ALG_ECDSA(PSA_ALG_SHA_256) is allowed at the moment
SRAM_SCRATCH2+ 0x04		
Bits [31:0]	Number	Random number from server
SRAM_SCRATCH2+ 0x08		
Bits [31:0]	Number	Number of memory regions in the list
SRAM_SCRATCH2+ 0x0C		
Bits [31:0]	Address	Pointer to list of memory region start addresses
SRAM_SCRATCH2+ 0x10		
Bits [31:0]	Address	Pointer to list of memory region sizes
SRAM_SCRATCH2+ 0x14		
Bits [31:0]	Address	Pointer to array for hashes of the memory regions
SRAM_SCRATCH2+ 0x18		
Bits [31:0]	Number	Size in bytes of the array for hashes

Return

Address	Return Value	Description
SRAM_SCRATCH		
Bits [31:28]	0xA = SUCCESS 0xF = ERROR	Status code (see System Call Status for details).
Bits [27:0]	Error status	See System Call Status for details.
SRAM_SCRATCH + 0x04		

Address	Return Value	Description
Bits [31:0]	SRAM_SCRATCH2_ADDR	Address of SRAM where the extended set of API's response is stored, SRAM address where the API parameters are stored. This must be a 32-bit aligned address.
SRAM_SCRATCH2		
Bits [31:0]	Number	Random number from the syscall
SRAM_SCRATCH2+ 0x04		
Bits [31:0]	Number	Size in bytes of each memory region Hash
SRAM_SCRATCH2+ 0x08		
Bits [31:0]	Number	Size in bytes of the Signature
SRAM_SCRATCH2+ 0x0C		
Bits [31:0]	Address	Address where Signature is stored

16.3.3 System Call Status

At the end of every system call, a status code is written over the arguments in the IPC_DATA register or the SRAM address pointed by IPC_DATA. A success status is 0xAXXXXXXX, where X indicates don't care values or return data for system calls that return a value. A failure status is indicated by F<0/1/7>XXXXXX, where XXXXXX is the failure code.

Flashboot also returns several status code in DAP IPC_DATA at the end of boot. Debug probes needs to preserve them before invoking system calls.

Value	Description
0xAXXXXXXX	PASS, X - don't care
0xF0000000	Reject the system call when CPUSS_PROTECTION is not NORMAL
0xF0000001	Invalid flash page latch address
0xF0000002	Invalid flash address
0xF0000003	Row is write protected
0xF0000004	Wrong or out-of-bound flash address
0xF0000005	IPC structure is already locked by another process
0xF0000006	Input parameters passed to Flash API are not valid
0xF0000008	Returned by all APIs when client doesn't have access to region it is using for passing arguments.
0xF0000009	Command in progress.
0xF000000A	Checksum of flash resulted in non-zero.
0xF000000B	The opcode is not a valid API opcode.
0xF000000F	Returned when invalid arguments are passed to the API. For example, calling Silicon ID API with ID type of 0x5.
0xF0000013	An argument to the system call is not in a valid location.
0xF0000021	Sector erase requested on SFlash region.
0xF0000041	Bulk erase failed.
0xF0000042	Sector erase failed.
0xF0000043	Subsector erase failed.
0xF0000044	Verification of Bulk, Sector, or Subsector fails after program/erase.
0xF7000000	The fail status of the Secure FlashBoot return value
0xF700DEAD	Returned when device is in DEAD state.
0xF7000001	Returned if a device is provisioned and re-provisioning is not supported
0xF7000002	Returned when write to flash operation fails.

Value	Description
0xF7000003	Returned when signature verification of the JWT packet, which should be signed by Cypress Private Key is failed(cy_auth or image_cert).
0xF7000004	Returned when signature verification of the JWT packet, which should be signed by OEM Private Key is failed(rot_auth or prov_req).
0xF7000005	Returned when signature verification of the JWT packet, which should be signed by HSM Private Key is failed(prov_cmd).
0xF7000006	Returned when cy_auth token in prov_cmd packet doesn't contain valid hsm_pub_key.
0xF7000007	Returned when import of Cypress Public Key from JSON to key storage is failed or when validation of this key in cy_auth
0xF7000009	Returned when Device Key generation is failed.
0xF700000A	Returned if a master with PC > 4 tries to use PSA syscall and protection is enabled in the policy ("protect_flags" & 2) !=
0xF700000B	Returned when Secure Hash calculation is failed.
0xF700000C	Returned when compare of HSM public key in cy_auth and rot_auth packet is failed.
0xF700000E	Returned when validation of prod_id in prov_req packet is failed.
0xF700000F	Returned when Efuse blow is failed.
0xF7000010	Returned when Protection Context change of the Crypto is failed
0xF7000011	Returned when validation of provision policy by schema failed.
0xF7000012	Returned by TransitionToRMA or SyscallDPControl when requested action (enable DP or RMA) is not permitted by the
0xF7000013	Returned when cryptographic operation failed.
0xF7000014	Returned when "prod_id" string is too long. Maximum allowed length is 64 bytes.
0xF7000015	Returned when calculated secure image hash doesn't match with a hash received in image_cert packet.
0xF7000017	Returned if generation of a device UDS failed or it already exists.
0xF7000018	Returned if import of Group key failed.
0xF7000019	Returned if encryption and storing of Device and Group private keys failed.
0xF700001A	Returned if device private keys already exist.
0xF700001B	Returned if the length of the Chain of trust string is too long. Maximum length is 5120 bytes.
0xF700001C	Returned if the length of the provisioning request string is too long. Maximum length is 10240 bytes.
0xF700001D	Returned if the size of the image certificate is too big. Maximum size is 1024 bytes.
0xF700001E	Returned if the image certificate signature validation failed.
0xF700001F	Returned if a structure of the image certificate is not correct.
0xF7000020	Returned if the device ID in the "cy_auth" packet does not match with the provisioned device.
0xF7000021	Returned if the die ID in the "cy_auth" packet does not match with the provisioned device.
0xF7000022	Returned if re-provisioning settings in the policy does not match previously provisioned ones.
0xF7000023	Returned if the re-provisioning certificate signature validation failed.
0xF7000024	Returned if invalid arguments are passed to a SysCall.
0xF7000025	Returned if a master with PC > 4 tries to use psa_asymmetric_sign API with an internal key and protection is enabled in
0xF7000026	Returned if EncryptedProgramming SysCall tries to program not permitted regions of the flash memory according to the
0xF7000027	Returned when provisioning key length in JSON format is equal to 0 or is greater than 172 bytes.
0xF70000FF	Returned when memory allocation failed.
0xF7000206	Returned when cy_auth token in prov_cmd packet doesn't contain valid cy_pub_key.
0xF7000306	Returned when rot_auth token in prov_cmd packet doesn't contain valid oem_pub_key.
0xF7000406	Returned when rot_auth token in prov_cmd packet doesn't contain valid hsm_pub_key.
0xF7000506	Returned when the body part of the prov_cmd packet is absent.
0xF7000606	Returned when base64_decode of the body part of the prov_cmd packet is failed.

Value	Description
0xF7000706	Returned when JSON parser failed with the decoded body part of the prov_cmd packet.
0xF7000806	Returned when "cy_auth" token is absent in the parsed body part of the prov_cmd packet.
0xF7000C06	Returned when "exp" token is absent in the parsed body part of the cy_auth packet.
0xF7001706	Returned when "oem_pub_key" token is absent in the parsed body part of the rot_auth packet.
0xF7001806	Returned when "hsm_pub_key" token is absent in the parsed body part of the rot_auth packet.
0xF7001906	Returned when "prod_id" token is absent in the parsed body part of the rot_auth packet.
0xFF000000	The status mask of the Secure FlashBoot return value

16.4 Amazon FreeRTOS PSA Integration

The PSoC 64 Standard Secure Amazon Web Services (AWS) product line provides a Platform Security Architecture (PSA) level-2 certified way of connecting to AWS. PSA is a set of threat models, security analysis, hardware and firmware architecture specifications, open source firmware reference implementation, and independent evaluation and certification.

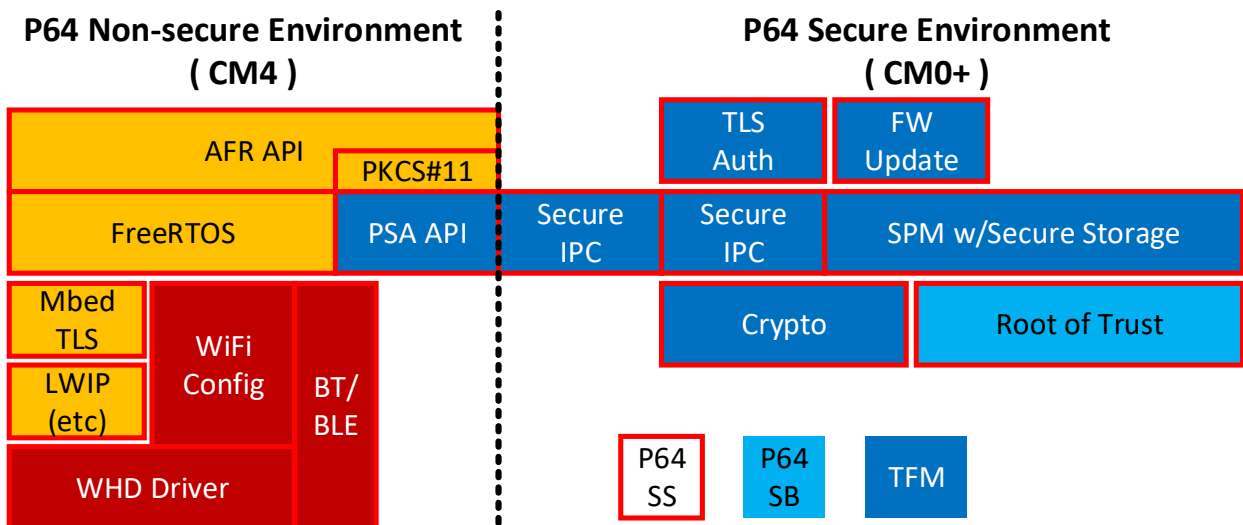
The goal of PSA is to provide the following:

- A hardware and firmware architecture via a set of specifications such as the Security Model (PSA-SM) and Firmware Framework (PSA-FF).
- Provides documents pertaining to boot and upgrade requirements.
- Defines isolation levels that compartmentalizes software and services.
- Outlines a standardized certification process that helps identify the level of security. Currently it has three levels of certification: L1, L2, and L3.
- Provides a standardized API for secure services such as Crypto, secure storage, and attestation.

Provides a reference implementation for Cortex-M class processors called Trusted Firmware M (TF-M).

Amazon FreeRTOS is an Amazon-supported embedded ecosystem that provides a development environment for IoT devices including WiFi connectivity. The PSoC 64 AWS Standard Secure provides a pre-built TF-M layer to Amazon FreeRTOS, as well as all the required AWS FreeRTOS enhancements. Figure 16-6 shows a block diagram of the PSoC 64 Standard Secure Features.

Figure 16-6. PSoC 64 Standard Secure Environment



Section C: System Resources Subsystem (SRSS)

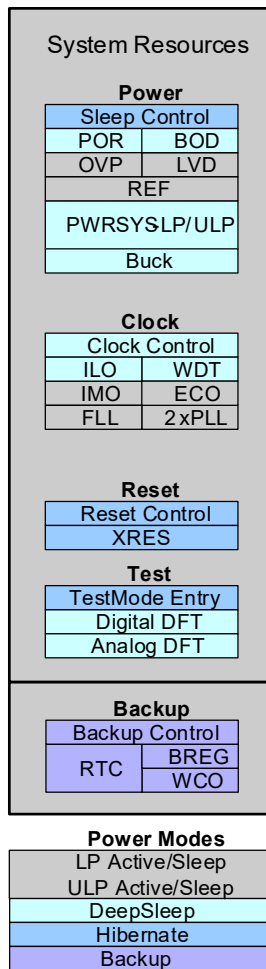


This section encompasses the following chapters:

- [Power Supply and Monitoring chapter on page 183](#)
- [Device Power Modes chapter on page 190](#)
- [Backup System chapter on page 200](#)
- [Clocking System chapter on page 207](#)
- [Reset System chapter on page 223](#)
- [I/O System chapter on page 226](#)
- [Watchdog Timer chapter on page 248](#)
- [Trigger Multiplexer Block chapter on page 260](#)
- [Profiler chapter on page 266](#)

Top Level Architecture

Figure C-1. System-Wide Resources Block Diagram



17. Power Supply and Monitoring



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SysPM
- Application notes
- Code examples

The PSoC 6 MCU family supports an operating voltage range of 1.7 V to 3.6 V. It integrates multiple regulators including an on-chip buck converter to power the blocks within the device in various power modes. The device supports multiple power supply rails – V_{DDD} , V_{DDA} , V_{DDIO} , and V_{BACKUP} – enabling the application to use a dedicated supply for different blocks within the device. For instance, V_{DDA} is used to power analog peripherals such as ADC and opamps.

The PSoC 6 MCU family supports power-on-reset (POR), brownout detection (BOD), over-voltage protection (OVP), and low-voltage-detection (LVD) circuit for power supply monitoring and failure protection purposes.


17.1 Features

The power supply subsystem of the PSoC 6 MCU supports the following features:

- Operating voltage range of 1.7 V to 3.6 V
- User-selectable core logic operation at either 1.1 V or 0.9 V
- Three independent supply rails (V_{DDD} , V_{DDA} , and V_{DDIO}) for PSoC core peripherals and one independent supply rail (V_{BACKUP}) for backup domain
- Multiple on-chip regulators
 - One low-dropout (LDO) regulator to power peripherals in Active power mode
 - One buck converter
 - Multiple low-power regulators to power peripherals operating in different low-power modes
- Two BOD circuit (V_{DDD} and V_{CCD}) in all power modes except Hibernate mode
- LVD circuit to monitor V_{DDD} , V_{AMUXA} , V_{AMUXB} , V_{BACKUP} , or V_{DDIO}
- One OVP block monitoring V_{CCD}

17.2 Architecture

Figure 17-1. Power System Block Diagram



See the [device datasheet](#) for the values to be used for the capacitors and inductor shown in [Figure 17-1](#).

The regulators and supply pins/rails shown in [Figure 17-1](#) power various blocks inside the device. The availability of various supply rails/pins for an application will depend on the device package selected. Refer to the [device datasheet](#) for details.

All the core regulators draw their input power from the V_{DD} supply pin. The on-chip buck uses the V_{DD_NS} supply pin as its input. V_{CCD} supply is used to power all active domain and high frequency peripherals. The output of the buck is connected to the V_{CCD} pin and in firmware the V_{CCD} supply can be switched to the buck output. A dedicated Deep Sleep regulator powers all the Deep Sleep peripherals. The Deep Sleep regulator switches its output to V_{CCD} when available and to its regulated output when V_{CCD} is not present (System Deep Sleep power mode). Hibernate domain does not implement any regulators and the peripherals available in that domain such as Low-Power comparator and ILO operate directly from V_{DD} .

When the V_{DDA} pin is not present, analog peripherals run directly from the V_{DD} .

The I/O cells operate from various V_{DDx} (V_{DDA} , V_{DD} , or V_{DDIO}) pins depending on the port where they are located. V_{CCD} supply is used to drive logic inside the I/O cells from core peripherals. V_{DDA} powers the analog logic such as analog mux switches inside the I/O cell. To know which I/Os operate from which supply, refer to the [device datasheet](#).

The device includes a V_{BACKUP} supply pin to power a small set of peripherals such as RTC and WCO, which run independent from other supply rails available in the device. When the V_{BACKUP} supply is not present, the device uses V_{DD} to power these peripherals.

In addition to the power rails and regulators, the device provides options to monitor supply rails and protection against potential supply failures. These include a POR circuit, a BOD circuit, an OVP circuit, and a LVD circuit.

17.3 Power Supply

17.3.1 Regulators Summary

17.3.1.1 Core Regulators

The device includes the following core regulators to power peripherals and blocks in various power modes.

Linear Core Regulator

The device includes a linear LDO regulator to power the Active and Sleep mode peripherals. This regulator generates the core voltage (V_{CCD}) required for Active mode operation of the peripherals from V_{DD} . The regulator is capable of providing 0.9 V and 1.1 V for core operation. See

[Core Operating Voltage](#) on [page 186](#). The regulator is available in Active and Sleep power modes. This regulator implements two sub-modes – high-current and low-current modes. LINREG_LPMODE bit [24] of the PWR_CTL register selects between the two modes of operation. The high-current mode is the normal operating mode, that is, the device operates to its full capacity in Active or Sleep power modes. In the low-current mode, the current output from the regulator is limited. This mode implements the Low-Power Active and Sleep power modes. The low-current mode sets a limitation on the capabilities and availability of resources in the Low-Power Active and Sleep modes. For details, see the [Device Power Modes](#) chapter on [page 190](#).

By default, the linear regulator is powered on reset. The regulator can be disabled by setting the LINREG_DIS bit [23] of the PWR_CTL register. Note that the linear regulator should be turned OFF only when the following conditions are satisfied:

- Switching buck regulator is ON.
- The load current requirement of the device from the V_{CCD} supply does not exceed 20 mA. This should be ensured by the firmware by disabling power consuming high-frequency peripherals and reducing the system clock frequency.

If the linear regulator is turned OFF without the above conditions satisfied, it will result in V_{CCD} brownout and the device will reset.

Switching (Buck) Core Regulator

The device includes a switching (buck) core regulator. The buck requires an inductor and a capacitor to generate the output. Note that the buck output is also available in the Deep Sleep device power mode.

The buck regulator can be enabled by setting the BUCK_EN bit [30] of the PWR_BUCK_CTL register. The buck output can be enabled/disabled using the BUCK_OUT1_EN bit [31] of the PWR_BUCK_CTLx registers. The buck output supports voltages from 0.85 V to 1.20 V. Use either 0.9 V or 1.1 V for V_{CCD} operation. The output selection can be made using BUCK_OUTx_SEL bits in the PWR_BUCK_CTLx registers.

When used to power the core peripherals, the buck regulator provides better power efficiency than the linear regulator, especially at higher V_{DD} . However, the buck regulator has less load current capability than the linear regulator. Therefore, when using the buck regulator, take care not to overload the regulator by running only the necessary peripherals at a lower frequency in firmware. Overload conditions can cause the buck output to drop and result in a brownout reset. Refer to the [device datasheet](#) for the load limitations on the buck regulator, recommended settings when used for V_{CCD} , and its efficiency characteristics. Follow these steps in firmware when

switching to the buck regulator for core (V_{CCD}) operation without causing a brownout.

1. Make sure the necessary inductor and capacitor connection for the buck as explained in [Figure 17-1](#) is present in the hardware.
2. Change the core supply voltage to 50 mV more than the final buck voltage. For instance, if the final buck voltage is 0.9 V, then set the LDO output to 0.95 V and 1.15 V for 1.1 V buck operation. Set the ACT_REG_TRIM bits[4:0] of the PWR_TRIM_PWRSYS_CTL register to '0x0B' to switch to 0.9 V and '0x1B' to switch to 1.1 V buck operation. This is discussed in [Core Operating Voltage](#).
3. Reduce the device current consumption by reducing clock frequency and switching off blocks to meet the buck regulator's load capacity.
4. Disable System Deep Sleep mode regulators. Because the buck regulator is available in the System Deep Sleep power mode, other deep-sleep regulators can be powered down. This is done by setting the following bits in the PWR_CTL register:
 - a. DPSLP_REG_DIS bit [20] – Disables the deep-sleep core regulator.
 - b. RET_REG_DIS bit [21] – Disables the logic retention regulator.
 - c. NWELL_REG_DIS bit [22] – Disables the nwell regulator.
5. Set the buck output to the desired value by writing '2' (for 0.9 V) or '5' (for 1.1 V) to the BUCK_OUT1_SEL bits[2:0] of the PWR_BUCK_CTL register.
6. Enable the buck regulator and output by setting the BUCK_EN bit[30] and BUCK_OUT1_EN bit[31] of the PWR_BUCK_CTL register.
7. Wait for the buck regulator to start up and settle.
8. Disable the linear regulator by setting the LINREG_DIS bit[23] of the PWR_CTL register.

After transitioning to the buck regulator, do not switch back to the linear regulator mode to ensure proper device operation. This should happen once during powerup.

Core Operating Voltage

PSoC 6 MCUs can operate at either 0.9 V LP mode (nominal) or 1.1 V ULP mode (nominal) core voltage. On reset, the core is configured to operate at 1.1 V by default. At 0.9 V, power consumption is less, but there are some limitations. The maximum operating frequency for all HFCLK paths should not exceed 50 MHz, whereas the peripheral and slow clock should not exceed 25 MHz.

Follow these steps to change the PSoC 6 MCU core voltage:

1. While transitioning to 0.9 V, reduce the operating frequency to be within the HFCLK and peri clock limits defined in the [device datasheet](#). Turn off blocks, if required, to be within the maximum current consumption limit of the linear regulator at 0.9 V.
2. Set the ACT_REG_TRIM bits[4:0] of the PWR_TRIM_PWRSYS_CTL register to '0x07' for 0.9 V or '0x17' for 1.1 V. For the buck regulator, set the BUCK_OUT1_SEL bits[2:0] of the PWR_BUCK_CTL register to '0x02' for 0.9 V and '0x05' for 1.1 V.
3. In the case of 1.1 V to 0.9 V transition, the time it takes to discharge or settle to the new voltage may depend on the load. So the system can continue to operate while the voltage discharges.
4. In the case of 0.9 V to 1.1 V transition, wait 9 μ s for the regulator to stabilize to the new voltage. The clock frequency can be increased after the settling delay.

Notes:

- Refer to the [device datasheet](#) for characterized numbers for the settling intervals and load current limitations.
- When changing clock frequencies, make sure to update wait states of RAM/ROM/FLASH. Refer to the [CPU Subsystem \(CPUSS\) chapter on page 33](#) for details on the wait states.

Other Low-power Regulators

In addition to the core active regulators, the device includes multiple low-power regulators for powering Deep Sleep peripherals/logic ($V_{CCDPSLP}$), digital retention logic/SRAM (V_{CCRET}), and N-wells (V_{NWELL}) in the device. Note that V_{NWELL} is not shown in [Figure 17-1](#) because this rail is used across the device for powering all the N-wells in the chip. These rails are shorted to V_{CCD} in Active and Sleep power modes. V_{CCRET} powers all the Active mode logic that needs to be in retention in Deep Sleep mode.

Note that none of these power rails are available in Hibernate mode. In Hibernate mode, all the hibernate logic and peripherals operate from V_{DDD} directly and a Hibernate wakeup resets the device. For details, refer to the [Device Power Modes chapter on page 190](#).

17.3.2 Power Pins and Rails

Table 17-1 lists all the power supply pins available in the device. The supply rails running inside the device ($V_{CCDPSLP}$, V_{CCRET} , V_{DDBAK} , and V_{NWELL}) are derived from these external supply pins/rails.

Table 17-1. Supply Pins

Supply Pin	Ground Pin	Voltage Range Supported ^a	Description
V_{DD} or V_{DDD}	V_{SS} or V_{SSD}	1.7 V to 3.6 V	V_{DD} is a single supply input for multiple supplies and V_{DDD} is a digital supply input. Either of the pins will be present in a given package.
V_{DD_NS}	V_{SS_NS}	1.7 V to 3.6 V	Supply input to the buck regulator on-chip.
V_{CCD}	V_{SS} or V_{SSD}	Capacitor (0.9 V or 1.1 V)	Core supply or bypass capacitor for the internal core regulator (LDO).
V_{DDA}	V_{SS} or V_{SSA} (if available)	1.7 V to 3.6 V	Analog supply voltage.
V_{DDIO}	V_{SS} or V_{SSIO} (if available)	1.7 V to 3.6 V	Additional I/O supply voltage.
V_{BACKUP}	V_{SS} or V_{SSD} (if available)	1.4 V to 3.6 V	Backup domain supply voltage.
V_{IND1}	—	Inductor	Inductor connection for the internal buck regulator.

a. Refer to the [device datasheet](#) for exact range and recommended connections.

17.3.3 Power Sequencing Requirements

V_{DDD} , V_{BACKUP} , V_{DDIO} , and V_{DDA} do not have any sequencing limitation and can establish in any order. The presence of V_{DDA} without V_{DD} or V_{DDD} can cause some leakage from V_{DDA} . However, it will not drive any analog or digital output. All the V_{DDA} pins in packages that offer multiple V_{DDA} supply pins, must be shorted externally (on the PCB). Note that the system will not exit reset until both V_{DDD} and V_{DDA} are established. However, it will not wait for other supplies to establish.

17.3.4 Backup Domain

The PSoC 6 MCU offers an independent backup supply option (V_{BACKUP}). This rail powers a small set of peripherals that includes an RTC, WCO, and a small number of retention registers. This rail is independent of all other rails and can exist even when other rails are absent. As [Figure 17-1](#) shows, this pin sources the V_{DDBAK} rail in the device. The V_{DDBAK} rail is connected to V_{DDD} when no V_{BACKUP} supply exists. For details on the backup domain, refer to the [Backup System chapter on page 200](#).

17.3.5 Power Supply Sources

The PSoC 6 MCU offers power supply options that support a wide range of application voltages and requirements. V_{DDD} input supports a voltage range of 1.7 V to 3.6 V. If the application voltage is in this range, then the PSoC 6 MCU (V_{DDD}) can be interfaced directly to the application voltage. In applications that have voltage beyond this range, a suitable PMIC (Buck or Boost or Buck-Boost) should be used to bring the voltage to this range.

Other supply rails and pins such as V_{DDA} , V_{DDIO} , and V_{BACKUP} exist independent of V_{DDD} and V_{CCD} .

17.4 Voltage Monitoring

The PSoC 6 MCU offers multiple voltage monitoring and supply failure protection options. This includes POR, BOD, LVD, and OVP.

17.4.1 Power-On-Reset (POR)

POR circuits provide a reset pulse during the initial power ramp. POR circuits monitor V_{DDD} voltage. Typically, the POR circuits are not very accurate about the trip-point.

POR circuits are used during initial chip power-up and then disabled. Refer to the [device datasheet](#) for details on the POR trip-point levels.

17.4.2 Brownout-Detect (BOD)

The BOD circuit protects the operating or retaining logic from possibly unsafe supply conditions by applying reset to the device. The PSoC 6 MCU offers two BOD circuits – high-voltage BOD (HVBOD) and low-voltage BOD (LVBOD). The HVBOD monitors the V_{DDD} voltage and LVBOD monitors the V_{CCD} voltage. Both BOD circuits generate a reset if a voltage excursion dips below the minimum V_{DDD}/V_{CCD} voltage required for safe operation (see the [device datasheet](#) for details). The system will not come out of RESET until the supply is detected to be valid again.

The HVBOD circuit guarantees a reset in System LP, ULP, and Deep Sleep power modes before the system crashes, provided the V_{DDD} supply ramp satisfies the datasheet

maximum supply ramp limits in that mode. There is no BOD support in Hibernate mode. Applications that require BOD support should not use Hibernate mode and should disable it. Refer to the [Device Power Modes chapter on page 190](#) for details.

The LVBOD, operating on V_{CCD} , is not as robust as the HVBOD. The limitation is because of the small voltage detection range available for LVBOD on the minimum allowed V_{CCD} .

For details on the BOD trip-points, supported supply ramp rate, and BOD detector response time, refer to the [device datasheet](#).

17.4.3 Low-Voltage-Detect (LVD)

An LVD circuit monitors external supply voltage and accurately detects depletion of the energy source. The LVD generates an interrupt to cause the system to take preventive measures.

The LVD can be configured to monitor V_{DD} , V_{AMUXA} , V_{AMUXB} , or V_{DDIO} . The HVLVD1_SRCSEL bits [6:4] of the PWR_LVD_CTL register selects the source of the LVD. The LVD support up to 15 voltage levels (thresholds) to monitor between 1.2 V to 3.1 V. The HVLVD1_TRIPSEL bits [3:0] of the PWR_LVD_CTL register select the threshold levels of the LVD. LVD should be disabled before selecting the threshold. The HVLVD1_EN bit [7] of the PWR_LVD_CTL register can be used to enable or disable the LVD.

Whenever the voltage level of the supply being monitored drops below the threshold, the LVD generates an interrupt. This interrupt status is available in the SRSS_INTR register. HVLVD1 bit [1] of the SRSS_INTR register indicates a pending LVD interrupt. The SRSS_INTR_MASK register decides whether LVD interrupts are forwarded to the CPU or not.

Note that the LVD circuit is available only in Active, LPACTIVE, Sleep, and LPSLEEP power modes. If an LVD is required in Deep Sleep mode, then the device should be configured to periodically wake up from Deep Sleep mode using a Deep Sleep wakeup source. This makes sure an LVD check is performed during Active/LPACTIVE mode.

When enabling the LVD circuit, it is possible to receive a false interrupt during the initial settling time. Firmware can mask this by waiting for 8 μ s after setting the HVLVD1_EN bit in the PWR_LVD_CTL register. The recommended firmware procedure to enable the LVD function is:

1. Ensure that the HVLVD1 bit in the SRSS_INTR_MASK register is 0 to prevent propagating a false interrupt.
2. Set the required trip-point in the HVLVD1_TRIPSEL field of the PWR_LVD_CTL register.
3. Configure the LVD edge (falling/rising/both) that triggers the interrupt by configuring HVLVD1_EDGE_SEL bits[1:0] of SRSS_INTR_CFG register. By default, the

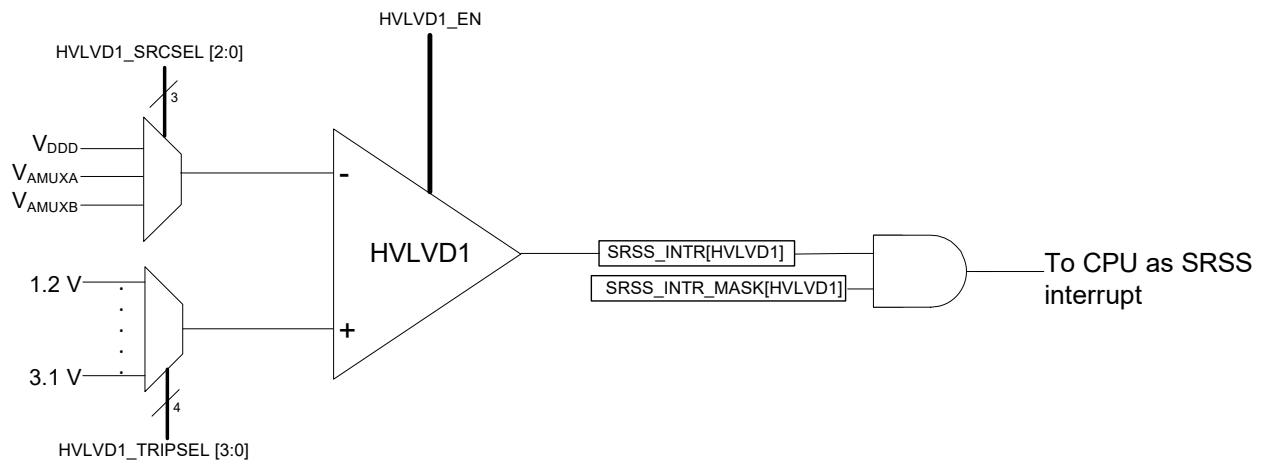
configuration disables the interrupt. **Note:** LVD logic may falsely detect a falling edge during Deep Sleep entry. This applies only when HVLVD1_EDGE_SEL is set to FALLING(2) or BOTH(3). Firmware can workaround this condition by disabling falling edge detection before entering Deep Sleep, and re-enabling it after exiting Deep Sleep.

4. Enable the LVD by setting the HVLVD1_EN bit in the PWR_LVD_CTL register. This may cause a false LVD event.
5. Wait at least 8 μ s for the circuit to stabilize.
6. Clear any false event by setting the HVLVD1 bit in the SRSS_INTR register. The bit will not clear if the LVD condition is truly present.
7. Unmask the interrupt by setting the HVLVD1 bit in SRSS_INTR_MASK.

For details on supported LVD thresholds, refer to the [device datasheet](#) and the PWR_LVD_CTL register definition in the [registers TRM](#).

Figure 17-2. PSoC 6 MCU LVD Block

PSoC 6 LVD block



17.4.4 Over-Voltage Protection (OVP)

The PSoC 6 MCU offers an over-voltage protection circuit that monitors the V_{CCD} supply. Similar to the BOD circuit, the OVP circuit protects the device from unsafe supply conditions by applying a reset. As the name suggests, the OVP circuit applies a device reset, when the V_{CCD} supply goes above the maximum allowed voltage. The OVP circuit can generate a reset in all device power modes except the Hibernate mode.

17.5 Register List

Name	Description
PWR_CTL	Power Mode Control register - controls the device power mode options and allows observation of current state
PWR_BUCK_CTL	Buck Control register - controls the buck output and master buck enable
PWR_BUCK_CTL2	Buck Control register 2 - controls the VRF output
PWR_LVD_CTL	LVD Configuration register
SRSS_INTR	SRSS Interrupt register - shows interrupt requests from the SRSS peripheral
SRSS_INTR_MASK	SRSS Interrupt Mask register - controls forwarding of the interrupt to CPU
SRSS_INTR_SET	SRSS Interrupt Set register - sets interrupts; this register is used for firmware testing
SRSS_INTR_MASKED	SRSS Interrupt Masked register - logical AND of corresponding SRSS interrupt request (SRSS Interrupt register) and mask bits (SRSS Interrupt Mask register)

18. Device Power Modes



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SysPM
- Application notes
- Code examples

The PSoC 6 MCU can operate in four system and three CPU power modes. These modes are intended to minimize the average power consumption in an application. The power modes supported by PSoC 6 MCUs, in the order of decreasing power consumption, are:

- System Low Power (LP) – All peripherals and CPU power modes are available at maximum speed and current
- System Ultra Low Power (ULP) – All peripherals and CPU power modes are available, but with limited speed and current
- CPU Active – CPU is executing code in system LP or ULP mode
- CPU Sleep – CPU code execution is halted in system LP or ULP mode
- CPU Deep Sleep – CPU code execution is halted and system deep sleep is requested while in system LP or ULP mode
- System Deep Sleep – Entered only after both CPUs enter CPU Deep Sleep mode. Only low-frequency peripherals are available
- System Hibernate – Device and I/O states are frozen and the device resets on wakeup

CPU Active, Sleep, and Deep Sleep are standard Arm-defined power modes supported by the Arm CPU instruction set architecture (ISA). System LP, ULP, Deep Sleep and Hibernate modes are additional low-power modes supported by PSoC 6. Hibernate mode is the lowest power mode in the PSoC 6 MCU and on wakeup, the CPU and all peripherals go through a reset.

18.1 Features

The PSoC 6 MCU power modes have the following features:

- Four system and three CPU power modes aimed at optimizing power consumption in an application
- System ULP mode with reduced operating current and clock frequency while supporting full device functionality
- System Deep Sleep mode with support for multiple wakeup sources and configurable amount of SRAM retention
- System Hibernate mode with wakeup from I/O, comparator, WDT, RTC, and timer alarms

The power consumption in different power modes is further controlled by using the following methods:

- Enabling and disabling clocks to peripherals
- Powering on/off clock sources
- Powering on/off peripherals and resources inside the PSoC 6 device

18.2 Architecture

The PSoC 6 device supports multiple power modes. Some modes only affect the CPUs (CPU power modes) and others affect the whole system (system power modes). The system and CPU power modes are used in combination to control the total system performance and power. CPU power modes are entered separately for each CPU on the device.

The SysPm Peripheral Driver Library (PDL) driver supports all device power mode transitions and is the recommended method of transition and configuration of PSoC 6 MCU power resources.

Table 18-1 summarizes the power modes available in PSoC 6 MCUs, their description, and details on their entry and exit conditions.

Table 18-1. PSoC 6 MCU Power Modes

System Power Mode	MCU Power Mode	Description	Entry Conditions	Wakeup Sources	Wakeup Action
LP	Active	System – Primary mode of operation. 1.1 V core voltage. All peripherals are available (programmable). Maximum clock frequencies CPU – Active mode	Reset from external reset, brownout, power on reset system and Hibernate mode. Manual register write from system ULP mode. Wakeup from CPU Sleep or CPU Deep Sleep while in system LP mode. Wakeup from system Deep Sleep after entered from LP mode.	Not applicable	N/A
	Sleep	1.1 V core voltage. One or more CPUs in Sleep mode (execution halted). All peripherals are available (programmable). Maximum clock frequencies	In system LP mode, CPU executes WFI/WFE instruction with Deep Sleep disabled	Any interrupt to CPU	Interrupt
	Deep Sleep	1.1 V core voltage. One CPU in Deep Sleep mode (execution halted). Other CPU in Active or Sleep mode. All peripherals are available (programmable). Maximum clock frequencies	In system LP mode, CPU executes WFI/WFE instruction with Deep Sleep enabled	Any interrupt to CPU	Interrupt
ULP	Active	0.9 V core voltage. All peripherals are available (programmable). Limited clock frequencies. No Flash write.	Manual register write from system LP mode. Wakeup from CPU Sleep or CPU Deep Sleep while in system ULP mode. Wakeup from system Deep Sleep after entered from ULP mode.	Not applicable	N/A
	Sleep	0.9 V core voltage. One or more CPUs in Sleep mode (execution halted). All peripherals are available (programmable). Limited clock frequencies. No Flash write.	In system ULP mode, CPU executes WFI/WFE instruction with Deep Sleep disabled	Any interrupt to CPU	Interrupt
	Deep Sleep	0.9 V core voltage. One CPU in Deep Sleep mode (execution halted). Other CPU in Active or Sleep mode. All peripherals are available (programmable). Limited clock frequencies. No Flash write.	In system ULP mode, CPU executes WFI/WFE instruction with Deep Sleep enabled	Any interrupt to CPU	Interrupt
Deep Sleep	Deep Sleep	All high-frequency clocks and peripherals are turned off. Low-frequency clock (32 kHz) and low-power analog and digital peripherals are available for operation and as wakeup sources. SRAM is retained (programmable).	Both CPUs simultaneously in CPU Deep Sleep mode.	GPIO interrupt, Low-Power comparator, SCB, watchdog timer, and RTC alarms	Interrupt
Hibernate	N/A	GPIO states are frozen. All peripherals and clocks in the device are completely turned off except optional low-power comparators and backup domain. Wakeup is possible through WAKEUP pins, XRES, low-power comparator (programmable), WDT, and RTC alarms (programmable). Device resets on wakeup.	Manual register write from LP or ULP modes.	WAKEUP pin, low-power comparator, watchdog timer ^a , and RTC ^b alarms	Reset

a. Watchdog timer is capable of generating a hibernate wakeup. See the [Watchdog Timer chapter on page 248](#) for details.

b. RTC (along with WCO) is part of the backup domain and is available irrespective of the device power mode. RTC alarms are capable of waking up the device from any power mode.

18.2.1 CPU Power Modes

The CPU Active, Sleep, and Deep Sleep modes are the standard Arm-defined power modes supported by both Cortex-M4 and Cortex-M0+ CPUs. All Arm CPU power modes are available in both system LP and ULP power modes. CPU power modes affect each CPU independently.

18.2.1.1 CPU Active Mode

In CPU Active mode, the CPU executes code and all logic and memory is powered. The firmware may decide to enable or disable specific peripherals and power domains depending on the application and power requirements. All the peripherals are available for use in Active mode. The device enters CPU Active mode upon any device reset or wakeup.

18.2.1.2 CPU Sleep Mode

In CPU Sleep mode, the CPU clock is turned off and the CPU halts code execution. Note that in the PSoC 6 MCU, Cortex-M4 and Cortex-M0+ both support their own CPU Sleep modes and each CPU can be in sleep independent of the other CPU state. All peripherals available in Active mode are also available in Sleep mode. Any peripheral interrupt, masked to the CPU, will wake the CPU to Active mode. Only the CPU(s) with the interrupt masked will wake.

18.2.1.3 CPU Deep Sleep Mode

In CPU Deep Sleep mode, the CPU requests the device to go into system Deep Sleep mode. When the device is ready, it enters Deep Sleep mode as detailed in [18.2.3 System Deep Sleep Mode](#).

Because PSoC 6 has more than one CPU, both CPUs must independently enter CPU Deep Sleep before the system will transition to system Deep Sleep.

18.2.2 System Power Modes

System power modes affect the whole device and may be combined with CPU power modes.

18.2.2.1 System Low Power Mode

System Low Power (LP) mode is the default operating mode of the device after reset and provides maximum system performance. In LP mode all resources are available for operation at their maximum power level and speed.

While in system LP mode the CPUs may operate in any of the standard Arm defined CPU modes detailed in [18.2.1 CPU Power Modes](#).

18.2.2.2 System Ultra Low Power Mode

System Ultra Low Power (ULP) mode is identical to LP mode with a performance tradeoff made to achieve lower system current. This tradeoff lowers the core operating

voltage, which then requires reduced operating clock frequency and limited high-frequency clock sources. Flash write operations are not available in ULP mode. [Table 18-5](#) provides the list of resources available in ULP mode along with limitations.

While in system ULP mode, the CPUs may operate in any of the standard Arm-defined CPU modes detailed in [18.2.1 CPU Power Modes](#).

Transitioning between LP and ULP modes is performed by reducing the core regulator voltage from the LP mode voltage to the ULP mode voltage. The lower voltage reduces system operating current and slows down signal speeds requiring a lower maximum operating frequency. Refer to Core Operating Voltage section in [Power Supply and Monitoring chapter page 183](#) for details on how to switch between LP and ULP modes.

18.2.3 System Deep Sleep Mode

In system Deep Sleep mode, all the high-speed clock sources are off. This in turn makes high-speed peripherals unusable in system Deep Sleep mode. However, low-speed clock sources and peripherals may continue to operate, if configured and enabled by the firmware. In addition, the peripherals that do not need a clock or receive a clock from their external interface (I²C or SPI slave) may continue to operate, if configured for system Deep Sleep operation. The PSoC 6 MCU provides an option to configure the amount of SRAM, in blocks of 32 KB, that are retained during Deep Sleep mode.

Both Cortex-M0+ and Cortex-M4 can enter CPU Deep Sleep mode independently. However, the entire device enters system Deep Sleep mode only when both the CPUs are in CPU Deep Sleep. On wakeup, the CPU that woke up enters CPU Active mode and the other CPU remains in CPU Deep Sleep mode. On wakeup, the system will return to LP or ULP mode based on what mode was Active before entering system Deep Sleep. Both CPUs may wake up to CPU Active simultaneously from the same wakeup source if so configured.

The device enters system Deep Sleep mode after the following conditions are met.

- The LPM_READY bit of the PWR_CTL register should read '1'. This ensures the system is ready to enter Deep Sleep. If the bit reads '0', then the device will wait in system LP or ULP mode instead of system Deep Sleep until the bit is set, at which instant the device automatically enters Deep Sleep mode, if requested.
- Both Cortex-M0+ and Cortex-M4 must be in CPU Deep Sleep. This is achieved by setting the SLEEPDEEP bit [2] of the SCR register of both Cortex-M0+ and Cortex-M4 and then executing WFI or WFE instruction.
- The HIBERNATE bit [31] of the PWR_HIBERNATE register should be cleared; otherwise, the device will enter system Hibernate mode.

In system Deep Sleep mode, the LP and ULP mode regulator is turned off and a lower power, Deep Sleep regulator sources all the peripherals enabled in system Deep Sleep mode. Alternatively, the buck regulator can be used to power the Deep Sleep peripherals. See the [Power Supply and Monitoring chapter on page 183](#) for details. [Table 18-5](#) provides the list of resources available in system Deep Sleep mode.

Interrupts from low-speed, asynchronous, or low-power analog peripherals can cause a CPU wakeup from system Deep Sleep mode. Note that when a debugger is running on either core, the device stays in system LP or ULP mode and the CPUs enter CPU Sleep mode instead of CPU Deep Sleep mode. PSoC 6 uses buffered writes. Therefore, writes to an MMIO register or memory can take few a cycles from the write instruction execution. The only way to ensure that the write operation is complete is by reading the same location after the write. It is required to follow a write by a read to the same location before executing a WFI/WFE instruction to enter CPU Deep Sleep mode.

18.2.4 System Hibernate Mode

System Hibernate mode is the lowest power mode of the device when external supplies are still present and XRES is deasserted. It is intended for applications in a dormant state. In this mode, both the Active LP/ULP mode regulator and Deep Sleep regulator are turned off and GPIO states are automatically frozen. Wakeup is facilitated through dedicated wakeup pins and a Low-Power comparator output. Low-Power comparator operation in Hibernate mode requires externally generated voltages for wakeup comparison. Internal references are not available in Hibernate mode. Optionally, an RTC alarm from the backup domain or a watchdog timer (16-bit free-running WDT) interrupt can generate a Hibernate wakeup signal. Set the MASK_HIBALARM bit [18] of the PWR_HIBERNATE register to enable the RTC alarm wakeup from Hibernate mode.

The device goes through a reset on wakeup from Hibernate. I/O pins remain in the configuration they were frozen before entering Hibernate mode. To differentiate between other system resets and a Hibernate mode wakeup, the TOKEN bits [7:0] of the PWR_HIBERNATE register can be used as described in the [Power Mode Transitions on page 195](#). The PWR_HIBERNATE (except the HIBERNATE bit [31]) register along with the PWR_HIB_DATA register are retained through the Hibernate wakeup sequence and can be used by the application for retaining some content. Note that these registers are reset by other reset events. On a Hibernate wakeup event, the HIBERNATE bit [31] of the PWR_HIBERNATE register is cleared.

The brownout detect (BOD) block is not available in Hibernate mode. As a result, the device does not recover from a brownout event in Hibernate mode. Do not enter Hibernate mode in applications that require brownout detection, that is, applications where the supply is not

stable. In addition, make sure the supply is stable for at least 250 μ s before the device enters Hibernate mode. To prevent accidental entry into Hibernate mode in applications that cannot meet these requirements, an option to disable the Hibernate mode is provided. Set the HIBERNATE_DISABLE bit [30] of the PWR_HIBERNATE register to disable Hibernate mode in the device. Note that this bit is a write-once bit during execution and is cleared only on reset. Debug functionality will be lost and the debugger will disconnect on entering Hibernate mode.

18.2.5 Other Operation Modes

In addition to the power modes discussed in the previous sections, there are three other states the device can be in – Reset, Off, and Backup states. These states are determined by the external power supply and XRES connections. No firmware action is required to enter these modes nor an interrupt or wakeup event to exit them.

18.2.5.1 Backup Domain

PSoC 6 offers an independent backup supply option that can be supplied through a separate Vbackup pin. For details on the backup domain and the powering options, refer to the [Backup System chapter on page 200](#). This domain powers a real-time clock (RTC) block, WCO, and a small set of backup registers. Because the power supply to these blocks come from a dedicated Vbackup pin, these blocks continue to operate in all CPU and system power modes, and even when the device power is disconnected or held in reset as long as a Vbackup supply is provided. The RTC present in the backup domain provides an option to wake up the device from any CPU or system power mode. The RTC can be clocked by an external crystal (WCO) or the internal low-speed oscillator (ILO). However, the ILO is available only if the device is powered – the device should not be in the off or reset state. Using ILO is not recommended for timekeeping purpose; however, it can be used for wakeup from Hibernate power mode.

18.2.5.2 Reset State

Reset is the device state when an external reset (XRES pin pulled low) is applied or when POR/BOD is asserted. Reset is not a power mode. During the reset state, all the components in the device are powered down and I/Os are tristated, keeping the power consumption to a minimum.

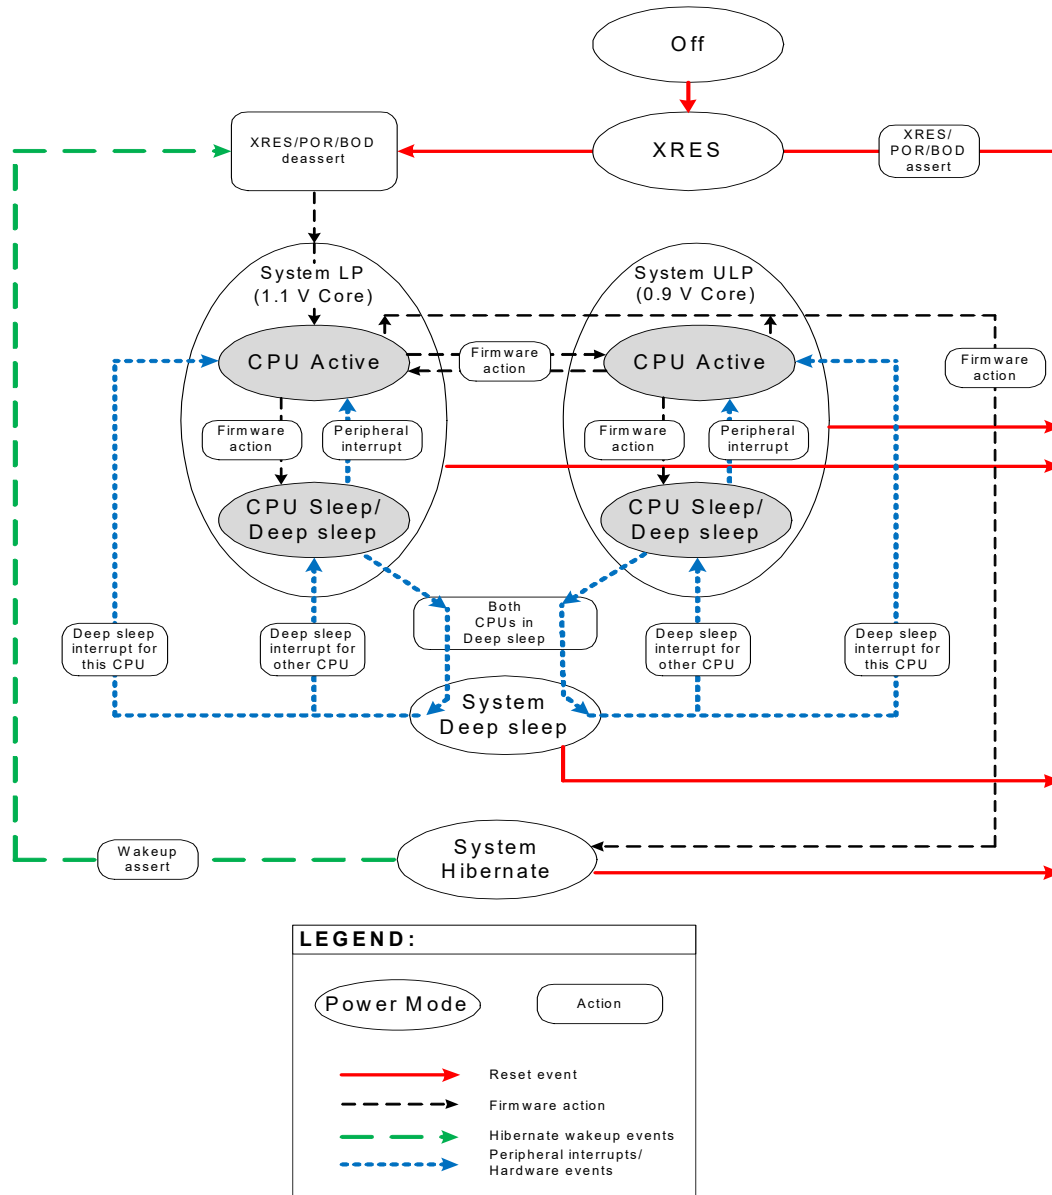
18.2.5.3 Off State

The off state simply represents the device state with no power applied. Even in the device off state, the backup domain can continue to receive power (Vbackup pin) and run the peripherals present in that domain. The reset and off states are discussed for completeness of all possible modes and states the device can be in. These states can be used in a system to further optimize power consumption. For instance, the system can control the supply of the PSoC 6 MCU by enabling or disabling the regulator output powering the device using the PMIC interface.

18.3 Power Mode Transitions

Figure 18-1 shows various states the device can be in along with possible power mode transition paths.

Figure 18-1. Power Mode Transitions in PSoC 6 MCU



18.3.1 Power-up Transitions

Table 18-2 summarizes various power-up transitions, their type, triggers, and actions.

Table 18-2. Power Mode Transitions

Initial State	Final State	Type	Trigger	Actions
Of	XRES	External	Power rail (V_{DD}) ramps up above POR voltage level with XRES pin asserted.	1. All high-voltage logic is reset
Off	Reset	External	Power rail (V_{DD}) ramps up above POR voltage level with XRES pin de-asserted.	1. All high-voltage logic is reset 2. Low-voltage (internal core and Deep Sleep mode) regulators and references are ramped up 3. All low-voltage logic (logic operating from internal regulators) is reset 4. IMO clock is started
XRES	Reset	External	XRES is de-asserted with V_{DD} present and above POR level.	1. Low-voltage regulators and references are ramped up 2. All low-voltage logic is reset 3. IMO clock is started
Reset	Active	Internal	Reset sequence completes. This transition can also be caused by internal resets.	1. Clock is released to the system 2. System reset is de-asserted 3. CPU starts execution

18.3.2 Power Mode Transitions

Table 18-3. Power Mode Transitions

Initial State	Final State	Type	Trigger	Actions
System LP	System ULP	Internal	Firmware action 1. Ensure the Clk_HF paths, peripheral, and slow clocks are less than the ULP clock speed limitations. 2. Flash/SRAM/ROM wait states values are increased to ULP values as detailed in the Nonvolatile Memory chapter on page 127 . 3. Configure the core regulator to 0.9 V.	1. Device is put into ULP mode with all peripherals available with limited speed. Flash write operation are not supported.
System ULP	System LP	Internal	Firmware action 1. Configure the core regulator to 1.1 V. 2. Wait 9 μ s to allow the core voltage to stabilize at the new value. 3. Flash/SRAM/ROM wait states values are decreased to LP values as detailed in the Nonvolatile Memory chapter on page 127 . 4. Increase the Clk_HF paths, peripheral, and slow clocks as desired up to the maximum LP clock speed specifications.	1. Device is put into LP mode with all peripherals available with maximum speed. Flash write operations are supported.

Table 18-3. Power Mode Transitions

Initial State	Final State	Type	Trigger	Actions
System LP/ ULP and CPU Active	System LP/ ULP and CPU Sleep	Internal	Firmware action 1. Clear the SLEEPDEEP bit [2] of the SCR register for both Cortex-M0+ and Cortex-M4. 2. Optionally, set the SLEEPONEXIT bit [1] of the SCR register, if the CPU runs only on interrupts. When this bit is set, the CPU does not return to application code after the WFI/WFE instruction is executed. The CPU wakes up on any enabled (masked to CPU) interrupt or event and enters CPU Sleep mode as soon as it exits the interrupt or services the event. 3. Optionally, set the SEVONPEND bit [4] of the SCR register if the application must wake up the CPU from any pending interrupt. If this bit is set, any interrupt that enters a pending state wakes up the CPU. This includes all the disabled (unmasked) interrupts to CPU. 4. Execute WFI/WFE instruction on both CPUs.	1. CPU clocks are gated off 2. CPU waits for an interrupt or event to wake it up.
System LP/ ULP and CPU Active	System LP/ UP and CPU Deep Sleep	Internal	Firmware action Perform these steps to enter Deep Sleep mode (LPM_READY bit [5] of the PWR_CTL register should read '1' before performing these steps): 1. Clear the HIBERNATE bit [31] of the PWR_HIBERNATE register. 2. Set the SLEEPDEEP bit [2] of the SCR register for one or both Cortex-M0+ and Cortex-M4. 3. Optionally, set the SLEEPONEXIT bit [1] of the SCR register if the CPU runs only on interrupts. When this bit is set, the CPU does not return to application code after the WFI/WFE instruction is executed. The CPU wakes up on any enabled (masked to CPU) interrupt or event and enters CPU Deep Sleep mode as soon as it exits the interrupt or services the event. 4. Optionally, set the SEVONPEND bit [4] of the SCR register if the application needs to wake up the CPU from any pending interrupt. If this bit is set, any interrupt that enters a pending state wakes up the CPU. This includes all the disabled (unmasked) interrupts to CPU. 5. Read the SCR register before executing a WFI/WFE instruction to ensure the write operation is complete. PSoC 6 uses buffered writes and any write transfer just before executing WFI/WFE instruction should be followed by a read to the same memory location. This ensures that the write operation has taken effect before entering Deep Sleep mode. Execute WFI/WFE instruction on both CPUs. 6. CPU in Deep Sleep mode generates a hardware request for the whole device to enter system Deep Sleep. A CPU waiting in CPU Deep Sleep state is functionally identical to CPU Sleep mode with the exception of the hardware request. 7. If only one CPU is in Deep Sleep mode the system will remain in system LP or ULP mode until the other CPU also enters CPU Deep Sleep. While waiting, a masked interrupt can wake the CPU to Active mode.	1. CPU clocks are gated off 2. CPU waits for a Deep Sleep interrupt to wake it up.
System LP/ ULP and CPU Deep Sleep	System Deep Sleep	Internal	Hardware action 1. When both CPUs enter CPU Deep Sleep mode and the LPM_READY bit [5] of the PWR_CTL register reads '1', the device will automatically transition to system Deep Sleep power mode.	1. High-frequency clocks are shut down. 2. Retention is enabled and non-retention logic is reset. 3. Active regulator is disabled and Deep Sleep regulator takes over.

Table 18-3. Power Mode Transitions

Initial State	Final State	Type	Trigger	Actions
System LP/ ULP and CPU Active	System Hibernate	Internal	<p>Firmware action</p> <p>Perform these steps to enter Hibernate mode (LPM_READY bit [5] of the PWR_CTL register should read '1' before performing these steps):</p> <ol style="list-style-type: none"> 1. Set the TOKEN bits [7:0] of the PWR_HIBERNATE register (optional) and PWR_HIB_DATA register to some application-specific branching data that can be used on a wakeup event from Hibernate mode. 2. Set the UNLOCK bits [8:15] of the PWR_HIBERNATE register to 0x3A, this ungates writes to FREEZE and HIBERNATE bits of the PWR_HIBERNATE register. 3. Configure wakeup pins polarity (POLARITY_HIBPIN bits [23:20]), wakeup pins mask (MASK_HIBPIN bits [27:24]) and wakeup alarm mask (MASK_HIBALARM bit [18]) in the PWR_HIBERNATE register based on the application requirement. 4. Optionally, set the FREEZE bit [17] of the PWR_HIBERNATE register to freeze the I/O pins. 5. Set the HIBERNATE bit [31] of the PWR_HIBERNATE register to enter Hibernate mode. 6. Ensure that the write operation to the PWR_HIBERNATE register is complete by reading the PWR_HIBERNATE register. Otherwise, instead of entering Hibernate mode, the CPU may start executing instructions after the write instruction. When the write is followed by a WFI/WFE instruction, the WFI/WFE can prevent the write from taking effect. On completion of the write operation, the device automatically enters Hibernate mode. 	<ol style="list-style-type: none"> 1. CPU enters low-power mode. 2. Both high-frequency and low-frequency clocks are shut down. 3. Retention is enabled and non-retention logic is reset. 4. Both Active and Deep Sleep regulators are powered down. The peripherals that are active in the Hibernate domain operate directly out of V_{DDD}.

18.3.3 Wakeup Transitions

Table 18-4. Wakeup Transitions

Initial State	Final State	Type	Trigger	Actions
CPU Sleep	CPU Active	Internal/External	Any peripheral interrupt masked to CPU	<ol style="list-style-type: none"> 1. Clock to CPU is ungated. 2. Peripheral interrupt is serviced by CPU. 3. Device remains in current system LP or ULP power mode.
System Deep Sleep	System LP/ ULP Active and CPU Active	Internal/External	Any Deep Sleep interrupt	<ol style="list-style-type: none"> 1. Active regulator and references are enabled. 2. Retention is disabled and non-retention reset is de-asserted. 3. High-frequency clocks are turned on. 4. CPU exits low-power mode and services the interrupt. 5. Returns to previous system LP or ULP power mode. <p>Note: If only one CPU wakes up from the system Deep Sleep interrupt, then the other CPU remains in the CPU Deep Sleep state until its Deep Sleep interrupt wakes it up. However, the system will wake up to LP or ULP mode. Note that a Deep Sleep interrupt can wake up either one or both CPUs depending on the WIC configuration for the CPU. See the Interrupts chapter on page 54.</p>
Hibernate	System LP and CPU Active	External	Wakeup pin, RTC alarm, WDT interrupt, or Low-Power comparator output asserts	<ol style="list-style-type: none"> 1. Device is reset and goes through a reset to active power-up transition. 2. Optionally, read the TOKEN bits [7:0] of the PWR_HIBERNATE and PWR_HIB_DATA registers for application-specific branching from hibernate wakeup. 3. Optionally, set the I/O drive modes by reading the I/O frozen output and setting the I/O output to the read value. 4. Unfreeze the I/O cells by clearing the FREEZE bit [17] of the PWR_HIBERNATE register.

18.4 Summary

Table 18-5. Resources Available in Different Power Modes

Component	Power Modes							
	LP		ULP		Deep Sleep	Hibernate	XRES	Power Off with Backup
	CPU Active	CPU Sleep/Deep Sleep	CPU Active	CPU Sleep/Deep Sleep				
Core functions								
CPU	On	Sleep	On	Sleep	Retention	Off	Off	Off
SRAM	On	On	On	On	Retention	Off	Off	Off
Flash	Read/Write	Read/Write	Read Only	Read Only	Off	Off	Off	Off
High-Speed Clock (IMO, ECO, PLL, FLL)	On	On	On	On	Off	Off	Off	Off
LVD	On	On	On	On	Off	Off	Off	Off
ILO	On	On	On	On	On	On	Off	Off
Peripherals								
SMIF/SHDC	On	On	On	On	Retention	Off	Off	Off
SAR ADC	On	On	On	On	Off	Off	Off	Off
LPCMP	On	On	On	On	On ^a	On ^a	Off	Off
TCPWM	On	On	On	On	Off	Off	Off	Off
CSD	On	On	On	On	Retention	Off	Off	Off
LCD	On	On	On	On	On	Off	Off	Off
SCB	On	On	On	On	Retention (I ² C/SPI wakeup available) ^b	Off	Off	Off
GPIO	On	On	On	On	On	Freeze	Off	Off
Watchdog timer	On	On	On	On	On	On	Off	Off
Multi-Counter WDT	On	On	On	On	On	Off	Off	Off
Resets								
XRES	On	On	On	On	On	On	On	Off
POR	On	On	On	On	On	On	Off	Off
BOD	On	On	On	On	On	Off	Off	Off
Watchdog reset	On	On	On	On	On	On ^c	Off	Off
Backup domain								
WCO, RTC, alarms	On	On	On	On	On	On	On	On

a. Low-Power comparator may be optionally enabled in the Hibernate mode to generate wakeup.

b. Only the SCB with system Deep Sleep support is available in the Deep Sleep power mode; other SCBs are not available in the Deep Sleep power mode.

c. Watchdog interrupt can generate a Hibernate wakeup. See the [Watchdog Timer chapter on page 248](#) for details.

18.5 Register List

Name	Description
PWR_CTL	Power Mode Control register – controls the device power mode options and allows observation of current state
PWR_HIBERNATE	Hibernate Mode register – controls various Hibernate mode entry/exit related options
PWR_HIB_DATA	Hibernate Mode Data register – data register that is retained through a hibernate wakeup sequence
CM4_SCS_SCR	Cortex-M4 System Control register – controls the CM4 CPU sleep and deep sleep decisions on the WFI/WFE instruction execution
CM0P_SCS_SCR	Cortex-M0+ System Control register – controls the CM0+ CPU sleep and deep sleep decisions on the WFI/WFE instruction execution

Name	Description
CPUSS_CM0_CTL	
CPUSS_CM0_STATUS	
CPUSS_CM4_PWR_CTL	
CPUSS_CM4_STATUS	

19. Backup System



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SysPM
- Application notes
- Code examples

The Backup domain adds an “always on” functionality to PSoC 6 MCUs using a separate power domain supplied by a backup supply (V_{BACKUP}) such as a battery or supercapacitor. It contains a real-time clock (RTC) with alarm feature, supported by a 32768-Hz watch crystal oscillator (WCO), and power-management IC (PMIC) control.


Backup is not a power mode; it is a power domain with its own power supply, which can be active during any of the device power modes. For more details, see the [Power Supply and Monitoring chapter on page 183](#) and [Device Power Modes chapter on page 190](#).

19.1 Features

- Fully-featured RTC
 - Year/Month/Date, Day-of-Week, Hour : Minute : Second fields
 - All fields binary coded decimal (BCD)
 - Supports both 12-hour and 24-hour formats
 - Automatic leap year correction
- Configurable alarm function
 - Alarm on Month/Date, Day-of-Week, Hour : Minute : Second fields
 - Two independent alarms
- 32768-Hz WCO with calibration
- Automatic backup power switching
- Built-in supercapacitor charger
- External PMIC control
- 32-byte backup registers

19.2 Architecture

Figure 19-1. Block Diagram



The Backup system includes an accurate WCO that can generate the clock required for the RTC with the help of an external crystal or external clock inputs. The RTC has a programmable alarm feature, which can generate interrupts to the CPU. An AHB-Lite interface provides firmware access to MMIO registers in the Backup domain.

An automatic backup power switch selects the V_{DDBAK} supply required to run the blocks in the Backup domain – either V_{DDD} (main power) or V_{BACKUP} (backup battery/supercapacitor power).

The domain also has backup registers that can store 32 bytes of data and retain its contents even when the main supply (V_{DDD}) is OFF as long as the backup supply (V_{BACKUP}) is present. The Backup system can also control an external PMIC that supplies V_{DDD} .

19.3 Power Supply

Power to the backup system (V_{DDBAK}) is automatically switched between V_{DDD} (main supply) and V_{BACKUP} (Backup domain supply). V_{BACKUP} is typically connected to an independent supply derived from a battery or supercapacitor (see the [Power Supply and Monitoring chapter on page 183](#) for more details).

There are no V_{BACKUP} versus V_{DDD} sequencing restrictions for the power selector switch. Either V_{BACKUP} or V_{DDD} may be removed during normal operation, and the Backup system will remain powered.

The V_{DDBAK_CTL} bitfield in the $BACKUP_CTL$ register controls the behavior of the power selector switch. See the [registers TRM](#) for details of this register. Possible options are:

- $V_{DDBAK_CTL} = 0$ (Default mode): Selects V_{DDD} when the brownout detector in the system resources is enabled and no brownout situation is detected (see the [Power Supply and Monitoring chapter on page 183](#) for more details). Otherwise, it selects the highest supply among V_{DDD} and V_{BACKUP} .
- $V_{DDBAK_CTL} = 1, 2, \text{ or } 3$: Always selects V_{BACKUP} for debug purposes.

If a supercapacitor is connected to V_{BACKUP} , the PSoC 6 MCU can charge the supercapacitor while V_{DDD} is available. Supercapacitor charging can be enabled by writing “3C” to the EN_CHARGE_KEY bitfield in the $BACKUP_CTL$ register. Note that this feature is for charging supercapacitors only and cannot safely charge a battery. Do not write this key when V_{BACKUP} is connected to a battery. Battery charging must be handled at the board level using external circuitry.

Note: If V_{DDD} and V_{BACKUP} are connected on the PCB, the Backup domain may require an explicit reset triggered by firmware using the $RESET$ bitfield in the $BACKUP_RESET$ register. This firmware reset is required if the V_{BACKUP} supply was invalid during a previous power supply ramp-up or brownout event. It is not necessary to reset the Backup domain if the RES_CAUSE register indicates a non-power-related reset as the reset cause, or if the PSoC 6 MCU just

woke from the Hibernate power mode and the supply is assumed to have been valid the entire time.

It is possible to monitor the V_{BACKUP} supply using an ADC attached to AMUXBUS-A by setting the $V_{\text{BACKUP_MEAS}}$ bit in the BACKUP_CTL register. Note that the V_{BACKUP} signal is scaled by 10 percent so it is within the supply range of the ADC. See the [SAR ADC chapter on page 453](#) for more details on how to connect the ADC to AMUXBUS-A.

19.4 Clocking

The RTC primarily runs from a 32768-Hz clock, after it is scaled down to one-second ticks. This clock signal can come from either of these internal sources:

- Watch-crystal oscillator (WCO). This is a high-accuracy clock generator that is suitable for RTC applications and requires a 32768-Hz external crystal populated on the application board. WCO can also operate without crystal, using external clock/sine wave inputs. These additional operating modes are explained later in this section. WCO is supplied by the Backup domain and can therefore run without V_{DD} present.
- Alternate Backup Clock (ALTBK): This option allows the use of LFCLK generated by the System Resources Subsystem (SRSS) as the Backup domain clock. Note that LFCLK is not available in all device power modes or when the V_{DD} is removed. (See the [Device Power Modes chapter on page 190](#) for more detail.)

Clock glitches can propagate into the Backup system when LFCLK is enabled or disabled by the SRSS. In addition, LFCLK may not be as accurate as WCO depending on the actual source of LFCLK. Because of these reasons, LFCLK is not recommended for RTC applications. Also, if the WCO is intended as the clock source then choose it directly instead of routing through LFCLK.

For more details on these clocks and calibration, see the [Clocking System chapter on page 207](#).

The RTC clock source can be selected using the CLK_SEL bitfield in the BACKUP_CTL register. The WCO_EN bit in the BACKUP_CTL register can be used to enable or disable the WCO. If the WCO operates with an external crystal, make sure the WCO_BYPASS bit in the BACKUP_CTL register is cleared before enabling the WCO. In addition, the PRESCALER bitfield in BACKUP_CTL must be configured for a prescaler value of 32768.

Note: External crystal and bypass capacitors of proper values must be connected to WCO_IN and WCO_OUT and pins. See the [device datasheet](#) for details of component values and electrical connections. In addition, GPIOs must be configured for WCO_OUT and WCO_IN signals. See the [I/O System chapter on page 226](#) to know how to configure the GPIOs.

19.4.1 WCO with External Clock/Sine Wave Input

The WCO can also operate from external clock/sine wave inputs. In these modes, WCO must be bypassed by setting the WCO_BYPASS bit in the BACKUP_CTL register before enabling the WCO. Also, GPIOs must be configured for WCO_OUT and WCO_IN signals (in Analog mode). The external clock/sine wave input modes, prescaler settings, and electrical connections are as follows:

- 32768-Hz external clock mode: In this mode, WCO_IN is floating and WCO_OUT is externally driven by a 32768-Hz square wave clock toggling between ground and V_{DD} supply levels. In this configuration, the WCO_OUT pin functions as a digital input pin for the external clock. The PRESCALER bitfield in BACKUP_CTL must be configured for a prescaler value of 32768.
- 60-Hz external clock mode: This mode can be used for deriving a clock from the 60-Hz AC mains supply. In this mode, WCO_OUT is floating and WCO_IN is driven with an external sine wave with zero DC offset, derived from the 60-Hz/120-V mains through a 100:1 capacitive divider. For example, a suitable capacitive divider can be formed by connecting a 220-pF/6-V capacitor between WCO_IN and ground, and a 2.2-pF/200-V capacitor between WCO_IN and the 60-Hz/120-V mains input. The PRESCALER bitfield in BACKUP_CTL must be configured for a prescaler value of 60.
- 50-Hz external clock mode: This mode is similar to the 60-Hz mode, and can be used for 50-Hz/220-V mains standard. The capacitive divider explained previously can be modified to fit this type of supply by having a 1-pF/250-V capacitor between WCO_IN and the mains input. The PRESCALER bitfield in BACKUP_CTL must be configured for a prescaler value of 50.

19.4.2 Calibration

The absolute frequency of the clock input can be calibrated using the BACKUP_CAL_CTL register. Calibration only works when the PRESCALER bitfield in BACKUP_CTL is set to 32768.

CALIB_VAL is a 6-bit field in the BACKUP_CAL_CTL register that holds the calibration value for absolute frequency (at a fixed temperature). One bit of this field translates into 128 ticks to be added or removed from the clock count. Therefore, each bit translates to a change of 1.085 ppm ($= 128/(60 \times 60 \times 32768)$).

The CALIB_SIGN field in the BACKUP_CAL_CTL register controls whether the ticks are added (it takes fewer clock ticks to count one second) or subtracted (it takes more clock ticks to count one second).

For more information, see the BACKUP_CAL_CTL register in the [registers TRM](#).

19.5 Reset

The PSoC 6 MCU reset sources that monitor device power supply such as power-on reset (POR) and brownout reset (BOD) cannot reset the backup system as long as the backup supply (V_{DDBAK}) is present. Moreover, internal and external resets such as watchdog timer (WDT) reset and XRES cannot reset the backup system. The backup system is reset only when:

- all the power supplies are removed from the Backup domain, also known as a “cold-start”.
- the firmware triggers a Backup domain reset using the RESET bitfield in the BACKUP_RESET register.

19.6 Real-Time Clock

The RTC consists of seven binary coded decimal (BCD) fields and one control bit as follows:

Table 19-1. RTC Fields

Bitfield Name	Number of Bits	Description
RTC_SEC	7	Calendar seconds in BCD, 0-59
RTC_MIN	7	Calendar minutes in BCD, 0-59
RTC_HOUR	6	Calendar hours in BCD; value depends on 12-hour or 24-hour format set in the BACKUP_RTC_TIME register. In 24-hour mode, bits RTC_HOUR[5:0] = 0–23 In 12-hour mode, bit RTC_HOUR[5] = 0 for AM and 1 for PM Bits RTC_HOUR[4:0] = 1–12
CTRL_12HR	1	Select the 12-hour or 24-hour mode: 1=12HR, 0=24HR
RTC_DAY	3	Calendar day of the week in BCD, 1-7 You should define the meaning of the values
RTC_DATE	6	Calendar day of the month in BCD, 1-31 Automatic leap year correction
RTC_MON	4	Calendar month in BCD, 1-12
RTC_YEAR	8	Calendar year in BCD, 0-99 (Can be used to represent years 2000 -2099)

BCD encoding indicates that each four-bit nibble represents one decimal digit. Constant bits are omitted in the RTC implementation. For example, the maximum RTC_SEC is 59, which can be represented as two binary nibbles 0101b 1001b. However, the most significant bit is always zero and is therefore omitted, making the RTC_SEC a 7-bit field.

The RTC supports both 12-hour format with AM/PM flag, and 24-hour format for “hours” field. The RTC also includes a “day of the week” field, which counts from 1 to 7. You should define which weekday is represented by a value of ‘1’.

The RTC implements automatic leap year correction for the Date field (day of the month). If the Year is divisible by 4, the month of February (Month=2) will have 29 days instead of 28. When the year reaches 2100 - the Year field rolls over from 99 to 00 - the leap year correction will be wrong (2100 is flagged as a leap year which it is not); therefore, an interrupt is raised to allow the firmware to take appropriate actions. This interrupt is called the century interrupt.

User registers containing these bitfields are BACKUP_RTC_TIME and BACKUP_RTC_DATE. See the corresponding register descriptions in the [registers TRM](#) for details. As the user registers are in the high-frequency bus-clock domain and the actual RTC registers run from the low-frequency 32768-Hz clock, reading and writing RTC

registers require special care. These processes are explained in the following section.

19.6.1 Reading RTC User Registers

To start a read transaction, the firmware should set the READ bit in the BACKUP_RTC_RW register. When this bit is set, the RTC registers will be copied to user registers and frozen so that a coherent RTC value can safely be read by the firmware. The read transaction is completed by clearing the READ bit.

The READ bit cannot be set if:

- RTC is still busy with a previous operation (that is, the RTC_BUSY bit in the BACKUP_STATUS register is set)
- WRITE bit in the BACKUP_RTC_RW register is set

The firmware should verify that the above bits are not set before setting the READ bit.

19.6.2 Writing to RTC User Registers

When the WRITE bit in the BACKUP_RTC_RW register is set, data can be written into the RTC user registers; otherwise, writes to the RTC user registers are ignored. When all the RTC writes are done, the firmware must clear the WRITE bit for the RTC update to take effect. After the

WRITE bit is cleared, the hardware will copy all the new data on one single WCO clock edge to ensure coherency to the actual RTC registers.

The WRITE bit cannot be set if:

- RTC is still busy with a previous operation (that is, the RTC_BUSY bit in the BACKUP_STATUS register is set)
- READ bit in the BACKUP_RTC_RW register is set

The firmware should make sure that the values written to the RTC fields form a coherent legal set. The hardware does not check the validity of the written values. Writing illegal values results in undefined behavior of the RTC.

When in the middle of an RTC update with the WRITE bit set, and a brownout, reset, or entry to Deep Sleep or Hibernate mode occurs, the write operation will not be complete. This is because the WRITE bit will be cleared by a reset, and the RTC update is triggered only when this bit is

cleared by a WRITE transaction. If the write operation is in progress (RTC_BUSY), data corruption can occur if the system is reset or enters Deep Sleep or Hibernate mode.

19.7 Alarm Feature

The Alarm feature allows the RTC to be used to generate an interrupt, which may be used to wake up the system from Sleep, Deep Sleep, and Hibernate power modes.

The Alarm feature consists of six fields corresponding to the fields of the RTC: Month/Date, Day-of-Week, and Hour : Minute : Second. Each Alarm field has an enable bit that needs to be set to enable matching; if the bit is cleared, then the field will be ignored for matching.

The Alarm bitfields are as follows:

Table 19-2. Alarm Bitfields

Bitfield Name	Number of Bits	Description
ALARM_SEC	7	Alarm seconds in BCD, 0-59
ALARM_SEC_EN	1	Alarm second enable: 0=disable, 1=enable
ALARM_MIN	7	Alarm minutes in BCD, 0-59
ALARM_MIN_EN	1	Alarm minutes enable: 0=disable, 1=enable
ALARM_HOUR	6	Alarm hours in BCD, value depending on the 12-hour or 24-hour mode In 12-hour mode, bit ALARM_HOUR[5] = 0 for AM and 1 for PM, bits ALARM_HOUR[4:0] = 1-12 In 24-hour mode, bits ALARM_HOUR[5:0] = 0-23
ALARM_HOUR_EN	1	Alarm hour enable: 0=disable, 1=enable
ALARM_DAY	3	Calendar day of the week in BCD, 1-7 You should define the meaning of the values
ALARM_DAY_EN	1	Alarm day of the week enable: 0=disable, 1=enable
ALARM_DATE	6	Alarm day of the month in BCD, 1-31 Leap year corrected
ALARM_DATE_EN	1	Alarm day of the month enable: 0=disable, 1=enable
ALARM_MON	4	Alarm month in BCD, 1-12
ALARM_MON_EN	1	Alarm month enable: 0=disable, 1=enable
ALM_EN	1	Master enable for alarm. 0: Alarm is disabled. Fields for date and time are ignored. 1: Alarm is enabled. If none of the date and time fields are enabled, then this alarm triggers once every second.

If the master enable (ALM_EN) is set, but all alarm fields for date and time are disabled, an alarm interrupt will be generated once every second. Note that there is no alarm field for Year because the life expectancy of a chip is about 20 years and thus setting an alarm for a certain year means that the alarm matches either once or never in the lifetime of the chip.

The PSoC 6 MCU has two independent alarms. See the BACKUP_ALM1_TIME, BACKUP_ALM1_DATE, BACKUP_ALM2_TIME, and BACKUP_ALM2_DATE registers in the [registers TRM](#) for details.

Note that the alarm user registers, similar to RTC user registers, require certain special steps before read/write operations, as explained in sections [Reading RTC User Registers on page 203](#) and [Writing to RTC User Registers on page 203](#).

Interrupts must be properly configured for the RTC to generate interrupts/wake up events. Also, to enable RTC interrupts to wake up the device from Hibernate mode, the MASK_HIBALARM bit in the PWR_HIBERNATE register must be set. See the [Device Power Modes chapter on page 190](#) and [Interrupts chapter on page 54](#) for details.

The BACKUP_INTR_MASK register can be used to disable certain interrupts from the backup system.

Table 19-3. Interrupt Mask Bits

Bit Name	Description
ALARM1	Mask bit for interrupt generated by ALARM1
ALARM2	Mask bit for interrupt generated by ALARM2
CENTURY	Mask bit for century interrupt (interrupt generated when the Year field rolls over from 99 to 00)

The RTC alarm can also control an external PMIC as explained in the following section.

19.8 PMIC Control

The backup system can control an external PMIC that supplies V_{DD} . PMIC enable is an active-high output that is available at a certain pin PMIC_Wakeup_Out. See the [device datasheet](#) for the location of this pin. This pin can be connected to the “enable” input of a PMIC.

Table 19-4 shows the bitfields in BACKUP_PMIC_CTL.

Table 19-4. PMIC Control Bits

Bitfield Name	Number of Bits	Description
UNLOCK	8	This byte must be set to 0x3A for PMIC to be disabled. Any other value in this field will cause writes to PMIC_EN to be ignored. Do not change PMIC_EN in the same write cycle as setting/clearing the UNLOCK code; do these in separate write cycles.
POLARITY	1	Reserved for future use. Keep this bit at ‘1’.
PMIC_EN_OUTEN	1	Output enable of the PMIC_EN pin. 0: PMIC_EN pin output is HI-Z. This allows the PMIC_EN pin to be used as a GPIO. The HI-Z condition is kept only if the UNLOCK key (0x3A) is present. 1: PMIC_EN pin output is enabled.
PMIC_ALWAYSSEN	1	Override normal PMIC controls to prevent errant firmware from accidentally turning off the PMIC. 0: Normal operation; PMIC_EN and PMIC_OUTEN work as explained in their bitfield descriptions. 1: PMIC_EN is forced set; PMIC_EN and PMIC_OUTEN are ignored. This bit is a write-once bit that cannot be cleared until the next Backup domain reset.
PMIC_EN	1	Enable external PMIC (hardware output available at pin PMIC_Wakeup_Out). This bit is enabled by default. This bit will only clear if the UNLOCK field was written correctly in a previous write operation and PMIC_ALWAYSSEN = 0.

Note that two writes to the BACKUP_PMIC_CTL register are required to change the PMIC_EN setting. The first write should update the desired settings (including the UNLOCK code) but should not change PMIC_EN or PMIC_EN_OUTEN. The second write must use the same bit values as the first one except desired PMIC_EN/PMIC_EN_OUTEN settings.

When the PMIC_EN bit is cleared by firmware, the external PMIC is disabled and the system functions normally until V_{DD} is no longer present (OFF with Backup mode). The firmware can set this bit if it does so before V_{DD} is actually removed. The time between firmware disabling the

PMIC_EN bit and the actual removal of V_{DD} depends on the external PMIC and supply-capacitor characteristics.

Additionally, PMIC can be turned on by one of these events:

- An RTC Alarm/Century Interrupt
- A logic high input at the PMIC_Wakeup_In pin. See the [device datasheet](#) for the location of this pin. This allows a mechanical button or an external input from another device to wake up the system and enable the PMIC. The POLARITY bit in the BACKUP_PMIC_CTL register must be set to ‘1’ to use this feature. The same wakeup pin PMIC_Wakeup_In can wake up the device from

Hibernate mode. See the [Device Power Modes chapter on page 190](#) for details.

Make sure that one or more of these events are configured properly, and a battery or supercapacitor is connected to V_{BACKUP} with sufficient charge to power the backup system until one of these events occur. Otherwise, the PMIC may continue to be in the disabled state with the PSoC 6 MCU unable to enable it again.

19.9 Backup Registers

The Backup domain has sixteen registers (BACKUP_BREG0 to BACKUP_BREG15), which can be used to store 32 bytes of important information/flags. These registers retain their contents even when the main supply (V_{DDD}) is off as long as backup supply (V_{BACKUP}) is present. These registers can also be used to store information that must be retained when the device enters Hibernate mode.

19.10 Register List

Table 19-5. Backup Registers

Register Name	Description
BACKUP_CTL	Main control register (including power and clock)
BACKUP_RTC_RW	RTC read/write control register
BACKUP_STATUS	Status register
BACKUP_RTC_TIME	Calendar seconds, minutes, hours, and day of week
BACKUP_RTC_DATE	Calendar day of month, month, and year
BACKUP_ALM1_TIME	Alarm 1 seconds, minute, hours, and day of week
BACKUP_ALM1_DATE	Alarm 1 day of month, and month
BACKUP_ALM2_TIME	Alarm 2 seconds, minute, hours, and day of week
BACKUP_ALM2_DATE	Alarm 2 day of month, and month
BACKUP_INTR	Interrupt request register
BACKUP_INTR_MASK	Interrupt mask register
BACKUP_INTR_MASKED	Interrupt masked request register
BACKUP_PMIC_CTL	PMIC control register
BACKUP_BREG0 to BACKUP_BREG15	Backup registers
BACKUP_RESET	Reset register for the backup domain

20. Clocking System



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SysClk
- Application notes
- Code examples

PSoC 6 MCU provides flexible clocking options with on-chip crystal oscillators, phase lock loop, frequency lock loop, and supports multiple external clock sources.

20.1 Features

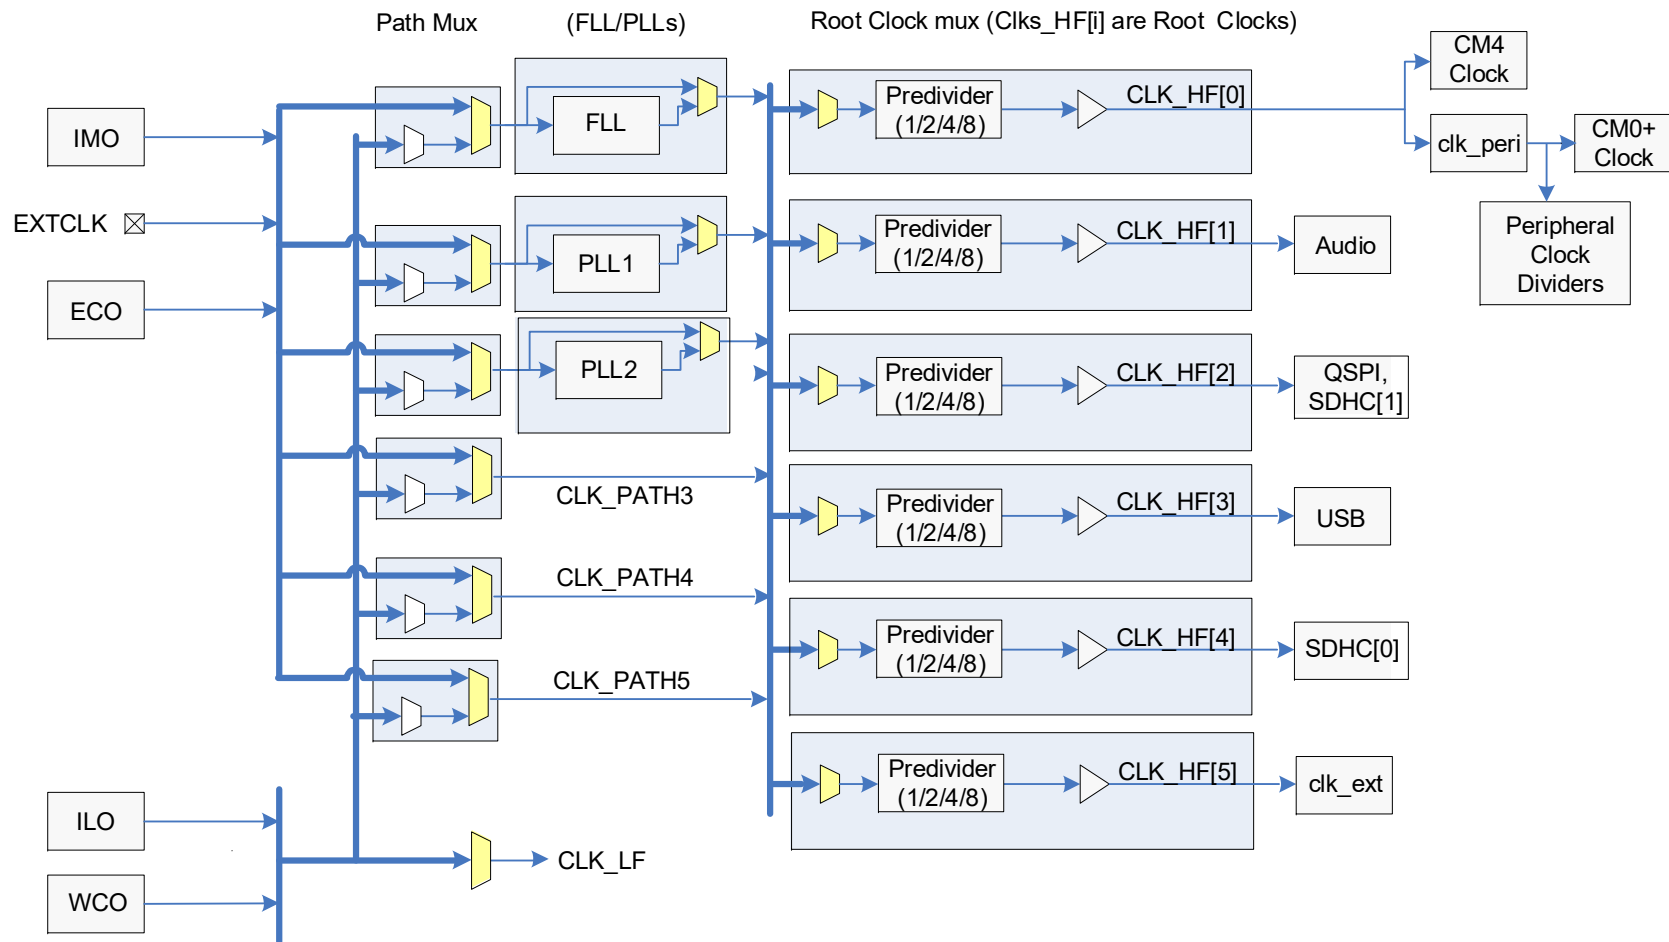
The PSoC 6 MCU clock system includes these resources:

- Two internal clock sources:
 - 8-MHz internal main oscillator (IMO)
 - 32-kHz internal low-speed oscillator (ILO)
- Three external clock sources
 - External clock (EXTCLK) generated using a signal from an I/O pin
 - External 4–35 MHz crystal oscillator (ECO)
 - External 32-kHz watch crystal oscillator (WCO)
- One frequency lock loop (FLL) with 24–100 MHz output range
- Two phase-locked loop (PLL) with 10.625–200 MHz output range

20.2 Architecture

Figure 20-1 gives a generic view of the clocking system in PSoC 6 MCUs.

Figure 20-1. Clocking System Block Diagram



20.3 Clock Sources

20.3.1 Internal Main Oscillator (IMO)

The IMO is an accurate, high-speed internal (crystal-less) oscillator that produces a fixed frequency of 8 MHz. The IMO output can be used by the PLL or FLL to generate a wide range of higher frequency clocks, or it can be used directly by the high-frequency root clocks.

When USB is present the USB Start-of-Frame (SOF) signal is used to trim the IMO to ensure that the IMO matches the accuracy of the USB SOF. The ENABLE_LOCK bitfield in the USBFS0_USBDEV_CR register of the USB block needs to be set for this feature to work. The driver for the USB block in the PDL does this automatically.

The IMO is available only in the system LP and ULP power modes.

20.3.2 External Crystal Oscillator (ECO)

The PSoC 6 MCU contains an oscillator to drive an external 4-MHz to 35-MHz crystal. This clock source is built using an oscillator circuit in PSoC. The circuit employs an external crystal that needs to be populated on the external crystal pins of the PSoC 6 MCU. See [AN218241 - PSoC 6 MCU Hardware Design Considerations](#) for more details.

The ECO can be enabled by using the CLK_ECO_CONFIG_ECO_EN register bitfields.

20.3.2.1 ECO Trimming

The ECO supports a wide variety of crystals and ceramic resonators with the nominal frequency range specification of $f = 4 \text{ MHz} - 35 \text{ MHz}$. The crystal manufacturer typically provides numerical values for parameters, namely the maximum drive level (DL), the equivalent series resistance (ESR), and the parallel load capacitance (C_L). These parameters can be used to calculate the transconductance (g_m) and the maximum peak to peak (V_{PP}).

$$\text{Max peak to peak value: } V_{PP} = 2 \times \sqrt{\frac{2 \times DL}{ESR}} \times \frac{1}{4\pi \times f \times C_L}$$

$$\text{Transconductance: } g_m = 4 \times 5 \times (2\pi \times f \times C_L)^2 \times ESR$$

ECO does not support V_{PP} less than 0.5 V. Similarly, g_m cannot be greater than or equal to 18 mA/V.

The following fields are found in the CLK_TRIM_ECO_CTL register. The Amplitude trim (ATRIM) and WDTRIM settings control the trim for amplitude of the oscillator output. ATRIM sets the crystal drive level when automatic gain control (AGC) is enabled (CLK_ECO_CONFIG.AGC_EN = 1). AGC must be enabled for $V_{PP} < 2 \text{ V}$ and disabled for all other cases.

WARNING: Take care when disabling AGC because driving a crystal beyond its rated limit can permanently damage the crystal.

Based on the V_{PP} value, the ATRIM and WDTRIM values are set as shown in [Table 20-1](#) and [Table 20-2](#).

Table 20-1. ATRIM Settings

V_{PP}	ATRIM
$V_{PP} < 0.6 \text{ V}$	0x00
$0.6 \text{ V} < V_{PP} < 0.7 \text{ V}$	0x02
$0.7 \text{ V} < V_{PP} < 0.8 \text{ V}$	0x04
$0.8 \text{ V} < V_{PP} < 0.9 \text{ V}$	0x06
$0.9 \text{ V} < V_{PP} < 1.025 \text{ V}$	0x08
$1.025 \text{ V} < V_{PP} < 1.15 \text{ V}$	0x0A
$1.15 \text{ V} < V_{PP} < 1.275 \text{ V}$	0x0C
$V_{PP} > 1.275 \text{ V}$	0x0E

Table 20-2. WDTRIM Settings

V_{PP}	WDTRIM
$V_{PP} < 0.6 \text{ V}$	0x00
$V_{PP} < 0.8 \text{ V}$	0x02
$V_{PP} < 1.0 \text{ V}$	0x04
$V_{PP} > 1.0 \text{ V}$	0x06

The GTRIM sets up the trim for amplifier gain based on the calculated g_m , as shown in [Table 20-3](#).

Table 20-3. GTRIM Settings

g_m	GTRIM
$g_m < 4.5 \text{ mA/V}$	0x01
$g_m = 4.5 \text{ mA/V}$	0x00
$g_m > 4.5 \text{ mA/V}$	INT ($g_m / 4.5 \text{ mA/V}$)

RTRIM sets up the trim for filter characteristics based on the frequency, as shown in [Table 20-4](#).

Table 20-4. RTRIM Settings

Nominal Frequency f (MHz)	RTRIM
$f > 28.6 \text{ MHz}$	0x00
$28.6 \text{ MHz} \geq f > 23.33 \text{ MHz}$	0x01
$23.33 \text{ MHz} \geq f > 16.5 \text{ MHz}$	0x02
$16.5 \text{ MHz} \geq f$	0x03

FTRIM sets up the trim for filter characteristics based on ATRIM, as shown in [Table 20-5](#).

Table 20-5. FTRIM Settings

Nominal Frequency f (MHz)	FTRIM
ATRIM < 2	0x00
4 > ATRIM ≥ 2	0x01
6 > ATRIM ≥ 4	0x02
8 > ATRIM ≥ 6	0x03
10 > ATRIM ≥ 8	0x04
12 > ATRIM ≥ 10	0x05
14 > ATRIM ≥ 12	0x06
ATRIM ≥ 14	0x07

First, set up the trim values based on [Table 20-1](#) through [Table 20-5](#) and then enable the ECO. After the ECO is enabled, the CLK_ECO_STATUS register can be checked to ensure it is ready. The SysClk driver of the PDL provides a function that automatically calculates all trim values based on the XTALs Vpp and gm.

20.3.3 External Clock (EXTCLK)

The external clock is a 0- to 100-MHz range clock that can be sourced from a signal on a designated I/O pin. This clock can be used as the source clock for either the PLL or FLL, or can be used directly by the high-frequency clocks.

When manually configuring a pin as an input to EXTCLK, the drive mode of the pin must be set to high-impedance digital to enable the digital input buffer. See the [I/O System chapter on page 226](#) for more details. Consult the [device datasheet](#) to determine the specific pin used for EXTCLK. See [KBA224493](#) for more details.

20.3.4 Internal Low-speed Oscillator (ILO)

The ILO operates with no external components and outputs a stable clock at 32.768 kHz nominal. The ILO is relatively low power and low accuracy. It is available in all power modes. If the ILO is to remain active in Hibernate mode, and across power-on-reset (POR) or brownout detect (BOD), the ILO_BACKUP bit must be set in the CLK_ILO_CONFIG register.

The ILO can be used as the clock source for: CLK_LF, which in turn can be used as a source for the backup domain (CLK_BAK). CLK_BAK runs the real-time clock (RTC). This can be useful if you do not wish to populate a WCO. Although the ILO is not suitable as an RTC due to its poor accuracy, it can be used as a HIBERNATE wakeup source using the wakeup alarm facility in the RTC. In this case, the V_{DD} rail must be supplied during HIBERNATE for the ILO to run, and the ILO_BACKUP bit must be set in the CLK_ILO_CONFIG register. CLK_LF is also the source of

the MCWDT timers; see the [Watchdog Timer chapter on page 248](#) for details.

The ILO is always the source of the watchdog timer (WDT).

The ILO is enabled and disabled with the ENABLE bit of the CLK_ILO_CONFIG register. Always leave the ILO enabled as it is the source of the WDT.

If the WDT is enabled, the only way to disable the ILO is to first clear the WDT_LOCK bit in the WDT_CTL register and then clear the ENABLE bit in the CLK_ILO_CONFIG register. If the WDT_LOCK bit is set, any register write to disable the ILO will be ignored. Enabling the WDT will automatically enable the ILO.

The calibration counters described in [Clock Calibration Counters on page 222](#) can be used to measure the ILO against a high-accuracy clock such as the ECO. This result can then be used to determine how the ILO must be adjusted. The ILO can be trimmed using the CLK_TRIM_ILO_CTL register.

20.3.5 Watch Crystal Oscillator (WCO)

The WCO is a highly accurate 32.768-kHz clock source. It is the primary clock source for the RTC. The WCO can also be used as a source for CLK_LF.


The WCO can be enabled and disabled by setting the WCO_EN bit in the CTL register for the backup domain. The WCO can also be bypassed and an external 32.768-kHz clock can be routed on a WCO output pin. This is done by setting the WCO_BYPASS bit in the CTL register for the backup domain. See [WCO with External Clock/Sine Wave Input on page 202](#) for more details.

20.4 Clock Generation

20.4.1 Phase-Locked Loop (PLL)

The PSoC 6 MCU contains two PLLs, which resides on CLK_PATH1 and CLK_PATH2. They are capable of generating a clock output in the range 10.625–200 MHz; the input frequency must be between 4 and 64 MHz. This makes it possible to use the IMO to generate much higher clock frequencies for the rest of the system.

Figure 20-2. PLL Block Diagram



The PLL is configured following these steps:

Note: f_{ref} is the input frequency of the PLL, that is, the frequency of input clock, such as 8 MHz for the IMO.

1. Determine the desired output frequency (f_{out}). Calculate the reference (REFERENCE_DIV), feedback (FEEDBACK_DIV), and output (OUTPUT_DIV) dividers subject to the following constraints:
 - a. PFD frequency (phase detector frequency). $f_{pfd} = f_{ref} / \text{REFERENCE_DIV}$. It must be in the range 4 MHz to 8 MHz. There may be multiple reference divider values that meet this constraint.
 - b. VCO frequency. $f_{vco} = f_{pfd} * \text{FEEDBACK_DIV}$. It must be in the range 170 MHz to 400 MHz. There may be multiple feedback divider values that meet this constraint with different REFERENCE_DIV choices.
 - c. Output frequency. $f_{out} = f_{vco} / \text{OUTPUT_DIV}$. It must be in the range 10.625 MHz to 200 MHz. Note that your device may not be capable of operating at this frequency; check the device datasheet. It may not be possible to get the desired frequency due to granularity; therefore, consider the frequency error of the two closest choices.
 - d. Choose the best combination of divider parameters depending on the application. Some possible decision-making factors are: minimum output frequency error, lowest power consumption (lowest f_{vco}), or lowest jitter (highest f_{vco}).
2. Program the divider settings in the appropriate CLK_PLL_CONFIG register. Do not enable the PLL on the same cycle as configuring the dividers. Do not change the divider settings while the PLL is enabled.
3. Enable the PLL (CLK_PLL_CONFIG.ENABLE = 1). Wait at least 1 μ s for PLL circuits to start.
4. Wait until the PLL is locked before using the output. By default, the PLL output is bypassed to its reference clock and will automatically switch to the PLL output when it is locked. This behavior can be changed using PLL_CONFIG.BYPASS_SEL. The status of the PLL can be checked by reading CLK_PLL_STATUS. This register contains a bit indicating the PLL has locked. It also contains a bit indicating if the PLL lost the lock status.

To disable the PLL, first set the PLL_CONFIG.BYPASS_SEL to PLL_REF. Then wait at least six PLL output clock cycles

before disabling the PLL by setting PLL_CONFIG.ENABLE to '0'.

20.4.2 Frequency Lock Loop (FLL)

The PSoC 6 MCU contains one frequency lock loop (FLL), which resides on Clock Path 0. The FLL is capable of generating a clock output in the range 24 MHz to 100 MHz; the input frequency must be between 0.001 and 100 MHz, and must be at least 2.5 times less than the CCO frequency. This makes it possible to use the IMO to generate much higher clock frequencies for the rest of the system.

The FLL is similar in purpose to a PLL but is not equivalent:

- FLL can start up (lock) much faster than the PLL.
- It consumes less current than the PLL.
- FLL does not lock the phase. At the heart of the FLL is a current-controlled oscillator (CCO). The output frequency of this CCO is controlled by adjusting the trim of the CCO; this is done in hardware and is explained in detail later in this section.
- FLL can produce up to 100-MHz clock with good duty cycle through its divided clock output.
- FLL reference clock can be the WCO (32 kHz), IMO (8 MHz), or any other periodic clock source.

Note: The CCO frequency must be at least 2.5 times greater than the reference frequency.

The CCO can output a stable frequency in the 48 MHz to 200 MHz range. This range is divided into five sub-ranges as shown by Table 20-6.

Table 20-6. CCO Frequency Ranges

CCO Range	0	1	2	3	4
f_{min}	48 MHz	64 MHz	85 MHz	113 MHz	150 MHz
f_{max}	64 MHz	85 MHz	113 MHz	150 MHz	200 MHz

Note: The output of the CCO has an option to enable a divide by two or not. For this device, the divide by two must always be enabled.

Within each range, the CCO output is controlled via a 9-bit trim field. This trim field is updated via hardware based on the control algorithm described below.

A reference clock must be provided to the FLL. This reference clock is typically the IMO, but could be many different clock sources. The FLL compares the reference clock against the CCO clock to determine how to adjust the CCO trim. Specifically, the FLL will count the number of CCO clock cycles inside a specified window of reference clock cycles. The number of reference clock cycles to count is set by the FLL_REF_DIV field in the CLK_FLL_CONFIG2 register.


After the CCO clocks are counted, they are compared against an ideal value and an error is calculated. The ideal value is programmed into the FLL_MULT field of the CLK_FLL_CONFIG register.

As an example, the reference clock is the IMO (8 MHz), the desired CCO frequency is 100 MHz, the value for FLL_REF_DIV is set to 146. This means that the FLL will

count the number of CCO clocks within 146 clock periods of the reference clock. In one clock cycle of the reference clock (IMO), there should be $100 / 8 = 12.5$ clock cycles of the CCO. Multiply this number by 146 and the value of FLL_MULT should be 1825.

If the FLL counts a value different from 1825, it attempts to adjust the CCO such that it achieves 1825 the next time it counts. This is done by scaling the error term with FLL_LF_IGAIN and FLL_LF_PGAIN found in CLK_FLL_CONFIG3. Figure 20-3 shows how the error (err) term is multiplied by FLL_LF_IGAIN and FLL_LF_PGAIN and then summed with the current trim to produce a new trim value for the CCO. The CCO_LIMIT field in the CLK_FLL_CONFIG4 can be used to put an upper limit on the trim adjustment; this is not needed for most situations.

Figure 20-3. FLL Error Correction Diagram



The FLL determines whether it is “locked” by comparing the error term with the LOCK_TOL field in the CLK_FLL_CONFIG2 register. When the error is less than LOCK_TOL the FLL is considered locked.

After each adjustment to the trim the FLL can be programmed to wait a certain number of reference clocks before doing a new measurement. The number of reference clocks to wait is set in the SETTLING_COUNT field of CLK_FLL_CONFIG3. Set this such that the FLL waits $\sim 1 \mu\text{s}$ before a new count. Therefore, if the 8-MHz IMO is used as the reference this field should be programmed to ‘8’.

When configuring the FLL there are two important factors that must be considered: lock time and accuracy. Accuracy is the closeness to the intended output frequency. These two numbers are inversely related to each other via the value of REF_DIV. Higher REF_DIV values lead to higher accuracy, whereas lower REF_DIV values lead to faster lock times.

In the example used previously the 8-MHz IMO was used as the reference, and the desired FLL output was 100 MHz. For that example, there are 12.5 CCO clocks in one reference clock. If the value for REF_DIV is set to ‘1’ then FLL_MULT must be set to either ‘13’ or ‘12’. This will result in a CCO output of either 96 MHz or 104 MHz, and an error of 4 percent from the desired 100 MHz. Therefore, the best way to improve this is to increase REF_DIV. However, the larger REF_DIV is, the longer each measurement cycle takes, thus increasing the lock time. In this example, REF_DIV was set to 146. This means each measurement cycle takes $146 * (1/8 \text{ MHz}) = 18.25 \mu\text{s}$, whereas when REF_DIV is set to 1, each measurement cycle takes $1 * (1/8 \text{ MHz}) = 0.125 \mu\text{s}$.

Another issue with lower REF_DIV values is that the minimum LOCK_TOL is 1, so the output of the CCO can have an error of ± 1 . In the example where REF_DIV = 1 and FLL_MULT = 13, the MULT value can really be 12, 13, or 14 and still be locked. This means the output of the FLL may vary between 96 and 112 MHz, which may not be desirable.

Thus, a choice must be made between faster lock times and more accurate FLL outputs. The biggest change to make for this is the value of REF_DIV. The following section describes how to configure all of the FLL registers and gives some example equations to set REF_DIV N for best accuracy.

20.4.2.1 Configuring the FLL

This section provides a guide to calculate FLL parameters. The following equations are tailored to achieve the best accuracy.

In the following equations:

$N = \text{REF_DIV}$

$M = \text{FLL_MULT}$

1. Set CCO_RANGE in the CLK_FLL_CONFIG4 register.
 - a. Determine the output frequency of the FLL.
 - b. Calculate the CCO frequency. The CCO frequency = $2 \times \text{FLL output frequency}$.
 - c. Use [Table 20-7](#) to determine the CCO range.
 - d. Write CCO range into CCO_RANGE in the CLK_FLL_CONFIG4 register.

Table 20-7. CCO Frequency Ranges

CCO Range	0	1	2	3	4
f_{\min}	48 MHz	64 MHz	85 MHz	113 MHz	150 MHz
f_{\max}	64 MHz	85 MHz	113 MHz	150 MHz	200 MHz

2. Set the FLL_OUTPUT_DIV in CLK_FLL_CONFIG. Set the output divider to '1'.
3. Set FLL_REF_DIV in CLK_FLL_CONFIG2.

FLL_REF_DIV divides the FLL input clock. The FLL counts the number of CCO clocks within one period of the divided reference clock. A general equation to calculate the reference divider is as follows:

$$N = \text{ROUNDUP} \left(\frac{2 \times f_{\text{ref}}}{\text{CCO}_{\text{TrimStep}} \times f_{\text{CCO}_{\text{Target}}}} \right)$$

Equation 20-1

The $\text{CCO}_{\text{TrimStep}}$ is found in [Table 20-8](#).

Table 20-8. CCO Trim Steps

CCO Range	0	1	2	3	4
CCO_Trim_Steps	0.0011	0.0011	0.0011	0.0011	0.00117

A larger N results in better precision on the FLL output, but longer lock times; a smaller N will result in faster lock times, but less precision.

Note: When the WCO is used as the reference clock, N must be set to 19.

4. Set FLL_MULT in CLK_FLL_CONFIG.

FLL_MULT is the ratio between the desired CCO frequency and the divided input frequency. This is the ideal value for the counter that counts the number of CCO clocks in one period of the divided input frequency.

$$M = f_{\text{CCO}_{\text{Target}}} \times \frac{N}{f_{\text{ref}}}$$

Equation 20-2

5. Set the FLL_LF_IGAIN and FLL_LF_PGAIN in CLK_FLL_CONFIG3.

Within each range of the CCO there are 512 steps of adjustment for the CCO frequency. These steps are controlled by CCO_FREQ in the CLK_FLL_CONFIG4 register. The FLL automatically adjusts CCO_FREQ based on the output of the FLL counter.

The output of the counter gives the number of CCO clocks, over one period of the divided reference clock, by which the FLL is off. This value is then multiplied by the sum of FLL_LF_IGAIN and FLL_LF_PGAIN. The result of this multiplication is then summed with the value currently in the CCO_FREQ register.

To determine the values for I_GAIN and P_GAIN use the following equation:

$$I_{GAIN} < \left(\frac{0.85}{K_{CCO} \times \frac{N}{f_{ref}}} \right)$$

Equation 20-3

Find the value of I_GAIN closest but not over the values in the gain row in [Table 20-9](#).

Table 20-9. I_GAIN and P_GAIN Register Values

Register Value	0	1	2	3	4	5	6	7	8	9	10	11
Gain Value	1/256	1/128	1/64	1/32	1/16	1/8	1/4	1/2	1	2	4	8

Program FLL_LF_IGAIN with the register value that corresponds to the chosen gain value.

Take the I_GAIN value from the register and use it in the following equation:

$$P_{GAIN} < I_{GAIN_{reg}} - \frac{0.85}{K_{CCO} \times \frac{N}{f_{ref}}}$$

Equation 20-4

Find the value of P_GAIN closest but not over the values in the gain row in [Table 20-9](#). Program FLL_PF_IGAIN with the register value that corresponds to the chosen gain value.

For best performance Pgain_reg + Igain_reg should be as close as possible to calculated I_GAIN without exceeding it. kCCO is the gain of the CCO; [Table 20-10](#) lists the kCCO for each CCO range.

Table 20-10. kCCO Values

CCO Range	0	1	2	3	4
kCCO	48109.38	64025.65	84920	113300	154521.85

6. Set SETTLING_COUNT in CLK_FLL_CONFIG3.

SETTLING_COUNT is the number of reference clocks to wait for the CCO to settle after it has changed. It is best to set this such that the settling time is around 1 μ s.

$$SETTLING_COUNT = 1 \mu s * f_{ref}$$

Do not set the settling time to anything less than 1 μ s, greater will lead to longer lock times.

1. Set LOCK_TOL in CLK_FLL_CONFIG2.

LOCK_TOL determines how much error the FLL can tolerate at the output of the counter that counts the number of CCO clocks in one period of the divided reference clock. A higher tolerance can be used to lock more quickly or track a less accurate source. The tolerance should be set such that the FLL does not unlock under normal conditions. A lower tolerance means a more accurate output, but if the input reference is unstable then the FLL may unlock.

The following equation can be used to help determine the value:

$$LOCK_TOL = M \times \left[\left(\frac{1 + CCO_{accuracy}}{1 - F_{refaccuracy}} \right) - 1 \right]$$

Equation 20-5

CCO (accuracy) = 0.25% or 0.0025

ref (accuracy) is the accuracy of the reference clock

2. Set CCO_FREQ in CLK_FLL_CONFIG4.

This field determines the frequency at which the FLL starts before any measurement. The nearer the FLL is to the desired frequency, the faster it will lock.

$$CCO_{FREQ} = \text{ROUPNDUP} \left(\frac{\text{LOG} \left(\frac{f_{CCO_{Target}}}{f_{CCO_{Base}}} \right)}{\text{LOG}(1 + CCO_{TrimStep})} \right)$$

Equation 20-6

$f_{CCO_{Base}}$ can be found in [Table 20-11](#).

Table 20-11. CCO Base Frequency

CCO Range	0	1	2	3	4
CCO_Base	43600	58100	77200	103000	132000

Table 20-12. CCO Trim Steps

CCO Range	0	1	2	3	4
CCO_Trim_Steps	0.0011	0.0011	0.0011	0.0011	0.00117

9. Calculating Precision, Accuracy, and Lock Time of FLL

To calculate the precision and accuracy of the FLL, the accuracy of the input source must be considered.

The precision is the larger of:

$$\text{Precision}_{FLL} = f_{ref} \times \left(\frac{1 + f_{ref}(\text{accuracy})}{N \times f_{CCO_{Target}}} \right)$$

Equation 20-7

Or $(CCO_Trim_Steps / 2)$

Example: The desired FLL output is 100 MHz, thus CCO target is 200 MHz. The 2 percent accurate 8 MHz IMO is used as the reference input. N is calculated to be 69.

$\text{Precision}_{FLL} = ((8 \text{ MHz} * 1.02) / (69 * 200 \text{ MHz})) = 0.059\%$, which is greater than the $(CCO_Trim_Steps / 2)$

The accuracy of the FLL output is the precision multiplied by the lock tolerance. If the CCO goes beyond this range, then the FLL will unlock.

$\text{Accuracy}_{FLL} = \text{Precision}_{FLL} * \text{LOCK_TOL}$

The value for the reference divider N should be tuned such that it achieves the best precision/accuracy versus lock time.

The lock time depends on the time for each adjustment step in the locking process;

$\text{Step_Time} = (N / f_{ref}) + (\text{SETTLING_COUNT} / f_{ref})$

Multiply this number by the number of steps it takes to lock, to determine lock time. Typically, the FLL locks within the first ~10 steps. ModusToolbox and the PDL provide drivers and calculators that automatically populate the FLL registers based on FLL input and output frequency requirements.

20.4.2.2 Enabling and Disabling the FLL

The FLL requires firmware sequencing when enabling, disabling, and entering/exiting DEEPSLEEP.

To enable the FLL, follow these steps:

1. Enable the CCO by writing `CLK_FLL_CONFIG4.CCO_ENABLE = 1` and wait until `CLK_FLL_STATUS.CCO_READY==1`.
2. Ensure the reference clock has stabilized and `CLK_FLL_CONFIG3.BYPASS_SEL = FLL_REF`.
3. Write `FLL_ENABLE = 1` and wait until `CLK_FLL_STATUS.LOCKED==1`.
4. Write `CLK_FLL_CONFIG3.BYPASS_SEL = FLL_OUT` to switch to the FLL output.

To disable the FLL, follow these steps:

1. Ensure the processor is operating from a different clock than `clk_path0`. If the muxes are changed, wait four FLL output clock cycles for it to complete.
2. Write `CLK_FLL_CONFIG3.BYPASS_SEL = FLL_REF` and read the same register to ensure the write completes.
3. Wait at least six FLL reference clock cycles and disable it with `FLL_ENABLE = 0`.
4. Disable the CCO by writing `CLK_FLL_CONFIG4.CCO_ENABLE = 0`.

Before entering DEEPSLEEP, either disable the FLL using the above sequence or use the following procedure to deselect/select it before/after DEEPSLEEP. Before entering DEEPSLEEP, write `CLK_FLL_CONFIG3.BYPASS_SEL = FLL_REF` to change the FLL to use its reference clock. After DEEPSLEEP wakeup, wait until `CLK_FLL_STATUS.LOCKED==1` and then write `CLK_FLL_CONFIG3.BYPASS_SEL = FLL_OUT` to switch to the FLL output.

Note: It is not recommended to use the `FLL_AUTO` option in the `BYPASS_SEL` field.

20.5 Clock Trees

The PSoC 6 MCU clocks are distributed throughout the device, as shown in [Figure 20-1](#). The clock trees are described in this section:

- Path Clocks
- High-Frequency Root Clocks (CLK_HF[i])
- Low-Frequency Clock (CLK_LF)
- Timer Clock (CLK_TIMER)
- Analog Pump Clock (CLK_PUMP)

20.5.1 Path Clocks

The PSoC 6 MCU has six clock paths: CLK_PATH0 contains the FLL, CLK_PATH1 and CLK_PATH2 contain the PLLs, and CLK_PATH3, CLK_PATH4, and CLK_PATH5 are a direct connection to the high-frequency root clocks. Note that the FLL and PLL(s) can be bypassed if they are not needed. These paths are the input sources for the high-frequency clock roots.

Each clock path has a mux to determine which source clock will clock that path. This configuration is done in CLK_PATH_SELECT[i] register.

Table 20-13. Clock Path Source Selections

Name	Description
PATH_MUX[2:0]	Selects the source for clk_path[i] 0: IMO 1: EXTCLK 2: ECO 3: Reserved 4: DSI_MUX 5-7: Reserved

The DSI mux is configured through the CLK_DSI_SELECT[i] register.

Table 20-14. DSI Mux Source Selection

Name	Description
DSI_MUX[4:0]	Selects the source for the DSI_MUX[i] 0-15: Reserved 16: ILO 17: WCO 18: Reserved 19: Reserved 20-31: Invalid

20.5.2 High-Frequency Root Clocks

The PSoC 6 MCU has six high-frequency root clocks (CLK_HF[i]). Each CLK_HF has a particular destination on the device, as shown in [Table 20-15](#).

Table 20-15. CLK_HF Destinations

Name	Description
CLK_HF[0]	Root clock for both CPUs, PERI, and AHB infrastructure
CLK_HF[1]	Root clock for the PDM/PCM and I ² S audio subsystem
CLK_HF[2]	Root clock for the Serial Memory Interface subsystem and SDHC[1] block
CLK_HF[3]	Root clock for USB communications
CLK_HF[4]	Root clock for SDHC[0] block
CLK_HF[5]	Clock output on clk_ext pin (when used as an output)

Each high-frequency root clock has a mux to determine its source. This configuration is done in the CLK_ROOT_SELECT[i] register.

Table 20-16. HFCLK Input Selection Bits

Name	Description
ROOT_MUX[3:0]	HFCLK input clock selection 0: Select CLK_PATH0 1: Select CLK_PATH1 2: Select CLK_PATH2 3: Select CLK_PATH3 4: Select CLK_PATH4 5: Select CLK_PATH5

Each CLK_HF has a pre-divider, which is set in the CLK_ROOT_SELECT register.

Table 20-17. HFCLK Divider Selection

Name	Description
ROOT_DVI[5:4]	Selects predivider value for the clock root and DSI input 0: No Divider 1: Divide clock by 2 2: Divide clock by 4 3: Divide clock by 8

CLK_HF[1-5] can be enabled and disabled. CLK_HF[0] is always enabled as it is the source of the CPU. To enable and disable CLK_HF[1-5] set the ENABLE bit in the CLK_ROOT_SELECT register.

20.5.3 Low-Frequency Clock

The low-frequency clock (CLK_LF) in the PSoC 6 MCU has two input options: ILO and WCO.

CLK_LF is the source for the multi-counter watchdog timers (MCWDT) and the RTC.

The source of CLK_LF is set in the LFCLK_SEL bits of the CLK_SELECT register.

Table 20-18. LFCLK Input Selection Bits LFCLK_SEL

Name	Description
LFCLK_SEL[1:0]	LFCLK input clock selection 0: ILO. Uses the internal local oscillator as the source of the LFCLK 1: WCO. Uses the watch crystal oscillator as the source of the LFCLK 2: Reserved 3: Reserved

20.5.4 Timer Clock

The timer clock (CLK_TIMER) can be used as a clock source for the profiler or the CPU SYSTICK timer. The source for CLK_TIMER can either be the IMO or CLK_HF[0]. This selection is made in the TIMER_SEL bitfield of the CLK_TIMER_CTL register. Several dividers can be applied to this clock, which are found in the CLK_TIMER_CTL register.

20.5.5 Group Clocks (clk_sys)

On the PSoC 6 platform, peripherals are grouped. Each group has a dedicated group clock (also referred to as clk_sys). The group clock sets the clock rate for the AHB interface on the peripheral; it also sets the clock rate for the trigger outputs and trigger input synchronization. Each group clock has an eight-bit divider located in the CLOCK_CTL register in the PERI_GROUP_STRUCT in the PERI register set. For a majority of applications these dividers should be left at default (divide by 1).


20.5.6 Backup Clock (clk_bak)

The backup clock is used to clock the backup domain, specifically the RTC. For more information see [WCO with External Clock/Sine Wave Input on page 202](#).

20.6 CLK_HF[0] Distribution

clk_hf[0] is the root clock for the CPU subsystem and for the peripheral clock dividers.

Figure 20-4. CLK_HF[0] Distribution



20.6.1 CLK_FAST

CLK_FAST clocks the Cortex-M4 processor. This clock is a divided version of CLK_HF[0]. The divider for this clock is set in the CM4_CLOCK_CTL register of the CPU subsystem.

20.6.2 CLK_PERI

CLK_PERI is the source clock for all programmable peripheral clock dividers and for the Cortex-M0+ processor. It is a divided version of CLK_HF[0]. The divider for this clock is set in the PERI_INT_DIV bitfields of the CM0_CLOCK_CTL register.

20.6.3 CLK_SLOW

CLK_SLOW is the source clock for the Cortex-M0+. This clock is a divided version of CLK_PERI. The divider for this clock is set in the SLOW_INT_DIV bitfields of the CM0_CLOCK_CTL register.

20.7.1 Fractional Clock Dividers

Fractional clock dividers allow the clock divisor to include a fraction of 0..31/32. For example, a 16.5-bit divider with an integer divide value of 3 generates a 16-MHz clock from a 48-MHz CLK_PERI. A 16.5-bit divider with an integer divide value of 4 generates a 12-MHz clock from a 48-MHz CLK_PERI. A 16.5-bit divider with an integer divide value of 3 and a fractional divider of 16 generates a $48 / (3 + 16/32) = 48 / 3.5 = 13.7$ -MHz clock from a 48-MHz CLK_PERI. Not all 13.7-MHz clock periods are equal in size; some will have a 16-MHz period and others will have a 12-MHz period, such that the average is 13.7 MHz.

Fractional dividers are useful when a high-precision clock is required (for example, for a UART/SPI serial interface). Fractional dividers are not used when a low jitter clock is required, because the clock periods have a jitter of one CLK_PERI cycle.

20.7.2 Peripheral Clock Divider Configuration

The peripheral clock dividers are configured using registers from the peripheral block; specifically DIV_CMD, DIV_8_CTL, DIV_16_CTL, DIV_16_5_CTL, DIV_24_5_CTL, and CLOCK_CTL registers.

First the clock divider needs to be configured. This is done via the DIV_8_CTL, DIV_16_CTL, DIV_16_5_CTL, and DIV_24_5_CTL registers. There is one register for each divider; for example, there are eight DIV_8_CTL registers as there are eight 8-bit dividers. In these registers, set the value of the integer divider; if it is a fractional divider then set the fraction portion as well.

20.7 Peripheral Clock Dividers

The PSoC 6 MCU peripherals such as SCBs and TCPWMs require a clock. These peripherals can be clocked only by a peripheral clock divider.


The PSoC 6 MCU has 29 peripheral clock dividers (PCLK). It has eight 8-bit dividers, sixteen 16-bit dividers, four fractional 16.5-bit dividers (16 integer bits, five fractional bits), and one 24.5-bit divider (24 integer bits, five fractional bits). The output of any of these dividers can be routed to any peripheral.

After the divider is configured use the DIV_CMD register to enable the divider. This is done by setting the DIV_SEL to the divider number you want to enable, and setting the TYPE_SEL to the divider type. For example, if you wanted to enable the 0th 16.5-bit divider, write '0' to DIV_SEL and '2' to TYPE_SEL. If you wanted to enable the tenth 16-bit divider, write '10' to DIV_SEL and '1' to TYPE_SEL. See the [registers TRM](#) for more details.

20.7.2.1 Phase Aligning Dividers


For specific use cases, you must generate clocks that are phase-aligned. For example, consider the generation of two gated clocks at 24 and 12 MHz, both of which are derived from a 48-MHz CLK_PERI. If phase alignment is not considered, the generated gated clocks appear as follows.

Figure 20-5. Non Phase-Aligned Clock Dividers




These clock signals may or may not be acceptable, depending on the logic functionality implemented on these two clocks. If the two clock domains communicate with each other, and the slower clock domain (12 MHz) assumes that each high/'1' pulse on its clock coincides with a high/'1' phase pulse in the higher clock domain (24 MHz), the phase misalignment is not acceptable. To address this, it is possible to have dividers produce clock signals that are phase-aligned with any of the other (enabled) clock dividers. Therefore, if (enabled) divider x is used to generate the 24-MHz clock, divider y can be phase-aligned to divider x and used to generate the 12-MHz clock. The aligned clocks appear as follows.

Figure 20-6. Phase-Aligned Clock Dividers



Phase alignment also works for fractional divider values. If (enabled) divider x is used to generate the 38.4-MHz clock (divide by $1 \frac{8}{32}$), divider y can be phase-aligned to divider x and used to generate the 19.2-MHz clock (divide by $2 \frac{16}{32}$). The generated gated clocks appear as follows.

Figure 20-7. Phase-Aligned Fractional Dividers



Divider phase alignment requires that the divider to which it is phase-aligned is already enabled. This requires the dividers to be enabled in a specific order.

Phase alignment is implemented by controlling the start moment of the divider counters in hardware. When a divider is enabled, the divider counters are set to '0'. The divider counters will only start incrementing from '0' to the

programmed integer and fractional divider values when the divider to which it is phase-aligned has an integer counter value of '0'.

Note that the divider and clock multiplexer control register fields are all retained during the Deep Sleep power mode. However, the divider counters that are used to implement the integer and fractional clock dividers are not. These counters are set to '0' during the Deep Sleep mode. Therefore, when transitioning from Deep Sleep to Active mode, all dividers (and clock signals) are enabled and phase-aligned by design.

Phase alignment is accomplished by setting the PA_DIV_SEL and PA_DIV_TYPE bits in the DIV_CMD register before enabling the clock. For example, to align the fourth 8-bit divider to the third 16-bit divider, set DIV_SEL to '4', TYPE_SEL to '0', PA_DIV_SEL to '3', and PA_TYPE_SEL to '1'.

20.7.2.2 Connecting Dividers to Peripheral

The PSoC 6 MCU has 54 peripherals, which can connect to one of the programmable dividers. [Table 20-19](#) lists those peripherals.

Table 20-19. Clock Dividers to Peripherals

Clock Number	Destination
0	scb[0].clock
1	scb[1].clock
2	scb[2].clock
3	scb[3].clock
4	scb[4].clock
5	scb[5].clock
6	scb[6].clock
7	scb[7].clock
8	scb[8].clock
9	scb[9].clock
10	scb[10].clock
11	scb[11].clock
12	scb[12].clock
13	smartio[8].clock
14	smartio[9].clock
15	tcpwm[0].clocks[0]
16	tcpwm[0].clocks[1]
17	tcpwm[0].clocks[2]
18	tcpwm[0].clocks[3]
19	tcpwm[0].clocks[4]
20	tcpwm[0].clocks[5]
21	tcpwm[0].clocks[6]
22	tcpwm[0].clocks[7]
23	tcpwm[1].clocks[0]

Table 20-19. Clock Dividers to Peripherals

Clock Number	Destination
24	tcpwm[1].clocks[1]
25	tcpwm[1].clocks[2]
26	tcpwm[1].clocks[3]
27	tcpwm[1].clocks[4]
28	tcpwm[1].clocks[5]
29	tcpwm[1].clocks[6]
30	tcpwm[1].clocks[7]
31	tcpwm[1].clocks[8]
32	tcpwm[1].clocks[9]
33	tcpwm[1].clocks[10]
34	tcpwm[1].clocks[11]
35	tcpwm[1].clocks[12]
36	tcpwm[1].clocks[13]
37	tcpwm[1].clocks[14]
38	tcpwm[1].clocks[15]
39	tcpwm[1].clocks[16]
40	tcpwm[1].clocks[17]
41	tcpwm[1].clocks[18]
42	tcpwm[1].clocks[19]
43	tcpwm[1].clocks[20]
44	tcpwm[1].clocks[21]
45	tcpwm[1].clocks[22]
46	tcpwm[1].clocks[23]
47	csd.clock
48	lcd.clock
49	profile.clock_profile
50	cpuss.clock_trace_in
51	pass.clock_pump_peri
52	pass.clock_sar
53	usb.clock_dev_brs

To connect a peripheral to a specific divider, the PERI_CLOCK_CTL register is used. There is one PERI_CLOCK_CTL register for each entry in [Table 20-19](#). For example, to select the twelfth 16-bit divider for tcpwm[1].clocks[2] write to the twenty-fifth CLOCK_CTL register, set the DIV_SEL to '12', and the TYPE_SEL to '1'.

20.8 Clock Calibration Counters

A feature of the clocking system in PSoC 6 MCUs is built-in hardware calibration counters. These counters can be used to compare the frequency of two clock sources against one another. The primary use case is to take a higher accuracy clock such as the ECO and use it to measure a lower accuracy clock such as the ILO or PILO. The result of this measurement can then be used to trim the ILO and PILO.

There are two counters: Calibration Counter 1 is clocked off of Calibration Clock 1 (generally the high-accuracy clock) and it counts down; Calibration Counter 2 is clocked off of Calibration Clock 2 and it counts up. When Calibration Counter 1 reaches 0, Calibration Counter 2 stops counting up and its value can be read. From that value the frequency of Calibration Clock 2 can be determined with the following equation.

$$\text{Calibration Clock 2 Frequency} = \frac{\text{Counter 2 Final Value}}{\text{Counter 1 Initial Value}} \times \text{Calibration Clock 1 Frequency}$$

For example, if Calibration Clock 1 = 8 MHz, Counter 1 = 1000, and Counter 2 = 5

Calibration Clock 1 Frequency = $(5/1000) \times 8 \text{ MHz} = 40 \text{ kHz}$.

Calibration Clock 1 and Calibration Clock 2 are selected with the CLK_OUTPUT_FAST register. All clock sources are available as a source for these two clocks. CLK_OUTPUT_SLOW is also used to select the clock source.

Calibration Counter 1 is programmed in CLK_CAL_CNT1. Calibration Counter 2 can be read in CLK_CAL_CNT2.

When Calibration Counter 1 reaches 0, the CAL_COUNTER_DONE bit is set in the CLK_CAL_CNT1 register.

21. Reset System



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SysLib
- Application notes
- Code examples

The PSoC 6 MCU family supports several types of resets that guarantee error-free operation during power up and allow the device to reset based on user-supplied external hardware or internal software reset signals. The PSoC 6 MCU also contains hardware to enable the detection of certain resets.

21.1 Features

The PSoC 6 MCU has these reset sources:

- Power-on reset (POR) to hold the device in reset while the power supply ramps up to the level required for the device to function properly
- Brownout reset (BOD) to reset the device if the power supply falls below the device specifications during normal operation
- External reset (XRES) to reset the device using an external input
- Watchdog timer (WDT) reset to reset the device if the firmware execution fails to periodically service the watchdog timer
- Software initiated reset to reset the device on demand using firmware
- Logic-protection fault resets to reset the device if unauthorized operating conditions occur
- Clock-supervision logic resets to reset the device when clock-related errors occur
- Hibernate wakeup reset to bring the device out of the Hibernate low-power mode

21.2 Architecture

The following sections provide a description of the reset sources available in the PSoC 6 MCU family.

Note: None of these sources can reset the Backup system. The Backup domain is reset only when all the power supplies are removed from it, also known as a “cold start” or if the firmware triggers a reset using the BACKUP_RESET register. For more details, see the [Backup System chapter on page 200](#).

21.2.1 Power-on Reset

Power-on reset is provided to keep the system in a reset state during power-up. POR holds the device in reset until the supply voltage, V_{DD} reaches the datasheet specification. The POR activates automatically at power-up. See the [Power Supply and Monitoring chapter on page 183](#) for more details.

POR events do not set a reset cause status bit, but can be partially inferred by the absence of any other reset source. If no other reset event is detected, then the reset is caused by POR, BOD, or XRES.

21.2.2 Brownout Reset

Brownout reset monitors the chip digital voltage supply V_{DD} and generates a reset if V_{DD} falls below the minimum logic operating voltage specified in the [device datasheet](#). See the [Power Supply and Monitoring chapter on page 183](#) for more details.

BOD events do not set a reset cause status bit, but in some cases they can be detected. In some BOD events, V_{DD} will fall below the minimum logic operating voltage specified by the datasheet, but remain above the minimum logic retention voltage. Thus, some BOD events may be distinguished from POR events by checking for logic retention. This is explained further in [Identifying Reset Sources on page 225](#).

21.2.3 Watchdog Timer Reset

Watchdog timer reset causes a reset if the WDT is not serviced by the firmware within a specified time limit. See the [Watchdog Timer chapter on page 248](#) for more details.

The RESET_HWWDT bit or RESET_SWWDTO to RESET_SWWDT3 status bits of the RES_CAUSE register is set when a watchdog reset occurs. This bit remains set until cleared by the firmware or until a POR, XRES, or BOD reset occurs. All other resets leave this bit unaltered.

For more details, see the [Watchdog Timer chapter on page 248](#).

21.2.4 Software Initiated Reset

Software initiated reset is a mechanism that allows the CPU to request a reset. The Cortex-M0+ and Cortex-M4 Application Interrupt and Reset Control registers (CM0_AIRCR and CM4_AIRCR, respectively) can request a reset by writing a '1' to the SYSRESETREQ bit of the respective registers.

Note that a value of 0x5FA should be written to the VECTKEY field of the AIRCR register before setting the SYSRESETREQ bit; otherwise, the processor ignores the write. See the [CPU Subsystem \(CPUSS\) chapter on page 33](#) for details.

The RESET_SOFT status bit of the RES_CAUSE register is set when a software reset occurs. This bit remains set until cleared by firmware or until a POR, XRES, or BOD reset occurs. All other resets leave this bit unaltered.

21.2.5 External Reset

External reset (XRES) is a reset triggered by an external signal that causes immediate system reset when asserted. The XRES pin is active low – a logic '1' on the pin has no effect and a logic '0' causes reset. The pin is pulled to logic '1' inside the device. XRES is available as a dedicated pin.

For detailed pinout, refer to the pinout section of the [device datasheet](#).

The XRES pin holds the device in reset as long as the pin input is '0'. When the pin is released (changed to logic '1'), the device goes through a normal boot sequence. The logical thresholds for XRES and other electrical characteristics are listed in the Electrical Specifications section of the [device datasheet](#). XRES is available in all power modes, but cannot reset the Backup system.

An XRES event does not set a reset cause status bit, but can be partially inferred by the absence of any other reset source. If no other reset event is detected, then the reset is caused by POR, BOD, or XRES.

21.2.6 Logic Protection Fault Reset

Logic protection fault reset detects any unauthorized protection violations and causes the device to reset if they occur. One example of a protection fault is reaching a debug breakpoint while executing privileged code.

The RESET_ACT_FAULT or RESET_DPSLP_FAULT bits of the RES_CAUSE register is set when a protection fault occurs in Active or Deep Sleep modes, respectively. These bits remain set until cleared or until a POR, XRES, or BOD reset. All other resets leave this bit unaltered.

21.2.7 Clock-Supervision Logic Reset

Clock-supervision logic initiates a reset due to the loss of a high-frequency clock or watch-crystal clock, or due to a high-frequency clock error.

The RESET_CSV_WCO_LOSS bit of the RES_CAUSE register is set when the clock supervision logic requests a reset due to the loss of a watch-crystal clock (if enabled).

The RESET_CSV_HF_LOSS is a 16-bit field in the RES_CAUSE2 register that can be used to identify resets caused by the loss of a high-frequency clock. Similarly, the RESET_CSV_HF_FREQ field can be used to identify resets caused by the frequency error of a high-frequency clock.

For more information on clocks, see the [Clocking System chapter on page 207](#).

21.2.8 Hibernate Wakeup Reset

Hibernate wakeup reset occurs when one of the Hibernate wakeup sources performs a device reset to return to the Active power mode. See the [Device Power Modes chapter on page 190](#) for details on Hibernate mode and available wakeup sources.

TOKEN is an 8-bit field in the PWR_HIBERNATE register that is retained through a Hibernate wakeup sequence. The firmware can use this bitfield to differentiate hibernate wakeup from a general reset event. Similarly, the PWR_HIB_DATA register can retain its contents through a

Hibernate wakeup reset, but is cleared when XRES is asserted.

21.3 Identifying Reset Sources

When the device comes out of reset, it is often useful to know the cause of the most recent or even older resets. This is achieved through the RES_CAUSE and RES_CAUSE2 registers. These registers have specific status bits allocated for some of the reset sources. These registers record the occurrences of WDT reset, software reset, logic-protection fault, and clock-supervision resets. However, these registers do not record the occurrences of POR, BOD, XRES, or

Hibernate wakeup resets. The bits in these registers are set on the occurrence of the corresponding reset and remain set after the reset, until cleared by the firmware or a loss of retention, such as a POR, XRES, or BOD.

Hibernate wakeup resets can be detected by examining the TOKEN field in the PWR_HIBERNATE register as described previously. Hibernate wakeup resets that occur as a result of an XRES cannot be detected. The other reset sources can be inferred to some extent by the status of the RES_CAUSE and RES_CAUSE2 registers, as shown in [Table 21-1](#).

Table 21-1. Reset Cause Bits to Detect Reset Source

Register	Bitfield	Number of Bits	Description
RES_CAUSE	RESET_HWWDT	1	A hardware WDT reset has occurred since the last power cycle.
RES_CAUSE	RESET_ACT_FAULT	1	Fault logging system requested a reset from its Active logic.
RES_CAUSE	RESET_DPSLP_FAULT	1	Fault logging system requested a reset from its Deep Sleep logic.
RES_CAUSE	RESET_CSV_WCO_LOSS	1	Clock supervision logic requested a reset due to loss of a watch-crystal clock.
RES_CAUSE	RESET_SOFT	1	A CPU requested a system reset through its SYSRESETREQ. This can be done via a debugger probe or in firmware.
RES_CAUSE	RESET_MCWDT0	1	Multi-counter WDT reset #0 has occurred since the last power cycle.
RES_CAUSE	RESET_MCWDT1	1	Multi-counter WDT reset #1 has occurred since the last power cycle.
RES_CAUSE	RESET_MCWDT2	1	Multi-counter WDT reset #2 has occurred since the last power cycle.
RES_CAUSE	RESET_MCWDT3	1	Multi-counter WDT reset #3 has occurred since the last power cycle.
RES_CAUSE2	RESET_CSV_HF_LOSS	16	Clock supervision logic requested a reset due to loss of a high-frequency clock. Each bit index K corresponds to a clk_hf<K>. Unimplemented clock bits return zero.
RES_CAUSE2	RESET_CSV_HF_FREQ	16	Clock supervision logic requested a reset due to frequency error of a high-frequency clock. Each bit index K corresponds to a clk_hf<K>. Unimplemented clock bits return zero.

For more information, see the RES_CAUSE and RES_CAUSE2 registers in the [registers TRM](#).

If these methods cannot detect the cause of the reset, then it can be one of the non-recorded and non-retention resets: BOD, POR, or XRES. These resets cannot be distinguished using on-chip resources.

21.4 Register List

Table 21-2. Reset System Register List

Register	Description
RES_CAUSE	Reset cause observation register
RES_CAUSE2	Reset cause observation register 2
PWR_HIBERNATE	Hibernate power mode control register. Contains a TOKEN field that can be used to detect the Hibernate wakeup reset
PWR_HIB_DATA	Retains its contents through Hibernate wakeup reset
CM4_AIRCR	Application interrupt and reset control register of Cortex-M4
CM0_AIRCR	Application interrupt and reset control register of Cortex-M0+

22. I/O System



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - GPIO
- Application notes
- Code examples

This chapter explains the PSoC 6 MCU I/O system, its features, architecture, operating modes, and interrupts. The I/O system provides the interface between the CPU core and peripheral components to the outside world. The flexibility of PSoC 6 MCUs and the capability of its I/O to route most signals to most pins greatly simplifies circuit design and board layout. The GPIO pins in the PSoC 6 MCU family are grouped into ports; a port can have a maximum of eight GPIO pins.

22.1 Features

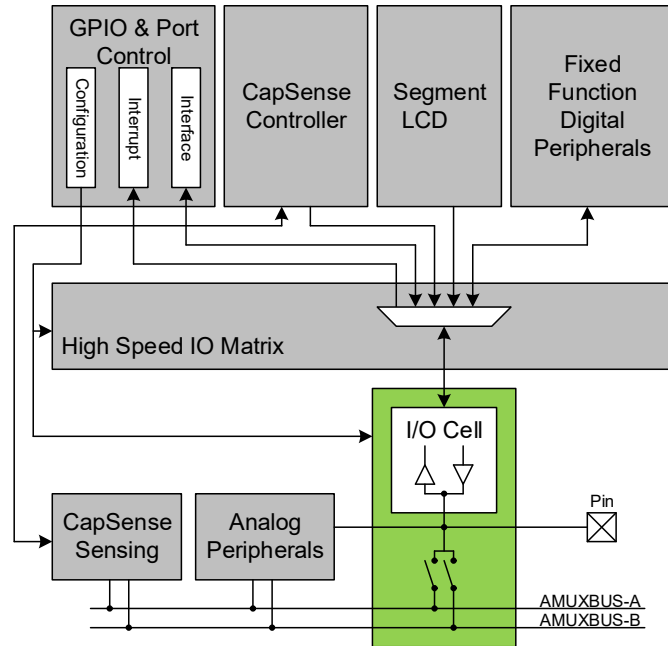
The PSoC 6 MCU GPIOs have these features:

- Analog and digital input and output capabilities
- Eight drive strength modes
- Separate port read and write registers
- Overvoltage tolerant (OVT-GPIO) pins
- Separate I/O supplies and voltages for up to six groups of I/O
- Edge-triggered interrupts on rising edge, falling edge, or on both edges, on all GPIO
- Slew rate control
- Frozen mode for latching previous state (used to retain the I/O state in System Hibernate Power mode)
- Selectable CMOS and low-voltage LVTTTL input buffer mode
- CapSense support
- Smart I/O provides the ability to perform Boolean functions in the I/O signal path
- Segment LCD drive support

22.2 Architecture

The PSoC 6 MCU is equipped with analog and digital peripherals. [Figure 22-1](#) shows an overview of the routing between the peripherals and pins.

Figure 22-1. GPIO Interface Overview

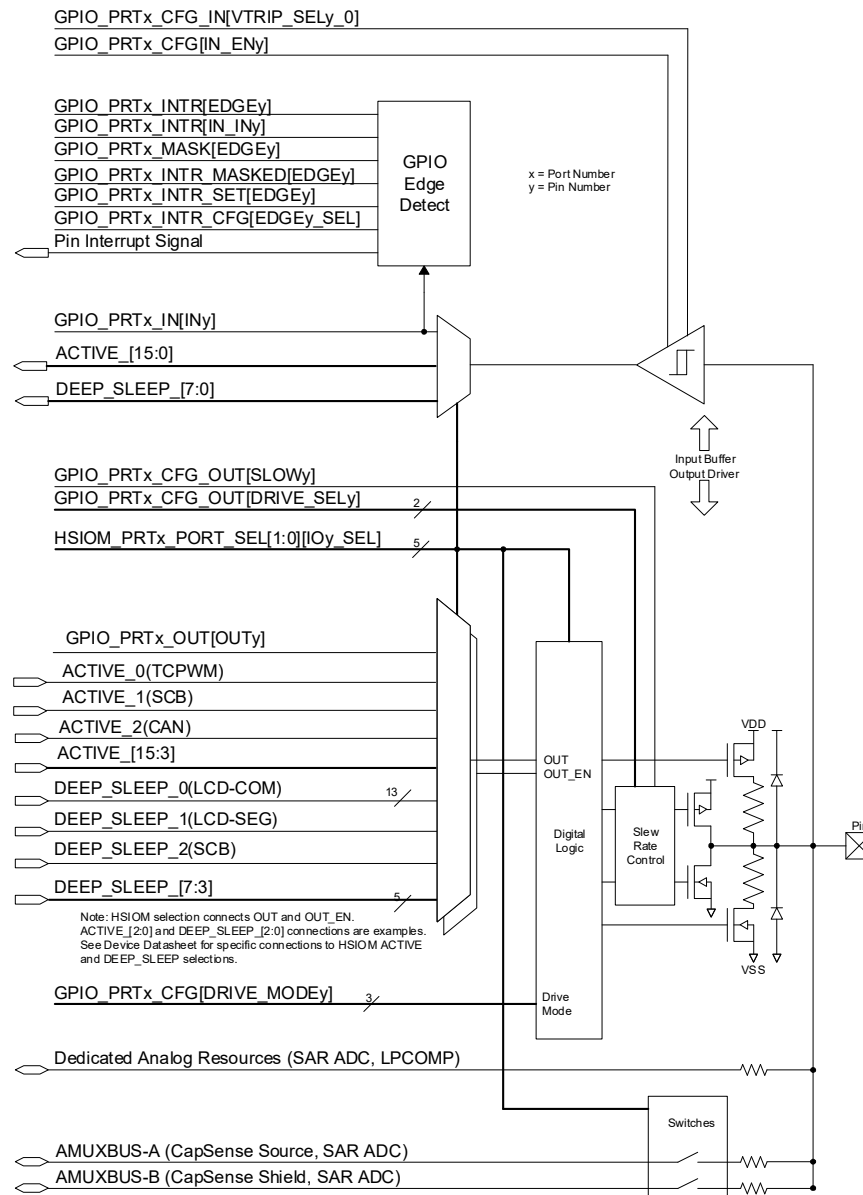


GPIO pins are connected to I/O cells. These cells are equipped with an input buffer for the digital input, providing high input impedance and a driver for the digital output signals. The digital peripherals connect to the I/O cells via the high-speed I/O matrix (HSIOM). The HSIOM for each pin contains multiplexers to connect between the selected peripheral and the pin. Analog peripherals such as SAR ADC, Low-Power comparator (LPCOMP), and CapSense are either connected to the GPIO pins directly or through the AMUXBUS.

22.2.1 I/O Cell Architecture

Figure 22-2 shows the I/O cell architecture present in every GPIO cell. It comprises an input buffer and an output driver that connect to the HSIOM multiplexers for digital input and output signals. Analog peripherals connect directly to the pin for point to point connections or use the AMUXBUS.

Figure 22-2. GPIO and GPIO_OVT Cell Architecture



22.2.2 Digital Input Buffer

The digital input buffer provides a high-impedance buffer for the external digital input. The buffer is enabled or disabled by the `IN_EN[7:0]` bit of the Port Configuration Register (`GPIO_PRTx_CFG`, where x is the port number).

The input buffer is connected to the HSIOM for routing to the CPU port registers and selected peripherals. Writing to the HSIOM port select register (`HSIOM_PORT_SELx`) selects the pin connection. See the [device datasheet](#) for the specific connections available for each pin.

If a pin is connected only to an analog signal, the input buffer should be disabled to avoid crowbar currents.

Each pin's input buffer trip point and hysteresis are configurable for the following modes:

- CMOS + I²C
- TTL

These buffer modes are selected by the VTRIP_SEL[7:0]_0 bit of the Port Input Buffer Configuration register. (GPIO_PRTx_CFG_IN).

22.2.3 Digital Output Driver

Pins are driven by the digital output driver. It consists of circuitry to implement different drive modes and slew rate control for the digital output signals. The HSIOM selects the control source for the output driver. The two primary types of control sources are CPU registers and fixed-function digital peripherals. A particular HSIOM connection is selected by writing to the HSIOM port select register (HSIOM_PORT_SELx).

I/O ports are powered by different sources. The specific allocation of ports to supply sources can be found in the Pinout section of the [device datasheet](#).

Each GPIO pin has ESD diodes to clamp the pin voltage to the I/O supply source. Ensure that the voltage at the pin does not exceed the I/O supply voltage $V_{DDIO}/V_{DD}/V_{DDA}$ or

drop below $V_{SSIO}/V_{SSD}/V_{SSA}$. For the absolute maximum and minimum GPIO voltage, see the [device datasheet](#).

The digital output driver can be enabled or disabled in hardware by using the DSI signal from a peripheral or the output data register (GPIO_PRTx_OUT) associated with the output pin. See [22.3 High-Speed I/O Matrix](#) for details on peripheral source selections supporting output enable control.

22.2.3.1 Drive Modes

Each I/O is individually configurable to one of eight drive modes by the DRIVE_MODE[7:0] field of the Port Configuration register, GPIO_PRTx_CFG. [Table 22-1](#) lists the drive modes. Drive mode '1' is reserved and should not be used in most designs. CPU register and AMUXBUS connections support seven discrete drive modes to maximize design flexibility. Fixed-function digital peripherals, such as SCB and TCPWM blocks, support modified functionality for the same seven drive modes compatible with fixed peripheral signaling. [Figure 22-3](#) shows simplified output driver diagrams of the pin view for CPU register control on each of the eight drive modes. [Figure 22-4](#) is a simplified output driver diagram that shows the pin view for fixed-function-based peripherals for each of the eight drive modes.

Table 22-1. Drive Mode Settings

Drive Mode	Value	CPU Register, AMUXBUS				Fixed-Function Digital Peripheral			
		OUT_EN = 1		OUT_EN = 0		OUT_EN = 1		OUT_EN = 0	
		OUT = 1	OUT = 0	OUT = 1	OUT = 0	OUT = 1	OUT = 0	OUT = 1	OUT = 0
High Impedance	0	HI-Z	HI-Z	HI-Z	HI-Z	HI-Z	HI-Z	HI-Z	HI-Z
Reserved	1	Strong 1	Strong 0	HI-Z	HI-Z	Strong 1	Strong 0	Weak 1 & 0	Weak 1 & 0
Resistive Pull Up	2	Weak 1	Strong 0	HI-Z	HI-Z	Strong 1	Strong 0	Weak 1	Weak 1
Resistive Pull Down	3	Strong 1	Weak 0	HI-Z	HI-Z	Strong 1	Strong 0	Weak 0	Weak 0
Open Drain, Drives Low	4	HI-Z	Strong 0	HI-Z	HI-Z	Strong 1	Strong 0	HI-Z	HI-Z
Open Drain, Drives High	5	Strong 1	HI-Z	HI-Z	HI-Z	Strong 1	Strong 0	HI-Z	HI-Z
Strong	6	Strong 1	Strong 0	HI-Z	HI-Z	Strong 1	Strong 0	HI-Z	HI-Z
Resistive Pull Up and Down	7	Weak 1	Weak 0	HI-Z	HI-Z	Strong 1	Strong 0	Weak 1	Weak 0

Figure 22-3. CPU I/O Drive Mode Block Diagram

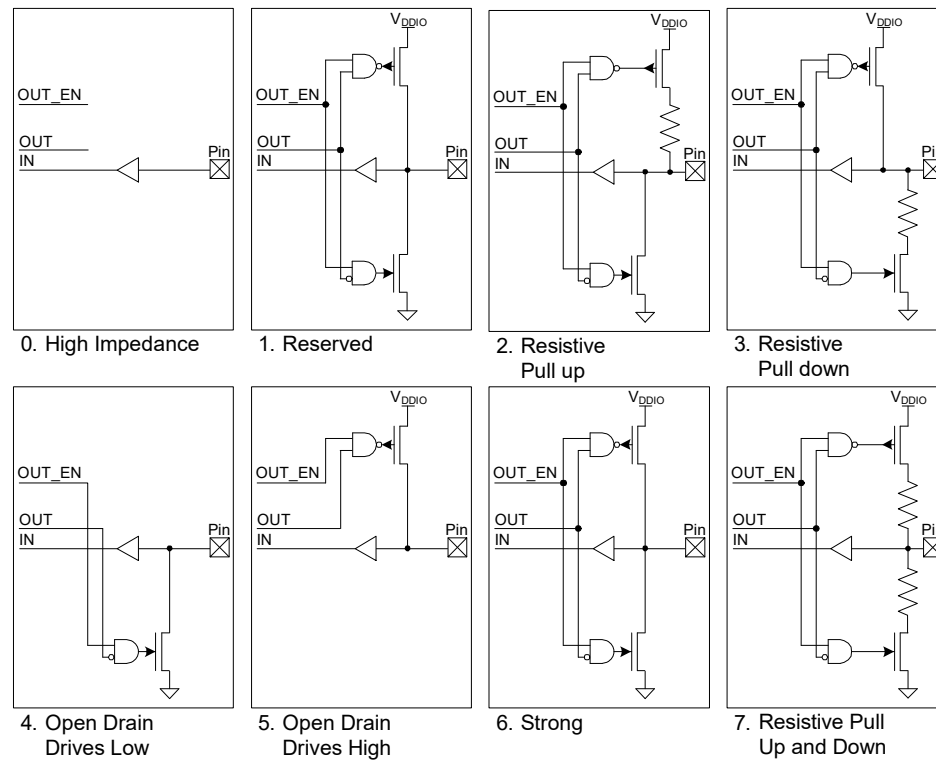



Figure 22-4. Peripheral I/O Drive Mode Block Diagrams



■ High-Impedance

This is the standard high-impedance (HI-Z) state recommended for analog and digital inputs. For digital signals, the input buffer is enabled; for analog signals, the input buffer is typically disabled to reduce crowbar current and leakage in low-power designs. To achieve the lowest device current, unused GPIOs must be configured to the high-impedance drive mode with input buffer disabled. High-impedance drive mode with input buffer disabled is also the default pin reset state.

■ Resistive Pull-Up or Resistive Pull-Down

Resistive modes provide a series resistance in one of the data states and strong drive in the other. Pins can be used for either digital input or digital output in these modes. If resistive pull-up is required, a '1' must be written to that pin's Data Register bit. If resistive pull-down is required, a '0' must be written to that pin's Data Register. Interfacing mechanical switches is a common application of these drive modes. The resistive modes are also used to interface PSoC with open drain drive lines. Resistive pull-up is used when the input is open drain low and resistive pull-down is used when the input is open drain high.

■ Open Drain Drives High and Open Drain Drives Low

Open drain modes provide high impedance in one of the data states and strong drive in the other. Pins are useful as digital inputs or outputs in these modes. Therefore, these modes are widely used in bidirectional digital communication. Open drain drive high mode is used when the signal is externally pulled down and open drain drive low is used when the signal is externally pulled high. A common application for the open drain drives low mode is driving I²C bus signal lines.

■ Strong Drive

The strong drive mode is the standard digital output mode for pins; it provides a strong CMOS output drive in both high and low states. Strong drive mode pins should not be used as inputs under normal circumstances. This mode is often used for digital output signals or to drive external devices.

■ Resistive Pull-Up and Resistive Pull-Down

In the resistive pull-up and pull-down mode, the GPIO will have a series resistance in both logic 1 and logic 0 output states. The high data state is pulled up while the low data state is pulled down. This mode is useful when the pin is driven by other signals that may cause shorts.

are recommended for low-frequency signals or signals without strict timing constraints.

When configured for fast slew rate, the drive strength can be set to one of four levels using the DRIVE_SEL field of the port output configuration register (GPIO_PRTx_CFG_OUT). The drive strength field determines the active portion of the output drivers used and can affect the slew rate of output signals. Drive strength options are full drive strength (default), one-half strength, one-quarter strength, and one-eighth strength. Drive strength must be set to full drive strength when the slow slew rate bit (SLOW) is set.

Note: For some devices in the PSoC 6 MCU family, simultaneous GPIO switching with unrestricted drive strengths and frequency can induce noise in on-chip subsystems affecting CapSense and ADC results. Refer to the Errata section in the respective device datasheet for details.

22.2.3.3 GPIO-OVT Pins

Select device pins are overvoltage tolerant (OVT) and are useful for interfacing to busses or other signals that may exceed the pin's V_{DDIO} supply, or where the whole device supply or pin V_{DDIO} may not be always present. They are identical to regular GPIOs with the additional feature of being overvoltage tolerant. GPIO-OVT pins have hardware to compare V_{DDIO} to the pin voltage. If the pin voltage exceeds V_{DDIO}, the output driver is disabled and the pin driver is tristated. This results in negligible current sink at the pin.

Note that in overvoltage conditions, the input buffer data will not be valid if the external source's specification of V_{OH} and V_{OL} do not match the trip points of the input buffer defined by the current V_{DDIO} voltage.

22.2.3.2 Slew Rate Control

GPIO pins have fast and slow output slew rate options for the strong drivers configured using the SLOW bit of the port output configuration register (GPIO_PRTx_CFG_OUT). By default, this bit is cleared and the port works in fast slew mode. This bit can be set if a slow slew rate is required. Slower slew rate results in reduced EMI and crosstalk and

22.3 High-Speed I/O Matrix

The high-speed I/O matrix (HSIOM) is a set of high-speed multiplexers that route internal CPU and peripheral signals to and from GPIOs. HSIOM allows GPIOs to be shared with multiple functions and multiplexes the pin connection to a user-selected peripheral. The HSIOM_PRTx_PORT_SEL[1:0] registers allow a single selection from up to 32 different connections to each pin as listed in [Table 22-2](#).

Table 22-2. HSIOM Connections

SELy_SEL	Name	Digital Driver Signal Source		Digital Input Signal Destination	Analog Switches		Description
		OUT	OUT_EN		AMUXA	AMUXB	
0	GPIO	OUT Register	1	IN Register	0	0	GPIO_PRTx_OUT register controls OUT
1	Reserved	–	–	–	–	–	–
2	Reserved	–	–	–	–	–	–
3	Reserved	–	–	–	–	–	–
4	AMUXA	OUT Register	1	IN Register	1	0	Analog mux bus A connected to pin
5	AMUXB	OUT Register	1	IN Register	0	1	Analog mux bus B connected to pin
6	Reserved	–	–	–	–	–	–
7	Reserved	–	–	–	–	–	–
8	ACT_0	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 0 - See the datasheet for specific pin connectivity
9	ACT_1	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 1 - See the datasheet for specific pin connectivity
10	ACT_2	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 2 - See the datasheet for specific pin connectivity
11	ACT_3	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 3 - See the datasheet for specific pin connectivity
12	DS_0	Deep Sleep Source OUT	Deep Sleep Source OUT_EN	Deep Sleep IN	0	0	Deep Sleep functionality 0 - See the datasheet for specific pin connectivity
13	DS_1	Deep Sleep Source OUT	Deep Sleep Source OUT_EN	Deep Sleep IN	0	0	Deep Sleep functionality 1 - See the datasheet for specific pin connectivity
14	DS_2	Deep Sleep Source OUT	Deep Sleep Source OUT_EN	Deep Sleep IN	0	0	Deep Sleep functionality 2 - See the datasheet for specific pin connectivity
15	DS_3	Deep Sleep Source OUT	Deep Sleep Source OUT_EN	Deep Sleep IN	0	0	Deep Sleep functionality 3 - See the datasheet for specific pin connectivity
16	ACT_4	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 4 - See the datasheet for specific pin connectivity
17	ACT_5	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 5 - See the datasheet for specific pin connectivity
18	ACT_6	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 6 - See the datasheet for specific pin connectivity
19	ACT_7	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 7 - See the datasheet for specific pin connectivity
20	ACT_8	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 8 - See the datasheet for specific pin connectivity
21	ACT_9	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 9 - See the datasheet for specific pin connectivity

Table 22-2. HSIOM Connections

SELy_SEL	Name	Digital Driver Signal Source		Digital Input Signal Destination	Analog Switches		Description
		OUT	OUT_EN		AMUXA	AMUXB	
22	ACT_10	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 10 - See the datasheet for specific pin connectivity
23	ACT_11	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 11 - See the datasheet for specific pin connectivity
24	ACT_12	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 12 - See the datasheet for specific pin connectivity
25	ACT_13	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 13 - See the datasheet for specific pin connectivity
26	ACT_14	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 14 - See the datasheet for specific pin connectivity
27	ACT_15	Active Source OUT	Active Source OUT_EN	Active Source IN	0	0	Active functionality 15 - See the datasheet for specific pin connectivity
28	DS_4	Deep Sleep Source OUT	Deep Sleep Source OUT_EN	Deep Sleep IN	0	0	Deep Sleep functionality 4 - See the datasheet for specific pin connectivity
29	DS_5	Deep Sleep Source OUT	Deep Sleep Source OUT_EN	Deep Sleep IN	0	0	Deep Sleep functionality 5 - See the datasheet for specific pin connectivity
30	DS_6	Deep Sleep Source OUT	Deep Sleep Source OUT_EN	Deep Sleep IN	0	0	Deep Sleep functionality 6 - See the datasheet for specific pin connectivity
31	DS_7	Deep Sleep Source OUT	Deep Sleep Source OUT_EN	Deep Sleep IN	0	0	Deep Sleep functionality 7 - See the datasheet for specific pin connectivity

Note: The Active and Deep Sleep sources are pin dependent. See the “Pinouts” section of the [device datasheet](#) for more details on the features supported by each pin.

22.4 I/O State on Power Up

During power up, all the GPIOs are in high-impedance analog state and the input buffers are disabled. During runtime, GPIOs can be configured by writing to the associated registers. Note that the pins supporting debug access port (DAP) connections (SWD lines) are always enabled as SWD lines during power up. The DAP connection does not provide pull-up or pull-down resistors; therefore, if left floating some crowbar current is possible. The DAP connection can be disabled or reconfigured for general-purpose use through the HSIOM only after the device boots and starts executing code.

22.5 Behavior in Low-Power Modes

Table 22-3 shows the status of GPIOs in low-power modes.


Table 22-3. GPIO in Low-Power Modes

Low-Power Mode	Status
CPU Sleep	<ul style="list-style-type: none"> Standard GPIO, GPIO-OVT, and SIO pins are active and can be driven by most peripherals such as CapSense, TCPWMs, and SCBs, which can operate in CPU Sleep mode. Inputs buffers are active; thus an interrupt on any I/O can be used to wake the CPU.
System Deep Sleep	<ul style="list-style-type: none"> GPIO, GPIO-OVT, and SIO pins, connected to System Deep Sleep domain peripherals, are functional. All other pins maintain the last output driver state and configuration. Pin interrupts are functional on all I/Os and can be used to wake the device.
System Hibernate	<ul style="list-style-type: none"> Pin output states and configuration are latched and remain in the frozen state. Pin interrupts are functional only on select IOs and can be used to wake the device. See the device datasheet for specific hibernate pin connectivity.

22.6 Interrupt

All port pins have the capability to generate interrupts. Figure 22-5 shows the routing of pin signals to generate interrupts.


Figure 22-5. Interrupt Signal Routing



- Pin signal through the “GPIO Edge Detect” block with direct connection to the CPU interrupt controller
- Pin signal through the port adapter and DSI to the CPU interrupt controller

Figure 22-6 shows the GPIO Edge Detect block architecture.

Figure 22-6. GPIO Edge Detect Block Architecture



An edge detector is present at each pin. It is capable of detecting rising edge, falling edge, and both edges without any reconfiguration. The edge detector is configured by writing into the EDGEv_SEL bits of the Port Interrupt Configuration register, GPIO_PRTx_INTR_CFG, as shown in [Table 22-4](#).

Table 22-4. Edge Detector Configuration

EDGE_SEL	Configuration
00	Interrupt is disabled
01	Interrupt on Rising Edge
10	Interrupt on Falling Edge
11	Interrupt on Both Edges

Writing '1' to the corresponding status bit clears the pin edge state. Clearing the edge state status bit is important; otherwise, an interrupt can occur repeatedly for a single trigger or respond only once for multiple triggers, which is explained later in this section. When the Port Interrupt Control Status register is read at the same time an edge is occurring on the corresponding port, it can result in the edge not being properly detected. Therefore, when using GPIO interrupts, read the status register only inside the corresponding interrupt service routine and not in any other part of the code.

Firmware and the debug interface are able to trigger a hardware interrupt from any pin by setting the corresponding bit in the GPIO_PRTx_INTR_SET register.

In addition to the pins, each port provides a glitch filter connected to its own edge detector. This filter can be driven by one of the pins of a port. The selection of the driving pin is done by writing to the FLT_SEL field of the GPIO_PRTx_INTR_CFG register as shown in [Table 22-5](#).

Table 22-5. Glitch Filter Input Selection

FLT_SEL	Selected Pin
000	Pin 0 is selected
001	Pin 1 is selected
010	Pin 2 is selected
011	Pin 3 is selected
100	Pin 4 is selected
101	Pin 5 is selected
110	Pin 6 is selected
111	Pin 7 is selected

When a port pin edge occurs, you can read the Port Interrupt Status register, GPIO_PRTx_INTR, to know which pin caused the edge. This register includes both the latched information on which pin detected an edge and the current pin status. This allows the CPU to read both information in a single read operation. This register has an additional use – to clear the latched edge state.

The GPIO_PRTx_INTR_MASK register enables forwarding of the GPIO_PRTx_INTR edge detect signal to the interrupt controller when a '1' is written to a pin's corresponding bitfield. The GPIO_PRTx_INTR_MASKED register can then be read to determine the specific pin that generated the interrupt signal forwarded to the interrupt controller. The masked edge detector outputs of a port are then ORed together and routed to the interrupt controller (NVIC in the CPU subsystem). Thus, there is only one interrupt vector per port.

The masked and ORed edge detector block output is routed to the Interrupt Source Multiplexer shown in [Figure 8-3 on page 57](#), which gives an option of Level and Rising Edge detection. If the Level option is selected, an interrupt is triggered repeatedly as long as the Port Interrupt Status register bit is set. If the Rising Edge detect option is selected, an interrupt is triggered only once if the Port Interrupt Status register is not cleared. Thus, the interrupt status bit must be cleared if the Edge Detect block is used.

Each port has a dedicated interrupt vector when the interrupt signal is routed through the fixed-function route. However, when the signal is routed through the DSI, interrupt vector connections are flexible and can occupy any of the DSI-connected interrupt lines of the NVIC. See the [Interrupts chapter on page 54](#) for details.

All of the port interrupt vectors are also ORed together into a single interrupt vector for use on devices with more ports than there are interrupt vectors available. To determine the port that triggered the interrupt, the GPIO_INTR_CAUSEx registers can be read. A '1' present in a bit location indicates that the corresponding port has a pending interrupt. The indicated GPIO_PRTx_INTR register can then be read to determine the pin source.

When the signal is routed through the DSI, bypassing the Edge Detect block, the edge detection is configurable in the Interrupt Source Multiplexer block. If the multiplexer is configured as Level, the interrupt is triggered repeatedly as long as the pin signal is high. Use the Rising Edge detect option when this route is selected to generate only one interrupt.

22.7 Peripheral Connections

22.7.1 Firmware-Controlled GPIO

For standard firmware-controlled GPIO using registers, the GPIO mode must be selected in the HSIOM_PORT_SELx register.

The GPIO_PRTx_OUT register is used to read and write the output buffer state for GPIOs. A write operation to this register changes the GPIO's output driver state to the written value. A read operation reflects the output data written to this register and the resulting output driver state. It does not return the current logic level present on GPIO pins,

which may be different. Using the GPIO_PRTx_OUT register, read-modify-write sequences can be safely performed on a port that has both input and output GPIOs.

In addition to the data register, three other registers – GPIO_PRTx_SET, GPIO_PRTx_CLR, and GPIO_PRTx_INV – are provided to set, clear, and invert the output data respectively on specific pins in a port without affecting other pins. This avoids the need for read-modify-write operations in most use cases. Writing '1' to these register bitfields will set, clear, or invert the respective pin; writing '0' will have no effect on the pin state.

GPIO_PRTx_IN is the port I/O pad register, which provides the actual logic level present on the GPIO pin when read. Writes to this register have no effect.

22.7.2 Analog I/O

Analog resources, such as LPCOMP, and SAR ADC, which require low-impedance routing paths have dedicated pins. Dedicated analog pins provide direct connections to specific analog blocks. They help improve performance and should be given priority over other pins when using these analog resources. See the [device datasheet](#) for details on these dedicated pins of the PSoC 6 MCU.

To configure a GPIO as a dedicated analog I/O, it should be configured in high-impedance analog mode (see [Table 22-1](#)) with input buffer disabled. The respective connection should be enabled via registers in the specific analog resource.

To configure a GPIO as an analog pin connecting to AMUXBUS, it should be configured in high-impedance analog mode with the input buffer disabled and then routed to the correct AMUXBUS using the HSIOM_PORT_SELx register.

While it is preferred for analog pins to disable the input buffer, it is acceptable to enable the input buffer if simultaneous analog and digital input features are required.

22.7.2.1 AMUXBUS Connection

Static connection of AMUXBUS A or B is made by selecting AMUXA or AMUXB in the HSIOM_PORT_SELx register. Dynamic hardware-controlled connection is made by selecting AMUXA_DSI or AMUXB_DSI in the HSIOM_PORT_SELx register, enabling you to implement hardware AMUXBUS switching.

To properly configure a pin as AMUXBUS input, follow these steps:

1. Configure the GPIO_PRTx_CFG register to set the pin in high-Impedance mode with input buffer disabled, enabling analog connectivity on the pin.
2. Configure the HSIOM_PRT_SELx register to connect the pin to AMUXBUS A or B. For static connections, select AMUXA or AMUXB. For dynamic connections, select AMUXA_DSI or AMUXB_DSI.

22.7.3 LCD Drive

GPIOs have the capability of driving an LCD common or segment line. HSIOM_PORT_SELx registers are used to select pins for the LCD drive. See the [LCD Direct Drive chapter on page 431](#) for details.

22.7.4 CapSense

The pins that support CapSense can be configured as CapSense widgets such as buttons, slider elements, touchpad elements, or proximity sensors. CapSense also requires external capacitors and optional shield lines. See the [PSoC 4 and PSoC 6 MCU CapSense Design Guide](#) for more details.

22.8 Smart I/O

The Smart I/O block adds programmable logic to an I/O port. This programmable logic integrates board-level Boolean logic functionality such as AND, OR, and XOR into the port. A graphical interface is provided with the ModusToolbox install for configuring the Smart I/O block. For more information about the configurator tool, see the [ModusToolbox Smart I/O Configurator Guide](#).

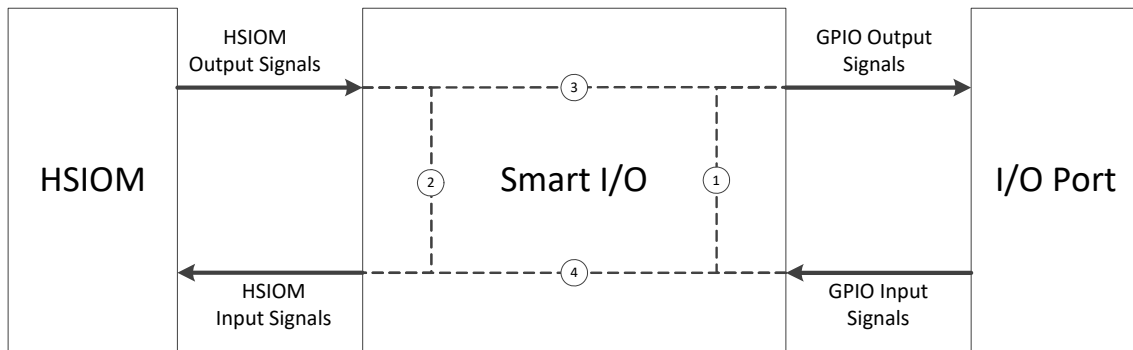
The Smart I/O block has these features:

- Integrate board-level Boolean logic functionality into a port
- Ability to preprocess HSIOM input signals from the GPIO port pins
- Ability to post-process HSIOM output signals to the GPIO port pins
- Support in all device power modes except Hibernate
- Integrate closely to the I/O pads, providing shortest signal paths with programmability

22.8.1 Overview

The Smart I/O block is positioned in the signal path between the HSIOM and the I/O port. The HSIOM multiplexes the output signals from fixed-function peripherals and CPU to a specific port pin and vice-versa. The Smart I/O block is placed on this signal path, acting as a bridge that can process signals between port pins and HSIOM, as shown in [Figure 22-7](#).

Figure 22-7. Smart I/O Interface



The signal paths supported through the Smart I/O block as shown in [Figure 22-7](#) are as follows:

1. Implement self-contained logic functions that directly operate on port I/O signals
2. Implement self-contained logic functions that operate on HSIOM signals
3. Operate on and modify HSIOM output signals and route the modified signals to port I/O signals
4. Operate on and modify port I/O signals and route the modified signals to HSIOM input signals

The following sections discuss the Smart I/O block components, routing, and configuration in detail. In these sections, GPIO signals (`io_data`) refer to the input/output signals from the I/O port; device or chip (`chip_data`) signals refer to the input/output signals from HSIOM.

22.8.2 Block Components

The internal logic of the Smart I/O includes these components:

- Clock/reset
- Synchronizers
- Three-input lookup table (LUT)
- Data unit

22.8.2.1 Clock and Reset

The clock and reset component selects the Smart I/O block's clock (`clk_block`) and reset signal (`rst_block_n`). A single clock and reset signal is used for all components in the block. The clock and reset sources are determined by the `CLOCK_SRC[4:0]` bitfield of the `SMARTIO_PRTx_CTL` register. The selected clock is used for the synchronous logic in the block components, which includes the I/O input synchronizers, LUT, and data unit components. The selected reset is used to asynchronously reset the synchronous logic in the LUT and data unit components.

Note that the selected clock (`clk_block`) for the block's synchronous logic is not phase-aligned with other synchronous logic in the device, operating on the same

clock. Therefore, communication between Smart I/O and other synchronous logic should be treated as asynchronous.

The following clock sources are available for selection:

- GPIO input signals "`io_data_in[7:0]`". These clock sources have no associated reset.
- HSIOM output signals "`chip_data[7:0]`". These clock sources have no associated reset.
- Smart I/O clock (`clk_smartio`). This is derived from the system clock (`clk_sys`) using a peripheral clock divider. See the [Clocking System chapter on page 207](#) for details on peripheral clock dividers. This clock is available only in System LP and ULP power modes. The clock can have one out of two associated resets: `rst_sys_act_n` and `rst_sys_dpslp_n`. These resets determine in which system power modes the block synchronous state is reset; for example, `rst_sys_act_n` is intended for Smart I/O synchronous functionality in the System LP and ULP power modes and reset is activated in the System Deep Sleep power mode.
- Low-frequency system clock (`clk_lf`). This clock is available in System Deep Sleep power mode. This clock has an associated reset, `rst_lf_dpslp_n`. Reset is activated if the system enters Hibernate, or is at POR.

When the block is enabled, the selected clock (`clk_block`) and associated reset (`rst_block_n`) are provided to the fabric components. When the fabric is disabled, no clock is released to the fabric components and the reset is activated (the LUT and data unit components are set to the reset value of '0').

The I/O input synchronizers introduce a delay of two `clk_block` cycles (when synchronizers are enabled). As a result, in the first two cycles, the block may be exposed to stale data from the synchronizer output. Hence, during the first two clock cycles, the reset is activated and the block is in bypass mode.


Table 22-6. Clock and Reset Register Control

Register[BIT_POS]	Bit Name	Description
SMARTIO_PRTn_CTL[12:8]	CLOCK_SRC[4:0]	<p>Clock (clk_block)/reset (rst_block_n) source selection:</p> <p>0: io_data_in[0]/1</p> <p>...</p> <p>7: io_data_in[7]/1</p> <p>8: chip_data[0]/1</p> <p>...</p> <p>15: chip_data[7]/1</p> <p>16: clk_smartio/rst_sys_act_n; asserts reset in any power mode other than System LP or ULP; that is, Smart I/O is active only in LP or ULP power modes with clock from the peripheral divider.</p> <p>17: clk_smartio/rst_sys_dpslp_n. Smart I/O is active in all power modes with a clock from the peripheral divider. However, the clock will not be active in System deep sleep power mode.</p> <p>19: clk_if/rst_if_dpslp_n. Smart I/O is active in all power modes with a clock from ILO.</p> <p>20-30: Clock source is a constant '0'. Any of these clock sources should be selected when Smart I/O is disabled to ensure low power consumption.</p> <p>31: clk_sys/1. This selection is not intended for clk_sys operation. However, for asynchronous operation, three clk_sys cycles after enabling, the Smart I/O is fully functional (reset is de-activated). To be used for asynchronous (clockless) block functionality.</p>

22.8.2.2 Synchronizer

Each GPIO input signal and device input signal (HSIOM input) can be used either asynchronously or synchronously. To use the signals synchronously, a double flip-flop synchronizer, as shown in Figure 22-8, is placed on both these signal paths to synchronize the signal to the Smart I/O clock (clk_block). The synchronization for each pin/input is enabled or disabled by setting or clearing the IO_SYNC_EN[i] bitfield for GPIO input signal and CHIP_SYNC_EN[i] for HSIOM signal in the SMARTIO_PRTx_SYNC_CTL register, where 'i' is the pin number.

Figure 22-8. Smart I/O Clock Synchronizer



22.8.2.3 Lookup Table (LUT)

Each Smart I/O block contains eight lookup table (LUT) components. The LUT component consists of a three-input LUT and a flip-flop. Each LUT block takes three input signals and generates an output based on the configuration set in the SMARTIO_PRTx_LUT_CTLy register (y denotes the LUT number). For each LUT, the configuration is determined by an 8-bit lookup vector LUT[7:0] and a 2-bit opcode OPC[1:0] in the SMARTIO_PRTx_LUT_CTLy register. The 8-bit vector is used as a lookup table for the three input signals. The 2-bit opcode determines the usage of the flip-flop. The LUT configuration for different opcodes is shown in Figure 22-9.

The SMARTIO_PRTx_LUT_SELy registers select the three input signals (tr0_in, tr1_in, and tr2_in) going into each LUT. The input can come from the following sources:


- Data unit output
- Other LUT output signals (tr_out)
- HSIOM output signals (chip_data[7:0])
- GPIO input signals (io_data[7:0])

LUT_TR0_SEL[3:0] bits of the SMARTIO_PRTx_LUT_SELy register selects the tr0_in signal for the yth LUT. Similarly, LUT_TR1_SEL[3:0] bits and LUT_TR2_SEL[3:0] bits select the tr1_in and tr2_in signals, respectively. See [Table 22-7](#) for details.

Table 22-7. LUT Register Control

Register[BIT_POS]	Bit Name	Description
SMARTIO_PRTx_LUT_CTLy[7:0]	LUT[7:0]	LUT configuration. Depending on the LUT opcode (LUT_OPC), the internal state, and the LUT input signals tr0_in, tr1_in, and tr2_in, the LUT configuration is used to determine the LUT output signal and the next sequential state.
SMARTIO_PRTx_LUT_CTLy[9:8]	LUT_OPC[1:0]	LUT opcode specifies the LUT operation as illustrated in Figure 22-9 .
SMARTIO_PRTx_LUT_SELy[3:0]	LUT_TR0_SEL[3:0]	LUT input signal “tr0_in” source selection: 0: Data unit output 1: LUT 1 output 2: LUT 2 output 3: LUT 3 output 4: LUT 4 output 5: LUT 5 output 6: LUT 6 output 7: LUT 7 output 8: chip_data[0] (for LUTs 0, 1, 2, 3); chip_data[4] (for LUTs 4, 5, 6, 7) 9: chip_data[1] (for LUTs 0, 1, 2, 3); chip_data[5] (for LUTs 4, 5, 6, 7) 10: chip_data[2] (for LUTs 0, 1, 2, 3); chip_data[6] (for LUTs 4, 5, 6, 7) 11: chip_data[3] (for LUTs 0, 1, 2, 3); chip_data[7] (for LUTs 4, 5, 6, 7) 12: io_data_in[0] (for LUTs 0, 1, 2, 3); io_data[4] (for LUTs 4, 5, 6, 7) 13: io_data_in[1] (for LUTs 0, 1, 2, 3); io_data[5] (for LUTs 4, 5, 6, 7) 14: io_data_in[2] (for LUTs 0, 1, 2, 3); io_data[6] (for LUTs 4, 5, 6, 7) 15: io_data_in[3] (for LUTs 0, 1, 2, 3); io_data[7] (for LUTs 4, 5, 6, 7)
SMARTIO_PRTx_LUT_SELy[11:8]	LUT_TR1_SEL[3:0]	LUT input signal “tr1_in” source selection: 0: LUT 0 output 1: LUT 1 output 2: LUT 2 output 3: LUT 3 output 4: LUT 4 output 5: LUT 5 output 6: LUT 6 output 7: LUT 7 output 8: chip_data[0] (for LUTs 0, 1, 2, 3); chip_data[4] (for LUTs 4, 5, 6, 7) 9: chip_data[1] (for LUTs 0, 1, 2, 3); chip_data[5] (for LUTs 4, 5, 6, 7) 10: chip_data[2] (for LUTs 0, 1, 2, 3); chip_data[6] (for LUTs 4, 5, 6, 7) 11: chip_data[3] (for LUTs 0, 1, 2, 3); chip_data[7] (for LUTs 4, 5, 6, 7) 12: io_data_in[0] (for LUTs 0, 1, 2, 3); io_data[4] (for LUTs 4, 5, 6, 7) 13: io_data_in[1] (for LUTs 0, 1, 2, 3); io_data[5] (for LUTs 4, 5, 6, 7) 14: io_data_in[2] (for LUTs 0, 1, 2, 3); io_data[6] (for LUTs 4, 5, 6, 7) 15: io_data_in[3] (for LUTs 0, 1, 2, 3); io_data[7] (for LUTs 4, 5, 6, 7)
SMARTIO_PRTx_LUT_SELy[19:16]	LUT_TR2_SEL[3:0]	LUT input signal “tr2_in” source selection. Encoding is the same as for LUT_TR1_SEL.

Figure 22-9. Smart I/O LUT Configuration



22.8.2.4 Data Unit (DU)

Each Smart I/O block includes a data unit (DU) component. The DU consists of a simple 8-bit datapath. It is capable of performing simple increment, decrement, increment/decrement, shift, and AND/OR operations. The operation performed by the DU is selected using a 4-bit opcode $DU_OPC[3:0]$ bitfield in the $SMARTIO_PRTx_DU_CTL$ register.

The DU component supports up to three input trigger signals ($tr0_in$, $tr1_in$, $tr2_in$) similar to the LUT component. These signals are used to initiate an operation defined by the DU opcode. In addition, the DU also includes two 8-bit data inputs ($data0_in[7:0]$ and $data1_in[7:0]$) that are used to initialize the 8-bit internal state ($data[7:0]$) or to provide a reference. The 8-bit data input source is configured as:

- Constant '0x00'
- $io_data_in[7:0]$
- $chip_data_in[7:0]$
- $DATA[7:0]$ bitfield of $SMARTIO_PRTx_DATA$ register

The trigger signals are selected using the DU_TRx_SEL[3:0] bitfield of the SMARTIO_PRTx_DU_SEL register. The DUT_DATAx_SEL[1:0] bits of the SMARTIO_PRTx_DU_SEL register select the 8-bit input data source. The size of the DU (number of bits used by the datapath) is defined by the DU_SIZE[2:0] bits of the SMARTIO_PRTx_DU_CTL register. See [Table 22-8](#) for register control details.

Table 22-8. Data Unit Register Control

Register[BIT_POS]	Bit Name	Description
SMARTIO_PRTx_DU_CTL[2:0]	DU_SIZE[2:0]	Size/width of the data unit (in bits) is DU_SIZE+1. For example, if DU_SIZE is 7, the width is 8 bits.
SMARTIO_PRTx_DU_CTL[11:8]	DU_OPC[3:0]	Data unit opcode specifies the data unit operation: 1: INCR 2: DECR 3: INCR_WRAP 4: DECR_WRAP 5: INCR_DECR 6: INCR_DECR_WRAP 7: ROR 8: SHR 9: AND_OR 10: SHR_MAJ3 11: SHR_EQL Otherwise: Undefined.
SMARTIO_PRTx_DU_SEL[3:0]	DU_TR0_SEL[3:0]	Data unit input signal "tr0_in" source selection: 0: Constant '0'. 1: Constant '1'. 2: Data unit output. 10–3: LUT 7–0 outputs. Otherwise: Undefined.
SMARTIO_PRTx_DU_SEL[11:8]	DU_TR1_SEL[3:0]	Data unit input signal "tr1_in" source selection. Encoding same as DU_TR0_SEL
SMARTIO_PRTx_DU_SEL[19:16]	DU_TR2_SEL[3:0]	Data unit input signal "tr2_in" source selection. Encoding same as DU_TR0_SEL
SMARTIO_PRTx_DU_SEL[25:24]	DU_DATA0_SEL[1:0]	Data unit input data "data0_in" source selection: 0: 0x00 1: chip_data[7:0]. 2: io_data[7:0]. 3: SMARTIO_PRTx_DATA.DATA[7:0] register field.
SMARTIO_PRTx_DU_SEL[29:28]	DU_DATA1_SEL[1:0]	Data unit input data "data1_in" source selection. Encoding same as DU_DATA0_SEL.
SMARTIO_PRTx_DATA[7:0]	DATA[7:0]	Data unit input data source.

The DU generates a single output trigger signal (tr_out). The internal state (du_data[7:0]) is captured in flip-flops and requires clk_block.

The following pseudo code describes the various datapath operations supported by the DU opcode. Note that "Comb" describes the combinatorial functionality – that is, functions that operate independent of previous output states. "Reg" describes the registered functionality – that is, functions that operate on inputs and previous output states (registered using flip-flops).

```
// The following is shared by all operations.
mask = (2 ^ (DU_SIZE+1) - 1)
data_eq1_data1_in = (data & mask) == (data1_in & mask));
data_eq1_0       = (data & mask) == 0);
data_incr        = (data + 1) & mask;
data_decr        = (data - 1) & mask;
```

```

data0_masked      = data_in0 & mask;

// INCR operation: increments data by 1 from an initial value (data0) until it reaches a
// final value (data1).
Comb:tr_out = data_eql_data1_in;
Reg:  data <= data;
      if (tr0_in)      data <= data0_masked; //tr0_in is reload signal - loads masked data0
                                      // into data
      else if (tr1_in) data <= data_eql_data1_in ? data : data_incr; //increment data until
                                      // it equals data1

// INCR_WRAP operation: operates similar to INCR but instead of stopping at data1, it wraps
// around to data0.
Comb:tr_out = data_eql_data1_in;
Reg:  data <= data;
      if (tr0_in)      data <= data0_masked;
      else if (tr1_in) data <= data_eql_data1_in ? data0_masked : data_incr;

// DECR operation: decrements data from an initial value (data0) until it reaches 0.
Comb:tr_out = data_eql_0;
Reg:  data <= data;
      if (tr0_in)      data <= data0_masked;
      else if (tr1_in) data <= data_eql_0      ? data : data_decr;

// DECR_WRAP operation: works similar to DECR. Instead of stopping at 0, it wraps around to
// data0.
Comb:tr_out = data_eql_0;
Reg:  data <= data;
      if (tr0_in)      data <= data0_masked;
      else if (tr1_in) data <= data_eql_0      ? data0_masked : data_decr;

// INCR_DECR operation: combination of INCR and DECR. Depending on trigger signals it either
// starts incrementing or decrementing. Increment stops at data1 and decrement stops at 0.
Comb:tr_out = data_eql_data1_in | data_eql_0;
Reg:  data <= data;
      if (tr0_in)      data <= data0_masked; // Increment operation takes precedence over
                                      // decrement when both signal are available
      else if (tr1_in) data <= data_eql_data1_in ? data : data_incr;
      else if (tr2_in) data <= data_eql_0      ? data : data_decr;

// INCR_DECR_WRAP operation: same functionality as INCR_DECR with wrap around to data0 on
// reaching the limits.
Comb:tr_out = data_eql_data1_in | data_eql_0;
Reg:  data <= data;
      if (tr0_in)      data <= data0_masked;
      else if (tr1_in) data <= data_eql_data1_in ? data0_masked : data_incr;
      else if (tr2_in) data <= data_eql_0      ? data0_masked : data_decr;

// ROR operation: rotates data right and LSb is sent out. The data for rotation is taken from
// data0.
Comb:tr_out = data[0];
Reg:  data <= data;
      if (tr0_in)      data <= data0_masked;
      else if (tr1_in) {
          data <= {0, data[7:1]} & mask; //Shift right operation
          data[du_size] <= data[0]; //Move the data[0] (LSb) to MSb
      }
  
```

```
// SHR operation: performs shift register operation. Initial data (data0) is shifted out and
// data on tr2_in is shifted in.
Comb:tr_out = data[0];
Reg:  data <= data;
      if (tr0_in)      data          <= data0_masked;
      else if (tr1_in) {
                          data          <= {0, data[7:1]} & mask; //Shift right operation
                          data[du_size] <= tr2_in; //tr2_in Shift in operation
      }

// SHR_MAJ3 operation: performs the same functionality as SHR. Instead of sending out the
// shifted out value, it sends out a '1' if in the last three samples/shifted-out values
// (data[0]), the signal high in at least two samples. otherwise, sends a '0'. This function
// sends out the majority of the last three samples.
Comb:tr_out = (data == 0x03)
              | (data == 0x05)
              | (data == 0x06)
              | (data == 0x07);
Reg:  data <= data;
      if (tr0_in)      data          <= data0_masked;
      else if (tr1_in) {
                          data          <= {0, data[7:1]} & mask;
                          data[du_size] <= tr2_in;
      }


// SHR_EQL operation: performs the same operation as SHR. Instead of shift-out, the output is
// a comparison result (data0 == data1).
Comb:tr_out = data_eql_data1_in;
Reg:  data <= data;
      if (tr0_in) data          <= data0_masked;
      else if (tr1_in) {
                          data          <= {0, data[7:1]} & mask;
                          data[du_size] <= tr2_in;
      }

// AND_OR operation: ANDs data1 and data0 along with mask; then, ORs all the bits of the
// ANDed output.
Comb:tr_out = | (data & data1_in & mask);
Reg:  data <= data;
      if (tr0_in) data <= data0_masked;
```

22.8.3 Routing

The Smart I/O block includes many switches that are used to route the signals in and out of the block and also between various components present inside the block. The routing switches are handled through the PRTGIO_PRTx_LUT_SELy and SMARTIO_PRTx_DU_SEL registers. Refer to the [registers TRM](#) for details. The Smart I/O internal routing is shown in [Figure 22-10](#). In the figure, note that LUT7 to LUT4 operate on io_data/chip_data[7] to io_data/chip_data[4] whereas LUT3 to LUT0 operate on io_data/chip_data[3] to io_data/chip_data[0].

Figure 22-10. Smart I/O Routing



22.8.4 Operation

The Smart I/O block should be configured and operated as follows:

1. Before enabling the block, all the components and routing should be configured as explained in ["Block Components" on page 237](#).
2. In addition to configuring the components and routing, some block level settings must be configured correctly for desired operation.
 - a. Bypass control: The Smart I/O path can be bypassed for a particular GPIO signal by setting the BYPASS[i] bitfield in the SMARTIO_PRTx_CTL register. When bit 'i' is set in the BYPASS[7:0] bitfield, the ith GPIO signal is bypassed to the HSIOM signal path directly – Smart I/O logic will not be present in that signal path. This is useful when the Smart I/O function is required only on select I/Os.
 - b. Pipelined trigger mode: The LUT input multiplexers and the LUT component itself do not include any combinatorial loops. Similarly, the data unit also does not include any combinatorial loops. However, when one LUT interacts with the other or to the data unit, inadvertent combinatorial loops are possible. To overcome this limitation, the PIPELINE_EN bitfield of the SMARTIO_PRTx_CTL register is used. When set, all the outputs (LUT and DU) are registered before branching out to other components.
3. After the Smart I/O block is configured for the desired functionality, the block can be enabled by setting the ENABLED bitfield of the SMARTIO_PRTx_CTL register. If disabled, the Smart I/O block is put in bypass mode, where the GPIO signals are directly controlled by the HSIOM signals and vice-versa. The Smart I/O block must be configured; that is, all register settings must be updated before enabling the block to prevent glitches during register updates.

Table 22-9. Smart I/O Block Controls

Register [BIT_POS]	Bit Name	Description
SMARTIO_PRTx_CTL[25]	PIPELINE_EN	Enable for pipeline register: 0: Disabled (register is bypassed). 1: Enabled
SMARTIO_PRTx_CTL[31]	ENABLED	Enable Smart I/O. Should only be set to '1' when the Smart I/O is completely configured: 0: Disabled (signals are bypassed; behavior as if BYPASS[7:0] is 0xFF). When disabled, the block (data unit and LUTs) reset is activated. If the block is disabled: - The PIPELINE_EN register field should be set to '1', to ensure low power consumption. - The CLOCK_SRC register field should be set to 20 to 30 (clock is constant '0'), to ensure low power consumption. 1: Enabled. When enabled, it takes three clk_block clock cycles until the block reset is de-activated and the block becomes fully functional. This action ensures that the I/O pins' input synchronizer states are flushed when the block is fully functional.
SMARTIO_PRTx_CTL[7:0]	BYPASS[7:0]	Bypass of the Smart I/O, one bit for each I/O pin: BYPASS[i] is for I/O pin i. When ENABLED is '1', this field is used. When ENABLED is '0', this field is not used and Smart I/O is always bypassed. 0: No bypass (Smart I/O is present in the signal path) 1: Bypass (Smart I/O is absent in the signal path)

22.9 Registers

Table 22-10. I/O Registers

Name	Description
GPIO_PRTx_OUT	Port output data register reads and writes the output driver data for I/O pins in the port.
GPIO_PRTx_OUT_CLR	Port output data clear register clears output data of specific I/O pins in the port.
GPIO_PRTx_OUT_SET	Port output data set register sets output data of specific I/O pins in the port.
GPIO_PRTx_OUT_INV	Port output data invert register inverts output data of specific I/O pins in the port.
GPIO_PRTx_IN	Port input state register reads the current pin state present on I/O pin inputs.
GPIO_PRTx_INTR	Port interrupt status register reads the current pin interrupt state.
GPIO_PRTx_INTR_MASK	Port interrupt mask register configures the mask that forwards pin interrupts to the CPU's interrupt controller.
GPIO_PRTx_INTR_MASKED	Port interrupt masked status register reads the masked interrupt status forwarded to the CPU interrupt controller.
GPIO_PRTx_INTR_SET	Port interrupt set register allows firmware to set pin interrupts.
GPIO_PRTx_INTR_CFG	Port interrupt configuration register selects the edge detection type for each pin interrupt.
GPIO_PRTx_CFG	Port configuration register selects the drive mode and input buffer enable for each pin.
GPIO_PRTx_CFG_IN	Port input buffer configuration register configures the input buffer mode for each pin.
GPIO_PRTx_CFG_OUT	Port output buffer configuration register selects the output driver slew rate for each pin.
HSIOM_PORT_SELx	High-speed I/O Mux (HSIOM) port selection register selects the hardware peripheral connection to I/O pins.

Note The 'x' in the GPIO register name denotes the port number. For example, GPIO_PTR1_OUT is the Port 1 output data register.

23. Watchdog Timer



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - WDT
- Application notes
- Code examples

The watchdog timer (WDT) is a hardware timer that automatically resets the device in the event of an unexpected firmware execution path. The WDT, if enabled, must be serviced periodically in firmware to avoid a reset. Otherwise, the timer elapses and generates a device reset. In addition, the WDT can be used as an interrupt source or a wakeup source in low-power modes.

The PSoC 6 MCU family includes one WDT and two multi-counter WDTs (MCWDT). The WDT has a 16-bit counter. Each MCWDT has two 16-bit counters and one 32-bit counter. Thus, the watchdog system has a total of seven counters – five 16-bit and two 32-bit. All 16-bit counters can generate a watchdog device reset. All seven counters can generate an interrupt on a match event.

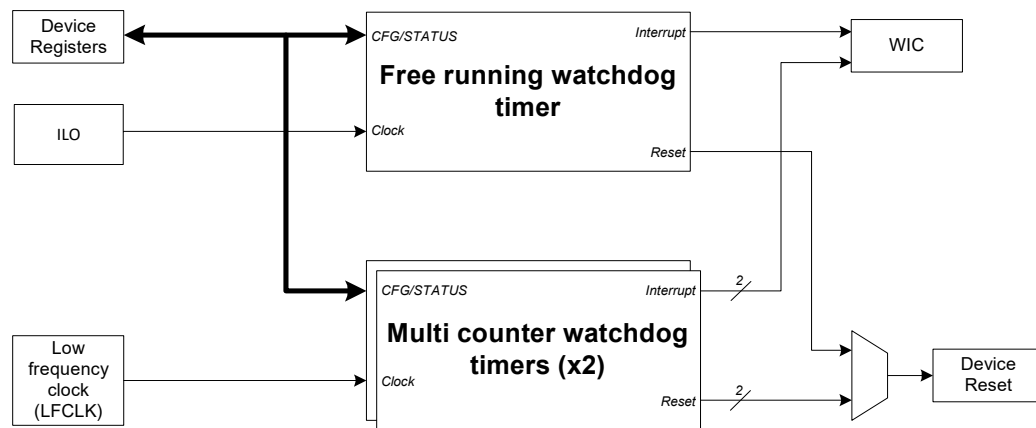
23.1 Features

The PSoC 6 MCU WDT supports these features:

- One 16-bit free-running WDT with:
 - ILO as the input clock source
 - Device reset generation if not serviced within a configurable interval
 - Interrupt/wakeup generation in LP/ULP Active, LP/ULP Sleep, Deep Sleep, and Hibernate power modes
- Two MCWDTs, each supporting:
 - Device reset generation if not serviced within a configurable interval
 - LFCLK (ILO or WCO) as the input clock source
 - Periodic interrupt/wake up generation in LP/ULP Active, LP/ULP Sleep, and Deep Sleep power modes
 - Two 16-bit and one 32-bit independent counters
 - One 64-bit or one 48-bit (with one 16-bit independent counter), or two 32-bit cascaded counters

23.2 Architecture

Figure 23-1. Watchdog Timer Block Diagram

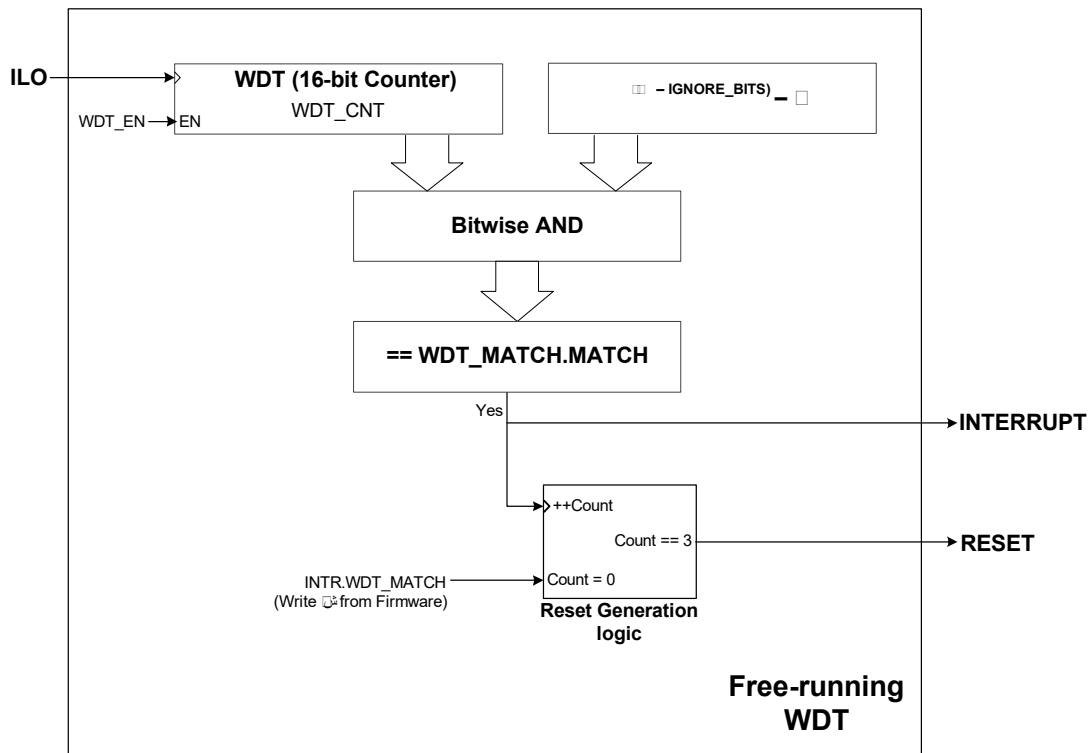


23.3 Free-running WDT

23.3.1 Overview

Figure 23-2 shows the functional overview of the WDT. The WDT has a free-running wraparound up-counter with a maximum of 16-bit resolution. The counter is clocked by the ILO. The timer has the capability to generate an interrupt on match and a reset event on the third unhandled interrupt. The number of bits used for a match comparison is configurable as depicted in Figure 23-2.

Figure 23-2. Free-running WDT Functional Diagram



When enabled, the WDT counts up on each rising edge of the ILO. When the counter value (WDT_CNT register) equals the match value stored in MATCH bits [15:0] of the WDT_MATCH register, an interrupt is generated. The match event does not reset the WDT counter and the WDT keeps counting until it reaches the 16-bit boundary (65535) at which point, it wraps around to 0 and counts up. The match interrupt is generated every time the counter value equals the match value.

The WDT_MATCH bit of the SRSS_INTR register is set whenever a WDT match interrupt occurs. This interrupt must be cleared by writing a '1' to the same bit (WDT_MATCH bit of SRSS_INTR). Clearing the interrupt resets the watchdog. If the firmware does not clear the interrupt for two consecutive occasions, the third interrupt generates a device reset.

In addition, the WDT provides an option to set the number of bits to be used for comparison. The IGNORE_BITS bits [19:16] of the WDT_MATCH register is used for this purpose. These bits configure the number of MSBs to ignore from the 16-bit count value while performing the match. For instance, when the value of these bits equals 3, the MSb 3

bits are ignored while performing the match and the WDT counter behaves similar to a 13-bit counter. Note that these bits do not reduce the counter size – the WDT_CNT register still counts from 0 to 65535 (16-bit).

The WDT can be enabled or disabled using the WDT_EN bit [0] of the WDT_CTL register. The WDT_CTL register provides a mechanism to lock the WDT configuration registers. The WDT_LOCK bits [31:30] control the lock status of the WDT registers. These bits are special bits, which can enable the lock in a single write; to release the lock, two different writes are required. The WDT_LOCK bits protect the WDT_EN bit, WDT_MATCH register, CLK_ILO_CONFIG register, and LFCLK_SEL bits [1:0] of the CLK_SELECT register. Note that the WDT_LOCK bits are not retained in Deep Sleep mode and reset to their default (LOCK) state after a deep sleep wakeup. As a result, to update any register protected by the WDT_LOCK bits after a deep sleep wakeup, a WDT UNLOCK sequence should be issued before the register update.

[Table 23-1](#) explains various registers and bitfields used to configure and use the WDT.

Table 23-1. Free-running WDT Configuration Options

Register [Bit_Pos]	Bit_Name	Description
WDT_CTL[0]	WDT_EN	Enable or disable the watchdog reset 0: WDT reset disabled 1: WDT reset enabled
WDT_CTL[31:30]	WDT_LOCK	Lock or unlock write access to the watchdog configuration and clock related registers. When the bits are set, the lock is enabled. 0: No effect 1: Clear bit 0 2: Clear bit 1 3: Set both bit 0 and 1 (lock enabled) WDT will lock on a reset. This field is not retained in Deep Sleep or Hibernate mode, so the WDT will be locked after wakeup from these modes.
WDT_CNT[15:0]	COUNTER	Current value of WDT counter
WDT_MATCH[15:0]	MATCH	Match value to a generate watchdog match event or interrupt
WDT_MATCH[19:16]	IGNORE_BITS	Number of MSBs of the WDT_CNT register to ignore for comparison with the MATCH value. Up to 12 MSBs can be ignored; settings above 12 act similar to a setting of 12.
SRSS_INTR[0]	WDT_MATCH	WDT interrupt request This bit is set whenever a watchdog match event happens. The WDT interrupt is cleared by writing a '1' to this bit
SRSS_INTR_MASK[0]	WDT_MATCH	Mask for the WDT interrupt 0: WDT interrupt is not masked to CPU 1: WDT interrupt is masked to CPU

23.3.2 Watchdog Reset

A watchdog is typically used to protect the device against firmware/system crashes or faults. When the WDT is used to protect against system crashes, the WDT interrupt bit should be cleared from a portion of the code that is not directly associated with the WDT interrupt. Otherwise, even if the main function of the firmware crashes or is in an endless loop, the WDT interrupt vector can still be intact and feed the WDT periodically.

The safest way to use the WDT against system crashes is to:

- Configure the watchdog reset period such that firmware is able to reset the watchdog at least once during the period, even along the longest firmware delay path.
- Reset (feed) the watchdog by clearing the interrupt bit regularly in the main body of the firmware code by writing a '1' to the WDT_MATCH bit in the SRSS_INTR register. Note that this does not reset the watchdog counter, it feeds only the watchdog so that it does not cause a reset for the next two match events.

Do not reset the watchdog (clear interrupt) in the WDT interrupt service routine (ISR) if WDT is being used as a reset source to protect the system against crashes. Therefore, do not use the WDT reset feature and ISR together.

The recommended steps to use WDT as a reset source are as follows:

1. Make sure the WDT configuration is unlocked by clearing the WDT_LOCK bits[31:30] of the WDT_CTL register. Note that clearing the bits requires two writes to the register with each write clearing one bit as explained in [Table 23-1](#).
2. Write the desired IGNORE_BITS in the WDT_MATCH register to set the counter resolution to be used for the match.
3. Write any match value to the WDT_MATCH register. The match value does not control the period of watchdog reset as the counter is not reset on a match event. This value provides an option to control only the first interrupt interval, after that the successive interrupts' period is defined by the IGNORE_BITS. Approximate watchdog period is given by the following equation:

$$\text{Watchdog reset period} = \text{ILO}_{\text{period}} \times (2 \times 2^{16 - \text{IGNORE_BITS}} + \text{WDT_MATCH})$$

Equation 23-1

4. Set the WDT_MATCH bit in the SRSS_INTR register to clear any pending WDT interrupt.
5. Enable ILO by setting the ENABLE bit [31] of the CLK_ILO_CONFIG register.
6. Enable the WDT by setting the WDT_EN bit in WDT_CTL register.
7. Lock the WDT and ILO configuration by writing '3' to the WDT_LOCK bits. This also locks the LFCLK_SEL bits of the CLK_SELECT register.
8. In the firmware, write '1' to the WDT_MATCH bit in the SRSS_INT register to feed (clear interrupt) the watchdog.

23.3.3 Watchdog Interrupt

In addition to generating a device reset, the WDT can be used to generate interrupts. The watchdog counter can send interrupt requests to the CPU in CPU Active power modes and to the wakeup interrupt controller (WIC) in CPU Sleep and Deep Sleep power modes. In addition, the watchdog is capable of waking up the device from Hibernate power mode. It works as follows:

- **Active Mode:** In Active power mode, the WDT can send the interrupt to the CPU. The CPU acknowledges the interrupt request and executes the ISR. Clear the interrupt in the ISR.
- **Sleep or Deep Sleep Mode:** In this mode, the CPU subsystem is powered down. Therefore, the interrupt request from the WDT is directly sent to the WIC, which then wakes up the CPU. The CPU acknowledges the interrupt request and executes the ISR. Clear the interrupt in the ISR firmware.
- **Hibernate Mode:** In this mode, the entire device except a few peripherals (such as WDT and LPCOMP) are powered down. Any interrupt to wake up the device in this mode results in a device reset. Hence, there is no interrupt service routine or mechanism associated with this mode.

For more details on device power modes, see the [Device Power Modes chapter on page 190](#).

Because of its free-running nature, the WDT should not be used for periodic interrupt generation. Use the MCWDT instead; see [23.4 Multi-Counter WDTs](#). The MCWDT counters can be used to generate periodic interrupts. If absolutely required, follow these steps to use the WDT as a periodic interrupt generator:

1. Unlock the WDT if this is the first update to the WDT registers after a deep sleep or hibernate wakeup, or a device reset.
2. Write the desired IGNORE_BITS in the WDT_MATCH register to set the counter resolution to be used for the match.

3. Write the desired match value to the WDT_MATCH register.
4. Set the WDT_MATCH bit in the SRSS_INTR register to clear any pending WDT interrupt.
5. Enable the WDT interrupt to CPU by setting the WDT_MATCH bit in SRSS_INTR_MASK.
6. Enable SRSS interrupt to the CPU by configuring the appropriate ISER register (See the [Interrupts chapter on page 54](#) for details).
7. In the ISR, unlock the WDT; clear the WDT interrupt and add the desired match value to the existing match value. By doing so, another interrupt is generated when the counter reaches the new match value (period).

Note that interrupt servicing and watchdog reset cannot be used simultaneously using the free-running WDT.


23.4 Multi-Counter WDTs

23.4.1 Overview

Figure 23-3 shows the functional overview of a single multi-counter WDT block. The PSoC 6 MCU includes two MCWDT blocks. Each MCWDT block includes two 16-bit counters (MCWDTx_WDT0 and MCWDTx_WDT1) and one 32-bit counter (MCWDTx_WDT2). These counters can be configured to work independently or in cascade (up to 64-bit). The 16-bit counters have the ability to generate an interrupt and reset the device. The 32-bit counter can only generate an interrupt. All the counters are clocked by LFCLK.

Note: Because the PSoC 6 MCU includes two CPUs (Cortex-M0+ and Cortex-M4), associate one MCWDT block to only one CPU during runtime. Although both the MCWDT blocks are available to both the CPUs, a single MCWDT is not intended to be used by multiple CPUs simultaneously.

Figure 23-3. Multi-Counter WDT Functional Diagram



23.4.1.1 WDT0 and WDT1 Counters Operation

MCWDTx_WDT0 and MCWDTx_WDT1 are 16-bit up counters, which can be configured to be a 16-bit free-running counter or a counter with any 16-bit period. These counters can be used to generate an interrupt or reset.

The WDT_CTR0 bits [15:0] and WDT_CTR1 bits [16:31] of the MCWDTx_CNTLOW register holds the current counter values of WDT0 and WDT1 respectively. The WDT_MATCH0 bits [15:0] and WDT_MATCH1 bits [16:31] of the MCWDTx_MATCH register store the match value for WDT0 and WDT1 respectively. The WDT_MODEx bits of the MCWDTx_CONFIG register configures the action the watchdog counter takes on a match event ($\text{WDT_MATCHx} == \text{WDT_CTRx}$). The WDT0/WDT1 counters perform the following actions:

- Assert interrupt (WDT_INTx) on match
- Assert a device reset on match
- Assert an interrupt on match and a device reset on the third unhandled interrupt

In addition to generating reset and interrupt, the match event can be configured to clear the counters. This is done by setting the WDT_CLEARx bit of the MCWDTx_CONFIG register. WDT0/WDT1 counter operation is shown in Figure 23-4.

Figure 23-4. WDT0/WDT1 Operation




Table 23-2. WDT0 and WDT1 Configuration Options

Register [Bit_Pos]	Bit_Name	Description
MCWDTx_CONFIG[1:0] MCWDTx_CONFIG[9:8]	WDT_MODE0 WDT_MODE1	WDT action on a match event (WDT_CTRx == WDT_MATCHx) 0: Do nothing 1: Assert interrupt (WDT_INTx) 2: Assert device reset 3: Assert interrupt on match and a device reset on the third unhandled interrupt
MCWDTx_CONFIG[2] MCWDTx_CONFIG[10]	WDT_CLEAR0 WDT_CLEAR1	Clear the WDTx counter on match. In other words, (WDT_MATCHx + 1) acts similar to a period for the WDTx counter. 0: Free-running counter 1: Clear WDT_CTRx bits on match
MCWDTx_CNTLOW[15:0] MCWDTx_CNTLOW[16:31]	WDT_CTR0 WDT_CTR1	Current watchdog counter value. Bits[15:0] contain the current value of WDT0 counter and bits[31:16] contain the current value of WDT1 counter.
MCWDTx_MATCH[15:0] MCWDTx_MATCH[16:31]	WDT_MATCH0 WDT_MATCH1	Watchdog match value Changing WDT_MATCH requires 1.5 LFCLK cycles to come into effect. After changing WDT_MATCH, do not enter the Deep Sleep mode for at least one LFCLK cycle to ensure the WDT updates to the new setting.

23.4.1.2 WDT2 Counter Operation

The MCWDTx_WDT2 is a 32-bit free-running counter, which can be configured to generate an interrupt. The MCWDTx_CNTHIGH register holds the current value of the WDT2 counter. WDT2 does not support a match feature. However, it can be configured to generate an interrupt when one of the counter bits toggle. The WDT_BITS2 bits [28:24] of the MCWDTx_CONFIG register selects the bit on which the WDT2 interrupt is asserted. WDT_MODE2 bit [16] of the MCWDTx_CONFIG register decides whether to assert an interrupt on bit toggle or not. [Figure 23-5](#) shows the WDT2 counter operation.

Figure 23-5. WDT2 Operation




Table 23-3. WDT2 Configuration Options

Register [Bit_Pos]	Bit_Name	Description
MCWDTx_CONFIG[16]	WDT_MODE2	WDT2 mode 0: Free-running counter 1: Free-running counter with interrupt (WDT_INTx) request generation when specified, but in MCWDTx_CNTHIGH register toggles
MCWDTx_CONFIG[28:24]	WDT_BITS2	Bit to monitor for WDT2 interrupt assertion 0: Asserts when bit [0] of MCWDTx_CNTHIGH register toggles (interrupt every tick) 31: Asserts when bit [0] of MCWDTx_CNTHIGH register toggles (interrupt every 2 ³¹ ticks)
MCWDTx_CNTHIGH[31:0]	WDT_CTR2	Current counter value of WDT2

23.4.2 Enabling and Disabling WDT

The WDT counters are enabled by setting the WDT_ENABLEx bit in the MCWDTx_CTL register and are disabled by clearing it. Enabling or disabling a WDT requires 1.5 LFCLK cycles to come into effect. Therefore, the WDT_ENABLEx bit value must not be changed more than once in that period and the WDT_ENABLEDx bit of the MCWDTx_CTL register can be used to monitor enabled/disabled state of the counter.

The WDT_RESETx bit of the MCWDTx_CTL register clears the corresponding WDTx counter when set in firmware. The hardware clears the bit after the WDTx counter resets. This option is useful when the WDT0 or WDT1 is configured to generate a device reset on a match event. In such cases, the device resets when the counter reaches the match value. Thus, setting the WDT_RESET0 or WDT_RESET1 bit resets the WDT0 or WDT1 counter respectively, preventing device reset.

After WDT is enabled, do not write to the WDT configuration (MCWDTx_CONFIG) and control (MCWDTx_CTL) registers. Accidental corruption of WDT registers can be prevented by setting the WDT_LOCK bits [31:30] of the MCWDTx_CTL register. If the application requires updating the match value (WDT_MATCH) when the WDT is running, the WDT_LOCK bits must be cleared. The WDT_LOCK bits require two different writes to clear both the bits. Writing a '1' to the bits clears bit 0. Writing a '2' clears bit 1. Writing a '3' sets both the bits and writing '0' does not have any effect. Note that the WDT_LOCK bits protects only MCWDTx_CTL (except the WDT_LOCK bits), MCWDTx_CONFIG, and MCWDTx_MATCH registers. The LFCLK select registers are protected by the free-running WDT lock bits.

Table 23-4. Watchdog Configuration Options

Register [Bit_Pos]	Bit_Name	Description
MCWDTx_CTL[0]	WDT_ENABLE0	Enable WDT counter x
MCWDTx_CTL[8]	WDT_ENABLE1	0: Counter is disabled
MCWDTx_CTL[16]	WDT_ENABLE2	1: Counter is enabled
MCWDTx_CTL[1]	WDT_ENABLED0	Indicates the actual enabled/disabled state of the counter. This bit should be monitored after changing the WDT_ENABLEx bit, to receive an acknowledgment of the change
MCWDTx_CTL[9]	WDT_ENABLED1	
MCWDTx_CTL[17]	WDT_ENABLED2	
MCWDTx_CTL[3]	WDT_RESET0	Reset WDT counter x to 0. Hardware clears the bit when the reset is complete
MCWDTx_CTL[11]	WDT_RESET1	0: Software - No action, Hardware - Counter is reset after software sets this bit
MCWDTx_CTL[19]	WDT_RESET2	1: Software - Resets the counter, Hardware - Counter is not reset after software sets this bit
MCWDTx_CTL[31:30]	WDT_LOCK	Locks or unlocks write access to the MCWDTx_CTL (except the WDT_LOCK bits), MCWDTx_CONFIG, and MCWDTx_MATCH registers. When the bits are set, the lock is enabled. 0: No effect 1: Clears bit 0 2: Clears bit 1 3: Sets both bit 0 and 1 (lock enabled)

Note: When the watchdog counters are configured to generate an interrupt every LFCLK cycle, make sure you read the MCWDTx_INTR register after clearing the watchdog interrupt (setting the WDT_INTx bit in the MCWDTx_INTR register). Failure to do this may result in missing the next interrupt. Hence, the interrupt cycle becomes LFCLK/2.

23.4.3 Watchdog Cascade Options

The cascade configuration shown in [Figure 23-3](#) provides an option to increase the WDT counter resolution. The WDT_CASCADE0_1 bit [3] of the MCWDTx_CONFIG register cascades WDT0 and WDT1 and the WDT_CASCADE1_2 bit [11] of the MCWDTx_CONFIG register cascades WDT1 and WDT2. Note that cascading two 16-bit counters does not provide a 32-bit counter; instead, you get a 16-bit period counter with a 16-bit prescaler. For example, when cascading WDT0 and WDT1, WDT0 acts as a prescaler for WDT1 and the prescaler value is defined by the WDT_MATCH0 bits [15:0] in the

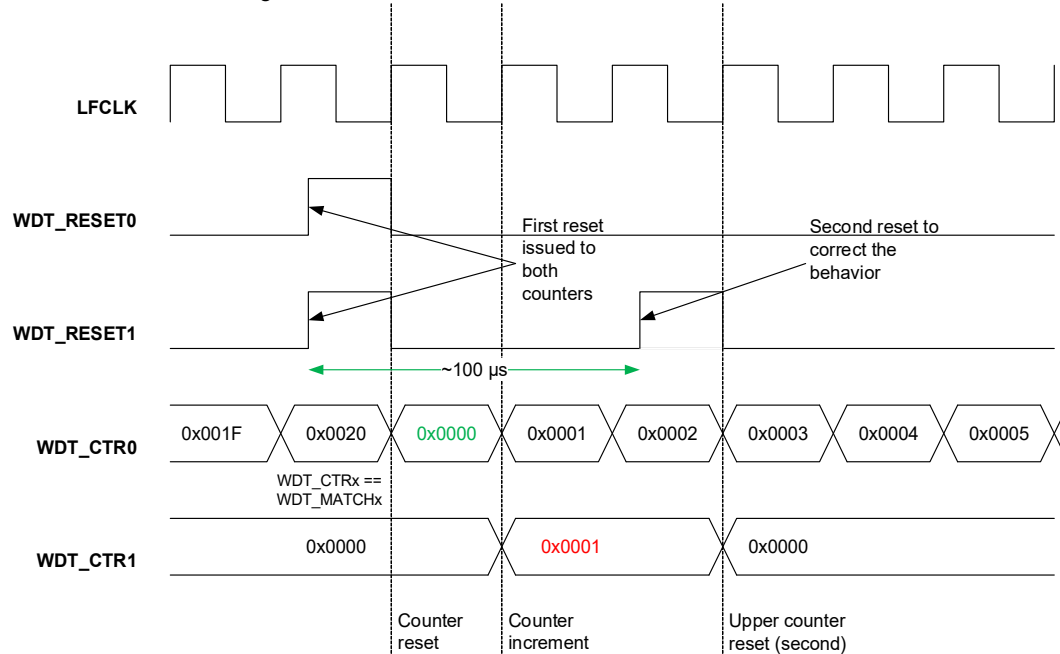
MCWDTx_MATCH register. The WDT1 has a period defined by WDT_MATCH1 bits [31:16] in the MCWDTx_MATCH register. The same logic applies to WDT1 and WDT2 cascading.

Table 23-5. Watchdog Cascade Options

Register [Bit_Pos]	Bit_Name	Description
MCWDTx_CONFIG[3]	WDT_CASCADE0_1	Cascade WDT0 and WDT1 0: WDT0 and WDT1 are independent counters 1: WDT1 increments two cycles after WDT_CTR0 == WDT_MATCH0
MCWDTx_CONFIG[11]	WDT_CASCADE1_2	Cascade WDT1 and WDT2 0: WDT1 and WDT2 are independent counters 1: WDT2 increments two cycles after WDT_CTR1 == WDT_MATCH1

When using cascade (WDT_CASCADE0_1 or WDT_CASCADE1_2 set), resetting the counters when the prescaler or lower counter is at its match value with the counter configured to clear on match, results in the upper counter incrementing to 1 instead of remaining at 0. This behavior can be corrected by issuing a second reset to the upper counter after approximately 100 μ s from the first reset. Note that the second reset is required only when the first reset is issued while the prescaler counter value is at its match value. Figure 23-6 illustrates the behavior when WDT0 and WDT1 are cascaded along with the second reset timing.

Figure 23-6. MCWDT Reset Behavior in Cascaded Mode



Other settings:

WDT_CASCADE0_1 = 1
 WDT_CLEAR0 = 1
 WDT_MATCH0 = 0x0020

In addition, the counters exhibit non-monotonicity in the following cascaded conditions:

- If WDT_CASCADE0_1 is set, then WDT_CTR1 does not increment the cycle after WDT_CTR0 = WDT_MATCH0.
- If WDT_CASCADE1_2 is set, then WDT_CTR2 does not increment the cycle after WDT_CTR1 = WDT_MATCH1.
- If both WDT_CASCADE0_1 and WDT_CASCADE1_2 are set, then WDT_CTR2 does not increment the cycle after WDT_CTR1 = WDT_MATCH1 and WDT_CTR1

does not increment the cycle after WDT_CTR0 = WDT_MATCH0.

When cascading is enabled, always read the WDT_CTR1 or WDT_CTR2 counter value only when the prescaler counter (WDT_CTR0 or WDT_CTR1) value is not 0. This makes sure the upper counter is incremented after a match event in the prescaler counter.

23.4.4 Watchdog Reset

WDT0 and WDT1 counters can be configured to generate a device reset similar to the free-running WDT reset. Follow these steps to use the WDT0 or WDT1 counter of a MCWDTx block to generate a system reset:

1. Configure the WDTx to generate a reset using the WDT_MODEx bits in MCWDTx_CONFIG. Configure the WDT_MODE0 or WDT_MODE1 bits in MCWDTx_CONFIG to '2' (reset on match) or '3' (interrupt on match and reset on the third unhandled interrupt).
2. Optionally, set the WDT_CLEAR0 or WDT_CLEAR1 bit in the MCWDTx_CONFIG register for WDT0 or WDT1 to reset the corresponding watchdog counter to '0' on a match event. Otherwise, the counters are free running. See [Table 23-2 on page 254](#) for details.
3. Calculate the watchdog reset period such that firmware is able to reset the watchdog at least once during the period, even along the longest firmware delay path. For WDT_MODEx == 2, match value is same as the watchdog period. For WDT_MODEx == 3, match value is one-third of the watchdog period. Write the calculated match value to the WDT_MATCH register for WDT0 or WDT1. Optionally, enable cascading to increase the interval. **Note:** The legal value for the WDT_MATCH field is 1 to 65535.
4. For WDT_MODEx == 2, set the WDT_RESETx bit in the MCWDTx_CONFIG register to reset the WDTx counter to 0. For WDT_MODEx == 3, set the WDT_INTx bit in MCWDTx_INTR to clear any pending interrupts.
5. Enable WDTx by setting the WDT_ENABLEx bit in the MCWDTx_CTL register. Wait until the WDT_ENABLEDx bit is set.
6. Lock the MCWDTx configuration by setting the WDT_LOCK bits of the MCWDTx_CTL register.
7. In the firmware, feed (reset) the watchdog as explained in step 4.

Do not reset watchdog in the WDT ISR. It is also not recommended to use the same watchdog counter to generate a system reset and interrupt. For example, if WDT0 is used to generate system reset against crashes, then WDT1 or WDT2 should be used for periodic interrupt generations.

23.4.5 Watchdog Interrupt

When configured to generate an interrupt, the WDT_INTx bits of the MCWDTx_INTR register provide the status of any pending watchdog interrupts. The firmware must clear the interrupt by setting the WDT_INTx. The WDT_INTx bits of the MCWDTx_INTR_MASK register mask the corresponding WDTx interrupt of the MCWDTx block to the CPU.

Follow these steps to use WDT as a periodic interrupt generator:

1. Write the desired match value to the WDT_MATCH register for WDT0/WDT1 or the WDT_BITS2 value to the MCWDTx_CONFIG register for WDT2. **Note:** The legal value for the WDT_MATCH field is 1 to 65535.
2. Configure the WDTx to generate an interrupt using the WDT_MODEx bits in MCWDTx_CONFIG. Configure the WDT_MODE0 or WDT_MODE1 bits in MCWDTx_CONFIG for WDT0 or WDT1 to '1' (interrupt on match) or '3' (interrupt on match and reset on third unhandled interrupt). For WDT2, set the WDT_MODE2 bit in the MCWDTx_CONFIG register.
3. Set the WDT_INT bit in MCWDTx_INTR to clear any pending interrupt.
4. Set the WDT_CLEAR0 or WDT_CLEAR1 bit in the MCWDTx_CONFIG register for WDT0 or WDT1 to reset the corresponding watchdog counter to '0' on a match event.
5. Mask the WDTx interrupt to the CPU by setting the WDT_INTx bit in the MCWDTx_INTR_MASK register.
6. Enable WDTx by setting the WDT_ENABLEx bit in the MCWDTx_CTL register. Wait until the WDT_ENABLEDx bit is set.
7. Enable MCWDTx interrupt to the CPU by configuring the appropriate ISER register. Refer to the [Interrupts chapter on page 54](#).
8. In the ISR, clear the WDTx interrupt by setting the WDT_INTx bit in the MCWDTx_INTR register.

Note that interrupts from all three WDTx counters of the MCWDT block are mapped as a single interrupt to the CPU. In the interrupt service routine, the WDT_INTx bits of the MCWDTx_INTR register can be read to identify the interrupt source. However, each MCWDT block has its own interrupt to the CPU. For details on interrupts, see the [Interrupts chapter on page 54](#).

The MCWDT block can send interrupt requests to the CPU in Active power mode and to the WIC in Sleep and Deep Sleep power modes. It works similar to the free-running WDT.

23.5 Reset Cause Detection

The RESET_WDT bit [0] in the RES_CAUSE register indicates the reset generated by the free-running WDT. The RESET_MCWDTx bit in the RES_CAUSE register indicates the reset generated by the MCWDTx block. These bits remain set until cleared or until a power-on reset (POR), brownout reset (BOD), or external reset (XRES) occurs. All other resets leave this bit unaltered.

For more details, see the [Reset System chapter on page 223](#).

23.6 Register List

Table 23-6. WDT Registers

Name	Description
WDT_CTL	Watchdog Counter Control Register
WDT_CNT	Watchdog Counter Count Register
WDT_MATCH	Watchdog Counter Match Register
MCWDTx_MCWDT_CNTLOW	Multi-counter WDT Sub-counters 0/1
MCWDTx_MCWDT_CNTHIGH	Multi-counter WDT Sub-counter 2
MCWDTx_MCWDT_MATCH	Multi-counter WDT Counter Match Register
MCWDTx_MCWDT_CONFIG	Multi-counter WDT Counter Configuration
MCWDTx_MCWDT_CTL	Multi-counter WDT Counter Control
MCWDTx_MCWDT_INTR	Multi-counter WDT Counter Interrupt Register
MCWDTx_MCWDT_INTR_SET	Multi-counter WDT Counter Interrupt Set Register
MCWDTx_MCWDT_INTR_MASK	Multi-counter WDT Counter Interrupt Mask Register
MCWDTx_MCWDT_INTR_MASKED	Multi-counter WDT Counter Interrupt Masked Register
CLK_SELECT	Clock Selection Register
CLK_ILO_CONFIG	ILO Configuration
SRSS_INTR	SRSS Interrupt Register
SRSS_INTR_SET	SRSS Interrupt Set Register
SRSS_INTR_MASK	SRSS Interrupt Mask Register
SRSS_INTR_MASKED	SRSS Interrupt Masked Register
RES_CAUSE	Reset Cause Observation Register

24. Trigger Multiplexer Block



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - Trigger Multiplexer
- Application notes
- Code examples


Every peripheral in the PSoC 6 MCU is interconnected using trigger signals. Trigger signals are means by which peripherals denote an occurrence of an event or a state. These triggers are used as means to affect or initiate some action in other peripherals. The trigger multiplexer block helps to route triggers from a source peripheral block to a destination.

24.1 Features

- Ability to connect any trigger signal from one peripheral to another
- Two-layer architecture with 15 trigger groups
- Supports a software trigger, which can trigger any signal in the block
- Supports multiplexing of triggers between peripherals
- One-to-one trigger paths for dedicated triggers that are more commonly routed
- Ability to configure a trigger multiplexer with trigger manipulation features in hardware such as inversion and edge/level detection
- Ability to block triggers in debug mode

24.2 Architecture

The trigger signals in the PSoC 6 MCU are digital signals generated by peripheral blocks to denote a state such as FIFO level, or an event such as the completion of an action. These trigger signals typically serve as initiator of other actions in other peripheral blocks. An example is an ADC peripheral block sampling three channels. After the conversion is complete, a trigger signal will be generated, which in turn triggers a DMA channel that transfers the ADC data to a memory buffer. This example is shown in [Figure 24-1](#).




A PSoC 6 MCU has multiple peripheral blocks; each of these blocks can be connected to other blocks through trigger signals, based on the system implementation. To support this, the PSoC 6 MCU has hardware, which is a series of multiplexers used to route the trigger signals from potential sources to destinations. This hardware is called the trigger multiplexer block. The

trigger multiplexer can connect to any trigger signal emanating out of any peripheral block in the PSoC 6 MCU and route it to any other peripheral to initiate or affect an operation at the destination peripheral block.

24.2.1 Trigger Multiplexer Group

The trigger multiplexer block is implemented using several trigger multiplexers. A trigger multiplexer selects a signal from a set of trigger output signals from different peripheral blocks to route it to a specific trigger input of another peripheral block. The multiplexers are grouped into trigger groups. All the trigger multiplexers in a trigger group have similar input options and are designed to feed similar destination signals. Hence the trigger group can be considered as a block that multiplexes multiple inputs to multiple outputs. This concept is illustrated in Figure 24-2. Both the trigger multiplexers and the one-to-one triggers offer trigger manipulation, which includes inversion or specifies the output signal as edge or level triggered. For edge triggering, the trigger is synchronized to the consumer block's (the peripheral block that is receiving the trigger) clock and a two-cycle pulse is generated on this clock.

Figure 24-2. Trigger Multiplexer Groups



24.2.2 One-to-one Trigger

In addition to trigger multiplexers, there are dedicated trigger paths called one-to-one triggers. These trigger paths/routes are between fixed source and destination peripherals and cannot be multiplexed. These paths can be enabled or disabled. A group of registers in the format `PERI_TR_1TO1_GR[X]_TR_CTL[Y]`, with X being the trigger group and Y being the output trigger number from the one-to-one trigger, can be used to enable or disable the triggers path.

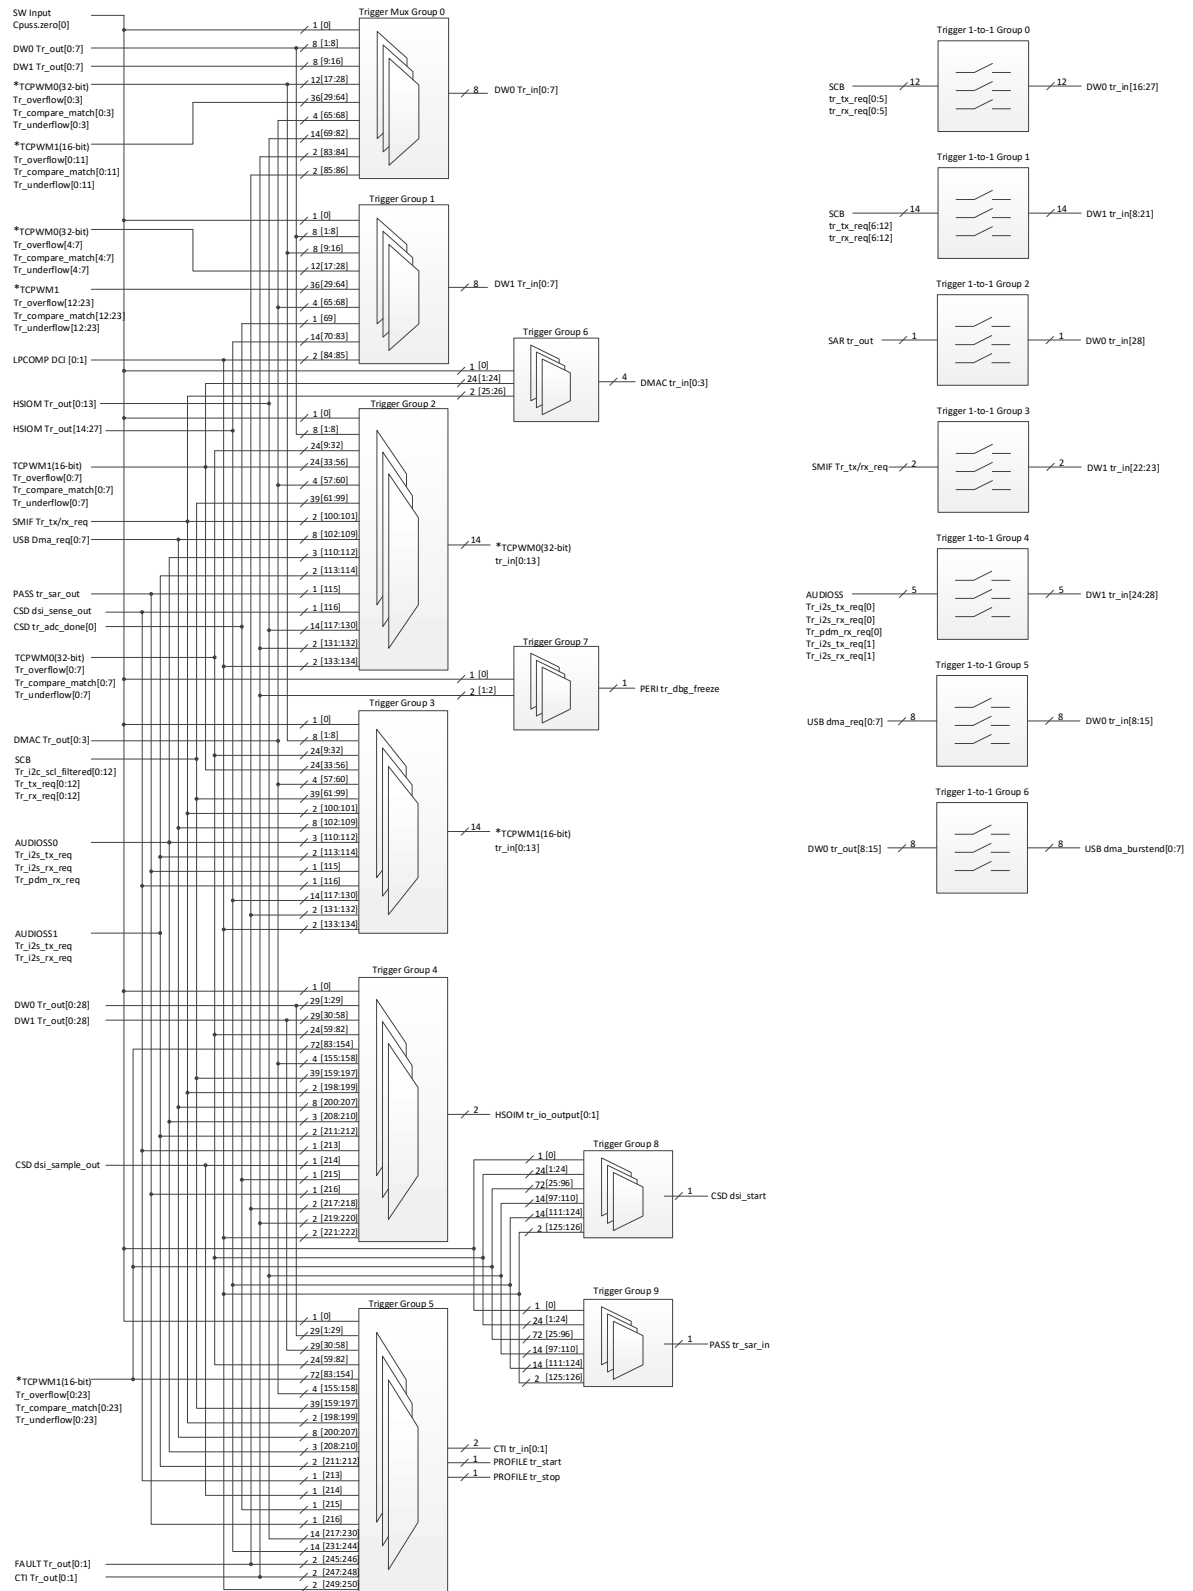
24.2.3 Trigger Multiplexer Block

The trigger multiplexer block has two types of trigger routings: trigger multiplexers that are capable of selecting from multiple trigger sources to a specific trigger destination and one-to-one triggers that are predetermined fixed trigger paths that can be enabled or disabled. See Table 24-2 for more information on each trigger group and its description.

Note: The triggers output into different peripherals, which may have more routing than is shown on the trigger routing diagram. For more information on this routing, go to the trigger destination peripheral block.


Figure 24-3 shows the trigger multiplexer architecture.

Figure 24-3. Trigger Multiplexer Block Architecture



Note: * indicates that Figure 24-4 shows details about the TCPWM and SCB inputs.

Figure 24-4. TCPWM and Multiplexer SCB Input Layout



Note: The format in these diagrams apply to all TCPWM and SCB connections in the trigger groups that have these connections. The trigger input line number will change between trigger groups and will not always start at 17 and 61 as shown in the TCPWM and SCB diagrams.

24.2.4 Software Triggers

All input and output signals to a trigger multiplexer can be triggered from software. This is accomplished by writing into the PERI_TR_CMD register. This register allows you to trigger the corresponding signal for a number of peripheral clock cycles.

The PERI_TR_CMD[GROUP_SEL] bitfield selects the trigger group of the signal being activated. The PERI_TR_CMD[OUT_SEL] bitfield determines whether the trigger signal is in output or input of the multiplexer. PERI_TR_CMD[TR_SEL] selects the specific line in the trigger group.

The PERI_TR_CMD[COUNT] bitfield sets up the number of peripheral clocks the trigger will be activated.

The PERI_TR_CMD[ACTIVATE] bitfield is set to '1' to activate the trigger line specified. Hardware resets this bit after the trigger is deactivated after the number of cycles set by the PERI_TR_CMD[COUNT].

Trigger signals that are software triggered may have some implications. If the output of a distribution multiplexer is being triggered, only the peripheral block trigger that the signal feeds is triggered. If an input of a reduction multiplexer is being triggered, then all the peripheral routes that connect to this signal through the trigger multiplexer block will be triggered. You should be aware of the routing that is being configured before deciding on the trigger signal. For distribution multiplexers, each trigger group will have its own clock_source, which will be used for trigger_manipulation and for generating software_trigger.

24.3 Register List


Table 24-1. Register List

Register Name	Description
PERI_TR_CMD	Trigger command register. The control enables software activation of a specific input trigger or output trigger of the trigger multiplexer structure.
PERI_TR_GRP[X]_TR_OUT_CTL[Y]	This register specifies the input trigger for a specific output trigger in a trigger group. It can also invert the signal and specify if the output signal should be treated as a level-sensitive or edge-sensitive trigger; only specific multiplexers can do this (see table). Every trigger multiplexer has a group of registers, the number of registers being equal to the output bus size from the multiplexer. In the register format, X is the trigger group and Y is the output trigger line number from the multiplexer.
PERI_TR_1TO1_GRP[X]_TR_OUT_CTL[Y]	This register specifies the input trigger for a specific output trigger in a one-to-one trigger group. It also includes trigger manipulation that can invert the signal and specify if the output signal should be treated as a level-sensitive or edge-sensitive trigger. Every one-to-one trigger multiplexer has a group of registers, the number of registers being equal to the output bus size from the multiplexer. The registers are formatted with X as the trigger group and Y as the output trigger line number from the multiplexer.

Table 24-2. Trigger Group

Trigger Group Number	Description
0	Routes all input trigger signals to DMA0
1	Routes all input trigger signals to DMA1
2	Routes all input trigger signals to TCPWM0(32-bit)
3	Routes all input trigger signals to TCPWM1(16-bit)
4	Routes all input trigger signals to HSIOM
5	Routes all input trigger signals to CPUSS CTI and the Profiler
6	Routes all input trigger signals to DMAC
7	Routes all input trigger signals to trigger freeze operation
8	Routes all input trigger signals to CSD ADC start
9	Routes all input trigger signals to PASS ADC start
Trigger 1-to-1 Group	
0	Connects SCB triggers to DMA0
1	Connects SCB triggers to DMA1
2	Connects SAR triggers to DMA0
3	Connects SMIF triggers to DMA1
4	Connects AUDIOSS triggers to DMA1
5	Connects USB triggers to DMA0
6	Connects DMA0 to USB

Figure 24-5. PSoC 6 MCU Trigger Multiplexer Block Architecture



25. Profiler



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - Profiler
- Application notes
- Code examples

The PSoC 6 MCU Profiler provides counters that can measure duration or number of events of a particular peripheral. Functions such as DMA transfers or buffered serial communications can happen asynchronously, and are not directly tied to the CPU or code execution. The profiler provides additional insight into the device so you can identify an asynchronous activity that could not be monitored previously.

The profiler manages a set of counters. You configure an available counter to monitor a particular source. Depending on the nature of the source, you count either duration (reference clock cycles) or the number of events.

The ability to monitor specific peripherals enables you to understand and optimize asynchronous hardware, including:

- Identify the activity of a particular peripheral
- Identify asynchronous activity

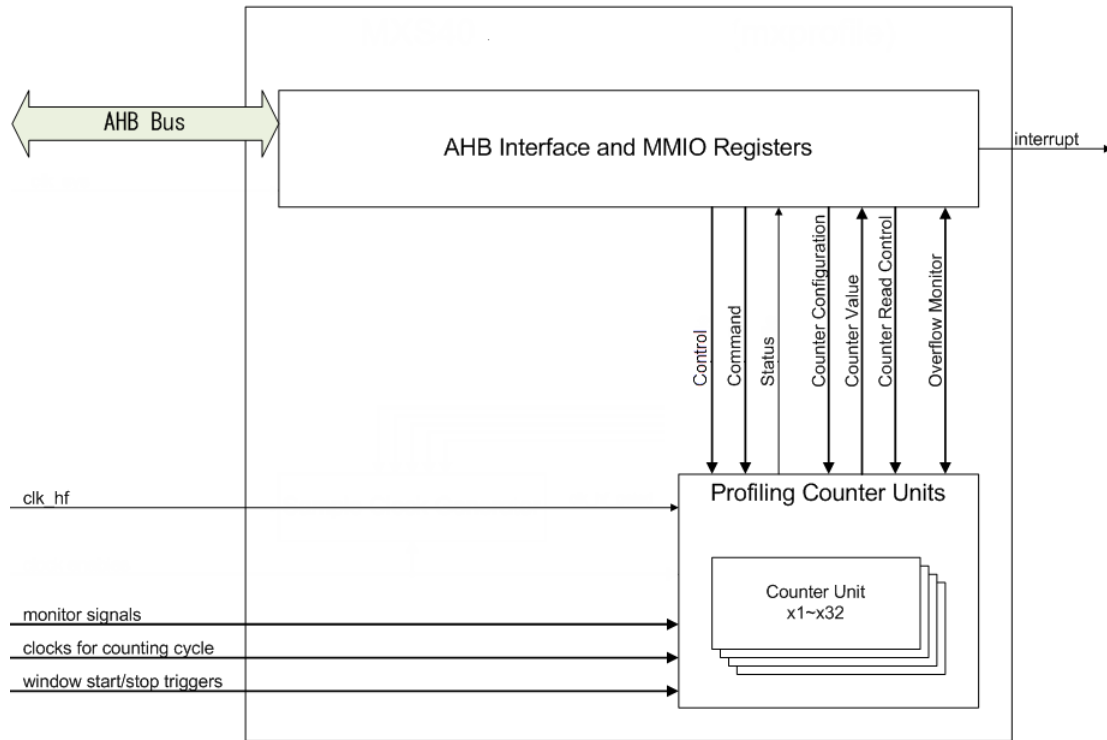
25.1 Features

The profiler has these features and capabilities:

- A variety of sources you can monitor (see [Table 25-1](#))
- Support for eight counters, to monitor up to eight sources simultaneously
- For each counter you specify:
 - ☐ what source to monitor (can be changed dynamically if required)
 - ☐ what to measure (duration or events)
 - ☐ what reference clock to use for the count (only affects duration)
- Provides ability to easily detect active peripherals that may be difficult to measure by external monitoring
- Provides an absolute count of events or reference clock cycles for each monitored source

25.2 Architecture

Figure 25-1. Profiler Block Diagram



The profiler supports up to 32 counters. The actual number of counters is hardware dependent. This device supports up to eight counters.

25.2.1 Profiler Design

With the profiler you can monitor:

- Peripherals that operate asynchronously to software
- CPU activity

This includes flash access, DMA, and other processes. See [Available Monitoring Sources on page 268](#).

You can measure total energy consumption directly using external hardware. The profiler can monitor internal hardware with no external probe or measuring device.

The profiler counts either the amount of time that a source is active (duration) or the number of events that occur. The profiler does not report the actual amount of energy consumed.

Your application may use a system resource or peripheral that is not among the profiler sources. A complete understanding of your application's energy consumption requires that you profile those parts of your code that use those other systems. For example, you could determine the time the peripheral is in use, and then derive energy use based on power consumption data from the data sheet, or external power monitoring.

All counters that are in use start and stop at the same time. The time between start and stop is called the profiling window, or the profiling session. During the profiling session, each counter increments at its own rate based on the monitored signal and (if measuring duration) the selected reference clock. See [Start and Stop Profiling on page 270](#).

Before starting a profiling session, configure the counters you want to use. See [Configure and Enable a Counter on page 270](#).

You can get interim results during a profiling session, or final results after you stop profiling. See [Get the Results on page 271](#).

25.2.2 Available Monitoring Sources

You specify which source a counter monitors. The available sources vary per device series. You can find constants in a header file named `<series>_config.h` (part of the [Peripheral Driver Library](#)). For example, [Table 25-1](#) lists the sources in `psoc63_config.h`.

Table 25-1. Available Signals to Monitor

Value	Symbol	Count	Description
0	PROFILE_ONE	Duration	Constant One
1	CPUSS_MONITOR_CM0	Events	CM0+ active cycle count
2	CPUSS_MONITOR_CM4	Events	CM4 active cycle count
3	CPUSS_MONITOR_FLASH	Events	Flash read count
4	CPUSS_MONITOR_DW0_AHB	Events	DW0 AHB transfer count (DMA transfer)
5	CPUSS_MONITOR_DW1_AHB	Events	DW1 AHB transfer count (DMA transfer)
6	CPUSS_MONITOR_CRYPT0	Events	Crypto memory access count
7	USB_MONITOR_AHB	Events	USB AHB transfer count
8	SCB0_MONITOR_AHB	Events	SCB 0 AHB transfer count
9	SCB1_MONITOR_AHB	Events	SCB 1 AHB transfer count
10	SCB2_MONITOR_AHB	Events	SCB 2 AHB transfer count
11	SCB3_MONITOR_AHB	Events	SCB 3 AHB transfer count
12	SCB4_MONITOR_AHB	Events	SCB 4 AHB transfer count
13	SCB5_MONITOR_AHB	Events	SCB 5 AHB transfer count
14	SCB6_MONITOR_AHB	Events	SCB 6 AHB transfer count
15	SCB7_MONITOR_AHB	Events	SCB 7 AHB transfer count
16	SCB8_MONITOR_AHB	Events	SCB 8 AHB transfer count
17–20	Reserved		
21	SMIF_MONITOR_SMIF_SPI_SELECT0	Duration	SPI select to memory 0 active
22	SMIF_MONITOR_SMIF_SPI_SELECT1	Duration	SPI select to memory 1 active
23	SMIF_MONITOR_SMIF_SPI_SELECT2	Duration	SPI select to memory 2 active
24	SMIF_MONITOR_SMIF_SPI_SELECT3	Duration	SPI select to memory 3 active
25	SMIF_MONITOR_SMIF_SPI_SELECT_ANY	Duration	SPI select to any of the memories 0-3 active

25.2.3 Reference Clocks

Each monitored source has its own clock reference, called the sample clock. [Table 25-2](#) lists the six choices for the sample clock. You may use one of two CLK_PROFILE sources, or one of four CLK_REF sources.

The counter units of the profiler are clocked using either CLK_HF or CLK_PERI as CLK_PROFILE. You configure CLK_HF in system resources and CLK_PERI in the CPU subsystem. See [Clocking System chapter on page 207](#) for details on clocks. In addition, there are four available reference clocks: CLK_TIMER, CLK_IMO, CLK_ECO, and CLK_LF. Any of the six clock sources can be used as the sample clock.

If you measure events, the sample clock is selected automatically to be either CLK_HF or CLK_PERI depending on the monitored source. The count represents the number of events that occurred during the profiling session. The profiler counts the edges of the monitored signal, using either CLK_HF or CLK_PERI to detect edges.

If you measure duration, you specify the sample clock from one of the six choices listed in [Table 25-2](#). When you measure duration, while the monitored source is active, the counter increments at each cycle of the sample clock. The profiler synchronizes reference clocks with CLK_PROFILE. As a result the four available reference clocks cannot exceed CLK_PROFILE/2.

To measure duration accurately, the sample clock should have a stable frequency throughout the profiling session. If your application remains in active power states, you can use CLK_HF as the sample clock. CLK_HF gives you the greatest resolution in your results.

The profiler can be enabled in CPU Active and Sleep modes. However, clock frequencies can change based on the power state of the application. The source clock frequency should be stable throughout the profiling session to ensure reliable data.

High-frequency clocks are not available in System Deep Sleep and Hibernate power modes, so the profiler is disabled. The configuration registers maintain state through Deep Sleep. Profiler data is lost (registers are not maintained) in Hibernate mode. See the [Device Power Modes chapter on page 190](#).

Table 25-2. Available Sample Clock Sources for the Profiler

Value	Clock Source	Symbol	Clock Type
0	Timer	CLK_TIMER	CLK_REF
1	Internal main oscillator	CLK_IMO	CLK_REF
2	External crystal oscillator	CLK_ECO	CLK_REF
3	Low frequency clock	CLK_LF	CLK_REF
4	High-frequency clock	CLK_HF	CLK_PROFILE
5	Peripheral clock	CLK_PERI	CLK_PROFILE

25.3 Using the Profiler

This section describes the steps required to use the profiler effectively. To use the profiler, you add instrumentation code to your application, to set or read values in particular registers. See the [registers TRM](#) for details.

At the highest level you perform these tasks:

- Enable the profiling block
- Configure and enable counters
- Start and stop profiling
- Handle counter overflow
- Get the results
- Exit gracefully

You can monitor as many sources as there are available counters. The actual number of available counters is hardware dependent. The value is defined in the symbol `PROFILE_PRFL_CNT_NR` in the `<series>_config.h` file. For example, the `psoc63_config.h` file defines this value as '8'.

Instead of manipulating registers directly, you can use the Peripheral Driver Library (PDL). The PDL provides a software API to manage the profiler and an interrupt to handle counter overflow. The software driver maintains an array of data for each counter, including the source you want to monitor, whether you are monitoring duration or events, the reference clock, and an overflow count for each counter. It provides function calls to configure and enable a counter, start and stop a profiling session, and calculate the results for each counter. The PDL API handles all register and bit access.

Refer to the PDL API Reference for more information. The PDL installer puts the documentation here:

`<PDL_directory>\doc\pdl_api_reference_manual.html`

25.3.1 Enable or Disable the Profiler

Before performing any operations, enable the profiling block. See [Table 25-3](#). This does not enable individual counters or start a profiling session.

Table 25-3. Enabling the Profiler

Task	Register	Bitfield	Value
Enable the profiler	PROFILE_CTL	ENABLED	0 = disabled; 1 = enabled

25.3.2 Configure and Enable a Counter

Each counter has a PROFILE_CTL register that holds configuration information. For each counter you specify:

- The source you want to monitor
- What you want to measure (events or duration)
- The reference clock for the counter (affects only duration)

You also enable each counter individually. Only the enabled counters will start when START_TR is set to one. See [Table 25-4](#).

Table 25-4. Configuring and Enabling a Counter

Task	Register	Bitfield	Value
Specify source	PROFILE_CTL	MON_SEL	Table 25-1 on page 268
Specify events or duration	PROFILE_CTL	CNT_DURATION	0 = events; 1 = duration. See Table 25-1 on page 268
Specify the reference clock	PROFILE_CTL	REF_CLK_SEL	Table 25-2 on page 269
Enable the counter	PROFILE_CTL	ENABLED	0 = disabled; 1 = enabled

The count value is a 32-bit register. If that is not a sufficiently large number for your purposes, you must also enable the profiling interrupt for the counter. See [Handle Counter Overflow on page 270](#).

When gathering data, you may want to know the actual number of clock cycles that occurred during the profiling session. To get this number, configure a counter to use:

- PROFILE_ONE as the monitored source
- duration (not events)
- a reference clock

The count for that counter is the actual number of reference clock cycles that occurred during the profiling session.

25.3.3 Start and Stop Profiling

After configuring and enabling the counters you want to use, you can start and stop a profiling session. You may wish to ensure that all counters are set to zero before beginning the profiling session. The PROFILE_CMD register applies to all counters, meaning that all enabled counters will be started at the same time. See [Table 25-5](#).

Table 25-5. Starting or Stopping a Profiling Session

Task	Register	Bitfield	Value
Clear all counters	PROFILE_CMD	CLEAR_ALL_CNT	1 = reset all counters to zero
Start profiling	PROFILE_CMD	START_TR	1 = start profiling
Stop profiling	PROFILE_CMD	STOP_TR	1 = stop profiling

By design, all counters that are in use start and stop simultaneously. During the profiling session, each counter increments at its own rate based on the monitored signal and (if measuring duration) the reference clock.

25.3.4 Handle Counter Overflow

Each profiling counter is a 32-bit number. If this is not sufficient for your purposes, then you must handle counter overflow. To do so you must:

- Enable the profiling interrupt for the particular counter
- Provide an interrupt handler

For example, your interrupt handler could maintain an overflow counter for each profiling counter. See the [Interrupts chapter on page 54](#) for details about interrupts.

Counter overflow involves the interaction of three registers, INTR, INTR_MASK, and INTR_MASKED. Each of these registers has one bit per profiling counter.

When an overflow occurs for a particular counter, hardware sets the corresponding bit in the INTR register. You typically do not read this register.

To enable the profiling interrupt for a particular counter, set the corresponding bit in the INTR_MASK register.

When an interrupt occurs, your interrupt handler reads the bits in the INTR_MASKED register. This register reflects a bitwise AND between the INTR register (an overflow interrupt has occurred for a particular counter) and the INTR_MASK register (this counter has the interrupt enabled). When a bit is set in the INTR_MASKED register, the corresponding counter is enabled and has experienced an overflow.

You can artificially trigger an overflow interrupt to test your code. The INTR_SET register also has one bit per counter. For debug purposes, software can set the appropriate bit to activate a specific overflow interrupt. This enables debug of the interrupt without waiting for hardware to cause the interrupt.

25.3.5 Get the Results

For each of your enabled counters, read the counter value in the CNT register for that counter. This is your absolute count. If you are tracking counter overflow, then the absolute count is $0x1\ 0000\ 0000 * \text{overflow count} + \text{counter value}$.

The common use case is to get results after you stop profiling. However, you can get a snapshot of the results without stopping the profiler. In this case, however, the profiler is still counting and the results are changing as you gather the data.

After completing a profiling session, you may wish to repeat the same profile. Clear all counters to zero, and then start and stop profiling. See [Start and Stop Profiling on page 270](#). If you are handling counter overflow, set your overflow counters to zero as well.

You can also reconfigure any or all counters to gather different data. See [Configure and Enable a Counter on page 270](#). Reconfigure the profiler, and start another profiling session.

25.3.6 Exit Gracefully

When finished, you should disable the profiler. To do this make sure that you:

- Stop profiling (see [Start and Stop Profiling on page 270](#))
- Clear any profiling configuration (see [Configure and Enable a Counter on page 270](#))
- Disable the profiling interrupt if you have set it up (see [Handle Counter Overflow on page 270](#))
- Disable the profiler itself (see [Enable or Disable the Profiler on page 269](#))

Section D: Digital Subsystem




This section encompasses the following chapters:

- Secure Digital Host Controller (SDHC) chapter on page 273
- Serial Communications Block (SCB) chapter on page 284
- Serial Memory Interface (SMIF) chapter on page 341
- Timer, Counter, and PWM (TCPWM) chapter on page 358
- Inter-IC Sound Bus chapter on page 394
- PDM-PCM Converter chapter on page 406
- Universal Serial Bus (USB) Device Mode chapter on page 415
- LCD Direct Drive chapter on page 431

Top Level Architecture

Figure D-1. Digital System Block Diagram



26. Secure Digital Host Controller (SDHC)




This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SDHC
- Application notes
- Code examples

The secure digital host controller (SDHC) in the PSoC 6 MCU allows interfacing with embedded multimedia card (eMMC)-based memory devices, secure digital (SD) cards, and secure digital input output (SDIO) cards. The block supports all three interfaces – SD, SDIO, and eMMC. The block can also work with devices providing SDIO interface, such as Cypress' Wi-Fi products (for example, CYW4343W). [Figure 26-1](#) illustrates a typical application using the SDHC block.

Figure 26-1. Typical SDHC Application



26.1 Features

- Complies with eMMC 5.1, SD 6.0, and SDIO 4.10 standards
- Supports host controller interface (HCI) 4.2 shared by eMMC and SD
- SD interface supports 1-bit and 4-bit bus interfaces, and the following speed modes. The specified data rate is for a 4-bit bus.
 - 3.3-V signal voltage: Default speed (12.5 MB/s at 25 MHz) and high speed (25 MB/s at 50 MHz)
 - UHS-I modes using 1.8-V signal voltage: SDR12 (12.5 MB/s at 25 MHz), SDR25 (25 MB/s at 50 MHz), SDR40 (40 MB/s at 80 MHz), and DDR40 (40 MB/s at 40 MHz)
- eMMC interface supports 1-bit and 4-bit bus interfaces, and the following speed modes. The specified data rate is for a 4-bit bus.
 - Legacy (13 MB/s at 26 MHz), high-speed SDR (26 MB/s at 52 MHz), and high-speed DDR (52 MB/s at 52 MHz)
- Supports three DMA modes – SDMA, ADMA2, and ADMA3 – through a dedicated DMA engine
- Provides 1KB SRAM for buffering up to two 512-byte blocks
- Provides I/O interfaces for bus interface voltage selection (3.3 V/1.8 V) and for power enable/disable
- Provides I/O interfaces for functions such as card detection, mechanical write protection, eMMC card reset, and LED control

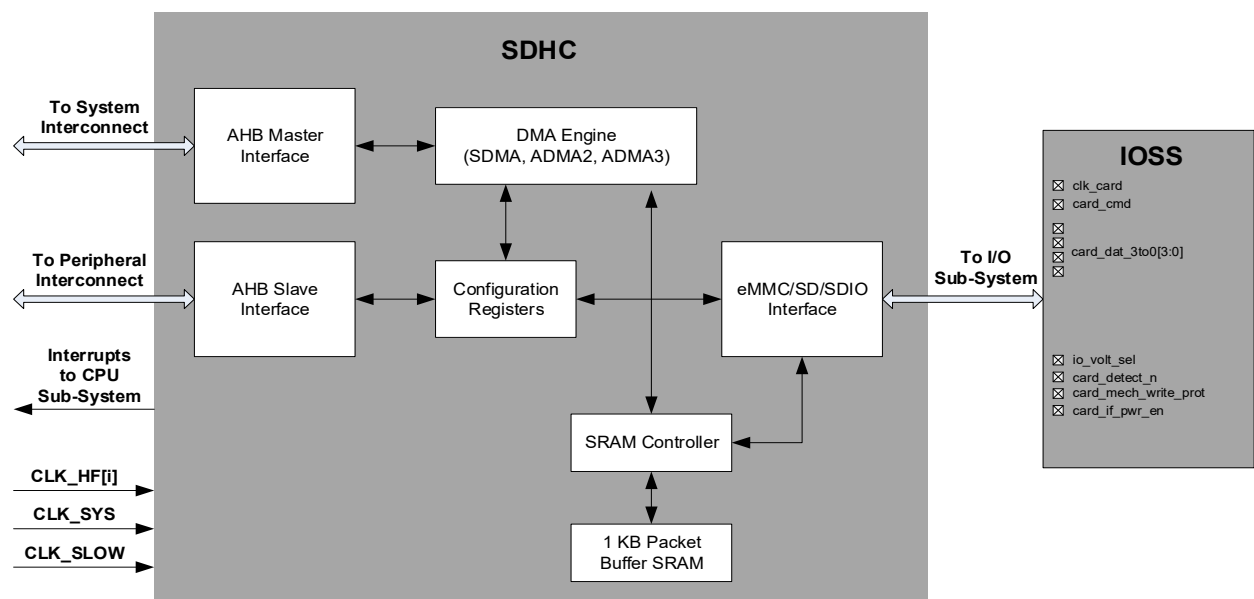
26.1.1 Features Not Supported

The SDHC block does not support the following features.

- SD/SDIO operation in UHS-II mode
- Command queuing engine (CQE)
- eMMC boot operation in dual data rate mode
- Read wait operation by DAT[2] signaling in an SDIO card
- Suspend/resume operation in an SDIO card
- Interrupt input pins for embedded SD systems
- SPI protocol mode of operation

26.2 Block Diagram

Figure 26-2. SDHC Block Diagram



The SDHC controller supports all three interfaces – SD, SDIO, and eMMC; it supports up to 4-bit bus width. The AHB master interface helps to transfer data to and from the system memory and the AHB slave interface provides access to the configuration registers. The register set comprises the standard SD host controller interface (HCI) registers as specified in the SD Specifications Part A2 SD Host Controller Standard Specification. These registers are described in the [registers TRM](#). The DMA engine handles direct data transfer between the SDHC logic and system memory. It supports SDMA, ADMA2, and ADMA3 modes based on the configuration.

The SDHC block complies with the following standards. Refer to the specifications documents for more information on the protocol and operations.

- SD Specifications Part 1 Physical Layer Specification Version 6.00
- SD Specifications Part A2 SD Host Controller Standard Specification Version 4.20
- SD Specifications Part E1 SDIO Specifications Version 4.10
- Embedded Multi-Media Card (eMMC) Electrical Standard 5.1

26.3 Clocking

Table 26-1 lists the different clocks used in the SDHC block. While configuring the clock for SDHC make sure that $\text{clk_slow} \geq \text{clk_sys} \geq \text{clk_card}$.

Table 26-1. Clocks in SDHC

Source	SDHC Clock	Function
CLK_SLOW	Core SDHC Clock	Used for core SDHC functions including the packet buffer SRAM; it is sourced from the slow clock (clk_slow); it must be \geq AHB slave clock.
	AHB Master Interface Clock	Used by the AHB master interface; it is sourced from the slow clock (clk_slow); it must be \geq AHB slave clock.
CLK_SYS	AHB Slave Interface Clock	Used by the AHB slave interface; it is clocked by the PERI group clock (clk_sys); it must be $\geq \text{clk_card}$. The group clock is derived from the PERI clock (clk_peri) using a divider. Because this divider can remain at the default value of '1' for most applications, clk_peri can be considered as clk_sys for SDHC. See Clocking System chapter on page 207 for information on clk_sys and clk_peri .
CLK_HF[i]	Base Clock/Card Clock	Used for sourcing the SD/eMMC interface clock (clk_card); it is derived from $\text{CLK_HF}[i]$; it must be set to 100 MHz to be compatible with the Capabilities register. See 26.3.2 Base Clock (CLK_HF[i]) Configuration for details. See High-Frequency Root Clocks on page 217 to know which $\text{CLK_HF}[i]$ drives an SDHC instance.
	Timer Clock	Used for command and data timeout functions; it is derived from $\text{CLK_HF}[i]$.

26.3.1 Clock Gating

All the clocks except the slave interface clock can be gated internally to enter standby mode (See [Power Modes on page 276](#)). In standby mode, you can also stop the clocks externally if required. The slave clock cannot be gated because it is used for wakeup logic (see [Interrupts to CPU on page 276](#)) during the standby mode.

The card clock is gated by clearing the SD_CLK_EN bit and other clocks are gated by clearing the INTERNAL_CLK_EN bit of the CLK_CTRL_R register. See [Clock Setup on page 282](#) for the sequence to be followed while modifying this bit.

26.3.2 Base Clock (CLK_HF[i]) Configuration

The HCI register (Capabilities) has a read-only field (BASE_CLK_FREQ) to indicate the base clock frequency so that an SD HCI-compatible driver can easily configure the divider for the required bus speed. This value is set to 0x64 (100 MHz) and hence $\text{CLK_HF}[i]$ must be set to 100 MHz. If this compatibility is not required, $\text{CLK_HF}[i]$ can be set to any value. See [26.3.4 Timeout \(TOUT\) Configuration](#).

26.3.3 Card Clock (SDCLK) Configuration

The SDCLK or card clock frequency is set by configuring the 10-bit divider in the CLK_CTRL_R register and selecting the 10-bit divided clock mode by clearing the CLK_GEN_SELECT bit of the same register. The default value of this bit is zero. The UPPER_FREQ_SEL field holds the upper two bits (9:8) and the FREQ_SEL field holds the lower eight bits (7:0) of the divider. Base clock frequency is

sourced from $\text{CLK_HF}[i]$ as explained in [Table 26-1](#). SDCLK frequency is equal to base clock frequency when the divider value is zero.

$$\text{SDCLK Frequency} = \text{Base Clock Frequency} / (2 \times 10\text{-bit divider value})$$

These fields are set automatically, based on the selected Bus Speed mode, to a value specified in one of the preset registers when $\text{HOST_CTRL2_R.PRESET_VAL_ENABLE}$ is set. The preset registers are selected according to [Table 26-2](#).

26.3.4 Timeout (TOUT) Configuration

An internal timer is used for command and data timeouts. The timeout value is specified through the $\text{TOUT_CTRL_R.TOUT_CNT}$ register field. The timer clock (TMCLK) frequency indicated by the read-only fields TOUT_CLK_FREQ and TOUT_CLK_UNIT of CAPABILITIES1_R register is 1 MHz. Timer clock is derived by dividing the $\text{CLK_HF}[i]$, which means that $\text{CLK_HF}[i]$ must be set to 100 MHz to be compatible with the Capabilities register.

26.4 Bus Speed Modes

The SDHC block can operate in either SD/SDIO mode or in eMMC mode. The SDHC block operates in eMMC mode when the EMMC_CTRL_R.CARD_IS_EMMC bit is set; otherwise, it operates in SD/SDIO mode. The speed mode is configured through HOST_CTRL1_R and HOST_CTRL2_R registers as per [Table 26-2](#). The HOST_CTRL2_R.UHS2_IF_ENABLE bit should remain at its default value of zero because the block does not support UHS-II mode. The card clock must be configured according to the selected speed mode through the CLK_CTRL_R register. See [26.3.3 Card Clock \(SDCLK\) Configuration](#) for more information.

Table 26-2. Bus Speed Mode Configuration

Bus Speed Mode	HOST_CTRL1_R Field	HOST_CTRL2_R Fields		Selected Preset Register
	HIGH_SPEED_EN	SIGNALING_EN	UHS_MODE_SEL	
SD Default Speed (DS)	0	0	Don't care	PRESET_DS_R
SD High Speed (HS)	1	0	Don't care	PRESET_HS_R
SDR12/eMMC Legacy	Don't care	1	000b	PRESET_SDR12_R
SDR25/eMMC High Speed	Don't care	1	001b	PRESET_SDR25_R
SDR50	Don't care	1	010b	PRESET_SDR50_R
DDR50/eMMC High Speed DDR	Don't care	1	100b	PRESET_DDR50_R

26.5 Power Modes

The block can operate during active and sleep system power modes. It does not support deep sleep mode and cannot wake up from events such as card insertion and removal when the system is in deep sleep. All the core registers except the packet buffer SRAM are retained when the system enters deep sleep mode and the SRAM is switched off to save power. Retention is performed so that the block can resume operation immediately after wakeup from deep sleep without requiring reconfiguration. Make sure that no AHB traffic (such as register read/write and DMA operation) is present, the SD/SDIO/eMMC bus interface is idle, and no data packets are pending in the packet buffer SRAM when the system transitions into deep sleep mode.

26.5.1 Standby Mode

The block can be put into standby mode to save power during the active and sleep system power modes by turning off the clocks. See [Clock Gating on page 275](#) for details. The block can detect wakeup interrupts (see [Interrupts to CPU on page 276](#)) in standby mode.

26.6 Interrupts to CPU

The block provides two interrupt signals to CPUSS:

- Wakeup Interrupt Signal – Triggered on events such as card insertion, removal, and SDIO card interrupt. This interrupt source cannot wake up the system from deep sleep mode and is provided so that a host driver can take appropriate action on those events. For example, resuming operation from standby mode on card

insertion. See [26.5.1 Standby Mode](#) for details. As card insertion and removal is not applicable to an embedded device, wakeup interrupt should not be used in this case. However it can still be used for SDIO card interrupt.

- General Interrupt Signal – Triggered on all other events, in either normal conditions or error conditions.

A host driver must not enable the wakeup and general interrupt signals at the same time.

To use only the wakeup interrupt signal, clear the NORMAL_INT_STAT_R and NORMAL_INT_SIGNAL_EN_R registers, and then set the enable bits of the required wakeup events in the WUP_CTRL_R and NORMAL_INT_STAT_EN registers.

To use only the general interrupt signal, clear the WUP_CTRL_R and NORMAL_INT_STAT_R registers. Then, set the required bits in NORMAL_INT_SIGNAL_EN_R and NORMAL_INT_STAT_EN registers.

These interrupts remain asserted until the CPU clears the interrupt status through one of the status registers – NORMAL_INT_STAT_R and ERROR_INT_STAT_R.

The SDIO card interrupt status bit, CARD_INTERRUPT, is a read-only bit. The host driver may clear the NORMAL_INT_STAT_EN_R.CARD_INTERRUPT_STAT_EN bit before servicing the SDIO card interrupt and may set this bit again after all interrupt requests from the card are cleared to prevent inadvertent interrupts.

Following is the list of registers used in interrupt configuration.

Table 26-3. Interrupt Control Registers

Register	Description
WUP_CTRL_R	Enables or disables different wakeup interrupts. Host driver must maintain voltage on the SD bus by setting PWR_CTRL_R.SD_BUS_PWR_VDD1 bit for these interrupts to occur. These interrupts cannot wakeup the device from deep sleep.
NORMAL_INT_STAT_R	Reflects the status of wakeup interrupts and non-error general interrupts. It also has a bit to indicate whether any of the bits in ERROR_INT_STAT_R is set.
ERROR_INT_STAT_R	Reflects the status of general interrupts that are triggered by error conditions.
NORMAL_INT_STAT_EN_R	Provides mask bits for wakeup interrupts and non-error general interrupts.
ERROR_INT_STAT_EN_R	Provides mask bits for general interrupts that are triggered by error conditions.
NORMAL_INT_SIGNAL_EN_R	Setting any of these bits to '1' enables interrupt generation for wakeup interrupts and non-error general interrupts.
ERROR_INT_SIGNAL_EN_R	Setting any of these bits to '1' enables interrupt generation for general interrupts that are triggered by error conditions.
FORCE_ERROR_INT_STAT_R	Forces an error interrupt to occur when the corresponding bit is set.

26.6.1 SDIO Interrupt

The SDIO interrupt function is supported on card_dat_3to0[1] line (SD pin 8). See the SDIO specifications for details on this feature. The CARD_INTERRUPT_STAT_EN bit in the NORMAL_INT_STAT_EN_R register and the CARD_INTERRUPT_SIGNAL_EN bit in the NORMAL_INT_SIGNAL_EN_R register must be set to enable this interrupt. To use this interrupt as wakeup interrupt, use WUP_CTRL_R.WUP_CARD_INT instead of NORMAL_INT_SIGNAL_EN_R.

26.7 I/O Interface

SDHC block provides the signals shown in [Table 26-4](#), which can be routed to pins through the I/O subsystem (IOSS). Refer to the [I/O System chapter on page 226](#) to configure the I/Os, and the [device datasheet](#) for specific pins available for each signal. SDHC also supports SDIO interrupt on DAT[1] line (card_data_3to0[1]). The output signals must be configured in strong drive mode, bi-directional signals in strong drive with the input buffer ON, and the input pins in high-impedance mode when an external pull-up resistor is available; otherwise, they must be configured in internal pull-up mode. Input buffer must be enabled for the input pins. The drive mode of the DAT lines must be set to high impedance after card removal. See [Card Detection on page 280](#) for details. In addition to configuring the drive mode and HSIOM registers in IOSS, the SDHC_CORE.GP_OUT_R register must be configured to enable the required signals. See [Table 26-4](#). The card_detect_n and card_write_prot should be connected to ground if an eMMC or an embedded SDIO device is connected.

Table 26-4. I/O Signal Interface

Signal	Function	Register Configuration
clk_card	Clock output	GP_OUT_R.CARD_CLOCK_OE
card_cmd	Command (bi-directional)	Always enabled
card_dat_3to0[3:0]	Data (bi-directional)	HOST_CTRL1_R.DAT_XFER_WIDTH
card_detect_n	Card detect signal input, Active low	GP_OUT_R.CARD_DETECT_EN
card_mech_write_prot	Mechanical write protect signal input, Active low	GP_OUT_R.CARD_MECH_WRITE_PROT_EN
io_volt_sel	Signaling voltage select output (see 26.7.1 Switching Signaling Voltage from 3.3 V to 1.8 V)	GP_OUT_R.IO_VOLT_SEL_OE
card_if_pwr_en	Card interface power enable output	GP_OUT_R.CARD_IF_PWR_EN_OE

26.7.1 Switching Signaling Voltage from 3.3 V to 1.8 V

The I/Os operate at the voltage level supplied through the external V_{DDIO} pin. The SD mode supports switching the signaling voltage from 3.3 V to 1.8 V after negotiation with the SD card. The block sets the HOST_CTRL2.R.SIGNALING_EN bit to indicate the switch. This value is reflected on the io_volt_sel pin, which can be connected to an external regulator powering V_{DDIO} to switch between 3.3 V and 1.8 V. Note that PSoC 6 does not provide an internal regulator to power the SD interface I/Os.

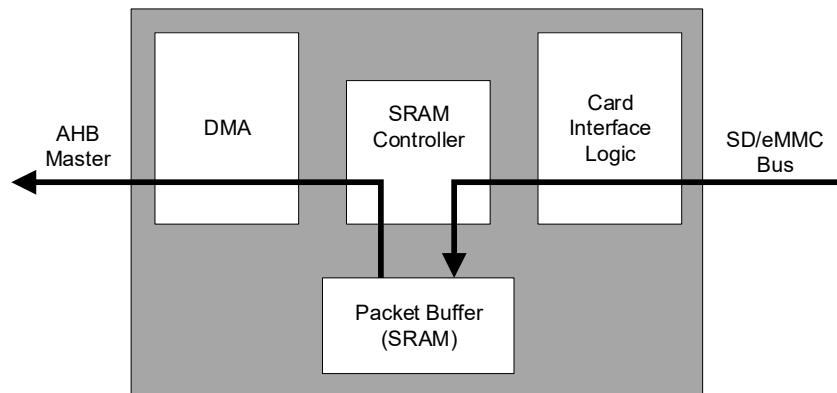
26.8 Packet Buffer SRAM

SRAM that is internal to the SDHC block is used as a packet buffer to store data packets while carrying out data transfer to and from the card. The size of the SRAM is 1KB to support buffering of two 512 bytes blocks. As write and read

transfers to the cards do not occur simultaneously, a single shared buffer is used for read and write operations. During the data transfer command handshake, the read/write bit of the command register is sampled and stored. This internal bit defines whether the SDHC is in read or write mode.

Figure 26-3 shows how data flows from the card interface to the AHB master interface through the packet buffer for a card read transfer. Received data from the card interface is written into packet buffer. When one block of data is received, DMA starts transmitting that data to the system by reading it from the packet buffer. For a card write transfer, data flows in the reverse direction. DMA writes data into a packet buffer that is subsequently read by the card interface logic. DMA and card interface logic can work simultaneously because read and write to packet buffer can be interleaved. For card read, DMA can send out the previous block while card interface logic is receiving the current block. For card write, DMA can write the current block into packet buffer while card interface logic is sending out the previous block.

Figure 26-3. Data Flow in a Read Transfer



26.8.1 Packet Buffer Full/Empty

When the packet buffer becomes full in card read, the clock to the card is stopped to prevent the card from sending the next data block. When packet buffer is empty, data block is not sent. In both cases, card interface logic is idle. SDHC does not support SDIO Read Wait signaling through DAT[2]. Therefore, the I/O command (CMD52) cannot be performed during a multiple read cycle because the card clock is stopped.

is moving data during task execution (for CMD46 and CMD47).

- Prefetches data for back-to-back eMMC write commands.
- Writes back the received data packets to system memory.

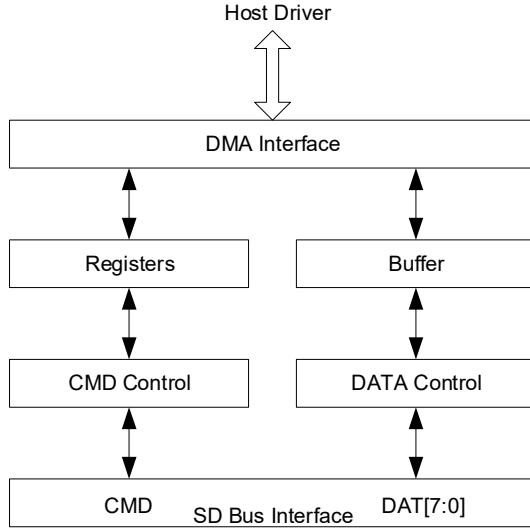
Figure 26-4 shows the data flow between the host driver and SD bus. The host driver can transfer data using either a programmed I/O (PIO) method in which the internal buffer is accessed through the buffer data port (BUF_DATA_R) register or using any of the defined DMA methods. PIO mode is much slower and burdens the processor. Do not use the PIO mode for large transfers.

26.9 DMA Engine

The DMA engine handles data transfer between SDHC and system memory. Following are the features of this unit:

- Supports SDMA, ADMA2, and ADMA3 modes based on the configuration.
- The same DMA engine is used to interleave data transfer and descriptor fetch. This enables new task descriptor fetches (for CMD44 and CMD45) while DMA

Figure 26-4. Data Flow



DMA supports both single block and multi-block transfers. The control bits in the block gap control (BGAP_CTRL_R) register is used to stop and restart a DMA operation. SDMA mode is used for short data transfer because it generates interrupts at page boundaries. These interrupts disturb the CPU to reprogram the new system address. Only one SD command transaction can be executed for every SDMA operation.

ADMA2 and ADMA3 are used for long data transfers. They adopt scatter gather algorithm so that higher data transfer speed is available. The host driver can program a list of data transfers between system memory and SD card to the descriptor table. ADMA2 performs one read/write SD command operation at a time. ADMA3 can program multiple read/write SD command operation in a descriptor table.

In SDMA and ADMA2 modes, writing the CMD_R register triggers the DMA operation. In ADMA3 mode, writing ADMA_ID_LOW_R register triggers the DMA operation.

SD mode commands are generated by writing into the following registers – system address (SDMASA_R), block size (BLOCKSIZE_R), block count (BLOCKCOUNT_R), transfer mode (XFER_MODE_R), and command (CMD_R). When HOST_CTRL2_R.HOST_VER4_EN = 0, SDMA uses SDMASA_R as system address register and hence Auto CMD23 cannot be used with SDMA because this register is assigned for Auto CMD23 as the 32-bit block count register. When HOST_CTRL2_R.HOST_VER4_EN = 1, SDMA uses ADMA_SA_LOW_R as system address register and SDMASA_R is reassigned to 32-bit block count and hence SDMA may use Auto CMD23.


To use the 32-bit block count register when HOST_CTRL2_R.HOST_VER4_EN = 1, it must be programmed with a non-zero value and the value of the 16-bit block count register BLOCK_COUNT_R must be zero.

Refer to the respective specifications documents listed in [Block Diagram on page 274](#) to learn more about the DMA operation.

26.10 Initialization Sequence

Figure 26-5 shows the sequence for initializing SDHC to work with SD/SDIO/eMMC cards. Subsequent sections describe each step. After initialization, SDHC is ready to communicate with the card. Refer to the corresponding specifications document for information on other sequences such as card initialization and identification, changing bus speed mode, signal voltage switch procedure, transaction generation, and error recovery.

Figure 26-5. SDHC Programming Sequence



26.10.1 Enabling SDHC

Ensure clk_sys is configured to be greater than or equal to clk_card and is running. Then, follow the sequence in [Figure 26-6](#) to enable the block. The internal clock can also be enabled later during clock setup. It must be enabled to detect card insertion or removal through general interrupts when SDHC is not in standby mode. See [26.10.2 Card Detection](#) for details.

Figure 26-6. SDHC Enable Sequence

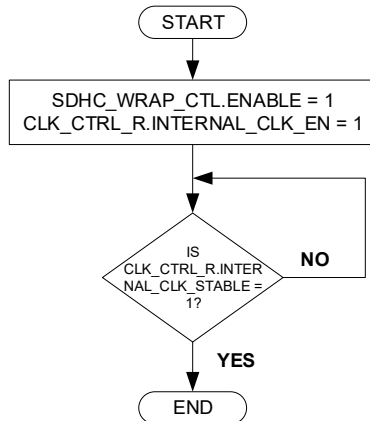



Figure 26-7. Card Status Check Sequence

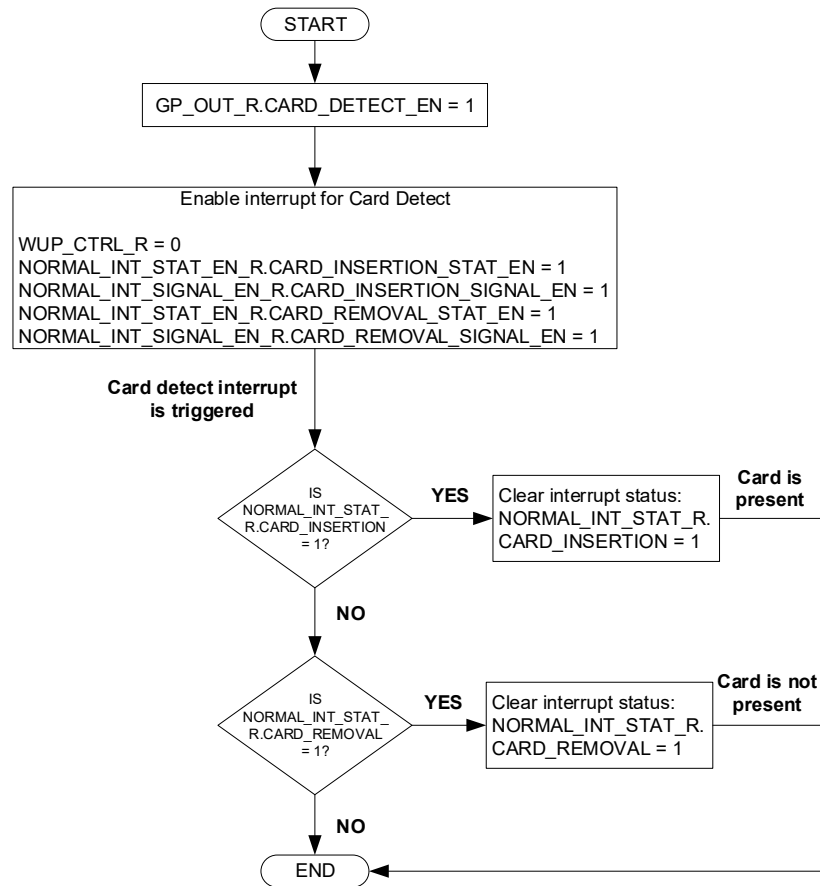


26.10.2 Card Detection

Check if the card is already inserted by following the sequence shown in [Figure 26-7](#). This step is required for a removable card. After the card is detected, the host driver can supply power and clock to the card. If the card is inserted, then proceed with SDHC initialization. To detect card insertion or removal through interrupt when the internal clock is already enabled, follow the sequence shown in [Figure 26-8](#). To detect the card status through interrupt when the internal clock is disabled (when SDHC is in standby mode), the bits in the WUP_CTRL_R register must be set and the NORMAL_INT_SIGNAL_EN register must be cleared. See [Interrupts to CPU on page 276](#) for details. To detect SDIO card interrupt on DAT[1] line, a separate bit is provided in these registers, which must be configured.

SDHC clears the PWR_CTRL_R.SD_BUS_PWR_VDD1 bit when the card is removed and drives the DAT lines low. Therefore, the drive mode of the DAT lines must be changed from strong (with input buffer ON) to HI-Z when the card is removed to keep the lines pulled high. After detecting card insertion, the drive mode must be configured back to strong (with input buffer ON) mode only after PWR_CTRL_R.SD_BUS_PWR_VDD1 is set to 1.

Figure 26-8. Card Detection Through Interrupt



26.10.3 SDHC Initialization

To initialize SDHC, configure the basic settings as shown in [Figure 26-9](#). This step can also be executed immediately after enabling SDHC.

Figure 26-9. SDHC Setup





Figure 26-10. Clock Setup



* See the CY8C62x8, CY8C62xA Registers TRM for details.

26.10.4 Clock Setup

Enable the internal clock followed by the card clock (SD clock) by following the sequence shown in [Figure 26-10](#). The SD clock frequency must be 100 kHz to 400 kHz during the card initialization. See [Card Clock \(SDCLK\) Configuration on page 275](#) for details. SD clock can be started and stopped by toggling the CLK_CTRL_R.SD_CLK_EN bit. The same sequence excluding the step of enabling the internal clock can be used to change the SD clock frequency. The SD clock must be stopped before changing its frequency. Note that GP_OUT_R.CARD_CLOCK_OE should have been set to '1' for the card clock to appear on the pin.

26.11 Error Detection

SDHC can detect different types of errors in SD and eMMC transactions. Error is detected in either the command or data portion of the transaction. When an error is detected, the ERR_INTERRUPT bit in the NORMAL_INT_STAT_R register is set. The exact error can then be identified through the ERROR_INT_STAT_R register. The Abort command is used to recover from an error detected during data transfer. In addition to these two registers, SDHC has two other error status registers – Auto CMD Error Status (AUTO_CMD_STAT_R) and ADMA Error Status (ADMA_ERR_STAT_R). The following table lists the errors detected by SDHC.

Table 26-5. Errors Detected by SHDC

Type	Error
Command Errors	Command Timeout Error Command CRC Error Command End Bit Error Command Index Error Command Conflict Error Response Error
Auto Command Errors	Command not issued by Auto CMD12 Error Auto Command Timeout Error Auto Command CRC Error Auto Command End Bit Error Auto Command Index Error Auto Command Conflict Error Auto CMD response Error
Data Errors	Data Timeout Error Data CRC Error Data End Bit Error ADMA Error Tuning Error

27. Serial Communications Block (SCB)



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SCB
- Application notes
- Code examples

The Serial Communications Block (SCB) supports three serial communication protocols: Serial Peripheral Interface (SPI), Universal Asynchronous Receiver Transmitter (UART), and Inter Integrated Circuit (I²C or IIC). Only one of the protocols is supported by an SCB at any given time. The number of SCBs in a PSoC 6 MCU varies by part number; consult the [device datasheet](#) to determine number of SCBs and the SCB pin locations. Not all SCBs support all three modes (SPI, UART, and I²C); consult the [device datasheet](#) to determine which modes are supported by which SCBs. Not all SCBs operate in deep sleep; consult the [device datasheet](#) to determine which SCBs operate in deep sleep.

27.1 Features

The SCB supports the following features:

- Standard SPI master and slave functionality with Motorola, Texas Instruments, and National Semiconductor protocols
- Standard UART functionality with SmartCard reader, Local Interconnect Network (LIN), and IrDA protocols
 - Standard LIN slave functionality with LIN v1.3 and LIN v2.1/2.2 specification compliance
- Standard I²C master and slave functionality
- Trigger outputs for connection to DMA
- Multiple interrupt sources to indicate status of FIFOs and transfers
- Features available only on Deep Sleep-capable SCB:
 - EZ mode for SPI and I²C slaves; allows for operation without CPU intervention
 - CMD_RESP mode for SPI and I²C slaves; allows for operation without CPU intervention
 - Low-power (Deep Sleep) mode of operation for SPI and I²C slaves (using external clocking)
 - Deep Sleep wakeup on I²C slave address match or SPI slave selection

27.2 Architecture

The operation modes supported by SCB are described in the following sections.

27.2.1 Buffer Modes

Each SCB has 256 bytes of dedicated RAM for transmit and receive operation. This RAM can be configured in three different modes (FIFO, EZ, or CMD_RESP). The following sections give a high-level overview of each mode. The sections on each protocol will provide more details.

- Masters can only use FIFO mode
- I²C and SPI slaves can use all three modes. **Note:** EZ Mode and CMD Response Mode are available only on the Deep Sleep-capable SCB
- UART only uses FIFO mode

Note: This document discusses hardware implementation of the EZ mode; for the firmware implementation, see the [PDL](#).

27.2.1.1 FIFO Mode

In this mode the RAM is split into two 128-byte FIFOs, one for transmit (TX) and one for receive (RX). The FIFOs can be configured to be 8 bits x 128 elements or 16 bits x 64 elements; this is done by setting the BYTE_MODE bit in the SCB control register.

FIFO mode of operation is available only in Active and Sleep power modes. However, the I²C address or SPI slave select can be used to wake the device from Deep Sleep on the Deep Sleep-capable SCB.

Statuses are provided for both the RX and TX FIFOs. There are multiple interrupt sources available, which indicate the status of the FIFOs, such as full or empty; see “[SCB Interrupts](#)” on page 332.

27.2.1.2 EZ Mode

In easy (EZ) mode the RAM is used as a single 256-byte buffer. The external master sets a base address and reads and writes start from that base address.

EZ Mode is available only for SPI slave and I²C slave. It is available only on the Deep Sleep capable SCB.

EZ mode is available in Active, Sleep, and Deep Sleep power modes.

Note: This document discusses hardware implementation of the EZ mode; for the firmware implementation, see the [PDL](#).

27.2.1.3 CMD_RESP Mode

Command Response (CMD_RESP) mode is similar to EZ mode except that the base address is provided by the CPU not the external master.

CMD_RESP mode is available only for SPI slave and I²C slave. It is available only on the Deep Sleep-capable SCB.

CMD_RESP mode operation is available in Active, Sleep, and Deep Sleep power modes.

27.2.2 Clocking Modes

The SCB can be clocked either by an internal clock provided by the peripheral clock dividers (referred to as `clk_scb` in this document), or it can be clocked by the external master.

- UART, SPI master, and I²C master modes must use `clk_scb`.
- Only SPI slave and I²C slave can use the clock from and external master, and only the Deep Sleep capable SCB supports this.

Internally- and externally-clocked slave functionality is determined by two register fields of the SCB CTRL register:

- `EC_AM_MODE` indicates whether SPI slave selection or I²C address matching is internally ('0') or externally ('1') clocked.
- `EC_OP_MODE` indicates whether the rest of the protocol operation (besides SPI slave selection and I²C address matching) is internally ('0') or externally ('1') clocked.

Notes:

- FIFO mode supports an internally- or externally-clocked address match (`EC_AM_MODE` is '0' or '1'); however, data transfer must be done with internal clocking. (`EC_OP_MODE` is '1').
- EZ and CMD_RESP modes are supported with externally clocked operation (`EC_OP_MODE` is '1').

[Table 27-1](#) provides an overview of the clocking and buffer modes supported for each communication mode.

Table 27-1. Clock Mode Compatibility

	Internally clocked (IC)			Externally clocked (EC) (Deep Sleep SCB only)		
	FIFO	EZ	CMD_RESP	FIFO	EZ	CMD_RESP
I ² C master	Yes	No	No	No	No	No
I ² C slave	Yes	Yes	No	Yes ^a	Yes	Yes
I ² C master-slave	Yes	No	No	No	No	No
SPI master	Yes	No	No	No	No	No
SPI slave	Yes	Yes	No	Yes ^b	Yes	Yes
UART transmitter	Yes	No	No	No	No	No
UART receiver	Yes	No	No	No	No	No

a. In Deep Sleep mode the external-clocked logic can handle slave address matching, it then triggers an interrupt to wake up the CPU. The slave can be programmed to stretch the clock, or NACK until internal logic takes over. This applies only to the Deep Sleep-capable SCB.

b. In Deep Sleep mode the external-clocked logic can handle slave selection detection, it then triggers an interrupt to wake up the CPU. Writes will be ignored and reads will return 0xFF until internal logic takes over. This applies only to the Deep Sleep-capable SCB.

Table 27-2. Clock Configuration and Mode support

Mode	EC_AM_MODE is '0'; EC_OP_MODE is '0'	'EC_AM_MODE is '1'; EC_OP_MODE is '0'	'EC_AM_MODE is '1'; EC_OP_MODE is '1'
FIFO mode	Yes	Yes	No
EZ mode	Yes	Yes	Yes
CMD_RESP mode	No	No	Yes

27.3 Serial Peripheral Interface (SPI)

The SPI protocol is a synchronous serial interface protocol. Devices operate in either master or slave mode. The master initiates the data transfer. The SCB supports single-master-multiple-slaves topology for SPI. Multiple slaves are supported with individual slave select lines.

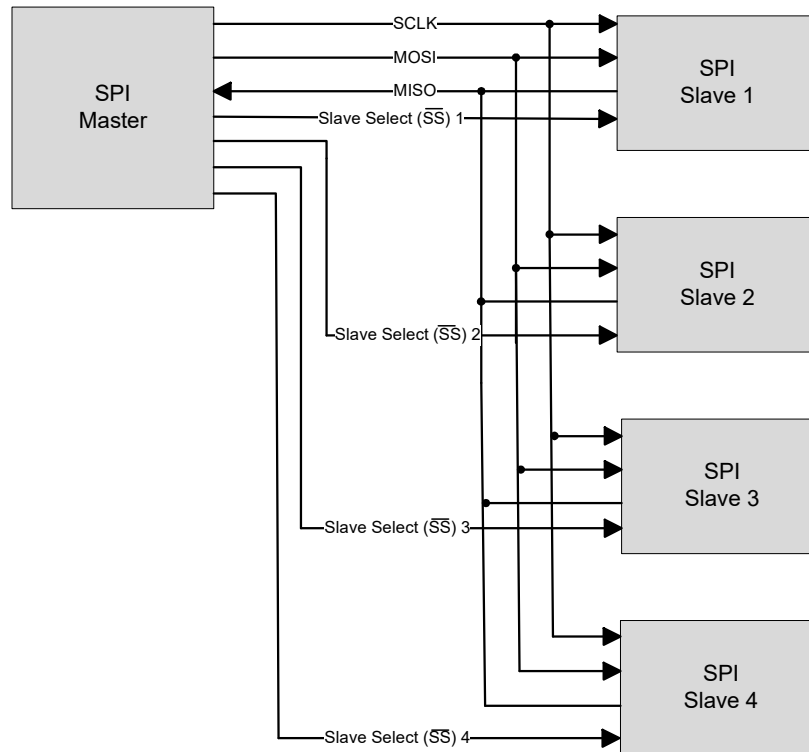
27.3.1 Features

- Supports master and slave functionality
- Supports three types of SPI protocols:
 - Motorola SPI – modes 0, 1, 2, and 3
 - Texas Instruments SPI, with coinciding and preceding data frame indicator – mode 1 only
 - National Semiconductor (MicroWire) SPI – mode 0 only
- Master supports up to four slave select lines
 - Each slave select has configurable active polarity (high or low)
 - Slave select can be programmed to stay active for a whole transfer, or just for each byte
- Master supports late sampling for better timing margin
- Master supports continuous SPI clock
- Data frame size programmable from 4 bits to 16 bits
- Programmable oversampling
- MSb or LSb first
- Median filter available for inputs
- Supports FIFO Mode
- Supports EZ Mode (slave only) and CMD_RESP mode (slave only) on the Deep Sleep-capable SCB

27.3.2 General Description

Figure 27-1 illustrates an example of SPI master with four slaves.

Figure 27-1. SPI Example



A standard SPI interface consists of four signals as follows.

- **SCLK:** Serial clock (clock output from the master, input to the slave).
- **MOSI:** Master-out-slave-in (data output from the master, input to the slave).
- **MISO:** Master-in-slave-out (data input to the master, output from the slave).
- **Slave Select (\overline{SS}):** Typically an active low signal (output from the master, input to the slave).

A simple SPI data transfer involves the following: the master selects a slave by driving its \overline{SS} line, then it drives data on the MOSI line and a clock on the SCLK line. The slave uses either of the edges of SCLK depending on the configuration to capture the data on the MOSI line; it also drives data on the MISO line, which is captured by the master.

By default, the SPI interface supports a data frame size of eight bits (1 byte). The data frame size can be configured to any value in the range 4 to 16 bits. The serial data can be transmitted either most significant bit (MSb) first or least significant bit (LSb) first.

Three different variants of the SPI protocol are supported by the SCB:

- **Motorola SPI:** This is the original SPI protocol.
- **Texas Instruments SPI:** A variation of the original SPI protocol, in which data frames are identified by a pulse on the \overline{SS} line.
- **National Semiconductors SPI:** A half-duplex variation of the original SPI protocol.

27.3.3 SPI Modes of Operation

27.3.3.1 Motorola SPI

The original SPI protocol was defined by Motorola. It is a full duplex protocol. Multiple data transfers may happen with the \overline{SS} line held at '0'. When not transmitting data, the \overline{SS} line is held at '1'.

Clock Modes of Motorola SPI

The Motorola SPI protocol has four different clock modes based on how data is driven and captured on the MOSI and MISO lines. These modes are determined by clock polarity (CPOL) and clock phase (CPHA).

Clock polarity determines the value of the SCLK line when not transmitting data. CPOL = '0' indicates that SCLK is '0' when not transmitting data. CPOL = '1' indicates that SCLK is '1' when not transmitting data.

Clock phase determines when data is driven and captured. CPHA = 0 means sample (capture data) on the leading (first) clock edge, while CPHA = 1 means sample on the trailing (second) clock edge, regardless of whether that clock edge is rising or falling. With CPHA = 0, the data must be stable for setup time before the first clock cycle.

- Mode 0: CPOL is '0', CPHA is '0': Data is driven on a falling edge of SCLK. Data is captured on a rising edge of SCLK.
- Mode 1: CPOL is '0', CPHA is '1': Data is driven on a rising edge of SCLK. Data is captured on a falling edge of SCLK.
- Mode 2: CPOL is '1', CPHA is '0': Data is driven on a rising edge of SCLK. Data is captured on a falling edge of SCLK.
- Mode 3: CPOL is '1', CPHA is '1': Data is driven on a falling edge of SCLK. Data is captured on a rising edge of SCLK.

Figure 27-2 illustrates driving and capturing of MOSI/MISO data as a function of CPOL and CPHA.

Figure 27-2. SPI Motorola, 4 Modes




Figure 27-3 illustrates a single 8-bit data transfer and two successive 8-bit data transfers in mode 0 (CPOL is '0', CPHA is '0').

Figure 27-3. SPI Motorola Data Transfer Example

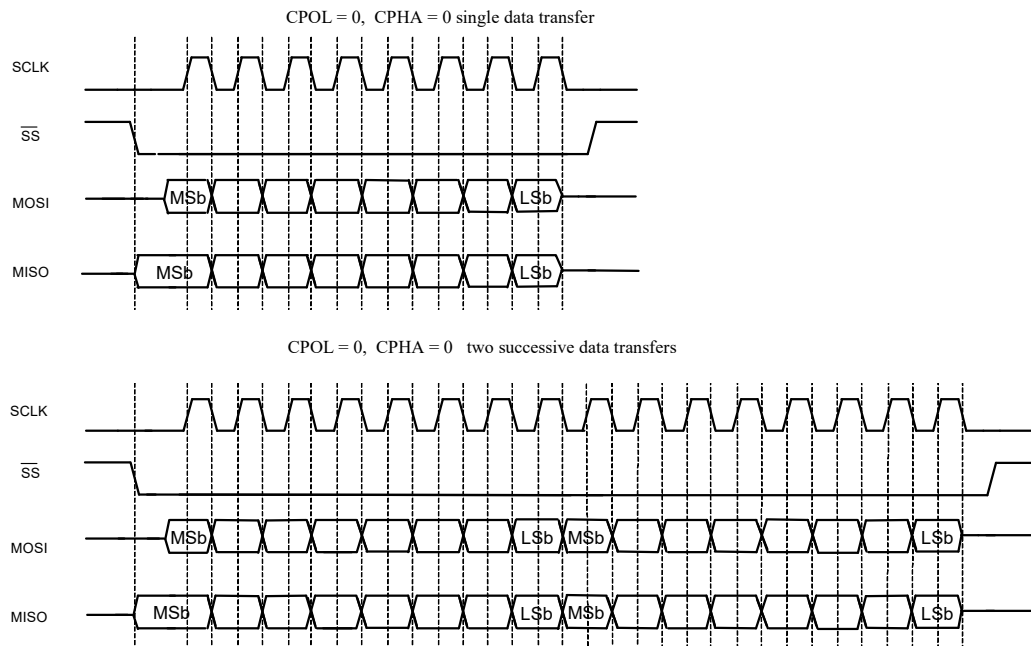



Figure 27-4. SELECT and SCLK Timing Correlation



For example above: OversamplingReg = 6 - 1 = 5.

$$\frac{3}{4} * SCLK = ((5 / 2) + 1) + (5 / 4 + 1)) * clk_scb = (3 + 2) * clk_scb = 5 * clk_scb.$$

$$\frac{1}{4} * SCLK = ((5 / 4) + 1) * clk_scb = 2 * clk_scb.$$

Note The value $\frac{3}{4} * SCLK$ is equal to $((\text{OversamplingReg} / 2) + 1) + (\text{OversamplingReg} / 4 + 1))$, where $\text{OversamplingReg} = \text{Oversampling} - 1$.

The value $\frac{1}{4} * SCLK$ is equal to $((\text{OversamplingReg} / 4) + 1)$.

The result of any division operation is rounded down to the nearest integer.

Note The provided timings are guaranteed by SCB block but do not take into account signal propagation time from SCB block to pins.

Configuring SCB for SPI Motorola Mode

To configure the SCB for SPI Motorola mode, set various register bits in the following order:

1. Select SPI by writing '01' to the MODE (bits [25:24]) of the SCB_CTRL register.
2. Select SPI Motorola mode by writing '00' to the MODE (bits [25:24]) of the SCB_SPI_CTRL register.
3. Select the clock mode in Motorola by writing to the CPHA and CPOL fields (bits 2 and 3 respectively) of the SCB_SPI_CTRL register.
4. Follow steps 2 to 4 mentioned in ["Enabling and Initializing SPI" on page 300](#).


For more information on these registers, see the [registers TRM](#).

27.3.3.2 Texas Instruments SPI

The Texas Instruments' SPI protocol redefines the use of the \overline{SS} signal. It uses the signal to indicate the start of a data transfer, rather than a low active slave select signal, as in the case of Motorola SPI. The start of a transfer is indicated by a high active pulse of a single bit transfer period. This pulse may occur one cycle before the transmission of the first data bit, or may coincide with the transmission of the first data bit. The TI SPI protocol supports only mode 1 (CPOL is '0' and CPHA is '1'): data is driven on a rising edge of SCLK and data is captured on a falling edge of SCLK.


[Figure 27-5](#) illustrates a single 8-bit data transfer and two successive 8-bit data transfers. The SELECT pulse precedes the first data bit. Note how the SELECT pulse of the second data transfer coincides with the last data bit of the first data transfer.

Figure 27-5. SPI TI Data Transfer Example



[Figure 27-6](#) illustrates a single 8-bit data transfer and two successive 8-bit data transfers. The SELECT pulse coincides with the first data bit of a frame.

Figure 27-6. SPI TI Data Transfer Example



Configuring SCB for SPI TI Mode

To configure the SCB for SPI TI mode, set various register bits in the following order:

1. Select SPI by writing '01' to the MODE (bits [25:24]) of the SCB_CTRL register.
2. Select SPI TI mode by writing '01' to the MODE (bits [25:24]) of the SCB_SPI_CTRL register.
3. Select the mode of operation in TI by writing to the SELECT_PRECEDE field (bit 1) of the SCB_SPI_CTRL register ('1' configures the SELECT pulse to precede the first bit of next frame and '0' otherwise).
4. Set the CPHA bit of the SCB_SPI_CONTROL register to '0', and the CPOL bit of the same register to '1'.
5. Follow steps 2 to 4 mentioned in ["Enabling and Initializing SPI" on page 300](#).

For more information on these registers, see the [registers TRM](#).


27.3.3.3 National Semiconductors SPI

The National Semiconductors' SPI protocol is a half-duplex protocol. Rather than transmission and reception occurring at the same time, they take turns. The transmission and reception data sizes may differ. A single 'idle' bit transfer period separates transfers from reception. However, successive data transfers are not separated by an idle bit transfer period.


The National Semiconductors SPI protocol supports only mode 0.

[Figure 27-7](#) illustrates a single data transfer and two successive data transfers. In both cases, the transmission data transfer size is eight bits and the reception data transfer size is four bits.

Figure 27-7. SPI NS Data Transfer Example
 CPOL=0, CPHA=0 single data transfer



CPOL=0, CPHA=0 two successive data transfers



Configuring SCB for SPI NS Mode

To configure the SCB for SPI NS mode, set various register bits in the following order:

1. Select SPI by writing '01' to the MODE (bits [25:24]) of the SCB_CTRL register.
2. Select SPI NS mode by writing '10' to the MODE (bits [25:24]) of the SCB_SPI_CTRL register.
3. Set the CPOL and CPHA bits of the SCB_SPI_CTRL register to '0'.
4. Follow steps 2 to 4 mentioned in ["Enabling and Initializing SPI" on page 300](#).

For more information on these registers, see the [registers TRM](#).

27.3.4 SPI Buffer Modes

SPI can operate in three different buffer modes – FIFO, EZ, and CMD_RESP modes. The buffer is used in different ways in each of these modes. The following subsections explain each of these buffer modes in detail.

27.3.4.1 FIFO Mode

The FIFO mode has a TX FIFO for the data being transmitted and an RX FIFO for the data received. Each FIFO is constructed out of the SRAM buffer. The FIFOs are either 64 elements deep with 16-bit data elements or 128 elements deep with 8-bit data elements. The width of a FIFO is configured using the BYTE_MODE bitfield of the SCB_CTRL register.

FIFO mode of operation is available only in Active and Sleep power modes, and not in the Deep Sleep mode.

Transmit and receive FIFOs allow write and read accesses. A write access to the transmit FIFO uses the TX_FIFO_WR register. A read access from the receive FIFO uses the RX_FIFO_RD register. For SPI master mode, data transfers are started when data is written into the TX FIFO. Note that when a master is transmitting and the FIFO becomes empty the slave is de-selected.

Transmit and receive FIFO status information is available through status registers, TX_FIFO_STATUS and RX_FIFO_STATUS, and through the INTR_TX and INTR_RX registers.

Each FIFO has a trigger output. This trigger output can be routed through the trigger mux to various other peripheral on the device such as DMA or TCPWMs. The trigger output of the SCB is controlled through the TRIGGER_LEVEL field in the RX_CTRL and TX_CTRL registers.

- For a TX FIFO a trigger is generated when the number of entries in the transmit FIFO is less than TX_FIFO_CTRL.TRIGGER_LEVEL.
- For the RX FIFO a trigger is generated when the number of entries in the FIFO is greater than the RX_FIFO_CTRL.TRIGGER_LEVEL.

Note that the DMA has a trigger deactivation setting. For the SCB this should be set to 16.

Active to Deep Sleep Transition

Before going to deep sleep ensure that all active communication is complete. For a master this can easily be done by checking the SPI_DONE bit in the INTR_M register, and ensuring the TX FIFO is empty.

For a slave this can be achieved by checking the BUS_BUSY bit in the SPI Status register. Also the RX FIFO should be empty before going to deep sleep. Any data in the FIFO will be lost during deep sleep.

Also before going to deep sleep the clock to the SCB needs to be disabled. This can be done by setting the SDA_IN_FILT_TRIM[1] bit in the I2C_CFG register to "0"

Lastly, when the device goes to deep sleep the SCB stops driving the GPIO lines. This leads to floating pins and can lead to undesirable current during deep sleep power modes. To avoid this condition before entering deep sleep mode change the HSIOM settings of the SCB pins to GPIO driven, then change the drive mode and drive state to the appropriate state to avoid floating pins. Consult the [I/O System chapter on page 226](#) for more information on pin drive modes.

Deep Sleep to Active Transition

EC_AM = 1, EC_OP = 0, FIFO Mode.

When the SPI Slave Select line is asserted the device will be awoken by an interrupt. After the device is awoken change the SPI pin drive modes and HSIOM settings back to what they were before deep sleep. When clk_hf[0] is at the desired frequency set SDA_IN_FILT_TRIM[1] to '1' to enable the clock to the SCB. Then write data into the TX FIFO. At this point the master can read valid data from the slave. Before that any data read by the master will be invalid.

27.3.4.2 EZSPI Mode

The easy SPI (EZSPI) protocol only works in the Motorola mode, with any of the clock modes. It allows communication between master and slave without the need for CPU intervention. In the PSoC 6 MCU, only the deep sleep-capable SCB supports EZSPI mode.

The EZSPI protocol defines a single memory buffer with an 8-bit EZ address that indexes the buffer (256-entry array of eight bit per entry) located on the slave device. The EZ address is used to address these 256 locations. All EZSPI data transfers have 8-bit data frames.

The CPU writes and reads to the memory buffer through the SCB_EZ_DATA registers. These accesses are word accesses, but only the least significant byte of the word is used.

EZSPI has three types of transfers: a write of the EZ address from the master to the slave, a write of data from the master to an addressed slave memory location, and a read by the master from an addressed slave memory location.

Note: When multiple bytes are read or written the master must keep SSEL low during the entire transfer.

EZ Address Write

A write of the EZ address starts with a command byte (0x00) on the MOSI line indicating the master's intent to write the EZ address. The slave then drives a reply byte on the MISO line to indicate that the command is acknowledged (0xFE) or not (0xFF). The second byte on the MOSI line is the EZ address.

Memory Array Write


A write to a memory array index starts with a command byte (0x01) on the MOSI line indicating the master's intent to write to the memory array. The slave then drives a reply byte on the MISO line to indicate that the command was registered (0xFE) or not (0xFF). Any additional bytes on the MOSI line are written to the memory array at locations indicated by the communicated EZ address. The EZ address is automatically incremented by the slave as bytes are written into the memory array. When the EZ address exceeds the maximum number of memory entries (256), it remains there and does not wrap around to 0. The EZ base address is reset to the address written in the EZ Address Write phase on each slave selection.

Memory Array Read

A read from a memory array index starts with a command byte (0x02) on the MOSI line indicating the master's intent to read from the memory array. The slave then drives a reply byte on the MISO line to indicate that the command was registered (0xFE) or not (0xFF). Any additional read data bytes on the MISO line are read from the memory array at locations indicated by the communicated EZ address. The EZ address is automatically incremented by the slave as bytes are read from the memory array. When the EZ address exceeds the maximum number of memory entries (256), it remains there and does not wrap around to 0. The EZ base address is reset to the address written in the EZ Address Write phase on each slave selection.

Figure 27-8 illustrates the write of EZ address, write to a memory array and read from a memory array operations in the EZSPI protocol.

Figure 27-8. EZSPI Example



Configuring SCB for EZSPI Mode

By default, the SCB is configured for non-EZ mode of operation. To configure the SCB for EZSPI mode, set the register bits in the following order:

1. Select EZ mode by writing '1' to the EZ_MODE bit (bit 10) of the SCB_CTRL register.
2. Set the EC_AM and EC_OP modes in the SCB_CTRL register as appropriate.
3. Set the BYTE_MODE bit of the SCB_CTRL register to '1'.
4. Follow the steps in ["Configuring SCB for SPI Motorola Mode" on page 290](#).
5. Follow steps 2 to 4 mentioned in ["Enabling and Initializing SPI" on page 300](#).

For more information on these registers, see the [registers TRM](#).

Active to Deep Sleep Transition

Before going to deep sleep ensure the master is not currently transmitting to the slave. This can be done by checking the BUS_BUSY bit in the SPI_STATUS register.

If the bus is not busy, disable the clock to the SCB by setting the SDA_IN_FILT_TRIM[1] bit to '0' in the I2C_CFG register.

Deep Sleep to Active Transition

- **EC_AM = 1, EC_OP = 0, EZ Mode.** MISO transmits 0xFF until the internal clock is enabled. Data on MOSI is ignored until the internal clock is enabled. Do not enable the internal clock until clk_hf[0] is at the desired frequency. After clk_hf[0] is at the desired frequency set the SDA_IN_FILT_TRIM[1] bit to '1' to enable the clock. The external master needs to be aware that when it reads 0xFF on MISO the device is not ready yet.
- **EC_AM = 1, EC_OP = 1, EZ Mode.** Do not enable the internal clock until clk_hf[0] is at the desired frequency. After clk_hf[0] is at the desired frequency set the SDA_IN_FILT_TRIM[1] bit to '1' to enable the clock.

27.3.4.3 Command-Response Mode

The command-response mode is defined only for an SPI slave. In the PSoC 6 MCU, only the deep sleep-capable SCB supports this mode. This mode has a single memory buffer, a base read address, a current read address, a base write address, and a current write address that are used to index the memory buffer. The base addresses are provided by the CPU. The current addresses are used by the slave to index the memory buffer for sequential accesses of the memory buffer. The memory buffer holds 256 8-bit data elements. The base and current addresses are in the range [0, 255]. This mode is only supported by the Motorola mode of operation.

The CPU writes and reads to the memory buffer through the SCB_EZ_DATA registers. These accesses are word accesses, but only the least significant byte of the word is used.

The slave interface accesses the memory buffer using the current addresses. At the start of a write transfer (SPI slave selection), the base write address is copied to the current write address. A data element write is to the current write address location. After the write access, the current address is incremented by '1'. At the start of a read transfer, the base read address is copied to the current read address. A data element read is to the current read address location. After the read data element is transmitted, the current read address is incremented by '1'.

If the current addresses equal the last memory buffer address (address equals 255), the current addresses are not incremented. Subsequent write accesses will overwrite any previously written value at the last buffer address. Subsequent read accesses will continue to provide the (same) read value at the last buffer address. The bus master should be aware of the memory buffer capacity in command-response mode.


The base addresses are provided through CMD_RESP_CTRL. The current addresses are provided through CMD_RESP_STATUS. At the end of a transfer (SPI slave de-selection), the difference between a base and current address indicates how many read/write accesses were performed. The block provides interrupt cause fields to identify the end of a transfer. Command-response mode operation is available in Active, Sleep, and Deep Sleep power modes.

The command-response mode has two phases of operation:

- **Write phase** – The write phase begins with a selection byte, which has its last bit set to '0' indicating a write. The master writes 8-bit data elements to the slave's memory buffer following the selection byte. The slave's current write address is set to the slave's base write address. Received data elements are written to the current write address memory location. After each memory write, the current write address is incremented.
- **Read phase** – The read phase begins with a selection byte, which has its last bit set to '1' indicating a read. The master reads 8-bit data elements from the slave's memory buffer. The slave's current read address is set to the slave's base read address. Transmitted data elements are read from the current address memory location. After each read data element is transferred, the current read address is incremented.

During the reception of the first byte, the slave (MISO) transmits either 0x62 (ready) or a value different from 0x62 (busy). When disabled or reset, the slave transmits 0xFF (busy). The byte value can be used by the master to determine whether the slave is ready to accept the SPI request.

Figure 27-9. Command-Response Mode Example



Note that a slave's base addresses are updated by the CPU and not by the master.

Active to Deep Sleep Transition

Before going to deep sleep ensure the master is not currently transmitting to the slave. This can be done by checking the BUS_BUSY bit in the SPI_STATUS register.

If the bus is not busy, disable the clock to the SCB by setting the SDA_IN_FILT_TRIM[1] bit to 0 in the I2C_CFG register.

Deep Sleep to Active Transition

EC_AM = 1, EC_OP = 1, CMD_RESP Mode.

Do not enable the internal clock until clk_hf[0] is at the desired frequency. When clk_hf[0] is at the desired frequency, set the SDA_IN_FILT_TRIM[1] bit to '1' to enable the clock.

Configuring SCB for CMD_RESP Mode

By default, the SCB is configured for non-CMD_RESP mode of operation. To configure the SCB for CMD_RESP mode, set the register bits in the following order:

1. Select the CMD_RESP mode by writing '1' to the CMD_RESP_MODE bit (bit 12) of the SCB_CTRL register.
2. Set the EC_AM and EC_OP modes to '1' in the SCB_CTRL register.
3. Set the BYTE_MODE bit in the SCB_CTRL register.
4. Follow the steps in ["Configuring SCB for SPI Motorola Mode" on page 290](#).
5. Follow steps 2 to 4 mentioned in ["Enabling and Initializing SPI" on page 300](#).

For more information on these registers, see the [registers TRM](#).

27.3.5 Clocking and Oversampling

27.3.5.1 Clock Modes

The SCB SPI supports both internally and externally clocked operation modes. Two bitfields (EC_AM_MODE and EC_OP_MODE) in the SCB_CTRL register determine the SCB clock mode. EC_AM_MODE indicates whether SPI slave selection is internally (0) or externally (1) clocked. EC_OP_MODE indicates whether the rest of the protocol operation (besides SPI slave selection) is internally (0) or externally (1) clocked.

An externally-clocked operation uses a clock provided by the external master (SPI SCLK). **Note:** In the PSoC 6 MCU only the Deep Sleep-capable SCB supports externally-clocked mode of operation and only for SPI slave mode.

An internally-clocked operation uses the programmable clock dividers. For SPI, an integer clock divider must be used for both master and slave. For more information on system clocking, see the [Clocking System chapter](#) on page 207.

The SCB_CTRL bitfields EC_AM_MODE and EC_OP_MODE can be configured in the following ways.

- EC_AM_MODE is '0' and EC_OP_MODE is '0': Use this configuration when only Active mode functionality is required.
 - FIFO mode: Supported.
 - EZ mode: Supported.
 - Command-response mode: Not supported. The slave (MISO) transmits a value different from a ready (0x62) byte during the reception of the first byte if command-response mode is attempted in this configuration.
- EC_AM_MODE is '1' and EC_OP_MODE is '0': Use this configuration when both Active and Deep Sleep functionality are required. This configuration relies on the externally-clocked functionality to detect the slave selection and relies on the internally-clocked functionality to access the memory buffer.

The "hand over" from external to internal functionality relies on a busy/ready byte scheme. This scheme relies on the master to retry the current transfer when it receives a busy byte and requires the master to support busy/ready byte interpretation. When the slave is selected, INTR_SPI_EC.WAKE_UP is set to '1'. The associated Deep Sleep functionality interrupt brings the system into Active power mode.

 - FIFO mode: Supported. The slave (MISO) transmits 0xFF until the CPU is awoken and the TX FIFO is populated. Any data on the MOSI line will be dropped until clk_scb is enabled see ["Deep Sleep to Active Transition"](#) on page 297 for more details
 - EZ mode: Supported. In Deep Sleep power mode, the slave (MISO) transmits a busy (0xFF) byte during

the reception of the command byte. In Active power mode, the slave (MISO) transmits a ready (0xFE) byte during the reception of the command byte.

- CMD_RESP mode: Not supported. The slave transmits (MISO) a value different from a ready (0x62) byte during the reception of the first byte.
- EC_AM_MODE is '1' and EC_OP_MODE is '1'. Use this mode when both Active and Deep Sleep functionality are required. When the slave is selected, INTR_SPI_EC.WAKE_UP is set to '1'. The associated Deep Sleep functionality interrupt brings the system into Active power mode. When the slave is deselected, INTR_SPI_EC.EZ_STOP and/or INTR_SPI_EC.EZ_WRITE_STOP are set to '1'.
 - FIFO mode: Not supported.
 - EZ mode: Supported.
 - CMD_RESP mode: Supported.

If EC_OP_MODE is '1', the external interface logic accesses the memory buffer on the external interface clock (SPI SCLK). This allows for EZ and CMD_RESP mode functionality in Active and Deep Sleep power modes.

In Active system power mode, the memory buffer requires arbitration between external interface logic (on SPI SCLK) and the CPU interface logic (on system peripheral clock). This arbitration always gives the highest priority to the external interface logic (host accesses). The external interface logic takes two serial interface clock/bit periods for SPI. During this period, the internal logic is denied service to the memory buffer. The PSoC 6 MCU provides two programmable options to address this "denial of service":

- If the BLOCK bitfield of SCB_CTRL is '1': An internal logic access to the memory buffer is blocked until the memory buffer is granted and the external interface logic has completed access. This option provides normal SCB register functionality, but the blocking time introduces additional internal bus wait states.
- If the BLOCK bitfield of SCB_CTRL is '0': An internal logic access to the memory buffer is not blocked, but fails when it conflicts with an external interface logic access. A read access returns the value 0xFFFF:FFFF and a write access is ignored. This option does not introduce additional internal bus wait states, but an access to the memory buffer may not take effect. In this case, the following failures are detected:
 - Read Failure: A read failure is easily detected because the returned value is 0xFFFF:FFFF. This value is unique as non-failing memory buffer read accesses return an unsigned byte value in the range 0x0000:0000-0x0000:00ff.
 - Write Failure: A write failure is detected by reading back the written memory buffer location, and confirming that the read value is the same as the written value.

For both options, a conflicting internal logic access to the memory buffer sets INTR_TX.BLOCKED field to '1' (for write accesses) and INTR_RX.BLOCKED field to '1' (for read accesses). These fields can be used as either status fields or as interrupt cause fields (when their associated mask fields are enabled).

If a series of read or write accesses is performed and CTRL.BLOCKED is '0', a failure is detected by comparing the "logical-or" of all read values to 0xFFFF:FFFF and checking the INTR_TX.BLOCKED and INTR_RX.BLOCKED fields to determine whether a failure occurred for a series of write or read operations.

Table 27-3. SPI Modes Compatibility

	Internally clocked (IC)			Externally clocked (EC) (Deep Sleep SCB only)		
	FIFO	EZ	CMD_RESP	FIFO	EZ	CMD_RESP
SPI master	Yes	No	No	No	No	No
SPI slave	Yes	Yes	No	Yes ^a	Yes	Yes

a. In SPI slave FIFO mode, the external-clocked logic does selection detection, then triggers an interrupt to wake up the CPU. Writes will be ignored and reads will return 0xFF until the CPU is ready and the FIFO is populated.

27.3.5.2 Using SPI Master to Clock Slave

In a normal SPI Master mode transmission, the SCLK is generated only when the SCB is enabled and data is being transmitted. This can be changed to always generate a clock on the SCLK line while the SCB is enabled. This is used when the slave uses the SCLK for functional operations other than just the SPI functionality. To enable this, write '1' to the SCLK_CONTINUOUS (bit 5) of the SCB_SPI_CTRL register.

27.3.5.3 Oversampling and Bit Rate

SPI Master Mode

The SPI master does not support externally clocked mode. In internally clocked mode, the logic operates under internal clock. The internal clock has higher frequency than the interface clock (SCLK), such that the master can oversample its input signals (MISO).

The OVS (bits [3:0]) of the SCB_CTRL register specify the oversampling. The oversampling rate is calculated as the value in OVS register + 1. In SPI master mode, the valid range for oversampling is 4 to 16, when MISO is used; if MISO is not used then the valid range is 2 to 16. The bit rate is calculated as follows.

$$\text{Bit Rate} = \text{clk_scb} / \text{OVS}$$

Equation 27-1

Hence, with clk_scb at 100 MHz, the maximum bit rate is 25 Mbps with MISO, or 50 Mbps without MISO.

The numbers above indicate how fast the SCB hardware can run SCLK. It does not indicate that the master will be able to correctly receive data from a slave at those speeds. To determine that, the path delay of MISO must be calculated. It can be calculated using the following equation:

$$\frac{1}{2}t_{SCLK} \geq t_{SCLK_PCB_D} + t_{DSO} + t_{SCLK_PCB_D} + t_{DSI}$$

Equation 27-2

Where:

t_{SCLK} is the period of the SPI clock

$t_{SCLK_PCB_D}$ is the SCLK PCB delay from master to slave


t_{DSO} is the total internal slave delay, time from SCLK edge at slave pin to MISO edge at slave pin

$t_{SCLK_PCB_D}$ is the MISO PCB delay from slave to master

t_{DSI} is the master setup time

Most slave datasheets will list t_{DSO} . It may have a different name; look for MISO output valid after SCLK edge. Most master datasheets will also list t_{DSI} , or master setup time. $t_{SCLK_PCB_D}$ and $t_{SCLK_PCB_D}$ must be calculated based on specific PCB geometries.

If after doing these calculations the desired speed cannot be achieved, then consider using the MISO late sample feature of the SCB. This can be done by setting the SPI_CTRL.LATE_MISO_SAMPLE register. MISO late sample tells the SCB to sample the incoming MISO signal on the next edge of SCLK, thus allowing for ½ SCLK cycle more timing margin, see Figure 27-10.



This changes the equation to:

$$t_{SCLK} \geq t_{SCLK_PCB_D} + t_{DSO} + t_{SCLK_PCB_D} + t_{DSI} \quad \text{Equation 27-3}$$

Because late sample allows for better timing, leave it enabled all the time. The t_{DSI} specification in the PSoC 6 MCU datasheet assumes late sample is enabled.

Note: The SPI_CTRL.LATE_MISO_SAMPLE is set to '0' by default.

SPI Slave Mode

In SPI slave mode, the OVS field (bits [3:0]) of SCB_CTRL register is not used. The data rate is determined by Equation 26-2 and Equation 26-3. Late MISO sample is determined by the external master in this case, not by SPI_CTRL.LATE_MISO_SAMPLE.

For PSoC 6 MCUs, t_{DSO} is given in the [device datasheet](#). For internally-clocked mode, it is proportional to the frequency of the internal clock. For example it may be $20 \text{ ns} + 3 * t_{CLK_SCB}$. Assuming 0 ns PCB delays, and a 0 ns external master t_{DSI} Equation 26-1 can be re-arranged to $t_{CLK_SCB} \leq ((t_{SCLK}) - 40 \text{ ns})/6$.

27.3.6 Enabling and Initializing SPI

The SPI must be programmed in the following order:

1. Program protocol specific information using the SCB_SPI_CTRL register. This includes selecting the sub-modes of the protocol and selecting master-slave functionality. EZSPI and CMD_RESP can be used with slave mode only.
2. Program the OVS field and configure clk_scb as appropriate. See the [Clocking System chapter on page 207](#) for more information on how to program clocks and connect it to the SCB.
3. Configure SPI GPIO by setting appropriate drive modes and HSIOM settings.
4. Select the desired Slave Select line and polarity in the SCB_SPI_CTRL register.
5. Program the generic transmitter and receiver information using the SCB_TX_CTRL and SCB_RX_CTRL registers:
 - a. Specify the data frame width. This should always be 8 for EZSPI and CMD_RESP.
 - b. Specify whether MSb or LSb is the first bit to be transmitted/received. This should always be MSb first for EZSPI and CMD_RESP.
6. Program the transmitter and receiver FIFOs using the SCB_TX_FIFO_CTRL and SCB_RX_FIFO_CTRL registers respectively, as shown in SCB_TX_FIFO_CTRL/SCB_RX_FIFO_CTRL registers. Only for FIFO mode.
 - a. Set the trigger level.
 - b. Clear the transmitter and receiver FIFO and Shift registers.
7. Enable the block (write a '1' to the ENABLED bit of the SCB_CTRL register). After the block is enabled, control bits should not be changed. Changes should be made after disabling the block; for example, to modify the operation mode (from Motorola mode to TI mode) or to go from externally clocked to internally clocked operation. The change takes effect only

after the block is re-enabled. **Note:** Re-enabling the block causes re-initialization and the associated state is lost (for example, FIFO content).

27.3.7 I/O Pad Connection

27.3.7.1 SPI Master

Figure 27-11 and Table 27-4 list the use of the I/O pads for SPI Master.

Figure 27-11. SPI Master I/O Pad Connections




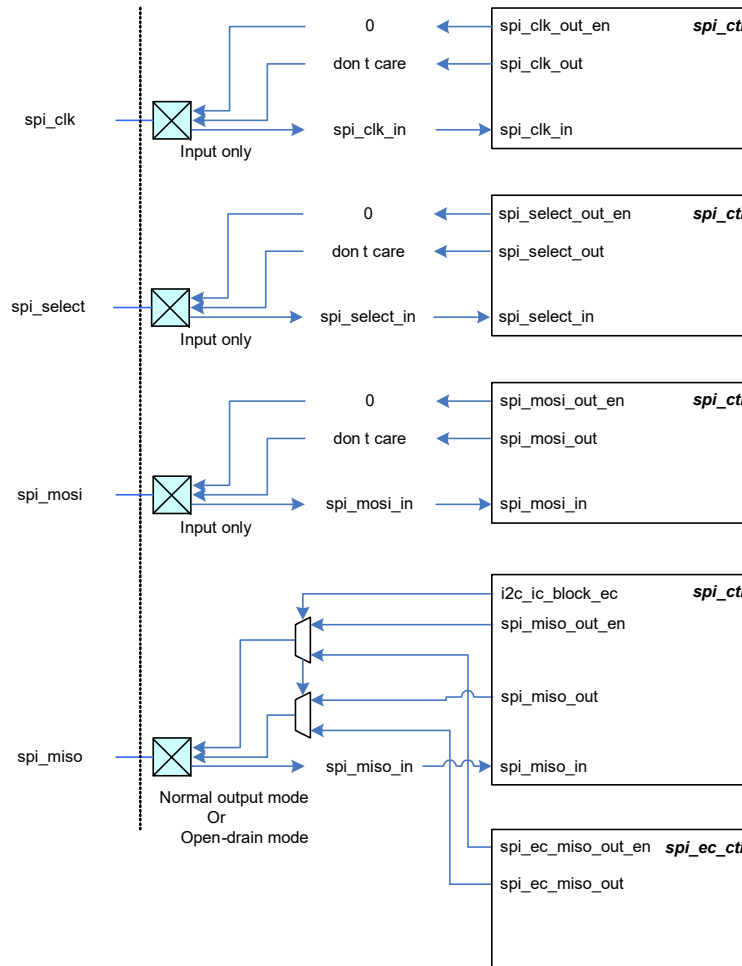
Table 27-4. SPI Master I/O Pad Connection Usage

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
spi_clk	Normal output mode	spi_clk_out_en spi_clk_out	Transmit a clock signal
spi_select	Normal output mode	spi_select_out_en spi_select_out	Transmit a select signal
spi_mosi	Normal output mode	spi_mosi_out_en spi_mosi_out	Transmit a data element
spi_miso	Input only	spi_miso_in	Receive a data element

27.3.7.2 SPI Slave

Figure 27-12 and Table 27-5 list the use of I/O pads for SPI Slave.

Figure 27-12. SPI Slave I/O Pad Connections



Open_Drain is set in the TX_CTRL register. In this mode the SPI MISO pin is actively driven low, and then high-z for driving high. This means an external pull-up is required for the line to go high. This mode is useful when there are multiple slaves on the same line. This helps to avoid bus contention issues.

Table 27-5. SPI Slave I/O Signal Description

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
spi_clk	Input mode	spi_clk_in	Receive a clock signal
spi_select	Input mode	spi_select_in	Receive a select signal
spi_mosi	Input mode	spi_mosi_in	Receive a data element
spi_miso	Normal output mode	spi_miso_out_en spi_miso_out	Transmit a data element

27.3.7.3 Glitch Avoidance at System Reset

The SPI outputs are in high-impedance digital state when the device is coming out of system reset. This can cause glitches on the outputs. This is important if you are concerned with SPI master SS0 – SS3 or SCLK output pins activity at either device startup or when coming out of Hibernate mode. External pull-up or pull-down resistor can be connected to the output pin to keep it in the inactive state.

27.3.8 SPI Registers

The SPI interface is controlled using a set of 32-bit control and status registers listed in [Table 27-6](#). Some of these registers are used specifically with the Deep Sleep SCB; for more information on these registers, see the [registers TRM](#).

Table 27-6. SPI Registers

Register Name	Operation
SCB_CTRL	Enables the SCB, selects the type of serial interface (SPI, UART, I ² C), and selects internally and externally clocked operation, and EZ and non-EZ modes of operation.
SCB_STATUS (Deep Sleep SCB only)	In EZ mode, this register indicates whether the externally clocked logic is potentially using the EZ memory.
SCB_SPI_CTRL	Configures the SPI as either a master or a slave, selects SPI protocols (Motorola, TI, National) and clock-based submodes in Motorola SPI (modes 0,1,2,3), selects the type of \overline{SS} signal in TI SPI.
SCB_SPI_STATUS	Indicates whether the SPI bus is busy and sets the SPI slave EZ address in the internally clocked mode.
SCB_TX_CTRL	Specifies the data frame width and specifies whether MSb or LSb is the first bit in transmission.
SCB_RX_CTRL	Performs the same function as that of the SCB_TX_CTRL register, but for the receiver. Also decides whether a median filter is to be used on the input interface lines.
SCB_TX_FIFO_CTRL	Specifies the trigger level, clears the transmitter FIFO and shift registers, and performs the FREEZE operation of the transmitter FIFO.
SCB_RX_FIFO_CTRL	Performs the same function as that of the SCB_TX_FIFO_CTRL register, but for the receiver.
SCB_TX_FIFO_WR	Holds the data frame written into the transmitter FIFO. Behavior is similar to that of a PUSH operation.
SCB_RX_FIFO_RD	Holds the data frame read from the receiver FIFO. Reading a data frame removes the data frame from the FIFO - behavior is similar to that of a POP operation. This register has a side effect when read by software: a data frame is removed from the FIFO.
SCB_RX_FIFO_RD_SILENT	Holds the data frame read from the receiver FIFO. Reading a data frame does not remove the data frame from the FIFO; behavior is similar to that of a PEEK operation.
SCB_TX_FIFO_STATUS	Indicates the number of bytes stored in the transmitter FIFO, the location from which a data frame is read by the hardware (read pointer), the location from which a new data frame is written (write pointer), and decides whether the transmitter FIFO holds the valid data.
SCB_RX_FIFO_STATUS	Performs the same function as that of the SCB_TX_FIFO_STATUS register, but for the receiver.
SCB_EZ_DATA (Deep Sleep SCB only)	Holds the data in EZ memory location

27.4 UART

The Universal Asynchronous Receiver/Transmitter (UART) protocol is an asynchronous serial interface protocol. UART communication is typically point-to-point. The UART interface consists of two signals:

- TX: Transmitter output
- RX: Receiver input

Additionally, two side-band signals are used to implement flow control in UART. Note that the flow control applies only to TX functionality.

- Clear to Send (CTS): This is an input signal to the transmitter. When active, the receiver signals to the transmitter that it is ready to receive.
- Ready to Send (RTS): This is an output signal from the receiver. When active, it indicates that the receiver is ready to receive data.

Not all SCBs support UART mode; refer to the [device datasheet](#) for details.


27.4.1 Features

- Supports UART protocol
 - Standard UART
 - Multi-processor mode
- SmartCard (ISO7816) reader
- IrDA
- Supports Local Interconnect Network (LIN)
 - Break detection
 - Baud rate detection
 - Collision detection (ability to detect that a driven bit value is not reflected on the bus, indicating that another component is driving the same bus)
- Data frame size programmable from 4 to 16 bits
- Programmable number of STOP bits, which can be set in terms of half bit periods between 1 and 4
- Parity support (odd and even parity)
- Median filter on RX input
- Programmable oversampling
- Start skipping
- Hardware flow control

27.4.2 General Description

Figure 27-13 illustrates a standard UART TX and RX.

Figure 27-13. UART Example



A typical UART transfer consists of a start bit followed by multiple data bits, optionally followed by a parity bit and finally completed by one or more stop bits. The start and stop bits indicate the start and end of data transmission. The parity bit is sent by the transmitter and is used by the receiver to detect single bit errors. Because the interface does not have a clock (asynchronous), the transmitter and receiver use their own clocks; thus, the transmitter and receiver need to agree on the baud rate.

By default, UART supports a data frame width of eight bits. However, this can be configured to any value in the range of 4 to 9. This does not include start, stop, and parity bits. The number of stop bits can be in the range of 1 to 4. The parity bit can be either enabled or disabled. If enabled, the type of parity can be set to either even parity or odd parity. The option of using the parity bit is available only in the Standard UART and SmartCard UART modes. For IrDA UART mode, the parity bit is automatically disabled.

Note: UART interface does not support external clocking operation. Hence, UART operates only in the Active and Sleep system power modes. UART also supports only the FIFO buffer mode.

27.4.3 UART Modes of Operation

27.4.3.1 Standard Protocol

A typical UART transfer consists of a start bit followed by multiple data bits, optionally followed by a parity bit and finally completed by one or more stop bits. The start bit value is always '0', the data bits values are dependent on the data transferred, the parity bit value is set to a value guaranteeing an even or odd parity over the data bits, and the stop bit value is '1'. The parity bit is generated by the transmitter and can be used by the receiver to detect single bit transmission errors. When not transmitting data, the TX line is '1' – the same value as the stop bits.


Because the interface does not have a clock, the transmitter and receiver must agree upon the baud rate. The transmitter and receiver have their own internal clocks. The receiver clock runs at a higher frequency than the bit transfer frequency, such that the receiver may oversample the incoming signal.

The transition of a stop bit to a start bit is represented by a change from '1' to '0' on the TX line. This transition can be used by the receiver to synchronize with the transmitter clock. Synchronization at the start of each data transfer allows error-free transmission even in the presence of frequency drift between transmitter and receiver clocks. The required clock accuracy is dependent on the data transfer size.

The stop period or the amount of stop bits between successive data transfers is typically agreed upon between transmitter and receiver, and is typically in the range of 1 to 3-bit transfer periods.


Figure 27-14 illustrates the UART protocol.

Figure 27-14. UART, Standard Protocol Example




The receiver oversamples the incoming signal; the value of the sample point in the middle of the bit transfer period (on the receiver's clock) is used. Figure 27-15 illustrates this.

Figure 27-15. UART, Standard Protocol Example (Single Sample)



Alternatively, three samples around the middle of the bit transfer period (on the receiver's clock) are used for a majority vote to increase accuracy; this is enabled by enabling the MEDIAN filter in the SCB_RX_CTRL register. Figure 27-16 illustrates this.

Figure 27-16. UART, Standard Protocol (Multiple Samples)



Parity

This functionality adds a parity bit to the data frame and is used to identify single-bit data frame errors. The parity bit is always directly after the data frame bits.

The transmitter calculates the parity bit (when `UART_TX_CTRL.PARITY_ENABLED` is 1) from the data frame bits, such that data frame bits and parity bit have an even (`UART_TX_CTRL.PARITY` is 0) or odd (`UART_TX_CTRL.PARITY` is 1) parity. The receiver checks the parity bit (when `UART_RX_CTRL.PARITY_ENABLED` is 1) from the received data frame bits, such that data frame bits and parity bit have an even (`UART_RX_CTRL.PARITY` is 0) or odd (`UART_RX_CTRL.PARITY` is 1) parity.


Parity applies to both TX and RX functionality and dedicated control fields are available.

- Transmit functionality: `UART_TX_CTRL.PARITY` and `UART_TX_CTRL.PARITY_ENABLED`.
- Receive functionality: `UART_RX_CTRL.PARITY` and `UART_RX_CTRL.PARITY_ENABLED`.

When a receiver detects a parity error, the data frame is either put in RX FIFO (`UART_RX_CTRL.DROP_ON_PARITY_ERROR` is 0) or dropped (`UART_RX_CTRL.DROP_ON_PARITY_ERROR` is 1).

The following figures illustrate the parity functionality (8-bit data frame).

Figure 27-17. UART Parity Examples



Start Skipping

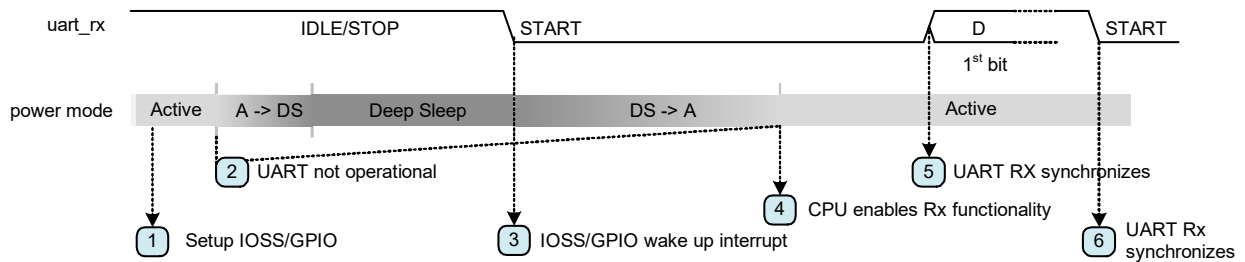
Start skipping applies only to receive functionality. The standard UART mode supports “start skipping”. Regular receive operation synchronizes on the START bit period (a 1 to 0 transition on the UART RX line), start skipping receive operation synchronizes on the first received data frame bit, which must be a ‘1’ (a 0 to 1 transition on UART RX).

Start skipping is used to allow for wake up from system Deep Sleep mode using UART. The process is described as follows:

1. Before entering Deep Sleep power mode, UART receive functionality is disabled and the GPIO is programmed to set an interrupt cause to ‘1’ when UART RX line has a ‘1’ to ‘0’ transition (START bit).
2. While in Deep Sleep mode, the UART receive functionality is not functional.
3. The GPIO interrupt is activated on the START bit and the system transitions from Deep Sleep to Active power mode.
4. The CPU enables UART receive functionality, with `UART_RX_CTRL.SKIP_START` bitfield set to ‘1’.
5. The UART receiver synchronizes data frame receipt on the next ‘0’ to ‘1’ transition. If the UART receive functionality is enabled in time, this is the transition from the START bit to the first received data frame bit.
6. The UART receiver proceeds with normal operation; that is, synchronization of successive data frames is on the START bit period.

Figure 27-18 illustrates the process.

Figure 27-18. UART Start Skip and Wakeup from Deep Sleep



Note that the above process works only for lower baud rates. The Deep Sleep to Active power mode transition and CPU enabling the UART receive functionality should take less than 1-bit period to ensure that the UART receiver is active in time to detect the '0' to '1' transition.

In step 4 of the above process, the firmware takes some time to finish the wakeup interrupt routine and enable the UART receive functionality before the block can detect the input rising edge on the UART RX line.

If the above steps cannot be completed in less than 1 bit time, first send a “dummy” byte to the device to wake it up before sending real UART data. In this case, the SKIP_START bit can be left as 0. For more information on how to perform this in firmware, visit the UART section of the [PDL](#).

Break Detection


Break detection is supported in the standard UART mode. This functionality detects when UART RX line is low (0) for more than UART_RX_CTRL.BREAK_WIDTH bit periods. The break width should be larger than the maximum number of low (0) bit periods in a regular data transfer, plus an additional 1-bit period. The additional 1-bit period is a minimum requirement and preferably should be larger. The additional bit periods account for clock inaccuracies between transmitter and receiver.

For example, for an 8-bit data frame with parity support, the maximum number of low (0) bit periods is 10 (START bit, 8 '0' data frame bits, and one '0' parity bit). Therefore, the break width should be larger than $10 + 1 = 11$ (UART_RX_CTRL.BREAK_WIDTH can be set to 11).

Note that the break detection applies only to receive functionality. A UART transmitter can generate a break by temporarily increasing TX_CTRL.DATA_WIDTH and transmitting an all zeroes data frame. A break is used by the transmitter to signal a special condition to the receiver. This condition may result in a reset, shut down, or initialization sequence at the receiver.

Break detection is part of the LIN protocol. When a break is detected, the INTR_RX.BREAK_DETECT interrupt cause is set to '1'. [Figure 27-19](#) illustrates a regular data frame and break frame (8-bit data frame, parity support, and a break width of 12-bit periods).

Figure 27-19. UART – Regular Frame and Break Frame



Flow Control

The standard UART mode supports flow control. Modem flow control controls the pace at which the transmitter transfers data to the receiver. Modem flow control is enabled through the `UART_FLOW_CTRL.CTS_ENABLED` register field. When this field is '0', the transmitter transfers data when its TX FIFO is not empty. When '1', the transmitter transfers data when UART CTS line is active and its TX FIFO is not empty.


Note that the flow control applies only to TX functionality. Two UART side-band signal are used to implement flow control:

- **UART RTS (uart_rts_out):** This is an output signal from the receiver. When active, it indicates that the receiver is ready to receive data (RTS: Ready to Send).
- **UART CTS (uart_cts_in):** This is an input signal to the transmitter. When active, it indicates that the transmitter can transfer data (CTS: Clear to Send).

The receiver's `uart_rts_out` signal is connected to the transmitter's `uart_cts_in` signal. The receiver's `uart_rts_out` signal is derived by comparing the number of used receive FIFO entries with the `UART_FLOW_CTRL.TRIGGER_LEVEL` field. If the number of used receive FIFO entries are less than `UART_FLOW_CTRL.TRIGGER_LEVEL`, `uart_rts_out` is activated.

Typically, the UART side-band signals are active low. However, sometimes active high signaling is used. Therefore, the polarity of the side-band signals can be controlled using bitfields `UART_FLOW_CTRL.RTS_POLARITY` and `UART_FLOW_CTRL.CTS_POLARITY`. Figure 27-20 gives an overview of the flow control functionality.


Figure 27-20. UART Flow Control Connection



UART Multi-Processor Mode

The `UART_MP` (multi-processor) mode is defined with single-master-multi-slave topology, as Figure 27-21 shows. This mode is also known as UART 9-bit protocol because the data field is nine bits wide. `UART_MP` is part of Standard UART mode.


Figure 27-21. UART MP Mode Bus Connections



The main properties of `UART_MP` mode are:

- Single master with multiple slave concept (multi-drop network).
- Each slave is identified by a unique address.
- Using 9-bit data field, with the ninth bit as address/data flag (MP bit). When set high, it indicates an address byte; when set low it indicates a data byte. A data frame is illustrated in Figure 27-22.
- Parity bit is disabled.

Figure 27-22. UART MP Address and Data Frame



The SCB can be used as either master or slave device in UART_MP mode. Both SCB_TX_CTRL and SCB_RX_CTRL registers should be set to 9-bit data frame size. When the SCB works as UART_MP master device, the firmware changes the MP flag for every address or data frame. When it works as UART_MP slave device, the MP_MODE field of the SCB_UART_RX_CTRL register should be set to '1'. The SCB_RX_MATCH register should be set for the slave address and address mask. The matched address is written in the RX_FIFO when ADDR_ACCEPT field of the SCB_CTRL register is set to '1'. If the received address does not match its own address, then the interface ignores the following data, until next address is received for compare.

Configuring the SCB as Standard UART Interface

To configure the SCB as a standard UART interface, set various register bits in the following order:


1. Configure the SCB as UART interface by writing '10b' to the MODE field (bits [25:24]) of the SCB_CTRL register.
2. Configure the UART interface to operate as a Standard protocol by writing '00' to the MODE field (bits [25:24]) of the SCB_UART_CTRL register.
3. To enable the UART MP Mode or UART LIN Mode, write '1' to the MP_MODE (bit 10) or LIN_MODE (bit 12) respectively of the SCB_UART_RX_CTRL register.
4. Follow steps 2 to 4 described in ["Enabling and Initializing the UART" on page 314](#).

For more information on these registers, see the [registers TRM](#).

27.4.3.2 UART Local Interconnect Network (LIN) Mode

The LIN protocol is supported by the SCB as part of the standard UART. LIN is designed with single-master-multi-slave topology. There is one master node and multiple slave nodes on the LIN bus. The SCB UART supports both LIN master and slave functionality. The LIN specification defines both physical layer (layer 1) and data link layer (layer 2). [Figure 27-23](#) illustrates the UART_LIN and LIN transceiver.

Figure 27-23. UART_LIN and LIN Transceiver




LIN protocol defines two tasks:

- Master task: This task involves sending a header packet to initiate a LIN transfer.
- Slave task: This task involves transmitting or receiving a response.

The master node supports master task and slave task; the slave node supports only slave task, as shown in [Figure 27-24](#).

Figure 27-24. LIN Bus Nodes and Tasks




LIN Frame Structure

LIN is based on the transmission of frames at pre-determined moments of time. A frame is divided into header and response fields, as shown in Figure 27-25.

- The header field consists of:
 - Break field (at least 13 bit periods with the value '0').
 - Sync field (a 0x55 byte frame). A sync field can be used to synchronize the clock of the slave task with that of the master task.
 - Identifier field (a frame specifying a specific slave).
- The response field consists of data and checksum.

Figure 27-25. LIN Frame Structure



In LIN protocol communication, the least significant bit (LSb) of the data is sent first and the most significant bit (MSb) last. The start bit is encoded as zero and the stop bit is encoded as one. The following sections describe all the byte fields in the LIN frame.

Break Field

Every new frame starts with a break field, which is always generated by the master. The break field has logical zero with a minimum of 13 bit times and followed by a break delimiter. The break field structure is as shown in Figure 27-26.

Figure 27-26. LIN Break Field



Sync Field

This is the second field transmitted by the master in the header field; its value is 0x55. A sync field can be used to synchronize the clock of the slave task with that of the master task for automatic baud rate detection. Figure 27-27 shows the LIN sync field structure.

Figure 27-27. LIN Sync Field




Protected Identifier (PID) Field

A protected identifier field consists of two sub-fields: the frame identifier (bits 0-5) and the parity (bit 6 and bit 7). The PID field structure is shown in [Figure 27-28](#).

- Frame identifier: The frame identifiers are divided into three categories
 - Values 0 to 59 (0x3B) are used for signal carrying frames
 - 60 (0x3C) and 61 (0x3D) are used to carry diagnostic and configuration data
 - 62 (0x3E) and 63 (0x3F) are reserved for future protocol enhancements
- Parity: Frame identifier bits are used to calculate the parity

[Figure 27-28](#) shows the PID field structure.

Figure 27-28. PID Field



Data. In LIN, every frame can carry a minimum of one byte and maximum of eight bytes of data. Here, the LSb of the data byte is sent first and the MSb of the data byte is sent last.

Checksum. The checksum is the last byte field in the LIN frame. It is calculated by inverting the 8-bit sum along with carryover of all data bytes only or the 8-bit sum with the carryover of all data bytes and the PID field. There are two types of checksums in LIN frames. They are:

- Classic checksum: the checksum calculated over all the data bytes only (used in LIN 1.x slaves).
- Enhanced checksum: the checksum calculated over all the data bytes along with the protected identifier (used in LIN 2.x slaves).

LIN Frame Types

The type of frame refers to the conditions that need to be valid to transmit the frame. According to the LIN specification, there are five different types of LIN frames. A node or cluster does not have to support all frame types.

Unconditional Frame. These frames carry the signals and their frame identifiers (of 0x00 to 0x3B range). The subscriber will receive the frames and make it available to the application; the publisher of the frame will provide the response to the header.

Event-Triggered Frame. The purpose of an event-triggered frame is to increase the responsiveness of the LIN cluster without assigning too much of the bus bandwidth to polling of multiple slave nodes with seldom occurring events. Event-triggered frames carry the response of one or more

unconditional frames. The unconditional frames associated with an event triggered frame should:

- Have equal length
- Use the same checksum model (either classic or enhanced)
- Reserve the first data field to its protected identifier
- Be published by different slave nodes
- Not be included directly in the same schedule table as the event-triggered frame

Sporadic Frame. The purpose of the sporadic frames is to merge some dynamic behavior into the schedule table without affecting the rest of the schedule table. These frames have a group of unconditional frames that share the frame slot. When the sporadic frame is due for transmission, the unconditional frames are checked whether they have any updated signals. If no signals are updated, no frame will be transmitted and the frame slot will be empty.

Diagnostic Frames. Diagnostic frames always carry transport layer, and contains eight data bytes.

The frame identifier for diagnostic frame is:

- Master request frame (0x3C), or
- Slave response frame (0x3D)

Before transmitting a master request frame, the master task queries its diagnostic module to see whether it will be transmitted or whether the bus will be silent. A slave response frame header will be sent unconditionally. The slave tasks publish and subscribe to the response according to their diagnostic modules.

Reserved Frames. These frames are reserved for future use; their frame identifiers are 0x3E and 0x3F.

LIN Go-To-Sleep and Wake-Up

The LIN protocol has the feature of keeping the LIN bus in Sleep mode, if the master sends the go-to-sleep command. The go-to-sleep command is a master request frame (ID = 0x3C) with the first byte field is equal to 0x00 and rest set to 0xFF. The slave node application may still be active after the go-to-sleep command is received. This behavior is application specific. The LIN slave nodes automatically enter Sleep mode if the LIN bus inactivity is more than four seconds.

Wake-up can be initiated by any node connected to the LIN bus – either LIN master or any of the LIN slaves by forcing the bus to be dominant for 250 μs to 5 ms. Each slave should detect the wakeup request and be ready to process headers within 100 ms. The master should also detect the wakeup request and start sending headers when the slave nodes are active.

To support LIN, a dedicated (off-chip) line driver/receiver is required. Supply voltage range on the LIN bus is 7 V to 18 V.

Typically, LIN line drivers will drive the LIN line with the value provided on the SCB TX line and present the value on the LIN line to the SCB RX line. By comparing TX and RX lines in the SCB, bus collisions can be detected (indicated by the SCB_UART_ARB_LOST field of the SCB_INTR_TX register).

27.4.3.3 SmartCard (ISO7816)

ISO7816 is asynchronous serial interface, defined with single-master-single slave topology. ISO7816 defines both Reader (master) and Card (slave) functionality. For more information, refer to the [ISO7816 Specification](#). Only the master (reader) function is supported by the SCB. This block provides the basic physical layer support with asynchronous character transmission. The UART_TX line is connected to SmartCard I/O line by internally multiplexing between UART_TX and UART_RX control modules.


The SmartCard transfer is similar to a UART transfer, with the addition of a negative acknowledgement (NACK) that

may be sent from the receiver to the transmitter. A NACK is always '0'. Both master and slave may drive the same line, although never at the same time.

A SmartCard transfer has the transmitter drive the start bit and data bits (and optionally a parity bit). After these bits, it enters its stop period by releasing the bus. Releasing results in the line being '1' (the value of a stop bit). After one bit transfer period into the stop period, the receiver may drive a NACK on the line (a value of '0') for one bit transfer period. This NACK is observed by the transmitter, which reacts by extending its stop period by one bit transfer period. For this protocol to work, the stop period should be longer than one bit transfer period. Note that a data transfer with a NACK takes one bit transfer period longer, than a data transfer without a NACK. Typically, implementations use a tristate driver with a pull-up resistor, such that when the line is not transmitting data or transmitting the Stop bit, its value is '1'.

[Figure 27-29](#) illustrates the SmartCard protocol.

Figure 27-29. SmartCard Example



The communication Baud rate while using SmartCard is given as:

$$\text{Baud rate} = \text{Fscbclk} / \text{Oversample}$$

Configuring SCB as UART SmartCard Interface

To configure the SCB as a UART SmartCard interface, set various register bits in the following order; note that ModusToolbox does all this automatically with the help of GUIs. For more information on these registers, see the [registers TRM](#).

1. Configure the SCB as UART interface by writing '10b' to the MODE (bits [25:24]) of the SCB_CTRL register.
2. Configure the UART interface to operate as a SmartCard protocol by writing '01' to the MODE (bits [25:24]) of the SCB_UART_CTRL register.
3. Follow steps 2 to 4 described in ["Enabling and Initializing the UART" on page 314](#).

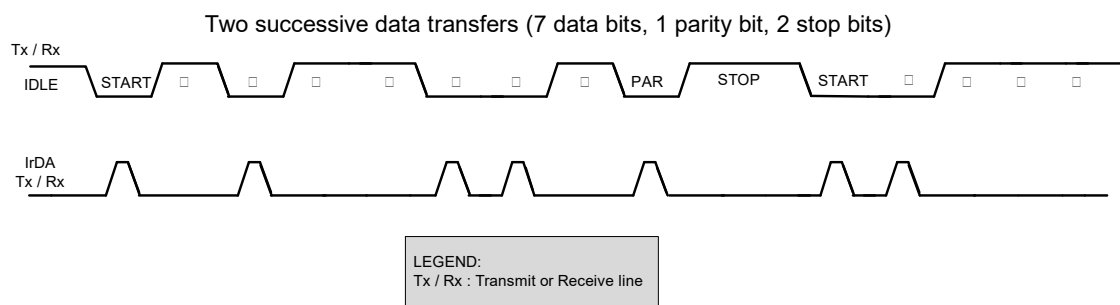
27.4.3.4 Infrared Data Association (IrDA)

The SCB supports the IrDA protocol for data rates of up to 115.2 kbps using the UART interface. It supports only the basic physical layer of IrDA protocol with rates less than 115.2 kbps. Hence, the system instantiating this block must consider how to implement a complete IrDA communication system with other available system resources.

The IrDA protocol adds a modulation scheme to the UART signaling. At the transmitter, bits are modulated. At the receiver, bits are demodulated. The modulation scheme uses a Return-to-Zero-Inverted (RZI) format. A bit value of '0' is signaled by a short '1' pulse on the line and a bit value of '1' is signaled by holding the line to '0'. For these data rates (≤ 115.2 kbps), the RZI modulation scheme is used and the pulse duration is 3/16 of the bit period. The sampling clock frequency should be set 16 times the selected baud rate, by configuring the SCB_OVS field of the SCB_CTRL register. In addition, the PSoC 6 MCU SCB supports a low-power IrDA receiver mode, which allows it to detect pulses with a minimum width of 1.41 μ s.

Different communication speeds under 115.2 kbps can be achieved by configuring clk_scb frequency. Additional allowable rates are 2.4 kbps, 9.6 kbps, 19.2 kbps, 38.4 kbps, and 57.6 kbps. [Figure 27-30](#) shows how a UART transfer is IrDA modulated.

Figure 27-30. IrDA Example



Configuring the SCB as a UART IrDA Interface

To configure the SCB as a UART IrDA interface, set various register bits in the following order; note that ModusToolbox does all this automatically with the help of GUIs. For more information on these registers, see the [registers TRM](#).

1. Configure the SCB as a UART interface by writing '10b' to the MODE (bits [25:24]) of the SCB_CTRL register.
2. Configure the UART interface to operate as IrDA protocol by writing '10' to the MODE (bits [25:24]) of the SCB_UART_CTRL register.
3. Enable the Median filter on the input interface line by writing '1' to MEDIAN (bit 9) of the SCB_RX_CTRL register.
4. Configure the SCB as described in ["Enabling and Initializing the UART" on page 314](#).

27.4.4 Clocking and Oversampling

The UART protocol is implemented using `clk_scb` as an oversampled multiple of the baud rate. For example, to implement a 100-kHz UART, `clk_scb` could be set to 1 MHz and the oversample factor set to '10'. The oversampling is set using the `SCB_CTRL.OVS` register field. The oversampling value is `SCB_CTRL.OVS + 1`. In the UART standard sub-mode (including LIN) and the SmartCard sub-mode, the valid range for the OVS field is [7, 15].

In UART transmit IrDA sub-mode, this field indirectly specifies the oversampling. Oversampling determines the interface clock per bit cycle and the width of the pulse. This sub-mode has only one valid OVS value—16 (which is a value of 0 in the OVS field of the `SCB_CTRL` register); the pulse width is roughly 3/16 of the bit period (for all bit rates).

In UART receive IrDA sub-mode (1.2, 2.4, 9.6, 19.2, 38.4, 57.6, and 115.2 kbps), this field indirectly specifies the oversampling. In normal transmission mode, this pulse is approximately 3/16 of the bit period (for all bit rates). In low-power transmission mode, this pulse is potentially smaller (down to 1.62 μ s typical and 1.41 μ s minimal) than 3/16 of the bit period (for less than 115.2 kbps bit rates).

Pulse widths greater than or equal to two SCB input clock cycles are guaranteed to be detected by the receiver. Pulse widths less than two clock cycles and greater than or equal to one SCB input clock cycle may be detected by the receiver. Pulse widths less than one SCB input clock cycle will not be detected by the receiver. Note that the `SCB_RX_CTRL.MEDIAN` should be set to '1' for IrDA receiver functionality.

The SCB input clock and the oversampling together determine the IrDA bit rate. Refer to the [registers TRM](#) for more details on the OVS values for different baud rates.

27.4.5 Enabling and Initializing the UART

The UART must be programmed in the following order:

1. Program protocol specific information using the `UART_TX_CTRL`, `UART_RX_CTRL`, and `UART_FLOW_CTRL` registers. This includes selecting the submodes of the protocol, transmitter-receiver functionality, and so on.
2. Program the generic transmitter and receiver information using the `SCB_TX_CTRL` and `SCB_RX_CTRL` registers.
 - a. Specify the data frame width.
 - b. Specify whether MSb or LSb is the first bit to be transmitted or received.
3. Program the transmitter and receiver FIFOs using the `SCB_TX_FIFO_CTRL` and `SCB_RX_FIFO_CTRL` registers, respectively.
 - a. Set the trigger level.
 - b. Clear the transmitter and receiver FIFO and Shift registers.
4. Enable the block (write a '1' to the `ENABLE` bit of the `SCB_CTRL` register). After the block is enabled, control bits should not be changed. Changes should be made after disabling the block; for example, to modify the operation mode (from SmartCard to IrDA). The change takes effect only after the block is re-enabled. Note that re-enabling the block causes re-initialization and the associated state is lost (for example FIFO content).

27.4.6 I/O Pad Connection

27.4.6.1 Standard UART Mode

Figure 27-31 and Table 27-7 list the use of the I/O pads for the Standard UART mode.

Figure 27-31. Standard UART Mode I/O Pad Connections




Table 27-7. UART I/O Pad Connection Usage

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
uart_tx	Normal output mode	uart_tx_out_en uart_tx_out	Transmit a data element
uart_rx	Input only	uart_rx_in	Receive a data element

27.4.6.2 SmartCard Mode

Figure 27-32 and Table 27-8 list the use of the I/O pads for the SmartCard mode.

Figure 27-32. SmartCard Mode I/O Pad Connections




Table 27-8. SmartCard Mode I/O Pad Connections

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
uart_tx	Open drain with pull-up	uart_tx_in	Used to receive a data element. Receive a negative acknowledgement of a transmitted data element
		uart_tx_out_en uart_tx_out	Transmit a data element. Transmit a negative acknowledgement to a received data element.

27.4.6.3 LIN Mode

Figure 27-33 and Table 27-9 list the use of the I/O pads for LIN mode.

Figure 27-33. LIN Mode I/O Pad Connections

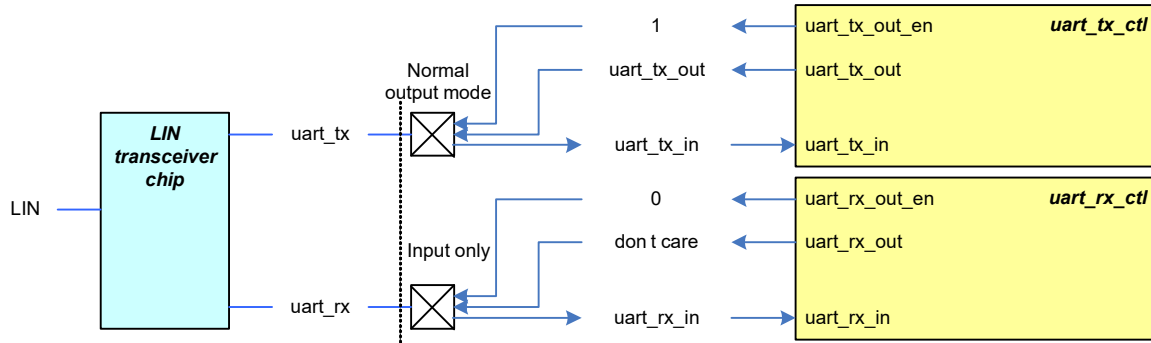


Table 27-9. LIN Mode I/O Pad Connections

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
uart_tx	Normal output mode	uart_tx_out_en uart_tx_out	Transmit a data element.
uart_rx	Input only	uart_rx_in	Receive a data element.

27.4.6.4 IrDA Mode

Figure 27-34 and Table 27-10 list the use of the I/O pads for IrDA mode.

Figure 27-34. IrDA Mode I/O Pad Connections

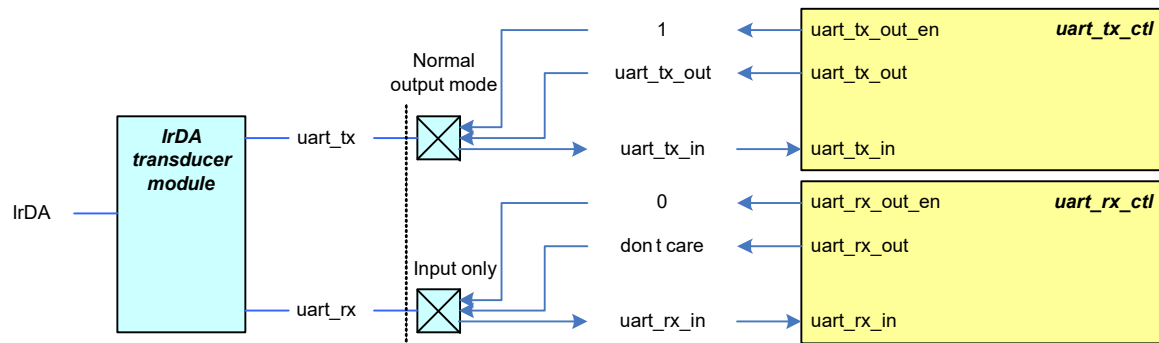


Table 27-10. IrDA Mode I/O Pad Connections

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
uart_tx	Normal output mode	uart_tx_out_en uart_tx_out	Transmit a data element.
uart_rx	Input only	uart_rx_in	Receive a data element.

27.4.7 UART Registers

The UART interface is controlled using a set of 32-bit registers listed in [Table 27-11](#). For more information on these registers, see the [registers TRM](#).

Table 27-11. UART Registers

Register Name	Operation
SCB_CTRL	Enables the SCB; selects the type of serial interface (SPI, UART, I ² C)
SCB_UART_CTRL	Used to select the sub-modes of UART (standard UART, SmartCard, IrDA), also used for local loop back control.
SCB_UART_RX_STATUS	Used to specify the BR_COUNTER value that determines the bit period. This is used to set the accuracy of the SCB clock. This value provides more granularity than the OVS bit in SCB_CTRL register.
SCB_UART_TX_CTRL	Used to specify the number of stop bits, enable parity, select the type of parity, and enable retransmission on NACK.
SCB_UART_RX_CTRL	Performs same function as SCB_UART_TX_CTRL but is also used for enabling multi processor mode, LIN mode drop on parity error, and drop on frame error.
SCB_TX_CTRL	Used to specify the data frame width and to specify whether MSb or LSb is the first bit in transmission.
SCB_RX_CTRL	Performs the same function as that of the SCB_TX_CTRL register, but for the receiver. Also decides whether a median filter is to be used on the input interface lines.
SCB_UART_FLOW_CONTROL	Configures flow control for UART transmitter.

27.5 Inter Integrated Circuit (I²C)

This section explains the I²C implementation in the PSoC 6 MCU. For more information on the I²C protocol specification, refer to the I²C-bus specification available on the [NXP website](#).

27.5.1 Features


This block supports the following features:

- Master, slave, and master/slave mode
- Standard-mode (100 kbps), fast-mode (400 kbps), and fast-mode plus (1000 kbps) data-rates
- 7-bit slave addressing
- Clock stretching
- Collision detection
- Programmable oversampling of I²C clock signal (SCL)
- Auto ACK when RX FIFO not full, including address
- General address detection
- FIFO Mode
- EZ and CMD_RESP modes

27.5.2 General Description

Figure 27-35 illustrates an example of an I²C communication network.

Figure 27-35. I²C Interface Block Diagram



The standard I²C bus is a two wire interface with the following lines:

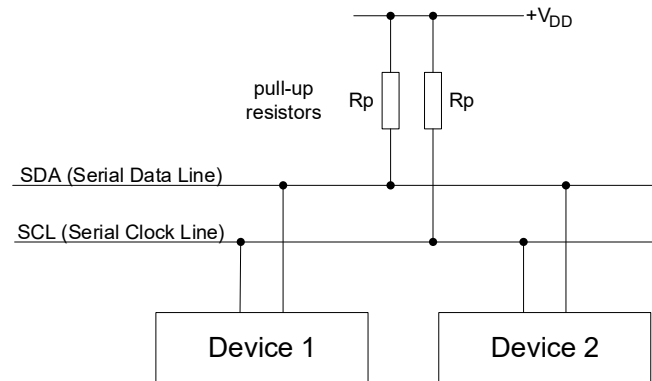
- Serial Data (SDA)
- Serial Clock (SCL)

I²C devices are connected to these lines using open collector or open-drain output stages, with pull-up resistors (Rp). A simple master/slave relationship exists between devices. Masters and slaves can operate as either transmitter or receiver. Each slave device connected to the bus is software addressable by a unique 7-bit address.

27.5.3 External Electrical Connections

As shown in Figure 27-36, the I²C bus requires external pull-up resistors. The pull-up resistors (R_P) are primarily determined by the supply voltage, bus speed, and bus capacitance. For detailed information on how to calculate the optimum pull-up resistor value for your design Cypress recommends using the UM10204 I²C-bus specification and user manual Rev. 6, available from the NXP website at www.nxp.com.

Figure 27-36. Connection of Devices to the I²C Bus



For most designs, the default values shown in Table 27-12 provide excellent performance without any calculations. The default values were chosen to use standard resistor values between the minimum and maximum limits.

Table 27-12. Recommended Default Pull-up Resistor Values

Standard Mode (0 – 100 kbps)	Fast Mode (0 – 400 kbps)	Fast Mode Plus (0 – 1000 kbps)	Units
4.7 k, 5%	1.74 k, 1%	620, 5%	Ω

These values work for designs with 1.8 V to 5.0 V V_{DD} , less than 200 pF bus capacitance (C_B), up to 25 μ A of total input leakage (I_{IL}), up to 0.4 V output voltage level (V_{OL}), and a max V_{IH} of $0.7 \times V_{DD}$. Calculation of custom pull-up resistor values is required if your design does not meet the default assumptions, you use series resistors (R_S) to limit injected noise, or you want to maximize the resistor value for low power consumption. Calculation of the ideal pull-up resistor value involves finding a value between the limits set by three equations detailed in the NXP I²C specification. These equations are:

$$R_{P\text{MIN}} = (V_{DD}(\text{max}) - V_{OL}(\text{max})) / I_{OL}(\text{min}) \quad \text{Equation 27-4}$$

$$R_{P\text{MAX}} = T_R(\text{max}) / 0.8473 \times C_B(\text{max}) \quad \text{Equation 27-5}$$

$$R_{P\text{MAX}} = V_{DD}(\text{min}) - (V_{IH}(\text{min}) + V_{NH}(\text{min})) / I_{IH}(\text{max}) \quad \text{Equation 27-6}$$

Equation parameters:

- V_{DD} = Nominal supply voltage for I²C bus
- V_{OL} = Maximum output low voltage of bus devices
- I_{OL} = Low-level output current from I²C specification
- T_R = Rise time of bus from I²C specification
- C_B = Capacitance of each bus line including pins and PCB traces
- V_{IH} = Minimum high-level input voltage of all bus devices
- V_{NH} = Minimum high-level input noise margin from I²C specification
- I_{IH} = Total input leakage current of all devices on the bus

The supply voltage (V_{DD}) limits the minimum pull-up resistor value due to bus devices maximum low output voltage (V_{OL}) specifications. Lower pull-up resistance increases current through the pins and can therefore exceed the spec conditions of V_{OH} . Equation 26-4 is derived using Ohm's law to determine the minimum resistance that will still meet the V_{OL} specification at 3 mA for standard and fast modes, and 20 mA for fast mode plus at the given V_{DD} .

Equation 26-5 determines the maximum pull-up resistance due to bus capacitance. Total bus capacitance is comprised of all pin, wire, and trace capacitance on the bus. The higher the bus capacitance the lower the pull-up resistance required to meet the specified bus speeds rise time due to RC delays. Choosing a pull-up resistance higher than allowed can result in failing timing requirements resulting in communication errors. Most designs with five or fewer I²C devices and up to 20 centimeters of bus trace length have less than 100 pF of bus capacitance.

A secondary effect that limits the maximum pull-up resistor value is total bus leakage calculated in Equation 26-6. The primary source of leakage is I/O pins connected to the bus. If leakage is too high, the pull-ups will have difficulty maintaining an acceptable V_{IH} level causing communication errors. Most designs with five or fewer I²C devices on the bus have less than 10 μ A of total leakage current.

27.5.4 Terms and Definitions

Table 27-13 explains the commonly used terms in an I²C communication network.

Table 27-13. Definition of I²C Bus Terminology

Term	Description
Transmitter	The device that sends data to the bus
Receiver	The device that receives data from the bus
Master	The device that initiates a transfer, generates clock signals, and terminates a transfer
Slave	The device addressed by a master
Multi-master	More than one master can attempt to control the bus at the same time
Arbitration	Procedure to ensure that, if more than one master simultaneously tries to control the bus, only one is allowed to do so and the winning message is not corrupted
Synchronization	Procedure to synchronize the clock signals of two or more devices

27.5.4.1 Clock Stretching

When a slave device is not yet ready to process data, it may drive a '0' on the SCL line to hold it down. Due to the implementation of the I/O signal interface, the SCL line value will be '0', independent of the values that any other master or slave may be driving on the SCL line. This is known as clock stretching and is the only situation in which a slave drives the SCL line. The master device monitors the SCL line and detects it when it cannot generate a positive clock pulse ('1') on the SCL line. It then reacts by delaying the generation of a positive edge on the SCL line, effectively synchronizing with the slave device that is stretching the clock. The SCB on the PSoC 6 MCU can and will stretch the clock.

27.5.4.2 Bus Arbitration

The I²C protocol is a multi-master, multi-slave interface. Bus arbitration is implemented on master devices by monitoring the SDA line. Bus collisions are detected when the master observes an SDA line value that is not the same as the value it is driving on the SDA line. For example, when master 1 is driving the value '1' on the SDA line and master 2 is driving the value '0' on the SDA line, the actual line value will be '0' due to the implementation of the I/O signal interface. Master 1 detects the inconsistency and loses control of the bus. Master 2 does not detect any inconsistency and keeps control of the bus.

27.5.5 I²C Modes of Operation

I²C is a synchronous single master, multi-master, multi-slave serial interface. Devices operate in either master mode, slave mode, or master/slave mode. In master/slave mode, the device switches from master to slave mode when it is addressed. Only a single master may be active during a data transfer. The active master is responsible for driving the clock on the SCL line. Table 27-14 illustrates the I²C modes of operation.

Table 27-14. I²C Modes

Mode	Description
Slave	Slave only operation (default)
Master	Master only operation
Multi-master	Supports more than one master on the bus

Table 27-15 lists some common bus events that are part of an I²C data transfer. The [Write Transfer](#) and [Read Transfer](#) sections explain the I²C bus bit format during data transfer.

Table 27-15. I²C Bus Events Terminology

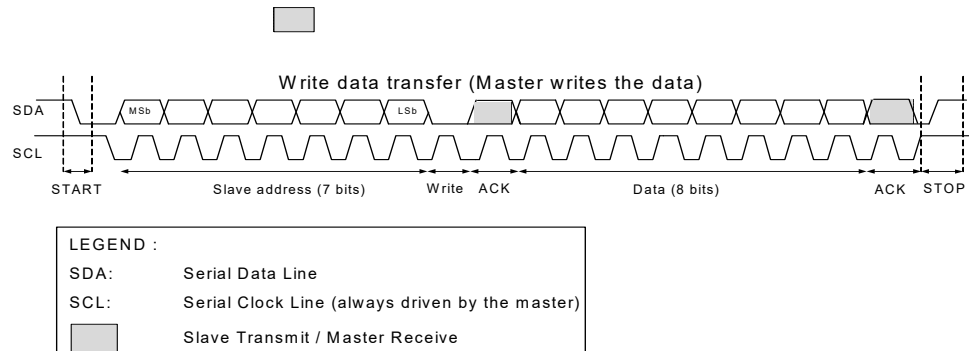
Bus Event	Description
START	A HIGH to LOW transition on the SDA line while SCL is HIGH
STOP	A LOW to HIGH transition on the SDA line while SCL is HIGH
ACK	The receiver pulls the SDA line LOW and it remains LOW during the HIGH period of the clock pulse, after the transmitter transmits each byte. This indicates to the transmitter that the receiver received the byte properly.
NACK	The receiver does not pull the SDA line LOW and it remains HIGH during the HIGH period of clock pulse after the transmitter transmits each byte. This indicates to the transmitter that the receiver did not receive the byte properly.
Repeated START	START condition generated by master at the end of a transfer instead of a STOP condition
DATA	SDA status change while SCL is LOW (data changing), and no change while SCL is HIGH (data valid)

With all of these modes, there are two types of transfer - read and write. In write transfer, the master sends data to slave; in read transfer, the master receives data from slave.

27.5.5.1 Write Transfer

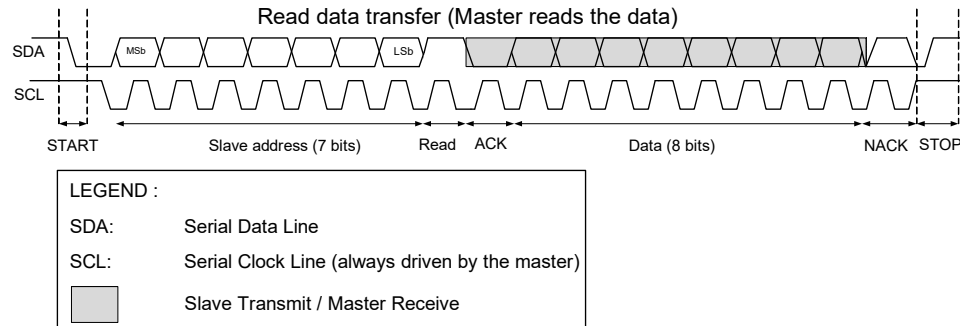
- A typical write transfer begins with the master generating a START condition on the I²C bus. The master then writes a 7-bit I²C slave address and a write indicator ('0') after the START condition. The addressed slave transmits an acknowledgment byte by pulling the data line low during the ninth bit time.
- If the slave address does not match any of the slave devices or if the addressed device does not want to acknowledge the request, it transmits a no acknowledgment (NACK) by not pulling the SDA line low. The absence of an acknowledgement, results in an SDA line value of '1' due to the pull-up resistor implementation.
- If no acknowledgment is transmitted by the slave, the master may end the write transfer with a STOP event. The master can also generate a repeated START condition for a retry attempt.
- The master may transmit data to the bus if it receives an acknowledgment. The addressed slave transmits an acknowledgment to confirm the receipt of every byte of data written. Upon receipt of this acknowledgment, the master may transmit another data byte.
- When the transfer is complete, the master generates a STOP condition.

Figure 27-37. Master Write Data Transfer



27.5.5.2 Read Transfer

Figure 27-38. Master Read Data Transfer



- A typical read transfer begins with the master generating a START condition on the I²C bus. The master then writes a 7-bit I²C slave address and a read indicator ('1') after the START condition. The addressed slave transmits an acknowledgment by pulling the data line low during the ninth bit time.
- If the slave address does not match with that of the connected slave device or if the addressed device does not want to acknowledge the request, a no acknowledgment (NACK) is transmitted by not pulling the SDA line low. The absence of an acknowledgment, results in an SDA line value of '1' due to the pull-up resistor implementation.
- If no acknowledgment is transmitted by the slave, the master may end the read transfer with a STOP event. The master can also generate a repeated START condition for a retry attempt.
- If the slave acknowledges the address, it starts transmitting data after the acknowledgment signal. The master transmits an acknowledgment to confirm the receipt of each data byte sent by the slave. Upon receipt of this acknowledgment, the addressed slave may transmit another data byte.
- The master can send a NACK signal to the slave to stop the slave from sending data bytes. This completes the read transfer.
- When the transfer is complete, the master generates a STOP condition.

elements or 128 elements deep with 8-bit data elements. The width of the data elements are configured using the CTRL.BYTE_MODE bitfield of the SCB. For I²C, put the FIFO in BYTE mode because all transactions are a byte wide.

The FIFO mode operation is available only in Active and Sleep power modes, not in the Deep Sleep power mode. However, on the Deep Sleep-capable SCB the slave address can be used to wake the device from sleep.

A write access to the transmit FIFO uses register TX_FIFO_WR. A read access from the receive FIFO uses register RX_FIFO_RD.

Transmit and receive FIFO status information is available through status registers TX_FIFO_STATUS and RX_FIFO_STATUS. When in debug mode, a read from this register behaves as a read from the SCB_RX_FIFO_RD_SILENT register; that is, data will not be removed from the FIFO.

Each FIFO has a trigger output. This trigger output can be routed through the trigger mux to various other peripheral on the device such as DMA or TCPWMs. The trigger output of the SCB is controlled through the TRIGGER_LEVEL field in the RX_CTRL and TX_CTRL registers.

- For a TX FIFO a trigger is generated when the number of entries in the transmit FIFO is less than TX_FIFO_CTRL.TRIGGER_LEVEL.
- For the RX FIFO a trigger is generated when the number of entries in the FIFO is greater than the RX_FIFO_CTRL.TRIGGER_LEVEL.

Note that the DMA has a trigger deactivation setting. For the SCB this should be set to 16.

Active to Deep Sleep Transition

Before going to deep sleep ensure that all active communication is complete. This can be done by checking the BUS_BUSY bit in the I2C_Status register.

Ensure that the TX and RX FIFOs are empty as any data will be lost during deep sleep.

27.5.6 I²C Buffer Modes

I²C can operate in three different buffered modes – FIFO, EZ, and CMD_RESP modes. The buffer is used in different ways in each of the modes. The following subsections explain each of these buffered modes in detail.

27.5.6.1 FIFO Mode

The FIFO mode has a TX FIFO for the data being transmitted and an RX FIFO for the data being received. Each FIFO is constructed out of the SRAM buffer. The FIFOs are either 64 elements deep with 16-bit data

Before going to deep sleep the clock to the SCB needs to be disabled. This can be done by setting the SDA_IN_FILT_TRIM[1] bit in the I2C_CFG register to '0'.

Deep Sleep to Active Transition

EC_AM = 1, EC_OP = 0, FIFO Mode.

The following descriptions only apply to slave mode.

Master Write:

- I2C_CTRL.S_NOT_READY_ADDR_NACK = 0, I2C_CTRL.S_READY_ADDR_ACK = 1. The clock is stretched until SDA_IN_FILT_TRIM[1] is set to '1'. After that bit is set to 1, the clock stretch will be released
- I2C_CTRL.S_NOT_READY_ADDR_NACK = 0, I2C_CTRL.S_READY_ADDR_ACK = 0. The clock is stretched until SDA_IN_FILT_TRIM[1] is set to '1', and S_ACK or S_NACK is set in the I2C_S_CMD register.
- I2C_CTRL.S_NOT_READY_ADDR_NACK = 1, I2C_CTRL.S_READY_ADDR_ACK = x. The incoming address is NACK'd until SDA_IN_FILT_TRIM[1] is set to '1'. After that bit is set to 1, the slave will respond with an ACK to a master address.

Master Read:

- I2C_CTRL.S_NOT_READY_ADDR_NACK = 0, I2C_CTRL.S_READY_ADDR_ACK = x. The incoming address is stretched until SDA_IN_FILT_TRIM[1] is set to '1'. After that bit is set to 1, the clock stretch will be released.
- I2C_CTRL.S_NOT_READY_ADDR_NACK = 1, I2C_CTRL.S_READY_ADDR_ACK = x. The incoming address is NACK'd until SDA_IN_FILT_TRIM[1] is set to '1'. After that bit is set to 1, the slave will ACK.

Note: When doing a repeated start after a write, wait until the UNDERFLOW interrupt status is asserted before setting the I2C_M_CMD.START bit and writing the new address into the TX_FIFO. Otherwise, the address in the FIFO will be sent as data and not as an address.

27.5.6.2 EZI2C Mode

The Easy I²C (EZI2C) protocol is a unique communication scheme built on top of the I²C protocol by Cypress. It uses a meta protocol around the standard I²C protocol to communicate to an I²C slave using indexed memory transfers. This removes the need for CPU intervention.

The EZI2C protocol defines a single memory buffer with an 8-bit address that indexes the buffer (256-entry array of 8-bit per entry is supported) located on the slave device. The EZ address is used to address these 256 locations. The CPU writes and reads to the memory buffer through the EZ_DATA registers. These accesses are word accesses, but only the least significant byte of the word is used.

The slave interface accesses the memory buffer using the current address. At the start of a transfer (I²C START/RESTART), the base address is copied to the current address. A data element write or read operation is to the current address location. After the access, the current address is incremented by '1'.

If the current address equals the last memory buffer address (255), the current address is not incremented. Subsequent write accesses will overwrite any previously written value at the last buffer address. Subsequent read accesses will continue to provide the (same) read value at the last buffer address. The bus master should be aware of the memory buffer capacity in EZ mode.

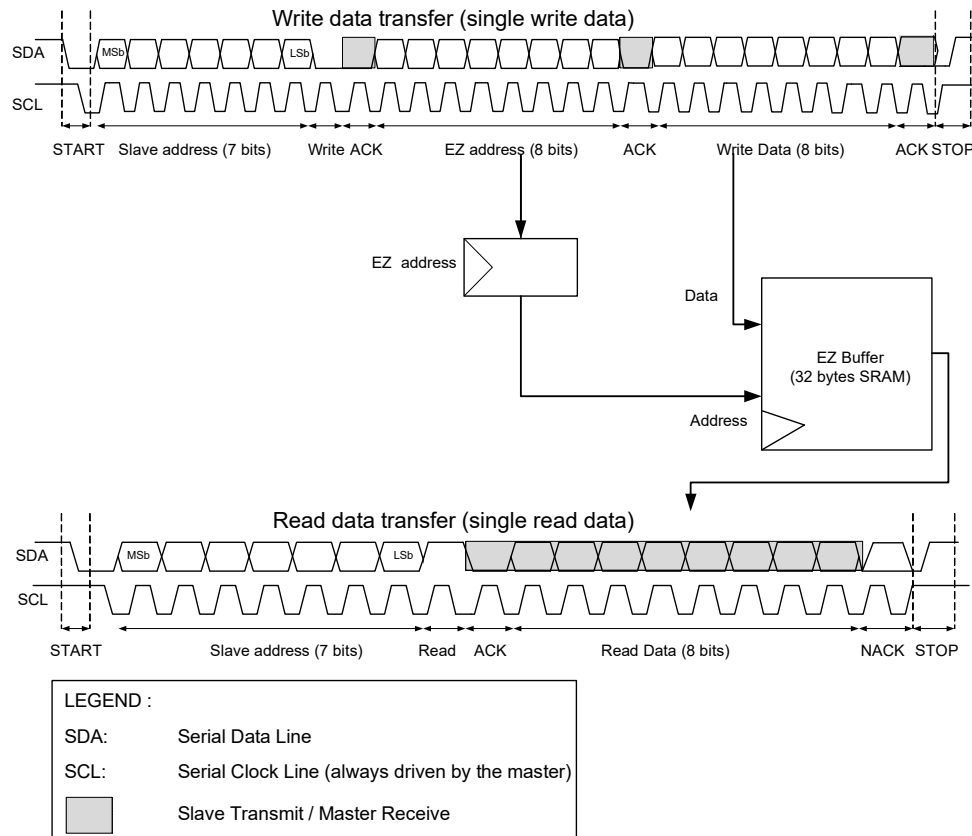
The I²C base and current addresses are provided through I2C_STATUS. At the end of a transfer, the difference between the base and current addresses indicates how many read or write accesses were performed. The block provides interrupt cause fields to identify the end of a transfer. EZI2C can be implemented through firmware or hardware. All SCBs can implement EZI2C through a firmware implementation in both Active and Sleep power modes. The Deep Sleep SCB can implement a hardware- and firmware-based EZI2C with a Deep Sleep power mode. This document focuses on hardware-implemented EZI2C; for more information on software implementation, see the [PDL](#).

EZI2C distinguishes three operation phases:

- Address phase: The master transmits an 8-bit address to the slave. This address is used as the slave base and current address.
- Write phase: The master writes 8-bit data element(s) to the slave's memory buffer. The slave's current address is set to the slave's base address. Received data elements are written to the current address memory location. After each memory write, the current address is incremented.
- Read phase: The master reads 8-bit data elements from the slave's memory buffer. The slave's current address is set to the slave's base address. Transmitted data elements are read from the current address memory location. After each memory read, the current address is incremented.

Note that a slave's base address is updated by the master and not by the CPU.

Figure 27-39. EZI2C Write and Read Data Transfer



Active to Deep Sleep Transition

Before going to deep sleep ensure that all active communication is complete. This can be done by checking the BUS_BUSY bit in the I2C_Status register.

Ensure that the TX and RX FIFOs are empty as any data will be lost during deep sleep.

Before going to deep sleep the clock to the SCB needs to be disabled. This can be done by setting the SDA_IN_FILT_TRIM[1] bit in the I2C_CFG register to '0'.

Deep Sleep to Active Transition

EC_AM = 1, EC_OP = 0, EZ Mode.

- S_NOT_READY_ADDR_NACK = 0, S_READY_ADDR_ACK = 1. The clock is stretched until SDA_IN_FILT_TRIM[1] is set to '1'. After that bit is set to 1, the clock stretch will be released.
- S_NOT_READY_ADDR_NACK = 1, S_READY_ADDR_ACK = x. The incoming address is NACK'd until SDA_IN_FILT_TRIM[1] is set to '1'. After that bit is set to 1, the slave will ACK.

27.5.6.3 Command-Response Mode

This mode has a single memory buffer, a base read address, a current read address, a base write address, and a current write address that are used to index the memory buffer. The base addresses are provided by the CPU. The current addresses are used by the slave to index the memory buffer for sequential accesses of the memory buffer. The memory buffer holds 256 8-bit data elements. The base and current addresses are in the range [0 to 255].

The CPU writes and reads to the memory buffer through the SCB_EZ_DATA registers. These are word accesses, but only the least significant byte of the word is used.

The slave interface accesses the memory buffer using the current addresses. At the start of a write transfer (I²C START/RESTART), the base write address is copied to the current write address. A data element write is to the current write address location. After the write access, the current address is incremented by '1'. At the start of a read transfer, the base read address is copied to the current read address. A data element read is to the current read address location. After the read data element is transmitted, the current read address is incremented by '1'.


If the current addresses equal the last memory buffer address (255), the current addresses are not incremented. Subsequent write accesses will overwrite any previously written value at the last buffer address. Subsequent read accesses will continue to provide the (same) read value at

the last buffer address. The bus master should be aware of the memory buffer capacity in command-response mode.

The base addresses are provided through CMD_RESP_CTRL. The current addresses can be viewed in CMD_RESP_STATUS. At the end of a transfer (I²C stop), the difference between a base and current address indicates how many read/write accesses were performed. This block provides interrupts to identify the end of a transfer, which can be found in SCB8_INTR_I2C_EC and SCB8_INTR_SPI_EC register sections. Command-response mode operation is available in Active, Sleep, and Deep Sleep power modes. The command-response mode has two phases of operation:

- **Write phase** - The write phase begins with a START/RESTART followed by the slave address with read/write bit set to '0' indicating a write. The slave's current write address is set to the slave's base write address. Received data elements are written to the current write address memory location. After each memory write, the current write address is incremented.
- **Read phase** - The read phase begins with a START/RESTART followed by the slave address with read/write bit set to '1' indicating a read. The slave's current read address is set to the slave's base read address. Transmitted data elements are read from the current address memory location. After each read data element is transferred, the current read address is incremented.

Figure 27-40. I²C Command-Response Mode



Note: A slave's base addresses are updated by the CPU and not by the master.

27.5.7 Clocking and Oversampling

The SCB I²C supports both internally and externally clocked operation modes. Two bitfields (EC_AM_MODE and EC_OP_MODE) in the SCB_CTRL register determine the SCB clock mode. EC_AM_MODE indicates whether I²C address matching is internally (0) or externally (1) clocked. I²C address matching comprises the first part of the I²C protocol. EC_OP_MODE indicates whether the rest of the protocol operation (besides I²C address matching) is internally (0) or externally (1) clocked. The externally clocked mode of operation is supported only in the I²C slave mode.

An internally-clocked operation uses the programmable clock dividers. For I²C, an integer clock divider must be used for both master and slave. For more information on system clocking, see the [Clocking System chapter on page 207](#). The internally-clocked mode does not support the command-response mode.

The SCB_CTRL bitfields EC_AM_MODE and EC_OP_MODE can be configured in the following ways.

- EC_AM_MODE is '0' and EC_OP_MODE is '0': Use this configuration when only Active mode functionality is required.
 - FIFO mode: Supported.
 - EZ mode: Supported.
 - Command-response mode: Not supported. The slave NACKs every slave address.
- EC_AM_MODE is '1' and EC_OP_MODE is '0': Use this configuration when both Active and Deep Sleep functionality are required. This configuration relies on the externally clocked functionality for the I²C address matching and relies on the internally clocked functionality to access the memory buffer. The "hand over" from external to internal functionality relies either on an ACK/NACK or clock stretching scheme. The former may result in termination of the current transfer and relies on a master retry. The latter stretches the current transfer after a matching address is received. This mode requires the master to support either NACK generation (and retry) or clock stretching. When the I²C address is matched, INTR_I2C_EC.WAKE_UP is set to '1'. The associated Deep Sleep functionality interrupt brings the system into Active power mode.
 - FIFO mode: See ["Deep Sleep to Active Transition" on page 323](#)
 - EZ mode: See ["Deep Sleep to Active Transition" on page 324](#).
 - CMD_RESP mode: Not supported. The slave NACKs every slave address.
- EC_AM_MODE is '1' and EC_OP_MODE is '1': Use this mode when both Active and Deep Sleep functionality are required. When the slave is selected, INTR_I2C_EC.WAKE_UP is set to '1'. The associated

Deep Sleep functionality interrupt brings the system into Active power mode. When the slave is deselected, INTR_I2C_EC.EZ_STOP and/or INTR_I2C_EC.EZ_WRITE_STOP are set to '1'.

- FIFO mode: Not supported.
- EZ mode: Supported.
- CMD_RESP mode: Supported.

An externally-clocked operation uses a clock provided by the serial interface. The externally clocked mode does not support FIFO mode. If EC_OP_MODE is '1', the external interface logic accesses the memory buffer on the external interface clock (I²C SCL). This allows for EZ and CMD_RESP mode functionality in Active and Deep Sleep power modes.

In Active system power mode, the memory buffer requires arbitration between external interface logic (on I²C SCL) and the CPU interface logic (on system peripheral clock). This arbitration always gives the highest priority to the external interface logic (host accesses). The external interface logic takes one serial interface clock/bit periods for the I²C. During this period, the internal logic is denied service to the memory buffer. The PSoC 6 MCU provides two programmable options to address this "denial of service":

- If the BLOCK bitfield of SCB_CTRL is '1': An internal logic access to the memory buffer is blocked until the memory buffer is granted and the external interface logic has completed access. For a 100-kHz I²C interface, the maximum blocking period of one serial interface bit period measures 10 μ s (approximately 208 clock cycles on a 48 MHz SCB input clock). This option provides normal SCB register functionality, but the blocking time introduces additional internal bus wait states.
- If the BLOCK bitfield of SCB_CTRL is '0': An internal logic access to the memory buffer is not blocked, but fails when it conflicts with an external interface logic access. A read access returns the value 0xFFFF:FFFF and a write access is ignored. This option does not introduce additional internal bus wait states, but an access to the memory buffer may not take effect. In this case, following failures are detected:
 - Read Failure: A read failure is easily detected, as the returned value is 0xFFFF:FFFF. This value is unique as non-failing memory buffer read accesses return an unsigned byte value in the range 0x0000:0000-0x0000:00FF.
 - Write Failure: A write failure is detected by reading back the written memory buffer location, and confirming that the read value is the same as the written value.

For both options, a conflicting internal logic access to the memory buffer sets INTR_TX.BLOCKED field to '1' (for write access-es) and INTR_RX.BLOCKED field to '1' (for read access-es). These fields can be used as either status fields

or as interrupt cause fields (when their associated mask fields are enabled).

If a series of read or write accesses is performed and CTRL.BLOCKED is '0', a failure is detected by comparing the logical OR of all read values to 0xFFFF:FFFF and checking the INTR_TX.BLOCKED and INTR_RX.BLOCKED fields to determine whether a failure occurred for a (series of) write or read operation(s).

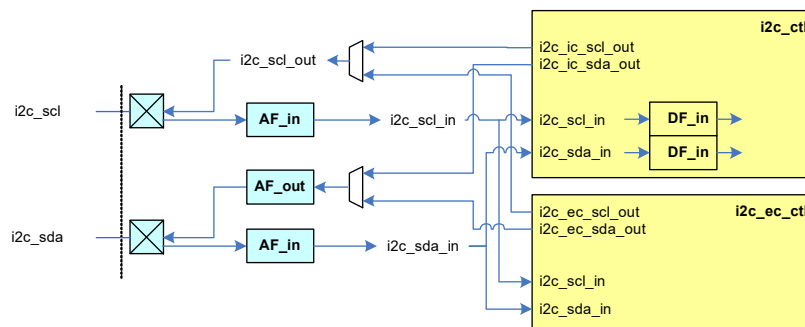
27.5.7.1 Glitch Filtering

The PSoC 6 MCU SCB I²C has analog and digital glitch filters. Analog glitch filters are applied on the i2c_scl_in and

i2c_sda_in input signals (AF_in) to filter glitches of up to 50 ns. An analog glitch filter is also applied on the i2c_sda_out output signal (AF_out). Analog glitch filters are enabled and disabled in the SCB.I2C_CFG register. Do not change the _TRIM bitfields; only change the _SEL bitfields in this register.

Digital glitch filters are applied on the i2c_scl_in and i2c_sda_in input signals (DF_in). The digital glitch filter is enabled in the SCB.RX_CTRL register via the MEDIAN bitfield.

Figure 27-41. I²C Glitch Filtering Connection



The following table lists the useful combinations of glitch filters.

Table 27-16. Glitch Filter Combinations

AF_in	AF_out	DF_in	Comments
0	0	1	Used when operating in internally-clocked mode and in master in fast-mode plus (1-MHz speed mode)
1	0	0	Used when operating in internally-clocked mode (EC_OP_MODE is '0')
1	1	0	Used when operating in externally-clocked mode (EC_OP_MODE is '1'). Only slave mode.

When operating in EC_OP_MODE = 1, the 100-kHz, 400-kHz, and 1000-kHz modes require the following settings for AF_out:

AF_in	AF_out	DF_in	
1	1	0	100-kHz mode: I2C_CFG.SDA_OUT_FILT_SEL = 3 400-kHz mode: I2C_CFG.SDA_OUT_FILT_SEL = 3 1000-kHz mode: I2C_CFG.SDA_OUT_FILT_SEL = 1

27.5.7.2 Oversampling and Bit Rate

Internally-clocked Master

The PSoC 6 MCU implements the I²C clock as an oversampled multiple of the SCB input clock. In master mode, the block determines the I²C frequency. Routing delays on the PCB, on the chip, and the SCB (including analog and digital glitch filters) all contribute to the signal interface timing. In master mode, the block operates off clk_scb and uses programmable oversampling factors for the SCL high (1) and low (0) times. For high and low phase oversampling, see

I2C_CTRL.LOW_PHASE_OVS and I2C_CTRL.HIGH_PHASE_OVS registers. For simple manipulation of the oversampling factor, see the SCB_CTRL.OVS register.

Table 27-17. I²C Frequency and Oversampling Requirements in I²C Master Mode

AF_in	AF_out	DF_in	Mode	Supported Frequency	LOW_PHASE_OVS	HIGH_PHASE_OVS	clk_scb Frequency
0	0	1	100 kHz	[62, 100] kHz	[9, 15]	[9, 15]	[1.98-3.2] MHz
			400 kHz	[264, 400] kHz	[13, 5]	[7, 15]	[8.45-10] MHz
			1000 kHz	[447, 1000] kHz	[8, 15]	[5, 15]	[14.32-25.8] MHz
1	0	0	100 kHz	[48, 100] kHz	[7, 15]	[7, 15]	[1.55-3.2] MHz
			400 kHz	[244, 400] kHz	[12, 15]	[7, 15]	[7.82-10] MHz
			1000 kHz	Not supported			

Table 27-17 assumes worst-case conditions on the I²C bus. The following equations can be used to determine the settings for your own system. This will involve measuring the rise and fall times on SCL and SDA lines in your system.

$$t_{CLK_SCB(Min)} = (t_{LOW} + t_F)/LOW_PHASE_OVS$$

If clk_scb is any faster than this, the t_{LOW} of the I²C specification will be violated. t_F needs to be measured in your system.

$$t_{CLK_SCB(Max)} = (t_{VD} - t_{RF} - 100 \text{ ns})/3 \text{ (When analog filter is enabled and digital disabled)}$$


$$t_{CLK_SCB(Max)} = (t_{VD} - t_{RF})/4 \text{ (When analog filter is disabled and analog filter is enabled)}$$

t_{RF} is the maximum of either the rise or fall time. If clk_scb is slower than this frequency, t_{VD} will be violated.

I²C Master Clock Synchronization

The HIGH_PHASE_OVS counter does not start counting until the SCB detects that the SCL line is high. This is not the same as when the SCB sets the SCL high. The differences are explained by three delays:

1. Delay from SCB to I/O pin
2. I²C bus t_R
3. Input delay (filters and synchronization)



If the above three delays combined are greater than one clk_scb cycle, then the high phase of the SCL will be extended. This may cause the actual data rate on the I²C bus to be slower than expected. This can be avoided by:

- Decreasing the pull-up resistor, or decreasing the bus capacitance to reduce t_R .
- Reducing the I2C_CTRL.HIGH_PHASE_OVS value.

Internally-clocked Slave

In slave mode, the I²C frequency is determined by the incoming I²C SCL signal. To ensure proper operation, clk_scb must be significantly higher than the I²C bus frequency. Unlike master mode, this mode does not use programmable oversampling factors.

Table 27-18. SCB Input Clock Requirements in I²C Slave Mode

AF_in	AF_out	DF_in	Mode	clk_scb Frequency Range
0	0	1	100 kHz	[1.98-12.8] MHz
			400 kHz	[8.45-17.14] MHz
			1000 kHz	[14.32-44.77] MHz
1	0	0	100 kHz	[1.55-12.8] MHz
			400 kHz	[7.82-15.38] MHz
			1000 kHz	[15.84-89.0] MHz

$$t_{CLK_SCB(Max)} = (t_{VD} - t_{RF} - 100 \text{ ns}) / 3 \text{ (When analog filter is enabled and digital disabled)}$$

$$t_{CLK_SCB(Max)} = (t_{VD} - t_{RF}) / 4 \text{ (When analog filter is disabled and analog filter is enabled)}$$

t_{RF} is the maximum of either the rise or fall time. If clk_scb is slower than this frequency, t_{VD} will be violated.

The minimum period of clk_scb is determined by one of the following equations:

$$t_{CLK_SCB(MIN)} = (t_{SU;DAT(min)} + t_{RF}) / 16$$

or

$$t_{CLK_SCB(Min)} = (0.6 * t_F - 50 \text{ ns}) / 2 \text{ (When analog filter is enabled and digital disabled)}$$

$$t_{CLK_SCB(Min)} = (0.6 * t_F) / 3 \text{ (When analog filter is disabled and digital enabled)}$$

The result that yields the largest period from the two sets of equations above should be used to set the minimum period of clk_scb.

Master-Slave

In this mode, when the SCB is acting as a master device, the block determines the I²C frequency. When the SCB is acting as a slave device, the block does not determine the I²C frequency. Instead, the incoming I²C SCL signal does.

To guarantee operation in both master and slave modes, choose clock frequencies that work for both master and slave using the tables above.

27.5.8 Enabling and Initializing the I²C

The following section describes the method to configure the I²C block for standard (non-EZ) mode and EZI²C mode.

27.5.8.1 Configuring for I²C FIFO Mode

The I²C interface must be programmed in the following order.

1. Program protocol specific information using the SCB_I2C_CTRL register. This includes selecting master - slave functionality.
2. Program the generic transmitter and receiver information using the SCB_TX_CTRL and SCB_RX_CTRL registers.
3. Set the SCB_CTRL.BYTE_MODE to '1' to enable the byte mode.
4. Program the SCB_CTRL register to enable the I²C block and select the I²C mode. For a complete description of the I²C registers, see the [registers TRM](#).

27.5.8.2 Configuring for EZ and CMD_RESP Modes

To configure the I²C block for EZ and CMD_RESP modes, set the following I²C register bits

- 1a. Select the EZI2C mode by writing '1' to the EZ_MODE bit (bit 10) of the SCB_CTRL register.
- 1b. Select CMD_RESP mode by writing a 1 to the CMD_RESP bit (bit 12) of the SCB_CTRL register.
2. Set the S_READY_ADDR_ACK (bit 12) and S_READY_DATA_ACK (bit 13) bits of the SCB_I2C_CTRL register.

Note: For all modes clk_scb must also be configured. For information on configuring a peripheral clock and connecting it to the SCB consult the [Clocking System chapter on page 207](#).

The GPIO must also be connected to the SCB; see the following section for more details.

27.5.9 I/O Pad Connections

Figure 27-43. I²C I/O Pad Connections




Table 27-19. I²C I/O Pad Descriptions

I/O Pads	Drive Mode	On-chip I/O Signals	Usage
i2c_scl	Open drain with external pull-up	i2c_scl_in	Receive a clock
		i2c_scl_out	Transmit a clock
i2c_sda	Open drain with external pull-up	i2c_sda_in	Receive data
		i2c_sda_out	Transmit data

When configuring the I²C SDA/SCL lines, the following sequence must be followed. If this sequence is not followed, the I²C lines may initially have overshoot and undershoot.

1. Set SCB_CTRL_MODE to '0'.
2. Configure HSIOM for SCL and SDA to connect to the SCB.
3. Set TX_CTRL.OPEN_DRAIN to '1'.
4. Configure I²C pins for high-impedance drive mode.
5. Configure SCB for I²C
6. Enable SCB
7. Configure I²C pins for Open Drain Drives Low.

27.5.10 I²C Registers

The I²C interface is controlled by reading and writing a set of configuration, control, and status registers, as listed in [Table 27-20](#).

Table 27-20. I²C Registers

Register	Function
SCB_CTRL	Enables the SCB block and selects the type of serial interface (SPI, UART, I ² C). Also used to select internally and externally clocked operation and EZ and non-EZ modes of operation.
SCB_I2C_CTRL	Selects the mode (master, slave) and sends an ACK or NACK signal based on receiver FIFO status.
SCB_I2C_STATUS	Indicates bus busy status detection, read/write transfer status of the slave/master, and stores the EZ slave address.
SCB_I2C_M_CMD	Enables the master to generate START, STOP, and ACK/NACK signals.
SCB_I2C_S_CMD	Enables the slave to generate ACK/NACK signals.
SCB_STATUS	Indicates whether the externally clocked logic is using the EZ memory. This bit can be used by software to determine whether it is safe to issue a software access to the EZ memory.
SCB_I2C_CFG	Configures filters, which remove glitches from the SDA and SCL lines.
SCB_TX_CTRL	Specifies the data frame width; also used to specify whether MSb or LSb is the first bit in transmission.
SCB_TX_FIFO_CTRL	Specifies the trigger level, clearing of the transmitter FIFO and shift registers, and FREEZE operation of the transmitter FIFO.
SCB_TX_FIFO_STATUS	Indicates the number of bytes stored in the transmitter FIFO, the location from which a data frame is read by the hardware (read pointer), the location from which a new data frame is written (write pointer), and decides if the transmitter FIFO holds the valid data.
SCB_TX_FIFO_WR	Holds the data frame written into the transmitter FIFO. Behavior is similar to that of a PUSH operation.
SCB_RX_CTRL	Performs the same function as that of the SCB_TX_CTRL register, but for the receiver. Also decides whether a median filter is to be used on the input interface lines.
SCB_RX_FIFO_CTRL	Performs the same function as that of the SCB_TX_FIFO_CTRL register, but for the receiver.
SCB_RX_FIFO_STATUS	Performs the same function as that of the SCB_TX_FIFO_STATUS register, but for the receiver.
SCB_RX_FIFO_RD	Holds the data read from the receiver FIFO. Reading a data frame removes the data frame from the FIFO; behavior is similar to that of a POP operation. This register has a side effect when read by software: a data frame is removed from the FIFO.
SCB_RX_FIFO_RD_SILENT	Holds the data read from the receiver FIFO. Reading a data frame does not remove the data frame from the FIFO; behavior is similar to that of a PEEK operation.
SCB_RX_MATCH	Stores slave device address and is also used as slave device address MASK.
SCB_EZ_DATA	Holds the data in an EZ memory location.

Note: Detailed descriptions of the I²C register bits are available in the [registers TRM](#).

27.6 SCB Interrupts

SCB supports interrupt generation on various events. The interrupts generated by the SCB block vary depending on the mode of operation.

Table 27-21. SCB Interrupts

Interrupt	Functionality	Active/Deep Sleep	Registers
interrupt_master	I ² C master and SPI master functionality	Active	INTR_M, INTR_M_SET, INTR_M_MASK, INTR_M_MASKED
interrupt_slave	I ² C slave and SPI slave functionality	Active	INTR_S, INTR_S_SET, INTR_S_MASK, INTR_S_MASKED
interrupt_tx	UART transmitter and TX FIFO functionality	Active	INTR_TX, INTR_TX_SET, INTR_TX_MASK, INTR_TX_MASKED
interrupt_rx	UART receiver and RX FIFO functionality	Active	INTR_RX, INTR_RX_SET, INTR_RX_MASK, INTR_RX_MASKED
interrupt_i2c_ec	Externally clocked I ² C slave functionality	Deep Sleep	INTR_I2C_EC, INTR_I2C_EC_MASK, INTR_I2C_EC_MASKED
interrupt_spi_ec	Externally clocked SPI slave functionality	Deep Sleep	INTR_SPI_EC, INTR_SPI_EC_MASK, INTR_SPI_EC_MASKED

Note: To avoid being triggered by events from previous transactions, whenever the firmware enables an interrupt mask register bit, it should clear the interrupt request register in advance.

Note: If the DMA is used to read data out of RX FIFO, the NOT_EMPTY interrupt may never trigger. This can occur when clk_peri (clocking DMA) is running much faster than the clock to the SCB. As a workaround to this issue, set the RX_FIFO_CTRL.TRIGGER_LEVEL to '1' (not 0); this will allow the interrupt to fire.

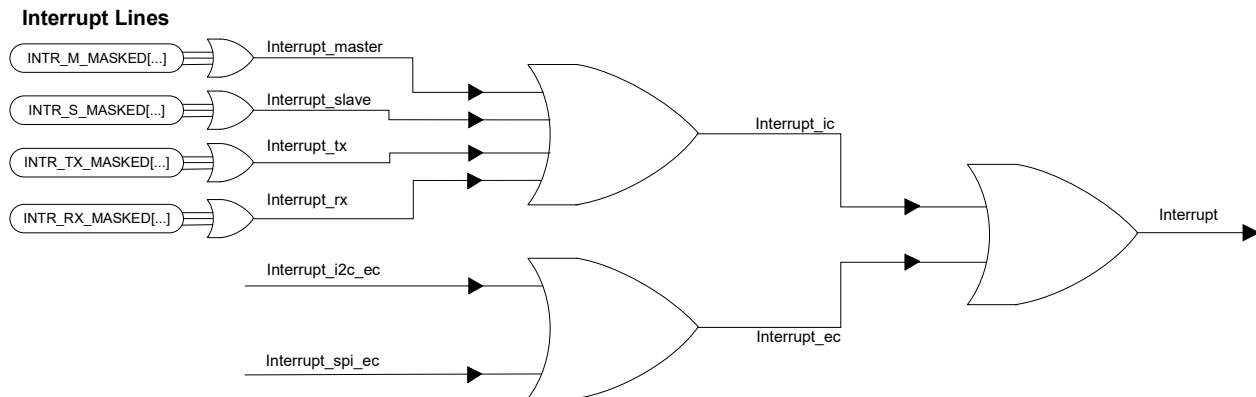
The following register definitions correspond to the SCB interrupts:

- **INTR_M:** This register provides the instantaneous status of the interrupt sources. A write of '1' to a bit will clear the interrupt.
- **INTR_M_SET:** A write of '1' into this register will set the interrupt.
- **INTR_M_MASK:** The bit in this register masks the interrupt sources. Only the interrupt sources with their masks enabled can trigger the interrupt.
- **INTR_M_MASKED:** This register provides the instantaneous value of the interrupts after they are masked. It provides logical and corresponding request and mask bits. This is used to understand which interrupt triggered the event.

Note: While registers corresponding to INTR_M are used here, these definitions can be used for INTR_S, INTR_TX, INTR_RX, INTR_I2C_EC, and INTR_SPI_EC.

Figure 27-44 shows the physical interrupt lines. All the interrupts are OR'd together to make one interrupt source that is the OR of all six individual interrupts. All the externally-clocked interrupts make one interrupt line called interrupt_ec, which is the OR'd signal of interrupt_i2c_ec and interrupt_spi_ec. All the internally-clocked interrupts make one interrupt line called interrupt_ic, which is the OR'd signal of interrupt_master, interrupt_slave, interrupt_tx, and interrupt_rx. The Active functionality interrupts are generated synchronously to clk_peri while the Deep Sleep functionality interrupts are generated asynchronously to clk_peri.

Figure 27-44. Interrupt Lines



27.6.1 SPI Interrupts

The SPI interrupts can be classified as master interrupts, slave interrupts, TX interrupts, RX interrupts, and externally clocked (EC) mode interrupts. Each interrupt output is the logical OR of the group of all possible interrupt sources classified under the section. For example, the TX interrupt output is the logical OR of the group of all possible TX interrupt sources. This signal goes high when any of the enabled TX interrupt sources are true. The SCB also provides an interrupt cause register (SCB_INTR_CAUSE) that can be used to determine interrupt source. The interrupt registers are cleared by writing '1' to the corresponding bitfield. Note that certain interrupt sources are triggered again as long as the condition is met even if the interrupt source was cleared. For example, the TX_FIFO_EMPTY is set as long as the transmit FIFO is empty even if the interrupt source is cleared. For more information on interrupt registers, see the [registers TRM](#). The SPI supports interrupts on the following events:

■ SPI Master Interrupts

- ❑ SPI master transfer done – All data from the TX FIFO are sent. This interrupt source triggers later than TX_FIFO_EMPTY by the amount of time it takes to transmit a single data element. TX_FIFO_EMPTY triggers when the last data element from the TX FIFO goes to the shifter register. However, SPI Done triggers after this data element is transmitted. This means SPI Done will be asserted one SCLK clock cycle earlier than the completion of data element reception.

■ SPI Slave Interrupts

- ❑ SPI Bus Error – Slave deselected at an unexpected time in the SPI transfer. The firmware may decide to clear the TX and RX FIFOs for this error.
- ❑ SPI slave deselected after any EZSPI transfer occurred.
- ❑ SPI slave deselected after a write EZSPI transfer occurred.

■ SPI TX

- ❑ TX FIFO has less entries than the value specified by TRIGGER_LEVEL in SCB_TX_FIFO_CTRL.
- ❑ TX FIFO not full – At least one data element can be written into the TX FIFO.
- ❑ TX FIFO empty – The TX FIFO is empty.
- ❑ TX FIFO overflow – Firmware attempts to write to a full TX FIFO.
- ❑ TX FIFO underflow – Hardware attempts to read from an empty TX FIFO. This happens when the SCB is ready to transfer data and EMPTY is '1'.
- ❑ TX FIFO trigger – Less entries in the TX FIFO than the value specified by TX_FIFO_CTRL.TRIGGER_LEVEL.

■ SPI RX

- ❑ RX FIFO has more entries than the value specified by TRIGGER_LEVEL in SCB_RX_FIFO_CTRL.
- ❑ RX FIFO full - RX FIFO is full.
- ❑ RX FIFO not empty - RX FIFO is not empty. At least one data element is available in the RX FIFO to be read.
- ❑ RX FIFO overflow - Hardware attempt to write to a full RX FIFO.
- ❑ RX FIFO underflow - Firmware attempts to read from an empty RX FIFO.

- ❑ RX FIFO trigger - More entries in the RX FIFO than the value specified by RX_FIFO_CTRL.TRIGGER_LEVEL.
- SPI Externally Clocked
 - ❑ Wake up request on slave select – Active on incoming slave request (with address match). Only set when EC_AM is '1'.
 - ❑ SPI STOP detection at the end of each transfer – Activated at the end of every transfer (I2C STOP). Only set for a slave request with an address match, in EZ and CMD_RESP modes, when EC_OP is '1'.
 - ❑ SPI STOP detection at the end of a write transfer – Activated at the end of a write transfer (I2C STOP). This event is an indication that a buffer memory location has been written to. For EZ mode, a transfer that only writes the base address does not activate this event. Only set for a slave request with an address match, in EZ and CMD_RESP modes, when EC_OP is '1'.
 - ❑ SPI STOP detection at the end of a read transfer – Activated at the end of a read transfer (I2C STOP). This event is an indication that a buffer memory location has been read from. Only set for a slave request with an address match, in EZ and CMD_RESP modes when EC_OP is '1'.

Figure 27-45 and Figure 27-46 show how each of the interrupts are triggered. Figure 27-45 shows the TX buffer and the corresponding interrupts while Figure 27-46 shows all the corresponding interrupts for the RX buffer. The FIFO has 256 split into 128 bytes for TX and 128 bytes for RX instead of the 8 bytes shown in the figures. For more information on how to implement and clear interrupts, see the SPI (SCB_SPI_PDL) datasheet and the PDL.

Figure 27-45. TX Interrupt Source Operation


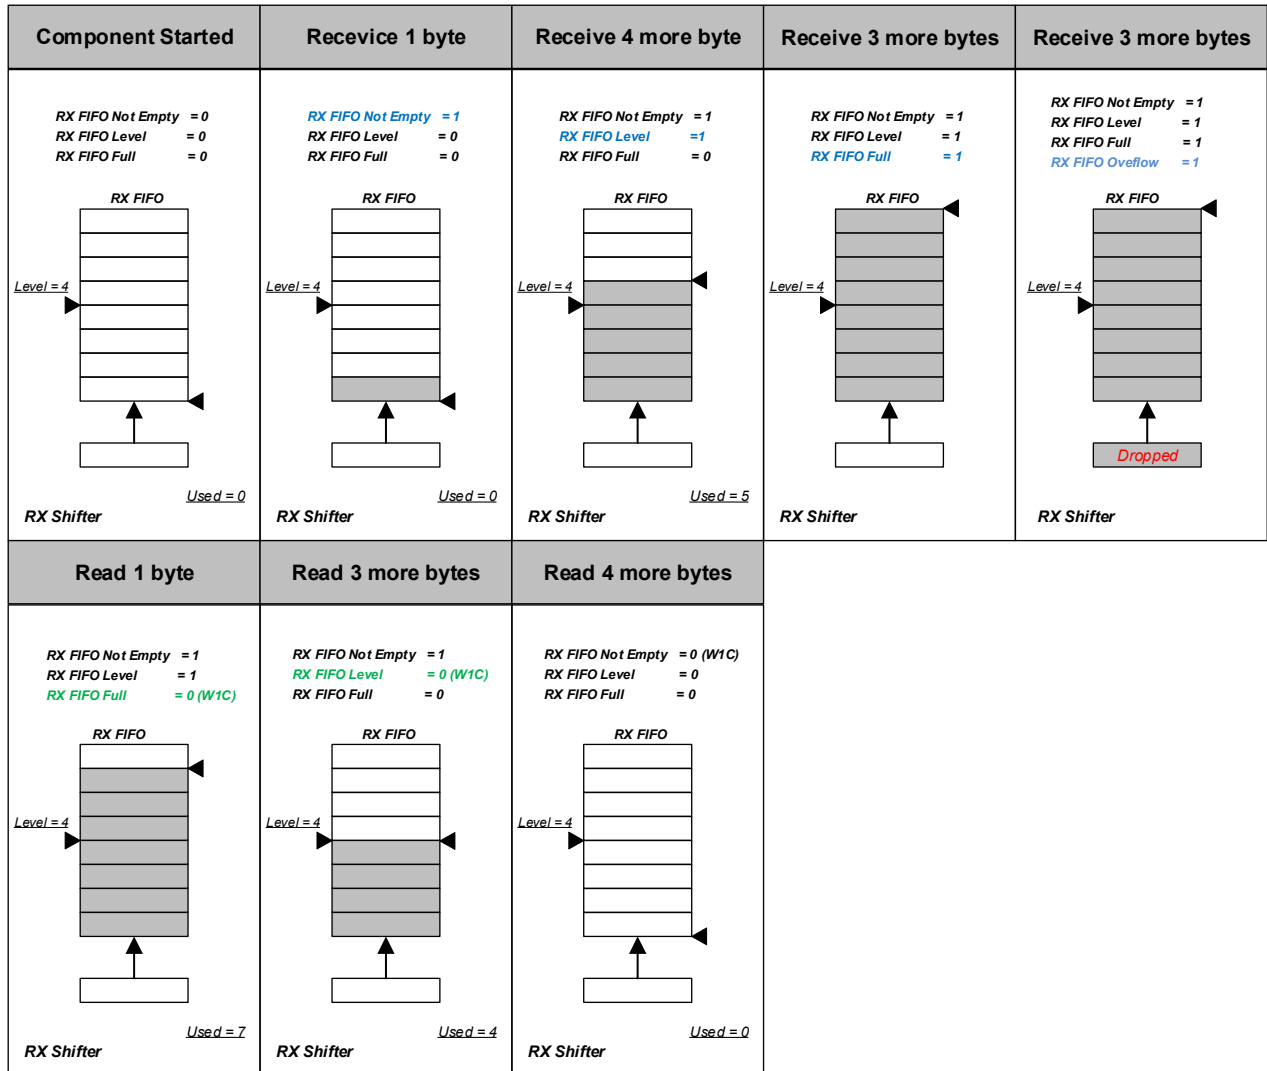


Figure 27-46. RX Interrupt Source Operation



27.6.2 UART Interrupts

The UART interrupts can be classified as TX interrupts and RX interrupts. Each interrupt output is the logical OR of the group of all possible interrupt sources classified under the section. For example, the TX interrupt output is the logical OR of the group of all possible TX interrupt sources. This signal goes high when any of the enabled TX interrupt sources are true. The SCB also provides an interrupt cause register (SCB_INTR_CAUSE) that can be used to determine interrupt source. The interrupt registers are cleared by writing '1' to the corresponding bitfield. Note that certain interrupt sources are triggered again as long as the condition is met even if the interrupt source was cleared. For example, the TX_FIFO_EMPTY is set as long as the transmit FIFO is empty even if the interrupt source is cleared. For more information on interrupt registers, see the [registers TRM](#). The UART block generates interrupts on the following events:

■ UART TX

- TX FIFO has fewer entries than the value specified by TRIGGER_LEVEL in SCB_TX_FIFO_CTRL.
- TX FIFO not full – TX FIFO is not full. At least one data element can be written into the TX FIFO.
- TX FIFO empty – The TX FIFO is empty.
- TX FIFO overflow – Firmware attempts to write to a full TX FIFO.

- ❑ TX FIFO underflow – Hardware attempts to read from an empty TX FIFO. This happens when the SCB is ready to transfer data and EMPTY is '1'.
- ❑ TX NACK – UART transmitter receives a negative acknowledgment in SmartCard mode.
- ❑ TX done – This happens when the UART completes transferring all data in the TX FIFO and the last stop field is transmitted (both TX FIFO and transmit shifter register are empty).
- ❑ TX lost arbitration – The value driven on the TX line is not the same as the value observed on the RX line. This condition event is useful when transmitter and receiver share a TX/RX line. This is the case in LIN or SmartCard modes.
- UART RX
 - ❑ RX FIFO has more entries than the value specified by TRIGGER_LEVEL in SCB_RX_FIFO_CTRL.
 - ❑ RX FIFO full – RX FIFO is full. Note that received data frames are lost when the RX FIFO is full.
 - ❑ RX FIFO not empty – RX FIFO is not empty.
 - ❑ RX FIFO overflow – Hardware attempts to write to a full RX FIFO.
 - ❑ RX FIFO underflow – Firmware attempts to read from an empty RX FIFO.
 - ❑ Frame error in received data frame – UART frame error in received data frame. This can be either a start of stop bit error:

Start bit error: After the beginning of a start bit period is detected (RX line changes from 1 to 0), the middle of the start bit period is sampled erroneously (RX line is '1'). **Note:** A start bit error is detected before a data frame is received.

Stop bit error: The RX line is sampled as '0', but a '1' was expected. A stop bit error may result in failure to receive successive data frames. **Note:** A stop bit error is detected after a data frame is received.
 - ❑ Parity error in received data frame – If UART_RX_CTL.DROP_ON_PARITY_ERROR is '1', the received frame is dropped. If UART_RX_CTL.DROP_ON_PARITY_ERROR is '0', the received frame is sent to the RX FIFO. In SmartCard sub mode, negatively acknowledged data frames generate a parity error. Note that firmware can only identify the erroneous data frame in the RX FIFO if it is fast enough to read the data frame before the hardware writes a next data frame into the RX FIFO.
 - ❑ LIN baud rate detection is completed – The receiver software uses the UART_RX_STATUS.BR_COUNTER value to set the clk_scb to guarantee successful receipt of the first LIN data frame (Protected Identifier Field) after the synchronization byte.
 - ❑ LIN break detection is successful – The line is '0' for UART_RX_CTRL.BREAK_WIDTH + 1 bit period. Can occur at any time to address unanticipated break fields; that is, "break-in-data" is supported. This feature is supported for the UART standard and LIN submodes. For the UART standard submodes, ongoing receipt of data frames is not affected; firmware is expected to take proper action. For the LIN submode, possible ongoing receipt of a data frame is stopped and the (partially) received data frame is dropped and baud rate detection is started. Set to '1', when event is detected. Write with '1' to clear bit.

Figure 27-47 and Figure 27-48 show how each of the interrupts are triggered. Figure 27-47 shows the TX buffer and the corresponding interrupts while Figure 27-48 shows all the corresponding interrupts for the RX buffer. The FIFO has 256 split into 128 bytes for TX and 128 bytes for RX instead of the 8 bytes shown in the figures. For more information on how to implement and clear interrupts see the UART (SCB_UART_PDL) datasheet and the PDL.

Figure 27-47. TX Interrupt Source Operation


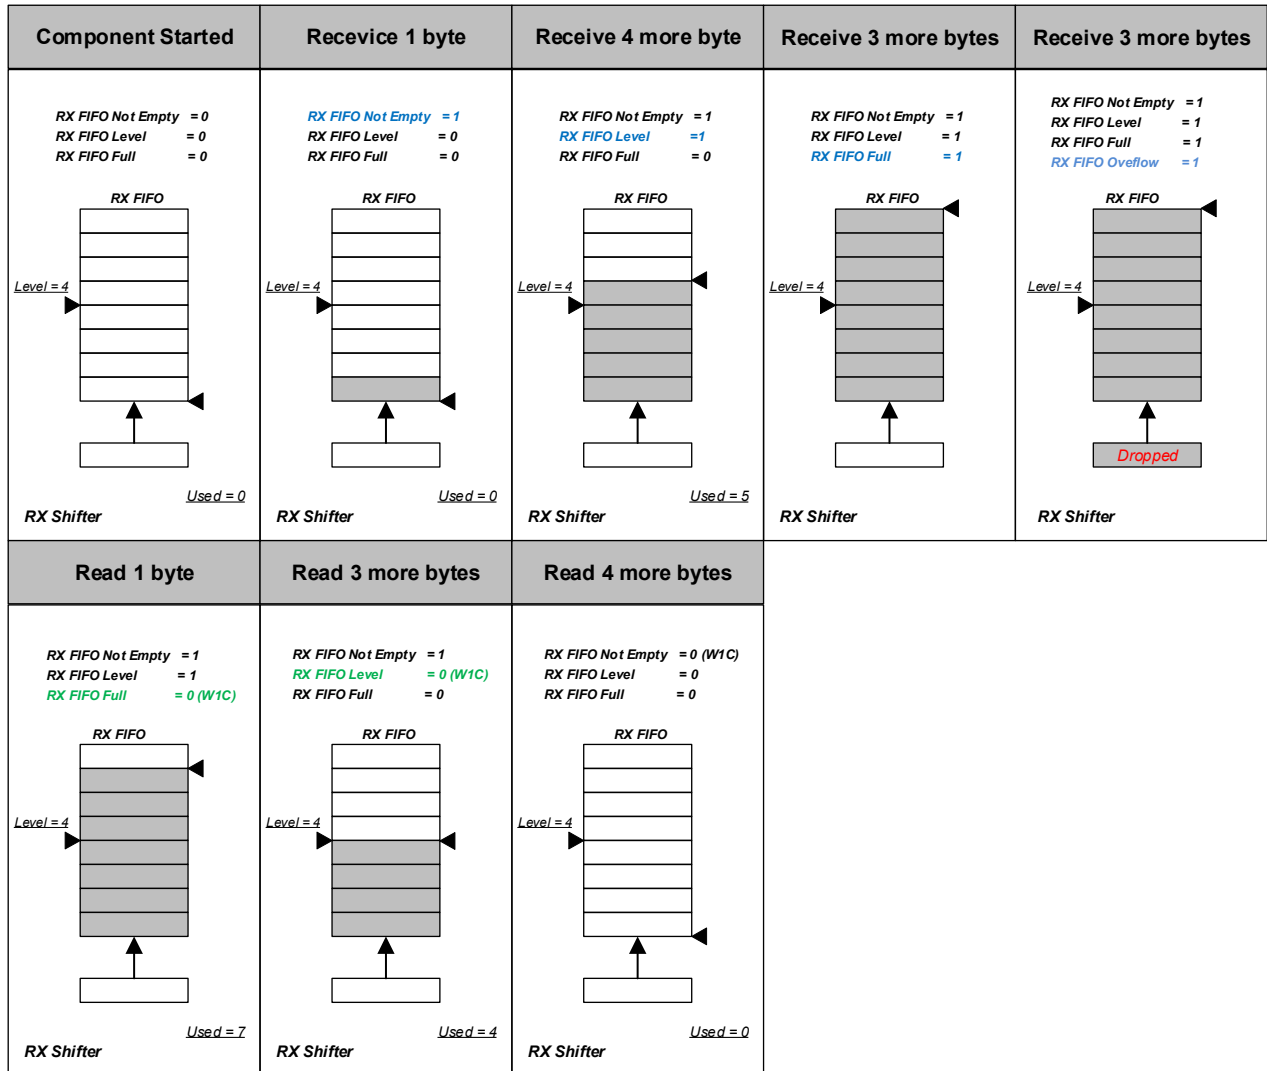


Figure 27-48. RX Interrupt Source Operation



27.6.3 I²C Interrupts

I²C interrupts can be classified as master interrupts, slave interrupts, TX interrupts, RX interrupts, and externally clocked (EC) mode interrupts. Each interrupt output is the logical OR of the group of all possible interrupt sources classified under the section. For example, the TX interrupt output is the logical OR of the group of all possible TX interrupt sources. This signal goes high when any of the enabled TX interrupt sources are true. The SCB also provides an interrupt cause register (SCB_INTR_CAUSE) that can be used to determine interrupt source. The interrupt registers are cleared by writing '1' to the corresponding bitfield. Note that certain interrupt sources are triggered again as long as the condition is met even if the interrupt source was cleared. For example, the TX_FIFO_EMPTY is set as long as the transmit FIFO is empty even if the interrupt source is cleared. For more information on interrupt registers, see the [registers TRM](#). The I²C block generates interrupts for the following conditions.

■ I²C Master

- I²C master lost arbitration – The value driven by the master on the SDA line is not the same as the value observed on the SDA line.
- I²C master received NACK – When the master receives a NACK (typically after the master transmitted the slave address or TX data).
- I²C master received ACK – When the master receives an ACK (typically after the master transmitted the slave address or TX data).
- I²C master sent STOP – When the master has transmitted a STOP.
- I²C bus error – Unexpected stop/start condition is detected.

■ I²C Slave

- I²C slave lost arbitration – The value driven on the SDA line is not the same as the value observed on the SDA line (while the SCL line is '1'). This should not occur; it represents erroneous I²C bus behavior. In case of lost arbitration, the I²C slave state machine aborts the ongoing transfer. Software may decide to clear the TX and RX FIFOs in case of this error.
- I²C slave received NACK – When the slave receives a NACK (typically after the slave transmitted TX data).
- I²C slave received ACK – When the slave receives an ACK (typically after the slave transmitted TX data).
- I²C slave received STOP – I²C STOP event for I²C (read or write) transfer intended for this slave (address matching is performed). When STOP or REPEATED START event is detected. The REPEATED START event is included in this interrupt cause such that the I²C transfers separated by a REPEATED START can be distinguished and potentially treated separately by the firmware. Note that the second I²C transfer (after a REPEATED START) may be to a different slave address.
The event is detected on any I²C transfer intended for this slave. Note that an I²C address intended for the slave (address matches) will result in an I2C_STOP event independent of whether the I²C address is ACK'd or NACK'd.
- I²C slave received START – When START or REPEATED START event is detected. In the case of an externally-clocked address matching (CTRL_EC_AM_MODE is '1') and clock stretching is performed (until the internally-clocked logic takes over) (I2C_CTRL.S_NOT_READY_ADDR_NACK is '0'), this field is not set. Firmware should use INTR_S_EC.WAKE_UP, INTR_S.I2C_ADDR_MATCH, and INTR_S.I2C_GENERAL.
- I²C slave address matched – I²C slave matching address received. If CTRL.ADDR_ACCEPT, the received address (including the R/W bit) is available in the RX FIFO. In the case of externally-clocked address matching (CTRL_EC_AM_MODE is '1') and internally-clocked operation (CTRL_EC_OP_MODE is '0'), this field is set when the event is detected.
- I²C bus error – Unexpected STOP/START condition is detected

■ I²C TX

- TX trigger – TX FIFO has fewer entries than the value specified by TRIGGER_LEVEL in SCB_TX_FIFO_CTRL.
- TX FIFO not full – At least one data element can be written into the TX FIFO.
- TX FIFO empty – The TX FIFO is empty.
- TX FIFO overflow – Firmware attempts to write to a full TX FIFO.
- TX FIFO underflow – Hardware attempts to read from an empty TX FIFO.

■ I²C RX

- RX FIFO has more entries than the value specified by TRIGGER_LEVEL in SCB_RX_FIFO_CTRL.

- ❑ RX FIFO is full – The RX FIFO is full.
- ❑ RX FIFO is not empty – At least one data element is available in the RX FIFO to be read.
- ❑ RX FIFO overflow – Hardware attempts to write to a full RX FIFO.
- ❑ RX FIFO underflow – Firmware attempts to read from an empty RX FIFO.
- I²C Externally Clocked
 - ❑ Wake up request on address match – Active on incoming slave request (with address match). Only set when EC_AM is '1'.
 - ❑ I²C STOP detection at the end of each transfer – Only set for a slave request with an address match, in EZ and CMD_RESP modes, when EC_OP is '1'.
 - ❑ I²C STOP detection at the end of a write transfer – Activated at the end of a write transfer (I²C STOP). This event is an indication that a buffer memory location has been written to. For EZ mode, a transfer that only writes the base address does not activate this event. Only set for a slave request with an address match, in EZ and CMD_RESP modes, when EC_OP is '1'.
 - ❑ I²C STOP detection at the end of a read transfer – Activated at the end of a read transfer (I²C STOP). This event is an indication that a buffer memory location has been read from. Only set for a slave request with an address match, in EZ and CMD_RESP modes, when EC_OP is '1'.

For more information on how to implement and clear interrupts see the I²C (SCB_I2C_PDL) datasheet and the PDL.

28. Serial Memory Interface (SMIF)



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SMIF
- Application notes
- Code examples

The SMIF block implements a single-SPI, dual-SPI, quad-SPI, or octal-SPI communication to interface with external memory chips. The SMIF block's primary use case is to set up the external memory and have it mapped to the PSoC 6 MCU memory space using the hardware. This mode of operation, called the XIP mode, allows the bus masters in the PSoC 6 MCU to directly interact with the SMIF for memory access to an external memory location.

28.1 Features

The Serial Memory Interface (SMIF) block provides a master interface to serial memory devices that supports the following functionality.

- Interfacing up to four memory devices (slaves) at a time
- SPI protocol
 - SPI mode 0: clock polarity (CPOL) and clock phase (CPHA) are both '0'
 - Support for single, dual, quad, and octal SPI protocols
 - Support for dual-quad SPI mode: the use of two quad SPI memory devices to increase data bandwidth for SPI read and write transfers
 - Support for configurable MISO sampling time and programmable receiver clock
- Support for device capacities in the range of 64 KB to 128 MB
- A memory-mapped mode of operation, eXecute In Place (XIP), which enables mapping the external memory into an internal memory address
- A command mode (MMIO mode), which enables using the SMIF block as a simple communication hardware
- Supports a 4-KB read cache in memory mapped (XIP) mode
- Supports on-the-fly 128-bit encryption and decryption

28.2 Architecture

Figure 28-1 shows a high-level block diagram of the SMIF hardware in PSoC 6 MCUs. Notice that the block is divided into multiple clock domains. This enables multiple domains to access the SMIF and still enable maintaining an asynchronous clock for the communication interface.


The SMIF block can also generate DMA triggers and interrupt signals. This allows events in the SMIF block to trigger actions in other parts of the system.

The SMIF interface is implemented using eight data lines, four slave select lines, and a clock line.

The access to the SMIF block can be by two modes: MMIO mode or XIP mode. The MMIO mode gives access to the SMIF's peripheral registers and the internal FIFOs. This mode is used when the user code is responsible for constructing the command structure for the external memory. Typically, this mode is used when the SMIF writes to an external flash memory. The MMIO interface is also used to configure the SMIF hardware block, including configuring the device registers that set up the XIP operation of the SMIF block.

The XIP mode of operation maps the external memory space to a range of addresses in the PSoC 6 MCU's address space. Refer to the [registers TRM](#) for details. When this address range is accessed, the hardware automatically generates the commands required to initiate the associated transfer from the external memory. The typical use case for the XIP mode is to execute code placed in external memory. Here, the memory will be mapped into the internal address space of the PSoC 6 MCU using the XIP mode. Thus code execution from external memory is seamless.

Figure 28-1. SMIF Hardware Block Diagram



The SMIF block has three AHB-Lite interfaces:

- An AHB-Lite interface to access the SMIF's MMIO registers.

- Two AHB-Lite interfaces to support execute-in-place (XIP).

All interfaces provide access to external memory devices. At any time, either the MMIO AHB-Lite interface or the two XIP AHB-Lite interfaces have access to the memory interface logic and external memory devices. The operation mode is specified by SMIFn_CTL.XIP_MODE. The operation mode should not be modified when the SMIF is busy (STATUS.BUSY is '1').

In the MMIO AHB-Lite interface, access is supported through software writes to transmit (Tx) FIFOs and software reads from receive (Rx) FIFOs. The FIFOs are mapped on SMIF registers. This interface provides the flexibility to implement any memory device transfer. For example, memory device transfers to setup, program, or erase the external memory devices.

In an XIP AHB-Lite interface, access is supported through XIP: AHB-Lite read and write transfers are automatically (by the hardware) translated in memory device read and write transfers. This interface provides efficient implementation of memory device read and write transfers, but does NOT support other types of memory device transfers. To improve XIP performance, the XIP AHB-Lite interface has a 4-KB read cache.

As mentioned, MMIO mode and XIP mode are mutually exclusive. The operation modes share Tx and Rx FIFOs and memory interface logic. In MMIO mode, the Tx and Rx FIFOs are accessed through the SMIF registers and under software control. In XIP mode, the Tx and Rx FIFOs are under hardware control. The memory interface logic is controlled through the Tx and Rx FIFOs and is agnostic of the operation mode.

28.2.1 Tx and Rx FIFOs

The SMIF block has two Tx FIFOs and one Rx FIFO. These FIFOs provide an asynchronous clock domain transfer between clk_hf logic and clk_if_tx/clk_if_rx memory interface logic. The memory interface logic is completely controlled through the Tx and Rx FIFOs.

- The Tx command FIFO transmits memory commands to the memory interface logic.
- The Tx data FIFO transmits write data to the memory interface transmit logic.
- The Rx data FIFO receives read data from the memory interface receive logic.

28.2.1.1 Tx Command FIFO

The Tx command FIFO consists of four 20-bit entries. Each entry holds a command. A memory transfer consists of a series of commands. In other words, a command specifies a phase of a memory transfer. Four different types of commands are supported:

- Tx command. A memory transfer must start with a Tx command. The Tx command includes a byte that is to be transmitted over the memory interface. The Tx command specifies the width of the data transfer (single, dual, quad, or octal data transfer). The Tx command specifies whether the command is for the last phase of the memory transfer (explicit "last command" indication). The Tx command specifies which of the four external devices are selected (multiple devices can be selected simultaneously); that is, the device selection as encoded by the Tx command is used for the complete memory transfer. The Tx command asserts the corresponding slave select lines. This is the reason every memory transfer should start with this command.
- TX_COUNT command. The TX_COUNT command specifies the number of bytes to be transmitted from the Tx data FIFO. This command relies on the Tx data FIFO to provide the bytes that are to be transmitted over the memory interface. The TX_COUNT command specifies the width of the data transfer and always constitutes the last phase of the memory transfer (implicit "last command" indication - de-asserts the slave select). Note that the TX_COUNT command does not assert the slave select lines. This must be done by a Tx command preceding it.
- RX_COUNT command. The RX_COUNT command specifies the number of bytes to be received from the Rx data FIFO. This command relies on the Rx data FIFO to accept the bytes that are received over the memory interface. The RX_COUNT command specifies the width of the data transfer and always constitutes the last phase of the memory transfer (implicit "last command" indication - de-asserts the slave select). Note that the RX_COUNT command does not assert the slave select lines. This must be done by a Tx command preceding it.
- DUMMY_COUNT command. The DUMMY_COUNT command specifies a number of dummy cycles. Dummy cycles are used to implement a turn-around (TAR) time in which the memory master changes from a transmitter driving the data lines to a receiver receiving on the same data lines. The DUMMY_COUNT command never constitutes the last phase of the memory transfer (implicit NOT "last command" indication - de-asserts the slave select); that is, it must be followed by another command. Note that the DUMMY_COUNT command does not assert the slave select lines. This must be done by a Tx command preceding it.

Together, the four command types can be used to construct any SPI transfer. The Tx command FIFO is used by both the memory interface transmit and receive logic. This ensures lockstep operation. The Tx command is a representation of a queue of commands that are to be processed.

The software will write the sequence of commands into the Tx command FIFO to generate a sequence responsible for the communication with slave device. The software can read

the number of used Tx command FIFO entries through the TX_CMD_FIFO_STATUS.USED[2:0] register field.

The software can write to the Tx command FIFO through the MMIO TX_CMD_FIFO_WR register. If software attempts to write to a full Tx command FIFO, the MMIO CTL.BLOCK field specifies the behavior:

- If CTL.BLOCK is '0', an AHB-Lite bus error is generated.
- If CTL.BLOCK is '1', the AHB-Lite write transfer is extended until an entry is available. This increases latency.

28.2.1.2 Tx Data FIFO

The Tx data FIFO consists of eight 8-bit entries. A Tx command FIFO TX_COUNT command specifies the number of bytes to be transmitted; that is, specifies the number of Tx data FIFO entries used. The Tx data FIFO is used by the memory interface transmit logic.

Software can read the number of used Tx data FIFO entries through the TX_DATA_FIFO_STATUS.USED[3:0] register field.

Software can write to the Tx data FIFO through the TX_DATA_FIFO_WR1, TX_DATA_FIFO_WR2, and TX_DATA_FIFO_WR4 registers:

- The TX_DATA_FIFO_WR1 register supports a write of a single byte to the FIFO.
- The TX_DATA_FIFO_WR2 register supports a write of two bytes to the FIFO.
- The TX_DATA_FIFO_WR4 register supports a write of four bytes to the FIFO. If software attempts to write more bytes than available entries in the Tx data FIFO, the MMIO CTL.BLOCK field specifies the behavior:
 - If CTL.BLOCK is '0', an AHB-Lite bus error is generated.
 - If CTL.BLOCK is '1', the AHB-Lite write transfer is extended until the required entries are available.

28.2.1.3 Rx Data FIFO

The Rx data FIFO consists of eight 8-bit entries. A Tx command FIFO RX_COUNT command specifies the number of bytes to be received; that is, specifies the number of Rx data FIFO entries used. The memory interface transmit logic will stop generating the SPI clock when the Rx data FIFO is full. This is how flow control is achieved.

Software can read the number of used Rx data FIFO entries through the RX_DATA_FIFO_STATUS.USED[3:0] register field.

Software can read from the Rx data FIFO through the MMIO RX_DATA_FIFO_RD1, RX_DATA_FIFO_RD2, and RX_DATA_FIFO_RD4 registers:

- The RX_DATA_FIFO_RD1 register supports a read of a single byte from the FIFO.

- The RX_DATA_FIFO_RD2 register supports a read of two bytes from the FIFO.
- The RX_DATA_FIFO_RD4 register supports a read of four bytes from the FIFO. If software attempts to read more bytes than available in the Rx data FIFO, the MMIO CTL.BLOCK field specifies the behavior:
 - If BLOCK is '0', an AHB-Lite bus error is generated and hard fault occurs.

If BLOCK is '1', the AHB-Lite read transfer is extended until the bytes are available.

Software can also read the first byte of the RX data FIFO without changing the status of the FIFO through the RX_DATA_FIFO_RD1_SILENT register.

28.2.2 MMIO Mode

If CTL.XIP_MODE is '0', the SMIF is in MMIO mode. Software generates SPI transfers by accessing the Tx FIFOs and Rx FIFO. Software writes to the Tx FIFOs and reads from the Rx FIFO. The Tx command FIFO has formatted commands (Tx, TX_COUNT, RX_COUNT, and DUMMY_COUNT) that are described in the [registers TRM](#).

Software should ensure that it generates correct memory transfers and accesses the FIFOs correctly. For example, if a memory transfer is generated to read four bytes from a memory device, software should read the four bytes from the Rx data FIFO. Similarly, if a memory transfer is generated to write four bytes to a memory device, software should write the four bytes to the Tx command FIFO or Tx data FIFO.

Incorrect software behavior can lock up the memory interface. For example, a memory transfer to read 32 bytes from a memory device, without software reading the Rx data FIFO will lock up the memory transfer as the memory interface cannot provide more than eight bytes to the Rx data FIFO (the Rx data FIFO has eight entries). This will prevent any successive memory transfers from taking place. Hence, the software should make sure that it read the FIFOs to avoid congestion. Note that a locked up memory transfer due to Tx or Rx FIFO states is still compliant to the memory bus protocol (but undesirable): the SPI protocol allows shutting down the interface clock in the middle of a memory transfer.

28.2.3 XIP Mode

If CTL.XIP_MODE is '1', the SMIF is in XIP mode. Hardware automatically (without software intervention) generates memory transfers by accessing the Tx FIFOs and Rx FIFO. Hardware supports only memory read and write transfers. Other functionality such as status reads are not supported. This means operations such as writing into a flash device may not be supported by XIP mode. This is because the writing operation into a flash memory involves not only a

write command transfer, but also a status check to verify the status of the operation.

- Hardware generates a memory read transfer for an AHB-Lite read transfer (to be precise: only for AHB-Lite read transfers that miss in the cache).
- Hardware generates a memory write transfer for an AHB-Lite write transfer.

Each slave device slot has a set of associated device configuration registers. To access a memory device in XIP mode, the corresponding device configuration registers (SMIFn_DEVICE_n) should be initialized. The device configuration register sets up the following parameters for memory:

- Write enable (WR_EN): Used to disable writes in XIP mode.
- Crypto enable (CRYPTO_EN): When enabled, all read access to the memory is decrypted and write access encrypted automatically.
- Data Select (DATA_SEL): Selects the data lines to be used ([Connecting SPI Memory Devices on page 348](#)).
- Base address and size: Sets up the mapped memory space. Any access in this space is converted to the access to the external memory automatically.
- Read and write commands: Used to communicate with the external memory device. These commands are defined by multiple settings.

As different memory devices support different types of memory read and write commands, you must provide the hardware with device specifics, such that it can perform the automatic translations. To this end, each memory device has a set of MMIO registers that specify its memory read and write transfers. This specification includes:

- Presence and value of the SPI command byte.
- Number of address bytes.
- Presence and value of the mode byte.
- Number of dummy cycles.
- Specified data transfer widths.

The XIP interface logic produces an AHB-Lite bus error under the following conditions:

- The SMIF is disabled (SMIFn_CTL.ENABLED is '0').
- The SMIF is not in XIP_MODE (SMIFn_CTL.XIP_MODE is '0').
- The transfer request is not in a memory region.
- The transfer is a write and the identified memory region does not support writes (SMIFn_DEVICE_n_CTL.WR_EN is '0').
- In XIP mode (CTL.XIP_MODE is '1') and dual quad SPI mode (ADDR_CTL.DIV2 is '1') or the transfer address is not a multiple of 2.

- In XIP mode (CTL.XIP_MODE is '1') and dual quad SPI mode (ADDR_CTL.DIV2 is '1'), the transfer size is not a multiple of 2.

28.2.4 Cache

To improve XIP performance, the XIP AHB-Lite interface has a cache. The cache is defined as follows:

- 4 KB capacity.
- Read-only cache. Write transfers bypass the cache. A write to an address, which is prefetched in the cache, invalidates the associated cache subsector. If there is a write to a memory in MMIO mode then you must invalidate the cache while switching back to XIP mode.
- Four-way set associative, with a least recently used (LRU) replacement scheme.

Each XIP interface implements a 4-KB cache memory, enabled by default. Any XIP access can be cached if a cache is enabled. There are separate cache registers for the slow cache (in the clk_slow domain) and fast cache (in the clk_fast domain). The cache can be enabled using the SLOW_CA_CTL[ENABLED] or FAST_CA_CTL[ENABLED] registers. Read transfers that "hit" are processed by the cache. Read transfers that "miss" result in a XIP memory read transfer.

If CA_CTL.PREF_EN is '1', prefetching is enabled and if CA_CTL.PREF_EN is '0', prefetching is disabled. If prefetch is enabled, a cache miss results in a 16 B (subsector) refill for the missing data AND a 16 B prefetch for the next sequential data (independent of whether this data is already in the cache or not).

Cache coherency is not supported by the hardware. For example, an XIP interface 0 write to an address in the XIP interface 0 cache invalidates (clears) the associated cache subsector in the XIP interface 0 cache, but not in the XIP interface 1 cache. This means XIP interface 1 cache now has outdated data. The user code can manually invalidate cache by using the SLOW_CA_CMD[INV] or FAST_CA_CMD[INV] register.

Caches should also be invalidated upon mode transitions. For example, in MMIO mode, a write to an address in the cache interface will cause the data in the cache interface to be outdated. The cache should be invalidated when transitioning to XIP mode to ensure that only valid data is used.

28.2.5 Arbitration

The SMIF provides two AHB-Lite slave interfaces to CPUSS (one fast interface and one slow interface). Both interfaces have a cache (as described in [28.2.4 Cache](#)) and can generate XIP requests to the external memory devices.

An arbitration component (as shown in [Figure 28-1](#)) arbitrates between the two interfaces. Arbitration is based

on the master identifier of the AHB-Lite transfer. The arbitration priority is specified by a system wide priority. Each master identifier has a 2-bit priority level ("0" is the highest priority level and "3" is the lowest priority level). Master identifiers with the same priority level are within the same priority group. Within a priority group, round-robin arbitration is performed.

28.2.6 Deselect Delay

The SMIF supports configuration of deselect delay between transfers. The SMIFn_CTL.DESELECT_DELAY field controls the minimum number of interface cycles to hold the chip select line inactive.

28.2.7 Cryptography

In XIP mode, a cryptography component supports on-the-fly encryption for write data and on-the-fly decryption for read data. The use of on-the-fly cryptography is determined by a device's MMIO CTL.CRYPTO_EN field. In MMIO mode, the cryptography component is accessible through a register interface to support offline encryption and decryption.

The usage scenario for cryptography is: data is encrypted in the external memory devices. Therefore, memory read and

write data transfers require decryption and encryption functionality respectively. By storing data encrypted in the external memory devices (nonvolatile devices), leakage of sensitive data is avoided.

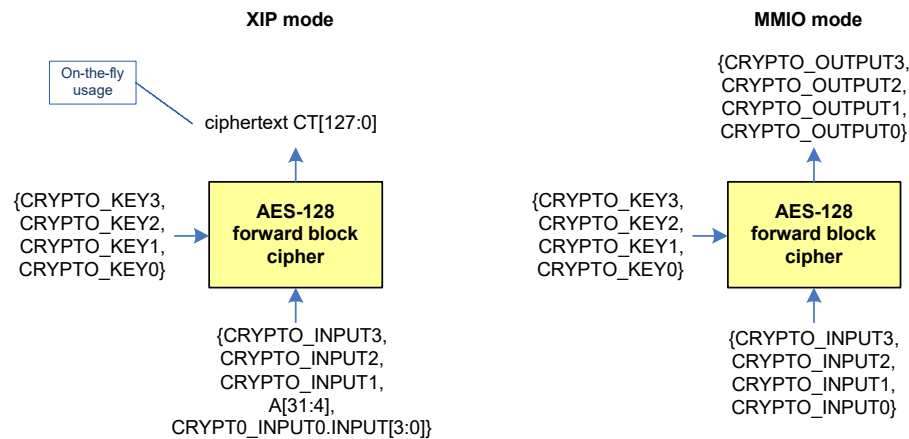
Encryption and decryption are based on the AES-128 forward block cipher: advanced encryption standard block cipher with a 128-bit key. KEY[127:0] is a secret (private) key programmed into the CRYPTO_KEY3, ..., CRYPTO_KEY0 registers. These registers are software write-only: a software read returns "0". In the SMIF hardware, by applying AES-128 with KEY[127:0] on a plaintext PT[127:0], we get a ciphertext CT[127:0].

In XIP mode, the XIP address is used as the plaintext PT[]. The resulting ciphertext CT[] is used on-the-fly and not software accessible. The XIP address is extended with the CRYPTO_INPUT3, ..., CRYPTO_INPUT0 registers.

In MMIO mode, the MMIO CRYPTO_INPUT3, ..., CRYPTO_INPUT0 registers provide the plaintext PT[]. The resulting ciphertext CT[] is provided through the MMIO CRYPTO_OUTPUT3, ..., CRYPTO_OUTPUT0 registers.

Figure 28-2 illustrates the functionality in XIP and MMIO modes.


Figure 28-2. Cryptography in XIP and MMIO Modes



In XIP mode, the resulting ciphertext CT[] (of the encrypted address) is XOR'd with the memory transfer's read data or write data. Note that the AES-128 block cipher is on the address of the data and not on the data itself. For memory read transfers, this means that as long as the latency of the memory transfer's read data is longer than the AES-128 block cipher latency, the on-the-fly decryption does not add any delay. Figure 28-3 illustrates the complete XIP mode functionality.

The XIP mode only encrypts the address and XORs with the data; to implement the same in MMIO, you must provide the address as the PT[] into the crypto_INPUTx registers.

Figure 28-3. XIP Mode Functionality



28.3 Memory Device Signal Interface

The SMIF acts as a master for SPI applications. SPI requires the definition of clock polarity and phase. In SPI mode, the SMIF supports a single clock polarity and phase configuration:

- Clock polarity (CPOL) is '0': the base value of the clock is 0.
- Clock phase (CPHA) is '0': driving of data is on the falling edge of the clock; capturing of data is specified by CTL.CLOCK_IF_RX_SEL.

The above configuration is also known as SPI configuration 0 and is supported by SPI memory devices.

28.3.1 Specifying Memory Devices

The SMIF requires that the memory devices are defined for their operation in XIP mode. The SMIF supports up to four memory devices. Each memory device is defined by a set of registers. The memory device specific register structure includes:

- The device base address and capacity. The ADDR register specifies the memory device's base address in the PSoC 6 MCU address space and the MASK register specifies the memory device's size/capacity. If a memory device is not present or is disabled, the ADDR and MASK registers specify a memory device with 0 B capacity. Typically, the devices' address regions in the PSoC 6 MCU address space are non-overlapping (except for dual-quad SPI mode) to ensure that the activation of select signals is mutually exclusive.
- The device data signal connections (as described in [Connecting SPI Memory Devices on page 348](#)).
- The definition of a read transfer to support XIP mode.
- The definition of a write transfer to support XIP mode.

Each memory device uses a dedicated device select signal: memory device 0 uses select[0], memory device 1 uses select[1], and so on. In other words, there is a fixed, one-to-one connection between memory device, register set, and select signal connection.

In XIP mode, the XIP AHB-Lite bus transfer address is compared with the device region. If the address is within the device region, the device select signal is activated. If an XIP AHB-Lite bus transfer address is within multiple regions (this is possible if the device regions overlap), all associated device select signals are activated. This overlap enables XIP in dual-quad SPI mode: the command, address, and mode bytes can be driven to two quad SPI devices simultaneously.

In XIP mode, dual quad SPI mode requires the ADDR_CTL.DIV2 field of the selected memory devices to be set to '1'. When this field is '1', the transfer address is divided by 2 and the divided by 2 address is provided to the memory devices.

In dual quad SPI mode, each memory device contributes a 4-bit nibble for each 8-bit byte. However, both memory devices are quad SPI memories with a byte interface. Therefore, the transfer size must be a multiple of 2.

The XIP_ALIGNMENT_ERROR interrupt cause is set under the following conditions (in XIP mode and when ADDR_CTL.DIV2 is '1'):

- The transfer address is not a multiple of 2. In this case the divided by 2 address for the memory devices is incorrect.
- The transfer size is not a multiple of 2. In this case, the memory devices contribute only a nibble of a byte. This is not supported as the memory devices have a byte interface.

28.3.2 Connecting SPI Memory Devices

Memory device I/O signals (SCK, $\overline{\text{CS}}$, SI/IO0, SO/IO1, IO2, IO3, IO4, IO5, IO6, IO7) are connected to the SMIF I/O signals (clk, select[3:0], and spi_data[7:0]). Not all memory devices provide the same number of I/O signals.

Table 28-1. Memory Device I/O Signals

Memory Device	IO Signals
Single SPI memory	SCK, $\overline{\text{CS}}$, SI, SO. This memory device has two data signals (SI and SO).
Dual SPI memory	SCK, $\overline{\text{CS}}$, IO0, IO1. This memory device has two data signals (IO0, IO1).
Quad SPI memory	SCK, $\overline{\text{CS}}$, IO0, IO1, IO2, IO3. This memory device has four data signals (IO0, IO1, IO2, IO3).
Octal SPI memory	SCK, $\overline{\text{CS}}$, IO0, IO1, IO2, IO3, IO4, IO5, IO6, IO7. This memory device has eight data signals (IO0, IO1, IO2, IO3, IO4, IO5, IO6, IO7).

Table 28-1 illustrates that each memory has a single clock signal SCK, a single (low active) select signal ($\overline{\text{CS}}$), and multiple data signals (IO0, IO1, ...).

Each memory device has a fixed select signal connection (to select[3:0]).

Each memory device has programmable data signal connections (to data[7:0]): the CTL.DATA_SEL[1:0] field specifies how a device's data signals are connected. The CTL.DATA_SEL[1:0] is responsible for configuring the selection of data lines to be used by a slave. This is not to be confused with the select lines that are used for addressing the four slaves of the SMIF master. This information is used by the SMIF interface to drive out data on the correct spi_data[] outputs and capture data from the correct spi_data[] inputs. If multiple device select signals are activated, the same data is driven to all selected devices simultaneously.

Not all data signal connections are legal/supported. Supported connections are dependent on the type of memory device.

Table 28-2. Data Signal Connections

DATA_SEL[1:0]	Single SPI Device	Dual SPI Device	Quad SPI Device	Octal SPI Device
0	spi_data[0] = SI spi_data[1] = SO	spi_data[0] = IO0 spi_data[1] = IO1	spi_data[0] = IO0 ... spi_data[3] = IO3	spi_data[0] = IO0 ... spi_data[7] = IO7
1	spi_data[2] = SI spi_data[3] = SO	spi_data[2] = IO0 spi_data[3] = IO1	Illegal	Illegal
2	spi_data[4] = SI spi_data[5] = SO	spi_data[4] = IO0 spi_data[5] = IO1	spi_data[4] = IO0 ... spi_data[7] = IO3	Illegal
3	spi_data[6] = SI spi_data[7] = SO	spi_data[6] = IO0 spi_data[7] = IO1	Illegal	Illegal

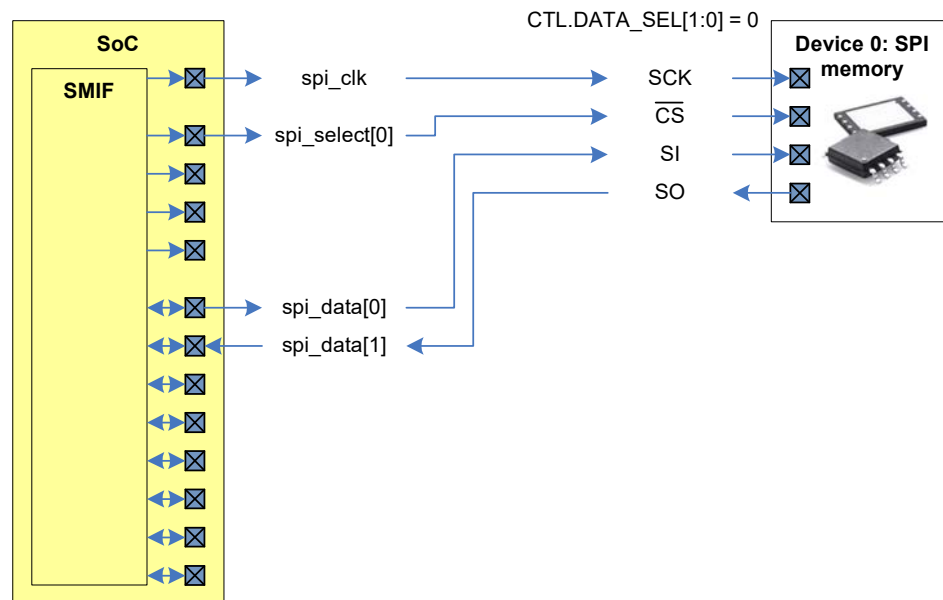
Memory devices can:

- Use shared data signal connections.
- Use dedicated data signal connections. This reduces the load on the data lines allowing faster signal level changes, which in turn allows for a faster I/O interface.

Note that dual-quad SPI mode requires dedicated data signals to enable read and/or write data transfer from and to two quad SPI devices simultaneously.

Figure 28-4 illustrates memory device 0, which is a single SPI memory with data signals connections to spi_data[1:0].

Figure 28-4. Single SPI Memory Device 0 Connected to spi_data[1:0]



Because of the pin layout, you might want to connect a memory device to specific data lines. Figure 28-5 illustrates memory device 0, which is a single SPI memory with data signals connections to spi_data[7:6].

Figure 28-5. Single SPI Memory Device 0 Connected to spi_data[7:6]




Figure 28-6 illustrates memory devices 0 and 1, both of which are single SPI memories. Each device uses dedicated data signal connections. The device address regions in the PSoC 6 MCU address space must be non-overlapping to ensure that the activation of select[0] and select[1] are mutually exclusive.

Figure 28-6. Single SPI Memory Devices 0 and 1 - Dedicated Data Signal




Figure 28-7 illustrates memory devices 0 and 1, both of which are single SPI memories. Both devices use shared data signal connections. The devices' address regions in the PSoC 6 MCU address space must be non-overlapping to ensure that the activation of `select[0]` and `select[1]` are mutually exclusive. Note that this solution increases the load on the data lines, which may result in a slower I/O interface.

Figure 28-7. Single SPI Memory Devices 0 and 1 - Shared Data Signal

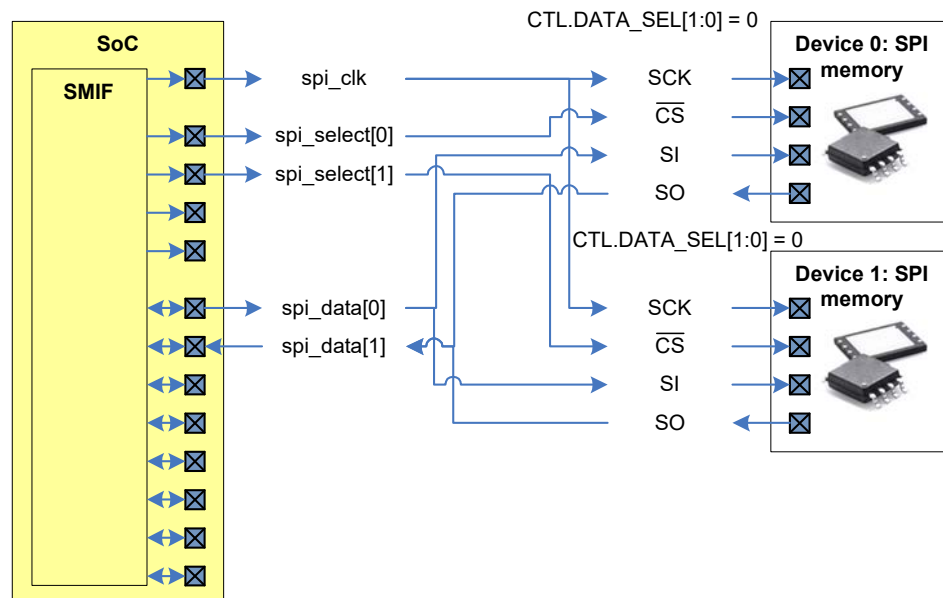


Figure 28-8 illustrates memory device 0, which is a quad SPI memory with data signals connections to `spi_data[7:4]`.

Figure 28-8. Quad SPI Memory Device 0

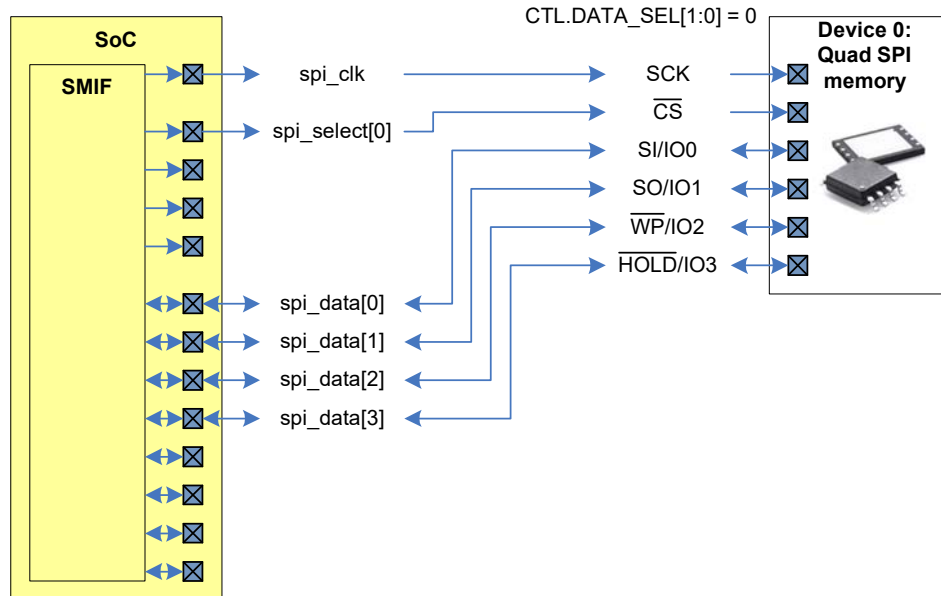


Figure 28-9 illustrates memory devices 0 and 1, device 0 is a single SPI memory and device 1 is a quad SPI memory. Each device uses dedicated data signal connections. The device address regions in the PSoC 6 MCU address space must be non-overlapping to ensure that the activation of select[0] and select[1] are mutually exclusive.

Figure 28-9. Single SPI Memory 0 and Quad SPI Memory 1 - Dedicated Data Signal

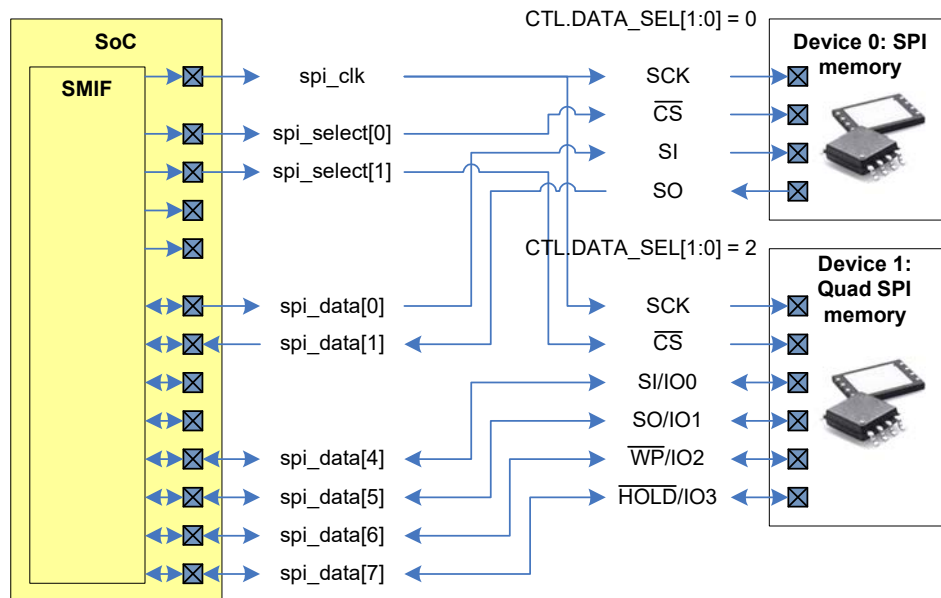


Figure 28-10 illustrates memory devices 0 and 1, device 0 is a single SPI memory and device 1 is a quad SPI memory. Both devices use shared data signal connections. The device address regions in the PSoC 6 MCU address space must be non-overlapping to ensure that the activation of select[0] and select[1] are mutually exclusive.

Figure 28-10. Single SPI Memory Device 0 and Quad SPI Memory Device 1 - Shared Data Signal

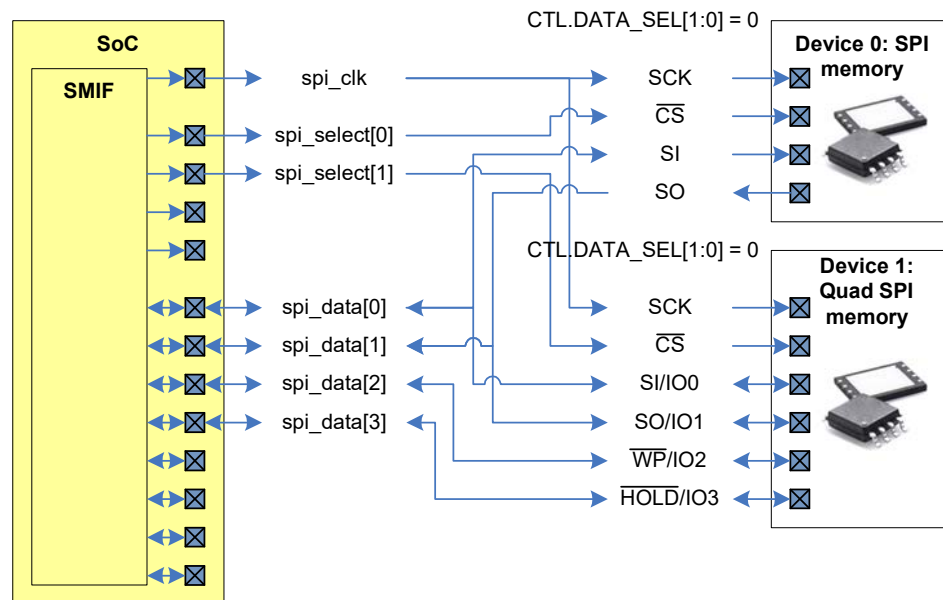


Figure 28-11 illustrates memory devices 0 and 1, both of which are quad SPI memories. Each device uses dedicated data signal connections. The device address regions in the PSoC 6 MCU address space are the same to ensure that the activation of select[0] and select[1] are the same (in XIP mode). This is known as a dual-quad configuration: during SPI read and write transfers, each device provides a nibble of a byte.

Figure 28-11. Quad SPI Memory Devices 0 and 1


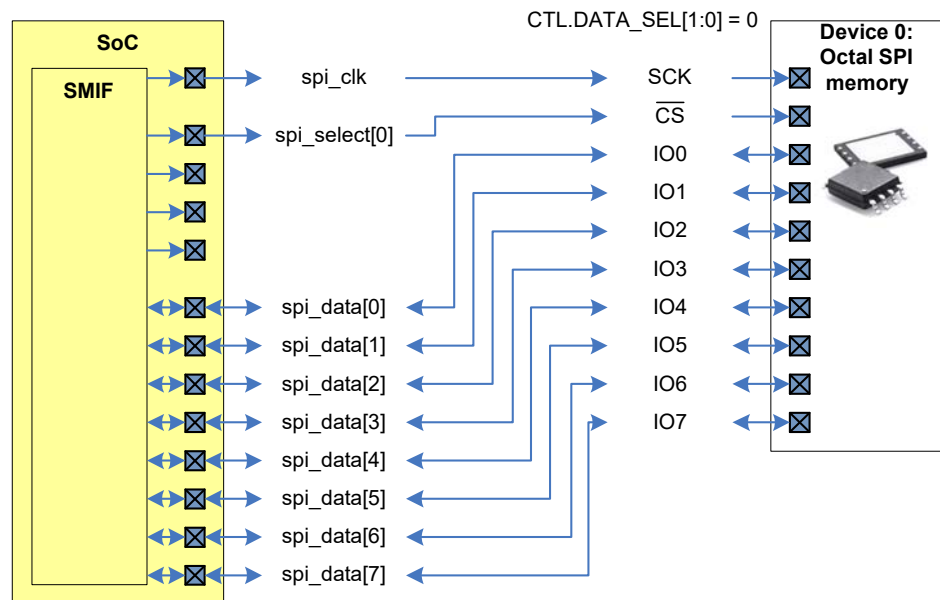


Figure 28-12 illustrates memory device 0, which is a octal SPI memory with data signals connections to spi_data[7:0].

Figure 28-12. Octal SPI Memory Device 0



28.3.3 SPI Data Transfer

SPI data transfer uses most-significant-bit (MSb) for the first data transfer. This means that for a byte B, consisting of bits b7, b6, ..., b0, bit b7 is transferred first, followed by bit b6, and so on. For dual, quad, dual quad, and octal SPI transfers, multiple bits are transferred per cycle. For a single SPI device and device data signal connections to spi_data[1:0] (DATA_SEL is "0"), Table 28-3 summarizes the transfer of byte B.

Table 28-3. Single Data Transfer

Cycle	Data Transfer
0	For a write transfer: b7 is transferred on data[0] and SI/IO0. For a read transfer: b7 is transferred on data[1] and SO/IO1.
1	For a write transfer: b6 is transferred on data[0] and SI/IO0. For a read transfer: b6 is transferred on data[1] and SO/IO1.
2	For a write transfer: b5 is transferred on data[0] and SI/IO0. For a read transfer: b5 is transferred on data[1] and SO/IO1.
3	For a write transfer: b4 is transferred on data[0] and SI/IO0. For a read transfer: b4 is transferred on data[1] and SO/IO1.
4	For a write transfer: b3 is transferred on data[0] and SI/IO0. For a read transfer: b3 is transferred on data[1] and SO/IO1.
5	For a write transfer: b2 is transferred on data[0] and SI/IO0. For a read transfer: b2 is transferred on data[1] and SO/IO1.
6	For a write transfer: b1 is transferred on data[0] and SI/IO0. For a read transfer: b1 is transferred on data[1] and SO/IO1.
7	For a write transfer: b0 is transferred on data[0] and SI/IO0. For a read transfer: b0 is transferred on data[1] and SO/IO1.

Note that in single SPI data transfer, the data signals are uni-directional: in the table, data[0] is exclusively used for write data connected to the device SI input signal and data[1] is exclusively used for read data connected to the device SO output signal.

For a dual SPI device and device data signal connections to data[1:0] (DATA_SEL is "0"), [Table 28-4](#) summarizes the transfer of byte B.

Table 28-4. Dual Data Transfer

Cycle	Data Transfer
0	b7, b6 are transferred on data[1:0] and IO1, IO0.
1	b5, b4 are transferred on data[1:0] and IO1, IO0.
2	b3, b2 are transferred on data[1:0] and IO1, IO0.
3	b1, b0 are transferred on data[1:0] and IO1, IO0.

For a quad SPI device and device data signal connections to data[3:0] (DATA_SEL is "0"), [Table 28-5](#) summarizes the transfer of byte B.

Table 28-5. Quad Data Transfer

Cycle	Data Transfer
0	b7, b6, b5, b4 are transferred on data[3:0] and IO3, IO2, IO1, IO0.
1	b3, b2, b1, b0 are transferred on data[3:0] and IO3, IO2, IO1, IO0.

For an octal SPI device and device data signal connections to data[7:0] (DATA_SEL is "0"), [Table 28-6](#) summarizes the transfer of byte B.

Table 28-6. Octal Data Transfer

Cycle	Data Transfer
0	b7, b6, b5, b4, b3, b2, b1, b0 are transferred on data[7:0] and IO7, IO6, IO5, IO4, IO3, IO2, IO1, IO0.

In dual-quad SPI mode, two quad SPI devices are used.

- The first device (the device with the lower device structure index) should have device data signal connections to data[3:0] (DATA_SEL is 0).
- The second device (the device with the higher device structure index) should have device data signal connection to data[7:4] (DATA_SEL is 2).

The command and data phases of the SPI transfer use different width data transfers:

- The command, address, and mode bytes use quad SPI data transfer.
- The read data and write data use octal data transfer. Each device provides a nibble of each data byte: the first device provides the lower nibble and the second device provides the higher nibble.

[Table 28-7](#) summarizes the transfer of a read data and write data byte B.


Table 28-7. Dual-quad SPI Mode, Octal Data Transfer

Cycle	Data transfer
0	b7, b6, b5, b4 are transferred on data[7:4] and second device IO3, IO2, IO1, IO0. b3, b2, b1, b0 are transferred on data[3:0] and first device IO3, IO2, IO1, IO0.

28.3.4 Example of Setting up SMIF

Devices 0 and 1 are used to implement the dual-quad SPI mode. Both devices are 1 MB / 8 Mb; the address requires 3 bytes. Device 0 has device data signal connections to data[3:0] and device 1 has device data signal connections to data[7:4].

Figure 28-13. Setting up SMIF



For dual quad SPI mode, the AHB-Lite bus transfer address is divided by two. Cryptography and write functionality are disabled in the following example.


```
#define MASK_1MB 0xffff00000;
DEV0_ADDR      = CPUSS_SMIF_BASE;
DEV0_MASK      = MASK_1MB      // MASK: 1 MB region
DEV0_CTL       =(0 << SMIF_DEVICE_CTL_DATA_SEL_Pos)      // DATA_SEL: data[3:0]
                | (0 << SMIF_DEVICE_CTL_CRYPTO_EN_Pos)   // CRYPTO_EN
                | (0 << SMIF_DEVICE_CTL_WR_EN_Pos));      // WR_EN
DEV0_ADDR_CTL  = (1 << SMIF_DEVICE_ADDR_CTL_DIV2_Pos)     // DIV2: enabled
                | ((3-1) << SMIF_DEVICE_ADDR_CTL_SIZE2_Pos)); // SIZE: 3 B address

DEV1_ADDR      = CPUSS_SMIF_BASE;
DEV1_MASK      = 0xffff00000; // MASK: 1 MB region
DEV1_CTL       = (2 << SMIF_DEVICE_CTL_DATA_SEL_Pos)     // DATA_SEL: data[7:4]
                | (0 << SMIF_DEVICE_CTL_CRYPTO_EN_Pos)   // CRYPTO_EN
                | (1 << SMIF_DEVICE_CTL_WR_EN_Pos));      // WR_EN
DEV1_ADDR_CTL  = (1 << SMIF_DEVICE_ADDR_CTL_DIV2_Pos)     // DIV2: enabled
                | ((3-1) << SMIF_DEVICE_ADDR_CTL_SIZE2_Pos)); // SIZE: 3 B address
```

For XIP read transfers, the 0xEB command/instruction is used (Figure 28-14 illustrates a two-byte transfer from devices 0 and 1 in dual quad SPI mode).

Figure 28-14. Two-Byte Transfer in Dual Quad SPI Mode

0xEB instruction, instruction 1 bit/cycle; address, mode, data 4 bits/cycle



The definition of a read transfer is as follows:

```
DEV0_RD_CMD_CTL = (1UL << SMIF_DEVICE_RD_CMD_CTL_PRESENT_Pos) // PRESENT
                  | (0 << SMIF_DEVICE_RD_CMD_CTL_WIDTH_Pos) // WIDTH: single data transfer
                  | 0xeb); // CODE
DEV0_RD_ADDR_CTL = (2 << SMIF_DEVICE_RD_ADDR_CTL_WIDTH_Pos)); // WIDTH: quad data transfer
DEV0_RD_MODE_CTL = (1UL << SMIF_DEVICE_RD_MODE_CTL_PRESENT_Pos) // PRESENT
                  | (2 << SMIF_DEVICE_RD_MODE_CTL_WIDTH_Pos) // WIDTH: quad data transfer
                  | 0x00); // CODE
DEV0_RD_DUMMY_CTL= (1UL << SMIF_DEVICE_RD_DUMMY_CTL_PRESENT_Pos) // PRESENT
                  | ((4-1) << SMIF_DEVICE_RD_DUMMY_CTL_SIZE5_Pos)); // SIZE: 4 dummy cycles
DEV0_RD_DATA_CTL = (3 << SMIF_DEVICE_RD_DATA_CTL_WIDTH_Pos)); // WIDTH: octal data transfer
```

Note that the command uses single data transfer, the address and mode byte use quad data transfer, and the read data byte uses octal data transfer.

28.4 Triggers

The SMIF has two level-sensitive triggers:

- tr_tx_req is associated with the Tx data FIFO.
- tr_rx_req is associated with the Rx data FIFO.

If the SMIF is enabled (CTL.ENABLED is '1') and MMIO operation mode is selected (CTL.XIP_MODE is '0'), the trigger functionality is enabled. If the SMIF is disabled (CTL.ENABLED is '0') or the XIP operation mode is selected (CTL.XIP_MODE is '1'), the triggers functionality is disabled. The trigger functionality is defined as follows:

- The MMIO TX_DATA_FIFO_CTL.TRIGGER_LEVEL field specifies a number of FIFO entries. The tr_tx_req trigger is active when the number of used Tx data FIFO entries is smaller or equal than the specified number; that is, TX_DATA_FIFO_STATUS.USED ≤ TRIGGER_LEVEL.
- The MMIO RX_DATA_FIFO_CTL.TRIGGER_LEVEL field specifies a number of FIFO entries. The tr_rx_req trigger is active when the number of used Rx data FIFO entries is greater than the specified number; that is, RX_DATA_FIFO_STATUS.USED > TRIGGER_LEVEL.

28.5 Interrupts

The SMIF has a single interrupt output with six interrupt causes:

- INTR.TR_TX_REQ. This interrupt cause is activated in MMIO mode when the tr_tx_req trigger is activated.
- INTR.TR_RX_REQ. This interrupt cause is activated in MMIO mode when the tr_rx_req trigger is activated.
- INTR.XIP_ALIGNMENT_ERROR. This interrupt cause is activated in XIP mode when the selected device's ADDR_CTL.DIV2 field is '1' and the AHB-Lite bus address is not a multiple of 2, or the requested transfer size is not a multiple of 2. This interrupt cause identifies erroneous behavior in dual-quad SPI mode (the selected device ADDR_CTL.DIV2 field is set to '1').
- INTR.TX_CMD_FIFO_OVERFLOW. This interrupt cause is activated in MMIO mode, on an AHB-Lite write transfer to the Tx command FIFO (TX_CMD_FIFO_WR) with insufficient free entries.
- INTR.TX_DATA_FIFO_OVERFLOW. This interrupt cause is activated in MMIO mode, on an AHB-Lite write transfer to the Tx data FIFO (TX_DATA_FIFO_WR1, TX_DATA_FIFO_WR2, and TX_DATA_FIFO_WR4) with insufficient free entries.
- INTR.RX_DATA_FIFO_OVERFLOW. This interrupt cause is activated in MMIO mode, on an AHB-Lite read transfer from the Rx data FIFO (RX_DATA_FIFO_RD1, RX_DATA_FIFO_RD2, and RX_DATA_FIFO_RD4) with insufficient free entries.

28.6 Sleep Operation

The SMIF hardware is operational in the Sleep and Active power modes. In the Sleep power mode, only DMA driven transactions can be performed because the CPU acting as the SPI master is not active.

28.7 Performance

Accesses to the external memory will have some latency, which is dependent upon the mode of SMIF operation, the amount of data being transferred, caching, and cryptography. In MMIO mode, the number of interface clock cycles per transfer is determined by the equation:

$$\text{Time to Transfer N bytes (in interface clock cycles)} = \left(\left(\frac{\text{opcode size}}{\text{opcode width}} \right) + \left(\frac{\text{address size}}{\text{address width}} \right) + \left(\frac{\text{mode size}}{\text{mode width}} \right) + \text{dummyscycles} + \left(\frac{\text{data size}}{\text{data width}} \right) \right)$$

For example, in [Figure 28-14](#), the equation to calculate the number of cycles would be (22 cycles = [8 bit instruction/1 single width] + [24 bit address/4 width] + [8 bit mode/4 width] + 4 dummy cycles + [16 bit data/8 width]).

In XIP Mode, the performance is affected by the cache. For data reads that hit in the cache, the read will not incur any interface cycles. Read operations that miss in the cache will occur as a normal SMIF read operation and, if prefetching is enabled, will result in two 16 B cache sub-sector refills. Writes to external memory in XIP mode will occur as a normal SMIF write operation.

Enabling cryptography may impact SMIF performance. The AES-128 block cipher has a typical latency of 13 clk_hf cycles. This means that for transfers that take more than 13 cycles, the on-the-fly decryption does not add any delay. If the transfer is less than 13 cycles, the transfer latency will be 13 cycles. When the cache is enabled, the 13-cycle latency for encryption is incurred only once for every 16 B fetched for the cache.

29. Timer, Counter, and PWM (TCPWM)



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - TCPWM
- Application notes
- Code examples

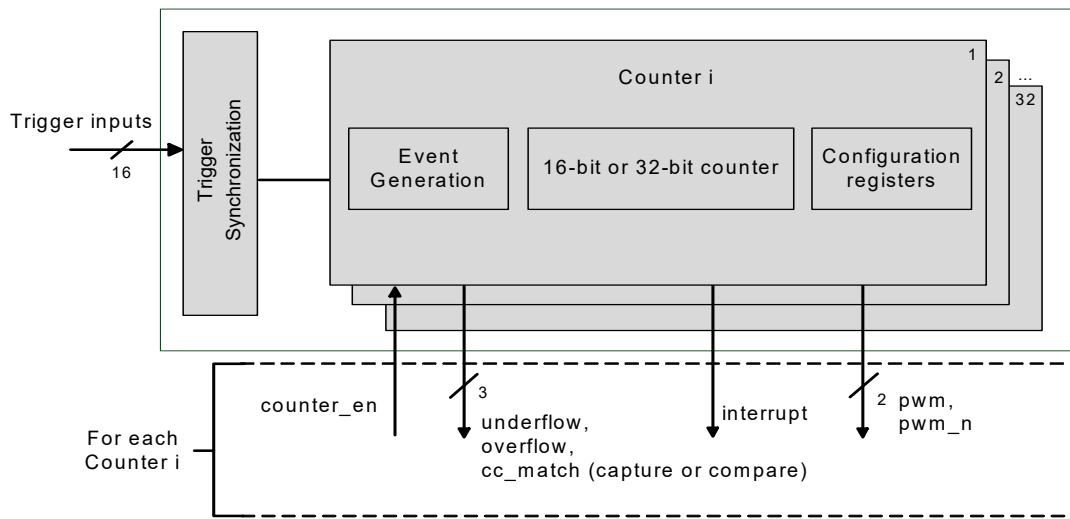
The Timer, Counter, and Pulse Width Modulator (TCPWM) block in the PSoC 6 MCU implements a 16-bit or 32-bit timer, counter, pulse width modulator (PWM), or quadrature decoder functionality. The block can be used to measure the period and pulse width of an input signal (timer), find the number of times a particular event occurs (counter), generate PWM signals, or decode quadrature signals. This chapter explains the features, implementation, and operational modes of the TCPWM block.

29.1 Features

- The TCPWM block supports the following operational modes:
 - Timer-counter with compare
 - Timer-counter with capture
 - Quadrature decoding
 - Pulse width modulation
 - Pseudo-random PWM
 - PWM with dead time
- Up, Down, and Up/Down counting modes.
- Clock prescaling (division by 1, 2, 4, ... 64, 128)
- Double buffering of compare/capture and period values
- Underflow, overflow, and capture/compare output signals
- Supports interrupt on:
 - Terminal count – Depends on the mode; typically occurs on overflow or underflow
 - Capture/compare – The count is captured to the capture register or the counter value equals the value in the compare register
- Complementary output for PWMs
- Selectable start, reload, stop, count, and capture event signals for each TCPWM with rising edge, falling edge, both edges, and level trigger options

29.2 Architecture

Figure 29-1. TCPWM Block Diagram



The TCPWM block in a PSoC 6 MCU can contain up to 32 counters. Each counter can be 16- or 32-bit wide. The number of 16- and 32-bit counters are device specific; refer to the [device datasheet](#) for details.

In this chapter, a TCPWM refers to the entire block and all the counters inside. A counter refers to the individual counter inside the TCPWM.

TCPWM has these interfaces:

- I/O signal interface: Consists of input triggers (such as reload, start, stop, count, and capture) and output signals (such as pwm, pwm_n, overflow (OV), underflow (UN), and capture/compare (CC)).
- Interrupts: Provides interrupt request signals from each counter, based on TC or CC conditions.

The TCPWM block can be configured by writing to the TCPWM registers. See [“TCPWM Registers” on page 393](#) for more information on all registers required for this block.

29.2.1 Enabling and Disabling Counters in TCPWM Block

A counter can be enabled by setting the corresponding bit of the TCPWM_CTRL_SET register. It can be disabled by setting the bit in the TCPWM_CTRL_CLR register. These registers are used to avoid race-conditions on read-modify-write attempts to the TCPWM_CTRL register, which controls the enable/disable fields of the counters.

Note: The counter must be configured before enabling it. Disabling the counter retains the values in the registers.

29.2.2 Clocking

Each TCPWM counter can have its own clock source. The only source for the clock is from the configurable peripheral clock dividers generated by the clocking system; see the [Clocking System chapter on page 207](#) for details. To select a clock divider for a particular counter inside a TCPWM, use the CLOCK_CTL register from the PERI register space. In this section the clock to the counter will be called clk_counter. Event generation is performed on the clk_counter. Another clock, clk_sys is used for the pulse width of the output triggers. clk_sys is synchronous to clk_peri (see [“CLK_PERI” on page 219](#)), but can be divided using CLOCK_CTL from the PERI_GROUP_STRUCT register space.

29.2.2.1 Clock Prescaling

clk_counter can be further divided inside each counter, with values of 1, 2, 4, 8...64, 128. This division is called prescaling. The prescaling is set in the GENERIC field of the TCPWM_CNT_CTLR register.

Note: Clock prescaling is not available in quadrature mode and pulse width modulation mode with dead time.

29.2.2.2 Count Event

The counter functionality is performed on an “active count” prescaled clock, which is gated by a “count event” signal as shown in [Figure 29-2](#). For example, a counter increments or decrements by ‘1’ every clk_counter cycle in which a count event is detected. Count event generation depends on edge detection modes on count inputs.


When the count input is configured as level, the count value is changed on each prescaled clk_counter edge in which the count input is high.

When the count input is configured as rising/falling the count value is changed on each prescaled `clk_counter` edge in which an edge is detected on the count input.

The next section contains additional details on edge detection modes configuration.

Note: Count events are not supported in quadrature and pulse-width modulation pseudo-random modes; the `clk_counter` is used in these cases instead of the active count prescaled clock.

Figure 29-2. Counter Clock Generation



29.2.3 Trigger Inputs

Each TCPWM block has 14 `Trigger_In` signals, which come from other on-chip resources such as other TCPWMs, SCBs, or DMA. The `Trigger_In` signals are shared with all counters inside of one TCPWM block. Use the Trigger Mux registers to configure which signals get routed to the `Trigger_In` for each TCPWM block. See the [Trigger Multiplexer Block chapter on page 260](#) for more details. Two constant trigger inputs '0' and '1' are available in addition to the 14 `Trigger_In`. For each counter, the trigger input source is selected using the `TCPWM_CNT_TR_CTRL0` register.

Each counter can select any of the 16 trigger signals to be the source for any of the following events:

- Reload
- Start
- Stop/Kill
- Count
- Capture/swap

The `TCPWM_CMD_RELOAD`, `TCPWM_CMD_STOP`, `TCPWM_CMD_START`, and `TCPWM_CMD_CAPTURE` registers can be used to trigger the reload, stop, start, and capture respectively from software.

The sections describing each TCPWM mode will describe the function of each input event in detail.

Typical operation uses the reload event once to initialize and start the counter and the stop event to stop the counter. When the counter is stopped, the start event can be used to start the counter with its counter value unmodified from when it was stopped.

If stop, reload, and start events coincide, the following precedence relationship holds:

- A stop event has higher priority than a reload event.
- A reload event has higher priority than a start event.

As a result, when a reload or start event coincides with a stop event, the reload or start event has no effect.


Before going to the counter each `Trigger_IN` can pass through a positive edge detector, negative edge detector, both edge detector, or pass straight through to the counter.

This is controlled using `TCPWM_CNT_TR_CTRL1`. In the quadrature mode, edge detection is done using `clk_counter`. For all other modes, edge detection is done using the `clk_peri`.

Multiple detected events are treated as follows:

- In the rising edge and falling edge modes, multiple events are effectively reduced to a single event. As a result, events may be lost.
- In the rising/falling edge mode, an even number of events are not detected and an odd number of events are reduced to a single event. This is because the rising/falling edge mode is typically used for capture events to determine the width of a pulse. The current functionality will ensure that the alternating pattern of rising and falling is maintained.

Figure 29-3. TCPWM Input Events


Notes:

- All trigger inputs are synchronized to `clk_peri`.
- When more than one event occurs in the same `clk_counter` period, one or more events may be missed. This can happen for high-frequency events (frequencies close to the counter frequency) and a timer configuration in which a pre-scaled (divided) `clk_counter` is used.

29.2.4 Trigger Outputs

Each counter can generate three trigger output events. These trigger output events can be routed through the trigger mux to other peripherals on the device. The three trigger outputs are:

- Overflow (OV): An overflow event indicates that in up-counting mode, COUNTER equals the PERIOD register, and is changed to a different value.
- Underflow (UN): An underflow event indicates that in a down-counting mode, COUNTER equals 0, and is changed to a different value.

- Compare/Capture (CC): This event is generated when the counter is running and one of the following conditions occur:
 - Counter equals the compare value. This event is either generated when the match is about to occur (COUNTER does not equal CC and is changed to CC) or when the match is not about to occur (COUNTER equals CC and is changed to a different value).
 - A capture event has occurred and the CC/CC_BUFF registers are updated.

Note: These signals remain high only for two cycles of `clk_sys`.

29.2.5 Interrupts

The TCPWM block provides a dedicated interrupt output for each counter. This interrupt can be generated for a TC or CC event. A TC is the logical OR of the OV and UN events.

Four registers are used to handle interrupts in this block, as shown in [Table 29-1](#).

Table 29-1. Interrupt Register

Interrupt Registers	Bits	Name	Description
TCPWM_CNT_INTR (Interrupt request register)	0	TC	This bit is set to '1', when a terminal count is detected. Write '1' to clear this bit.
	1	CC_MATCH	This bit is set to '1' when the counter value matches capture/compare register value. Write '1' to clear this bit.
TCPWM_CNT_INTR_SET (Interrupt set request register)	0	TC	Write '1' to set the corresponding bit in the interrupt request register. When read, this register reflects the interrupt request register status.
	1	CC_MATCH	Write '1' to set the corresponding bit in the interrupt request register. When read, this register reflects the interrupt request register status.
TCPWM_CNT_INTR_MASK (Interrupt mask register)	0	TC	Mask bit for the corresponding TC bit in the interrupt request register.
	1	CC_MATCH	Mask bit for the corresponding CC_MATCH bit in the interrupt request register.
TCPWM_CNT_INTR_MASKED (Interrupt masked request register)	0	TC	Logical AND of the corresponding TC request and mask bits.
	1	CC_MATCH	Logical AND of the corresponding CC_MATCH request and mask bits.

29.2.6 PWM Outputs

Each counter two outputs, pwm and pwm_n (complementary of pwm). Note that the OV, UN, and CC conditions are used to drive pwm and pwm_n, by configuring the TCPWM_CNT_TR_CTRL2 register ([Table 29-2](#)).

Table 29-2. Configuring Output for OV, UN, and CC Conditions

Field	Bit	Value	Event	Description
CC_MATCH_MODE Default Value = 3	1:0	0	Set pwm to '1	Configures output line on a compare match (CC) event
		1	Clear pwm to '0	
		2	Invert pwm	
		3	No change	
OVERFLOW_MODE Default Value = 3	3:2	0	Set pwm to '1	Configures output line on a overflow (OV) event
		1	Clear pwm to '0	
		2	Invert pwm	
		3	No change	
UNDERFLOW_- MODE Default Value = 3	5:4	0	Set pwm to '1	Configures output line on a underflow (UN) event
		1	Clear pwm to '0	
		2	Invert pwm	
		3	No change	

29.2.7 Power Modes

The TCPWM block works in Active and Sleep modes. The TCPWM block is powered from V_{CCD} . The configuration registers and other logic are powered in Deep Sleep mode to keep the states of configuration registers. See [Table 29-3](#) for details.

Table 29-3. Power Modes in TCPWM Block

Power Mode	Block Status
Active	This block is fully operational in this mode with clock running and power switched on.
Sleep	The CPU is in sleep but the block is still functional in this mode. All counter clocks are on.
Deep Sleep	Both power and clocks to the block are turned off, but configuration registers retain their states.
Hibernate	In this mode, the power to this block is switched off. Configuration registers will lose their state.

29.3 Operation Modes

The counter block can function in six operational modes, as shown in [Table 29-4](#). The MODE [26:24] field of the counter control register (TCPWM_CNTx_CTRL) configures the counter in the specific operational mode.

Table 29-4. Operational Mode Configuration

Mode	MODE Field [26:24]	Description
Timer	000	The counter increments or decrements by '1' at every clk_counter cycle in which a count event is detected. The Compare/Capture register is used to compare the count.
Capture	010	The counter increments or decrements by '1' at every clk_counter cycle in which a count event is detected. A capture event copies the counter value into the capture register.
Quadrature	011	Quadrature decoding. The counter is decremented or incremented based on two phase inputs according to an X1, X2, or X4 decoding scheme.
PWM	100	Pulse width modulation.
PWM_DT	101	Pulse width modulation with dead time insertion.
PWM_PR	110	Pseudo-random PWM using a 16- or 32-bit linear feedback shift register (LFSR) to generate pseudo-random noise.

The counter can be configured to count up, down, and up/down by setting the UP_DOWN_MODE[17:16] field in the TCPWM_CNT_CTRL register, as shown in [Table 29-5](#).

Table 29-5. Counting Mode Configuration

Counting Modes	UP_DOWN_MODE[17:16]	Description
UP Counting Mode	00	Increments the counter until the period value is reached. A Terminal Count (TC) and Overflow (OV) condition is generated when the counter changes from the period value.
DOWN Counting Mode	01	Decrements the counter from the period value until 0 is reached. A TC and Underflow (UN) condition is generated when the counter changes from a value of '0'.
UP/DOWN Counting Mode 1	10	Increments the counter until the period value is reached, and then decrements the counter until '0' is reached. TC and UN conditions are generated only when the counter changes from a value of '0'.
UP/DOWN Counting Mode 2	11	Similar to up/down counting mode 1 but a TC condition is generated when the counter changes from '0' and when the counter value changes from the period value. OV and UN conditions are generated similar to how they are generated in UP and DOWN counting modes respectively.

29.3.1 Timer Mode

The timer mode is commonly used to measure the time of occurrence of an event or to measure the time difference between two events. The timer functionality increments/decrements a counter between 0 and the value stored in the PERIOD register. When the counter is running, the count value stored in the COUNTER register is compared with the compare/capture register (CC). When the counter changes from a state in which COUNTER equals CC, the cc_match event is generated.

Timer functionality is typically used for one of the following:

- Timing a specific delay – the count event is a constant '1'.
- Counting the occurrence of a specific event – the event should be connected as an input trigger and selected for the count event.

Table 29-6. Timer Mode Trigger Input Description

Trigger Inputs	Usage
Reload	Sets the counter value and starts the counter. Behavior is dependent on UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: The counter is set to "0" and count direction is set to "up". ■ COUNT_DOWN: The counter is set to PERIOD and count direction is set to "down". ■ COUNT_UPDN1/2: The counter is set to "1" and count direction is set to "up". Can be used when the counter is running or not running.
Start	Starts the counter. The counter is not initialized by hardware. The current counter value is used. Behavior is dependent on UP_DOWN_MODE. When the counter is not running: <ul style="list-style-type: none"> ■ COUNT_UP: The count direction is set to "up". ■ COUNT_DOWN: The count direction is set to "down". ■ COUNT_UPDN1/2: The count direction is not modified. Note that when the counter is running, the start event has no effect. Can be used when the counter is running or not running.
Stop	Stops the counter.
Count	Count event increments/decrements the counter.
Capture	Not used.

Incrementing and decrementing the counter is controlled by the count event and the counter clock clk_counter. Typical operation will use a constant '1' count event and clk_counter without pre-scaling. Advanced operations are also possible; for example, the counter event configuration can decide to count the rising edges of a synchronized input trigger.

Table 29-7. Timer Mode Supported Features

Supported Features	Description
Clock pre-scaling	Pre-scales the counter clock clk_counter.
One-shot	Counter is stopped by hardware, after a single period of the counter: <ul style="list-style-type: none"> ■ COUNT_UP: on an overflow event. ■ COUNT_DOWN, COUNT_UPDN1/2: on an underflow event.
Auto reload CC	CC and CC_BUFF are exchanged on a cc_match event (when specified by CTRL.AUTO_RELOAD_CC)
Up/down modes	Specified by UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: The counter counts from 0 to PERIOD. ■ COUNT_DOWN: The counter counts from PERIOD to 0. ■ COUNT_UPDN1/2: The counter counts from 1 to PERIOD and back to 0.

Table 29-6 lists the trigger outputs and the conditions when they are triggered.

Table 29-8. Timer Mode Trigger Outputs

Trigger Outputs	Description
cc_match (CC)	Counter changes from a state in which COUNTER equals CC.
Underflow (UN)	Counter is decrementing and changes from a state in which COUNTER equals "0".
Overflow (OV)	Counter is incrementing and changes from a state in which COUNTER equals PERIOD.

Note: Each output is only two clk_sys wide and is represented by an arrow in the timing diagrams in this chapter, for example see [Figure 29-5](#).


Table 29-9. Timer Mode Interrupt Outputs

Interrupt Outputs	Description
tc	Specified by UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: The tc event is the same as the overflow event. ■ COUNT_DOWN: The tc event is the same as the underflow event. ■ COUNT_UPDN1: The tc event is the same as the underflow event. ■ COUNT_UPDN2: The tc event is the same as the logical OR of the overflow and underflow events.
cc_match (CC)	Counter changes from a state in which COUNTER equals CC.

Table 29-10. Timer Mode PWM Outputs

PWM Outputs	Description
pwm	Not used.
pwm_n	Not used.

Figure 29-4. Timer Functionality



Notes:

- The timer functionality uses only PERIOD (and not PERIOD_BUFF).
- Do not write to COUNTER when the counter is running.

[Figure 29-5](#) illustrates a timer in up-counting mode. The counter is initialized (to 0) and started with a software-based reload event.

Notes:

- PERIOD is 4, resulting in an effective repeating counter pattern of 4+1 = 5 counter clock periods. The CC register is 2, and sets the condition for a cc_match event.
- When the counter changes from a state in which COUNTER is 4, overflow and tc events are generated.
- When the counter changes from a state in which COUNTER is 2, a cc_match event is generated.
- A constant count event of '1' and clk_counter without prescaling is used in the following scenarios. If the count event is '0' and a reload event is triggered, the reload will be registered only on the first clock edge when the count event is '1'.

Figure 29-5. Timer in Up-counting Mode

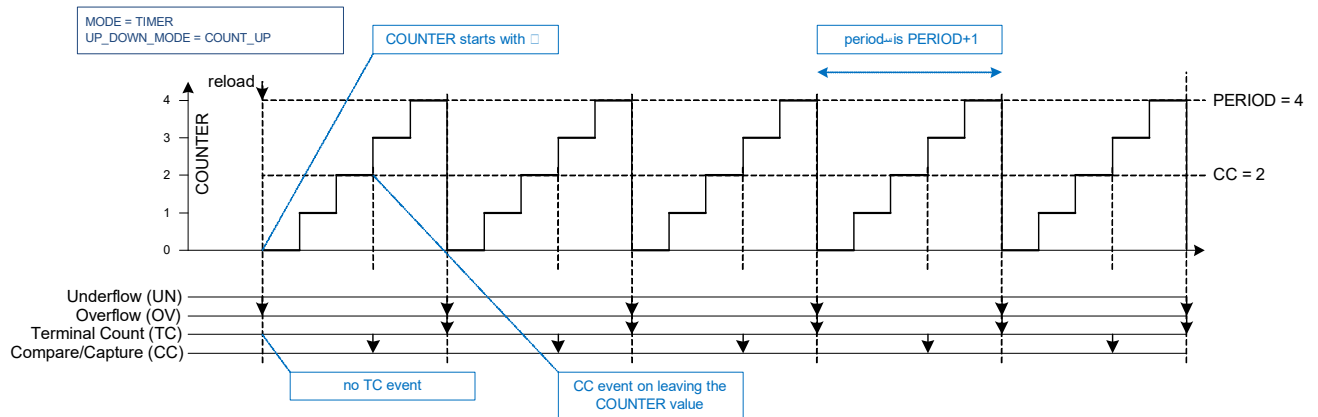


Figure 29-6 illustrates a timer in “one-shot” operation mode. Note that the counter is stopped on a tc event.

Figure 29-6. Timer in One-shot Mode

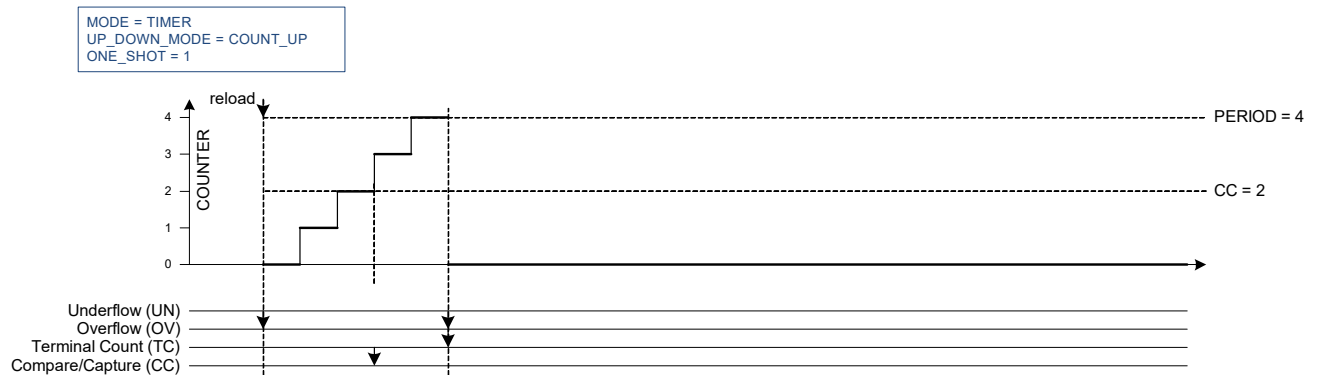


Figure 29-7 illustrates clock pre-scaling. Note that the counter is only incremented every other counter cycle.

Figure 29-7. Timer Clock Pre-scaling




Figure 29-8 illustrates a counter that is initialized and started (reload event), stopped (stop event), and continued/started (start event). Note that the counter does not change value when it is not running (STATUS.RUNNING).

Figure 29-8. Counter Start/Stopped/Continued




Figure 29-9 illustrates a timer that uses both CC and CC_BUFF registers. Note that CC and CC_BUFF are exchanged on a cc_match event.

Figure 29-9. Use of CC and CC_BUFF Register Bits




Figure 29-10 illustrates a timer in down counting mode. The counter is initialized (to PERIOD) and started with a software-based reload event.

Notes:

- When the counter changes from a state in which COUNTER is 0, a UN and TC events are generated.
- When the counter changes from a state in which COUNTER is 2, a cc_match event is generated.
- PERIOD is 4, resulting in an effective repeating counter pattern of $4+1 = 5$ counter clock periods.

Figure 29-10. Timer in Down-counting Mode

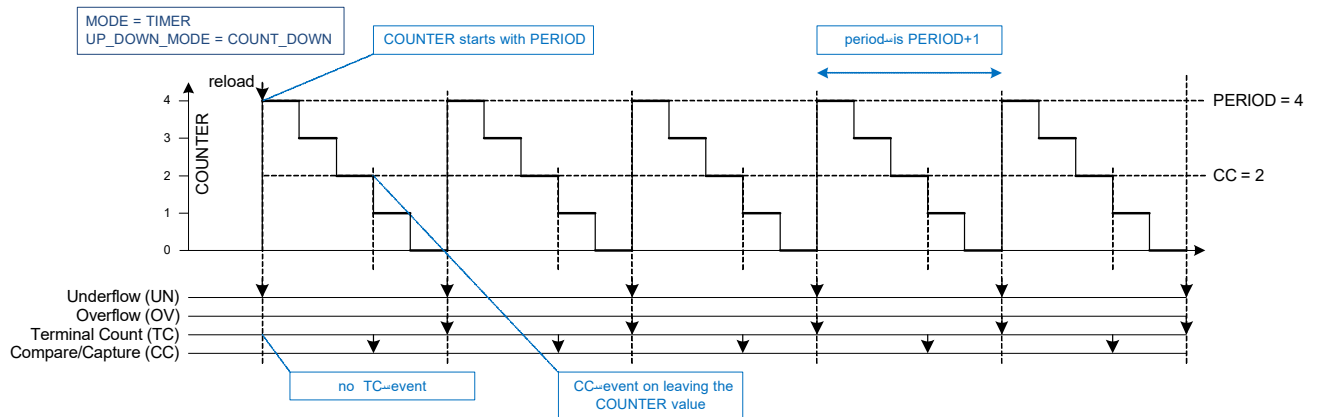


Figure 29-11 illustrates a timer in up/down counting mode 1. The counter is initialized (to 1) and started with a software-based reload event.

Notes:

- When the counter changes from a state in which COUNTER is 4, an overflow is generated.
- When the counter changes from a state in which COUNTER is 0, an underflow and tc event are generated.
- When the counter changes from a state in which COUNTER is 2, a cc_match event is generated.
- PERIOD is 4, resulting in an effective repeating counter pattern of $2 \times 4 = 8$ counter clock periods.

Figure 29-11. Timer in Up/Down Counting Mode 1




Figure 29-12 illustrates a timer in up/down counting mode 1, with different CC values.

Notes:

- When CC is 0, the cc_match event is generated at the start of the period (when the counter changes from a state in which COUNTER is 0).
- When CC is PERIOD, the cc_match event is generated at the middle of the period (when the counter changes from a state in which COUNTER is PERIOD).

Figure 29-12. Up/Down Counting Mode with Different CC Values

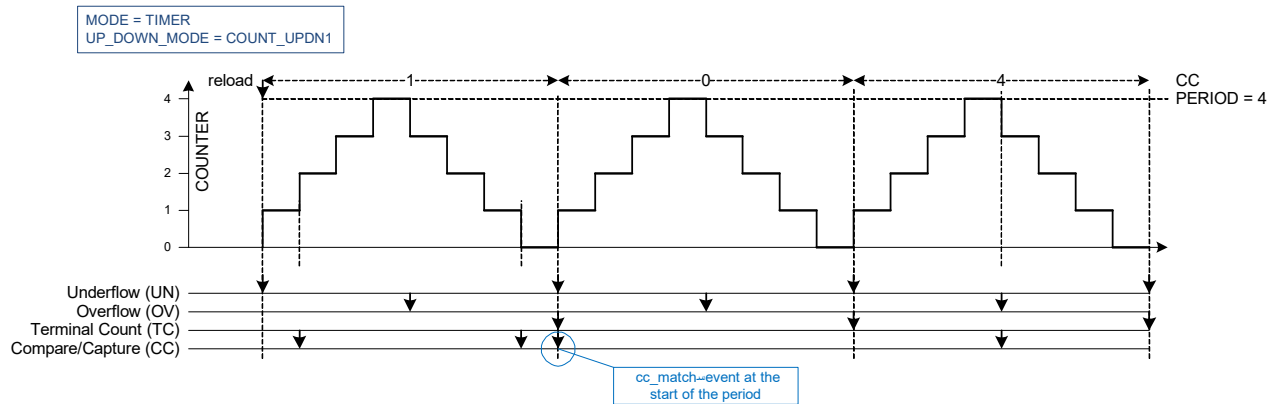



Figure 29-13 illustrates a timer in up/down counting mode 2. This mode is same as up/down counting mode 1, except for the TC event, which is generated when either underflow or overflow event occurs.

Figure 29-13. Up/Down Counting Mode 2



29.3.1.1 Configuring Counter for Timer Mode

The steps to configure the counter for Timer mode of operation and the affected register bits are as follows.

1. Disable the counter by writing '1' to the TCPWM_CTRL_CLR register.
2. Select Timer mode by writing '000' to the MODE[26:24] field of the TCPWM_CNT_CTRL register.
3. Set the required 16- or 32-bit period in the TCPWM_CNT_PERIOD register.
4. Set the 16- or 32-bit compare value in the TCPWM_CNT_CC register and the buffer compare value in the TCPWM_CNT_CC_BUFF register.
5. Set AUTO_RELOAD_CC field of the TCPWM_CNT_CTRL register, if required to swap values at every CC condition.
6. Set clock prescaling by writing to the GENERIC[15:8] field of the TCPWM_CNT_CTRL register.
7. Set the direction of counting by writing to the UP_DOWN_MODE[17:16] field of the TCPWM_CNT_CTRL register.
8. The timer can be configured to run either in continuous mode or one-shot mode by writing 0 or 1, respectively to the ONE_SHOT[18] field of the TCPWM_CNT_CTRL register.
9. Set the TCPWM_CNT_TR_CTRL0 register to select the trigger that causes the event (reload, start, stop, capture, and count).
10. Set the TCPWM_CNT_TR_CTRL1 register to select the edge of the trigger that causes the event (reload, start, stop, capture, and count).
11. If required, set the interrupt upon TC or CC condition.
12. Enable the counter by writing '1' to the TCPWM_CTRL_SET register. A start trigger must be provided through firmware (TCPWM_CMD_START register) to start the counter if the hardware start signal is not enabled.

29.3.2 Capture Mode

The capture functionality increments and decrements a counter between 0 and PERIOD. When the capture event is activated the counter value COUNTER is copied to CC (and CC is copied to CC_BUFF).

The capture functionality can be used to measure the width of a pulse (connected as one of the input triggers and used as capture event).

The capture event can be triggered through the capture trigger input or through a firmware write to command register (TCPWM_CMD_CAPTURE).

Table 29-11. Capture Mode Trigger Input Description

Trigger Inputs	Usage
reload	Sets the counter value and starts the counter. Behavior is dependent on UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: The counter is set to "0" and count direction is set to "up". ■ COUNT_DOWN: The counter is set to PERIOD and count direction is set to "down". ■ COUNT_UPDN1/2: The counter is set to "1" and count direction is set to "up". Can be used only when the counter is not running.
start	Starts the counter. The counter is not initialized by hardware. The current counter value is used. Behavior is dependent on UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: The count direction is set to "up". ■ COUNT_DOWN: The count direction is set to "down". ■ COUNT_UPDN1/2: The count direction is not modified. Can be used only when the counter is not running.
stop	Stops the counter.
count	Count event increments/decrements the counter.
capture	Copies the counter value to CC and copies CC to CC_BUFF.

Table 29-12. Capture Mode Supported Features

Supported Features	Description
Clock pre-scaling	Pre-scales the counter clock clk_counter.
One-shot	Counter is stopped by hardware, after a single period of the counter: <ul style="list-style-type: none"> ■ COUNT_UP: on an overflow event. ■ COUNT_DOWN, COUNT_UPDN1/2: on an underflow event.
Up/down modes	Specified by UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: The counter counts from 0 to PERIOD. ■ COUNT_DOWN: The counter counts from PERIOD to 0. ■ COUNT_UPDN1/2: The counter counts from 1 to PERIOD and back to 0.

Table 29-13. Capture Mode Trigger Output Description

Trigger Outputs	Description
cc_match (CC)	CC is copied to CC_BUFF and counter value is copied to CC (cc_match equals capture event).
Underflow (UN)	Counter is decrementing and changes from a state in which COUNTER equals "0".
Overflow (OV)	Counter is incrementing and changes from a state in which COUNTER equals PERIOD.

Table 29-14. Capture Mode Interrupt Outputs

Interrupt Outputs	Description
tc	Specified by UP_DOWN_MODE: ■ COUNT_UP: tc event is the same as the overflow event. ■ COUNT_DOWN: tc event is the same as the underflow event. ■ COUNT_UPDN1: tc event is the same as the underflow event. ■ COUNT_UPDN2: tc event is the same as the logical OR of the overflow and underflow events.
cc_match (CC)	CC is copied to CC_BUFF and counter value is copied to CC (cc_match equals capture event).

Table 29-15. Capture Mode PWM Outputs

PWM Outputs	Description
pwm	Not used.
pwm_n	Not used.

Figure 29-14. Capture Functionality





Figure 29-15 illustrates capture behavior in the up counting mode.

Notes:

- The capture event detection uses rising edge detection. As a result, the capture event is remembered until the next “active count” pre-scaled counter clock.
- When a capture event occurs, COUNTER is copied into CC. CC is copied to CC_BUFF.
- A cc_match event is generated when the counter value is captured.

Figure 29-15. Capture in Up Counting Mode




When multiple capture events are detected before the next “active count” pre-scaled counter clock, capture events are treated as follows:

- In the rising edge and falling edge modes, multiple events are effectively reduced to a single event.
- In the rising/falling edge mode, an even number of events is not detected and an odd number of events is reduced to a single event.

This behavior is illustrated by Figure 29-16, in which a pre-scaler by a factor of 4 is used.

Figure 29-16. Multiple Events Detected before Active-Count



29.3.2.1 Configuring Counter for Capture Mode

The steps to configure the counter for Capture mode operation and the affected register bits are as follows.

1. Disable the counter by writing '1' to the TCPWM_CTRL_CLR register.
2. Select Capture mode by writing '010' to the MODE[26:24] field of the TCPWM_CNT_CTRL register.
3. Set the required 16-bit period in the TCPWM_CNT_PERIOD register.
4. Set clock prescaling by writing to the GENERIC[15:8] field of the TCPWM_CNT_CTRL register.
5. Set the direction of counting by writing to the UP_DOWN_MODE[17:16] field of the TCPWM_CNT_CTRL register.
6. Counter can be configured to run either in continuous mode or one-shot mode by writing 0 or 1, respectively to the ONE_SHOT[18] field of the TCPWM_CNT_CTRL register.
7. Set the TCPWM_CNT_TR_CTRL0 register to select the trigger that causes the event (reload, start, stop, capture, and count).
8. Set the TCPWM_CNT_TR_CTRL1 register to select the edge that causes the event (reload, start, stop, capture, and count).
9. If required, set the interrupt upon TC or CC condition.
10. Enable the counter by writing '1' to the TCPWM_CTRL_SET register. A start trigger must be provided through firmware (TCPWM_CMD_START register) to start the counter if the hardware start signal is not enabled.

29.3.3 Quadrature Decoder Mode

Quadrature functionality increments and decrements a counter between 0 and 0xFFFF or 0xFFFFFFFF (32-bit mode). Counter updates are under control of quadrature signal inputs: index, phiA, and phiB. The index input is used to indicate an absolute position. The phiA and phiB inputs are used to determine a change in position (the rate of change in position can be used to derive speed). The quadrature inputs are mapped onto triggers (as described in [Table 29-16](#)).

Table 29-16. Quadrature Mode Trigger Input Description

Trigger Input	Usage
reload/index	This event acts as a quadrature index input. It initializes the counter to the counter midpoint 0x8000 (16-bit) or 0x8000000 (32-bit mode) and starts the quadrature functionality. Rising edge event detection or falling edge detection mode should be used.
start/phiB	This event acts as a quadrature phiB input. Pass through (no edge detection) event detection mode should be used.
stop	Stops the quadrature functionality.
count/phiA	This event acts as a quadrature phiA input. Pass through (no edge detection) event detection mode should be used.
capture	Not used.

Table 29-17. Quadrature Mode Supported Features

Supported Features	Description
Quadrature encoding	Three encoding schemes for the phiA and phiB inputs are supported (as specified by CTRL.QUADRA-TURE_MODE): X1 encoding. X2 encoding. X4 encoding.

Note: Clock pre-scaling is not supported and the count event is used as a quadrature input phiA. As a result, the quadrature functionality operates on the counter clock (clk_counter), rather than on an “active count” prescaled counter clock.

Table 29-18. Quadrature Mode Trigger Output Description

Trigger Outputs	Description
cc_match (CC)	Counter value COUNTER equals 0 or 0xFFFF or 0xFFFFFFFF (32-bit mode) or a reload/index event.
Underflow (UN)	Not used.
Overflow (OV)	Not used.

Table 29-19. Quadrature Mode Interrupt Outputs

Interrupt Outputs	Description
cc_match (CC)	Counter value COUNTER equals 0 or 0xFFFF or 0xFFFFFFFF (32-bit mode) or a reload/index event.
tc	Reload/index event.

Table 29-20. Quadrature Mode PWM Outputs

PWM Outputs	Description
pwm	Not used.
pwm_n	Not used.

```

graph LR
    subgraph Quadrature
        Reload[Reload/Index Start/phiB Stop Count/phiA] --> Counter[COUNTER]
        Counter --> CC[CC]
        CC --> CC_BUFF[CC_BUFF]
        Counter --> TC_Equal[==]
        CC_BUFF --> CC_Equal[==]
        TC_Equal -- tc --> Interrupt[Interrupt generation]
        TC_Equal -- tc --> Trigger[Trigger generation]
        CC_Equal -- cc_match --> Trigger
    end
    Interrupt --> interrupt_out[interrupt]
    Trigger --> tr_cc_match_out[tr_cc_match]
    clk_counter[clk_counter] --> Counter
    TC_Equal -- 0x0000 --> TC_Equal
    CC_Equal -- 0xFFFF --> CC_Equal

```

- A software-generated reload event starts quadrature operation. As a result, COUNTER is set to 0x8000 or 0x80000000, which is the counter midpoint (the COUNTER is set to 0x7FFF if the reload event coincides with a decrement event; the COUNTER is set to 0x8001 if the reload event coincides with an increment event). Note that a software-generated reload event is generated only once, when the counter is not running. All other reload/index events are hardware-generated reload events as a result of the quadrature index signal.

- The software interrupt handler uses the `tc` and `cc_match` interrupt cause fields to distinguish between a reload/index event and a situation in which a minimum/maximum counter value was reached (about to wrap around). The `CC` and `CC_BUFF` registers are used to determine when the interrupt causing event occurred.

Counter increments (incr1 event) and decrements (decr1 event) are determined by the quadrature encoding scheme as illustrated by [Figure 29-18](#).

Timing diagram illustrating the relationship between quadrature decoding modes and encoder events.

The diagram shows three quadrature decoding modes:

- Quadrature decoding QUADRATURE_MODE = X1:** Shows the relationship between ϕA , ϕB , $incr1$, and $decr1$ signals. Vertical dashed lines indicate events, with arrows pointing to the $incr1$ and $decr1$ signals.
- Quadrature decoding QUADRATURE_MODE = X2:** Shows the relationship between $incr1$ and $decr1$ signals. The events are twice as frequent as in X1 mode.
- Quadrature decoding QUADRATURE_MODE = X4:** Shows the relationship between $incr1$ and $decr1$ signals. The events are four times as frequent as in X1 mode.

Callouts indicate the frequency of events relative to X1 mode:

- Two times the events of X1 mode (for X2 mode)
- Four times the events of X1 mode (for X4 mode)

Figure 29-19 illustrates quadrature functionality as a function of the reload/index, incr1, and decr1 events. Note that the first reload/index event copies the counter value COUNTER to CC.

Figure 29-19. Quadrature Mode Reload/Index Timing


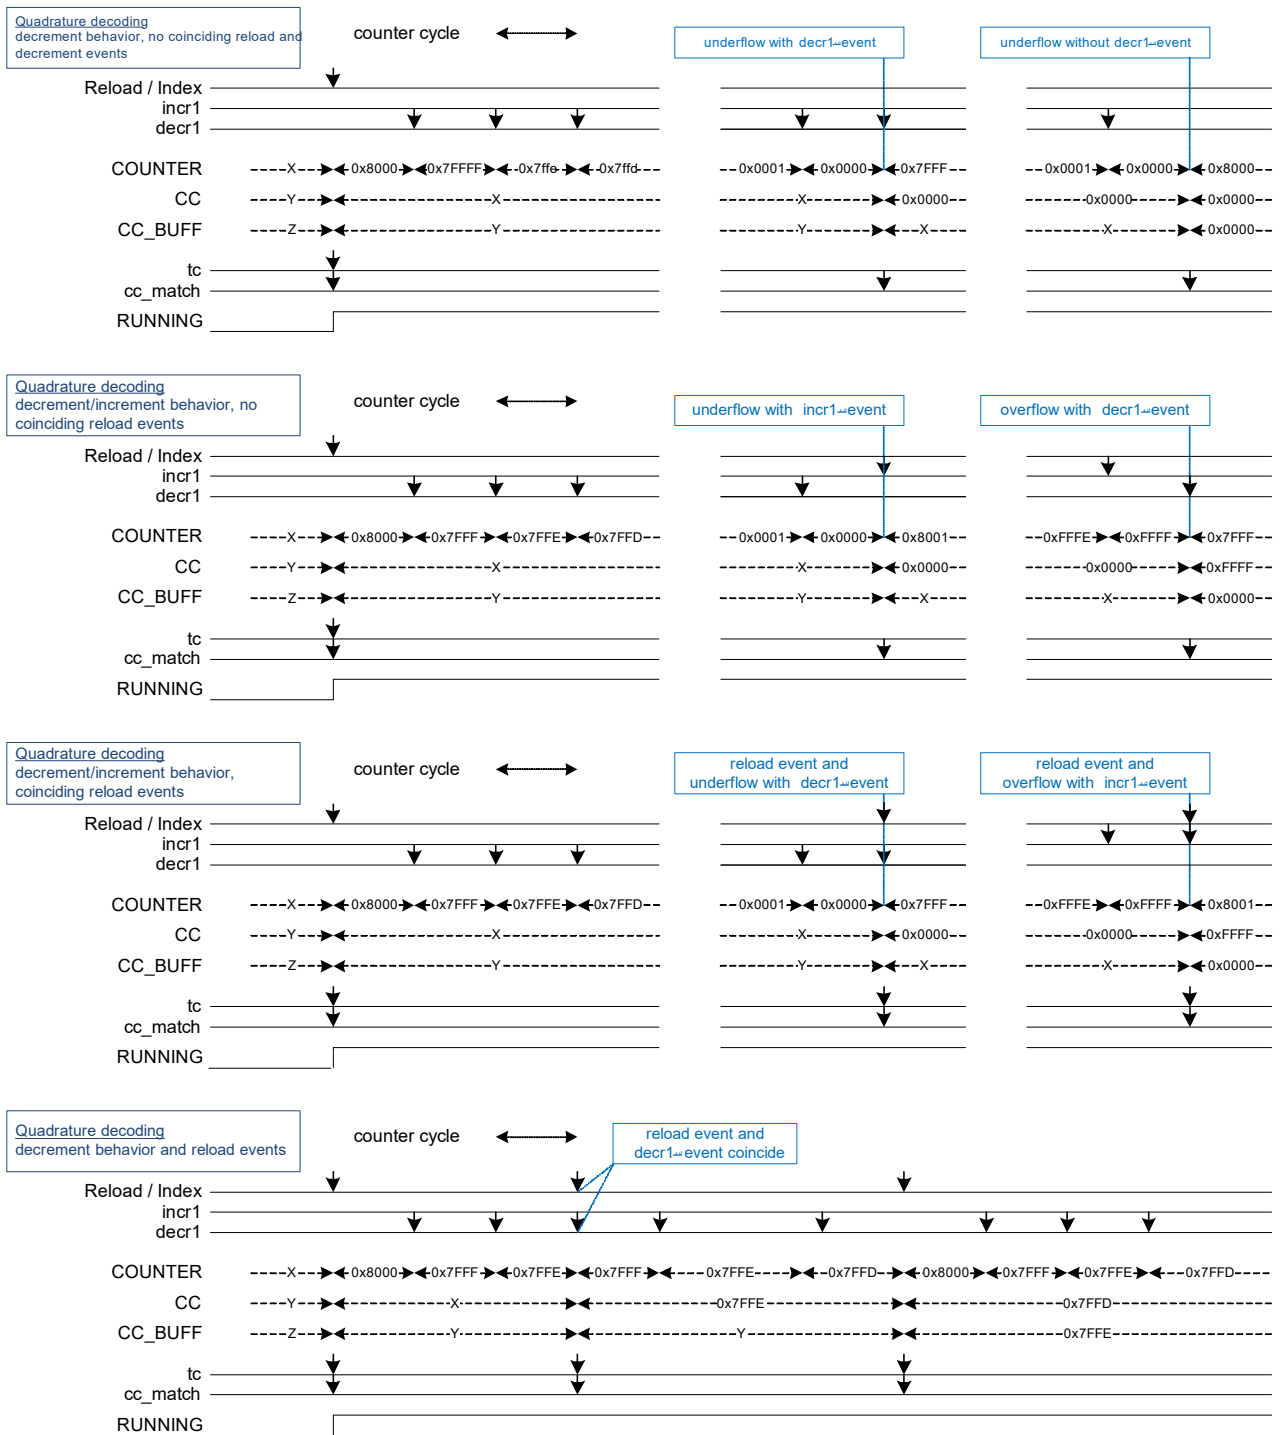


Figure 29-20 illustrate quadrature functionality for different event scenarios (including scenarios with coinciding events). In all scenarios, the first reload/index event is generated by software when the counter is not yet running.

Figure 29-20. Quadrature Mode Timing Cases



29.3.3.1 Configuring Counter for Quadrature Mode

The steps to configure the counter for quadrature mode of operation and the affected register bits are as follows.

1. Disable the counter by writing '1' to the TCPWM_CTRL_CLR register.
2. Select Quadrature mode by writing '011' to the MODE[26:24] field of the TCPWM_CNT_CTRL register.
3. Set the required encoding mode by writing to the QUADRATURE_MODE[21:20] field of the TCPWM_CNT_CTRL register.
4. Set the TCPWM_CNT_TR_CTRL0 register to select the trigger that causes the event (Index and Stop).
5. Set the TCPWM_CNT_TR_CTRL1 register to select the edge that causes the event (Index and Stop).
6. If required, set the interrupt upon TC or CC condition.
7. Enable the counter by writing '1' to the TCPWM_CTRL_SET register. A start trigger must be provided through firmware (TCPWM_CMD_START register) to start the counter if the hardware start signal is not enabled.

29.3.4 Pulse Width Modulation Mode

PWM functionality increments/decrements a counter between 0 and PERIOD. When the counter is running, the counter value COUNTER is compared with CC. When COUNTER equals CC, the cc_match event is generated. Additionally, on a counter overflow and counter underflow, the overflow and underflow events are generated. Combined, the cc_match, overflow and underflow events are used to generate a pulse-width modulated signal on the PWM "pwm" and "pwm_n" output signals. Left-aligned, right-aligned, and center-aligned PWM signals can be generated. Asymmetric PWM signals can be generated using the COUNT_UPDN2 mode.

Table 29-21. PWM Mode Trigger Input Description

Trigger Inputs	Usage
reload	Sets the counter value and starts the counter. Behavior is dependent on UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: The counter is set to "0" and count direction is set to "up". ■ COUNT_DOWN: The counter is set to PERIOD and count direction is set to "down". ■ COUNT_UPDN1/2: The counter is set to "1" and count direction is set to "up". Can be used only when the counter is not running.
start	Starts the counter. The counter is not initialized by hardware. The current counter value is used. Behavior is dependent on UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: The count direction is set to "up". ■ COUNT_DOWN: The count direction is set to "down". ■ COUNT_UPDN1/2: The count direction is set to "up". Can be used only when the counter is not running.
stop/kill	Stops the counter or suppresses the PWM output, depending on PWM_STOP_ON_KILL and PWM_SYN-C_KILL.
count	Count event increments/decrements the counter.
capture/swap	This event acts as a swap event. When this event is active, the CC/CC_BUFF and PERIOD/PERIOD_BUFF registers are exchanged on a tc event (when specified by CTRL.AUTO_RELOAD_CC and CTRL.AUTO_RELOAD_PERIOD). A swap event requires rising, falling, or rising/falling edge event detection mode. Pass-through mode is not supported, unless the selected event is a constant '0' or '1'. Note: When COUNT_UPDN2 mode exchanges PERIOD and PERIOD_BUFF at a tc event that coincides with an OV event, software should ensure that the PERIOD and PERIOD_BUFF values are the same. When a swap event is detected and the counter is running, the event is kept pending until the next tc event. When a swap event is detected and the counter is not running, the event is cleared by hardware.

Table 29-22. PWM Mode Supported Features

Supported Features	Description
Clock pre-scaling	Pre-scales the counter clock "clk_counter".
One-shot	Counter is stopped by hardware, after a single period of the counter: <ul style="list-style-type: none"> ■ COUNT_UP: on an overflow event. ■ COUNT_DOWN and COUNT_UPDN1/2: on an underflow event.
Compare Swap	CC and CC_BUFF are exchanged on a swap event and tc event (when specified by CTRL.AUTO_RELOAD_CC).
Period Swap	PERIOD and PERIOD_BUFF are exchanged on a swap event and tc event (when specified by CTRL.AUTO_RELOAD_PERIOD). Note: When COUNT_UPDN2/Asymmetric mode exchanges PERIOD and PERIOD_BUFF at a tc event that coincides with an overflow event, software should ensure that the PERIOD and PERIOD_BUFF values are the same.
Alignment (Up/Down modes)	Specified by UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: The counter counts from 0 to PERIOD. Generates a left-aligned PWM output. ■ COUNT_DOWN: The counter counts from PERIOD to 0. Generates a right-aligned PWM output. ■ COUNT_UPDN1/2: The counter counts from 1 to PERIOD and back to 0. Generates a center-aligned/asymmetric PWM output.
Kill modes	Specified by PWM_STOP_ON_KILL and PWM_SYNC_KILL: <ul style="list-style-type: none"> ■ PWM_STOP_ON_KILL = '1' (PWM_SYNC_KILL = don't care): Stop on Kill mode. This mode stops the counter on a stop/kill event. Reload or start event is required to start count running. ■ PWM_STOP_ON_KILL = '0' and PWM_SYNC_KILL = '0': Asynchronous kill mode. This mode keeps the counter running, but suppresses the PWM output signals and continues to do so for the duration of the stop/kill event. ■ PWM_STOP_ON_KILL = '0' and PWM_SYNC_KILL = '1': Synchronous kill mode. This mode keeps the counter running, but suppresses the PWM output signals and continues to do so until the next tc event without a stop/kill event.

Note that the PWM mode does not support dead time insertion. This functionality is supported by the separate PWM_DT mode.

Table 29-23. PWM Mode Trigger Output Description

Trigger Output	Description
cc_match (CC)	Specified by UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP and COUNT_DOWN: The counter changes to a state in which COUNTER equals CC. ■ COUNT_UPDN1/2: counter changes from a state in which COUNTER equals CC.
Underflow (UN)	Counter is decrementing and changes from a state in which COUNTER equals "0".
Overflow (OV)	Counter is incrementing and changes from a state in which COUNTER equals PERIOD.

Table 29-24. PWM Mode Interrupt Output Description


Interrupt Outputs	Description
tc	Specified by UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP: tc event is the same as the overflow event. ■ COUNT_DOWN: tc event is the same as the underflow event. ■ COUNT_UPDN1: tc event is the same as the underflow event. ■ COUNT_UPDN2: tc event is the same as the logical OR of the overflow and underflow events.
cc_match (CC)	Specified by UP_DOWN_MODE: <ul style="list-style-type: none"> ■ COUNT_UP and COUNT_DOWN: The counter changes to a state in which COUNTER equals CC. ■ COUNT_UPDN1/2: counter changes from a state in which COUNTER equals CC.

Table 29-25. PWM Mode PWM Outputs

PWM Outputs	Description
pwm	PWM output.
pwm_n	Complementary PWM output.


Note that the cc_match event generation in COUNT_UP and COUNT_DOWN modes are different from the generation in other functional modes or counting modes. This is to ensure that 0 percent and 100 percent duty cycles can be generated.

Figure 29-21. PWM Mode Functionality



The generation of the PWM output signals is a multi-step process and is illustrated as follows:

Figure 29-22. PWM Output Generation



PWM behavior depends on the PERIOD and CC registers. The software can update the PERIOD_BUFF and CC_BUFF registers, without affecting the PWM behavior. This is the main rationale for double buffering these registers.

Figure 29-23 illustrates a PWM in up counting mode. The counter is initialized (to 0) and started with a software-based reload event.

Notes:

- When the counter changes from a state in which COUNTER is 4, an overflow and tc event are generated.
- When the counter changes to a state in which COUNTER is 2, a cc_match event is generated.
- PERIOD is 4, resulting in an effective repeating counter pattern of $4+1 = 5$ counter clock periods.

Figure 29-23. PWM in Up Counting Mode

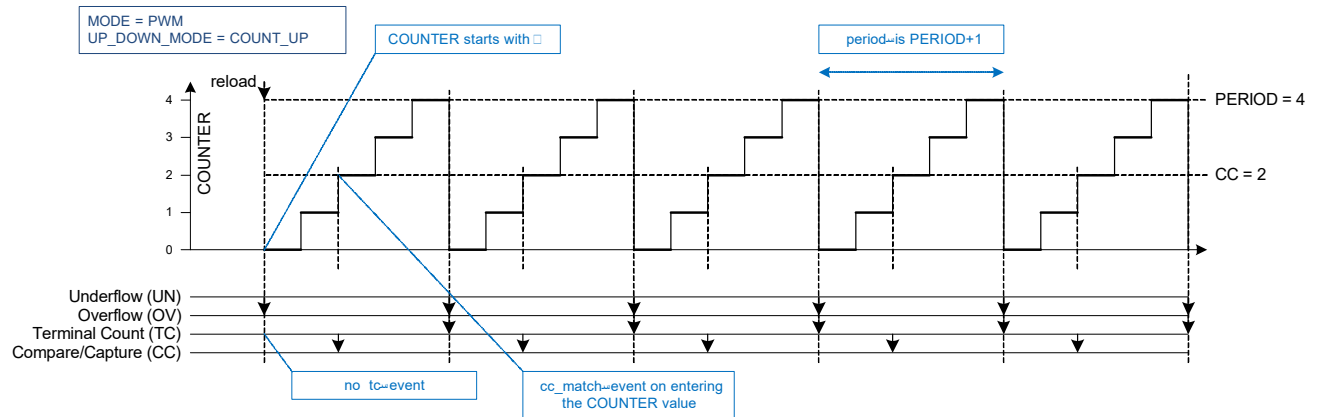


Figure 29-24 illustrates a PWM in up counting mode generating a left-aligned PWM. The figure also illustrates how a right-aligned PWM can be created using the PWM in up counting mode by inverting the OVERFLOW_MODE and CC_MATCH_MODE and using a CC value that is complementary ($\text{PERIOD}+1 - \text{pulse width}$) to the one used for left-aligned PWM. Note that CC is changed (to CC_BUFF, which is not depicted) on a tc event. The "pulse width" used as CC is the PWM output's duty cycle $\times (\text{PERIOD}+1)$.

Figure 29-24. PWM Left- and Right-Aligned Outputs

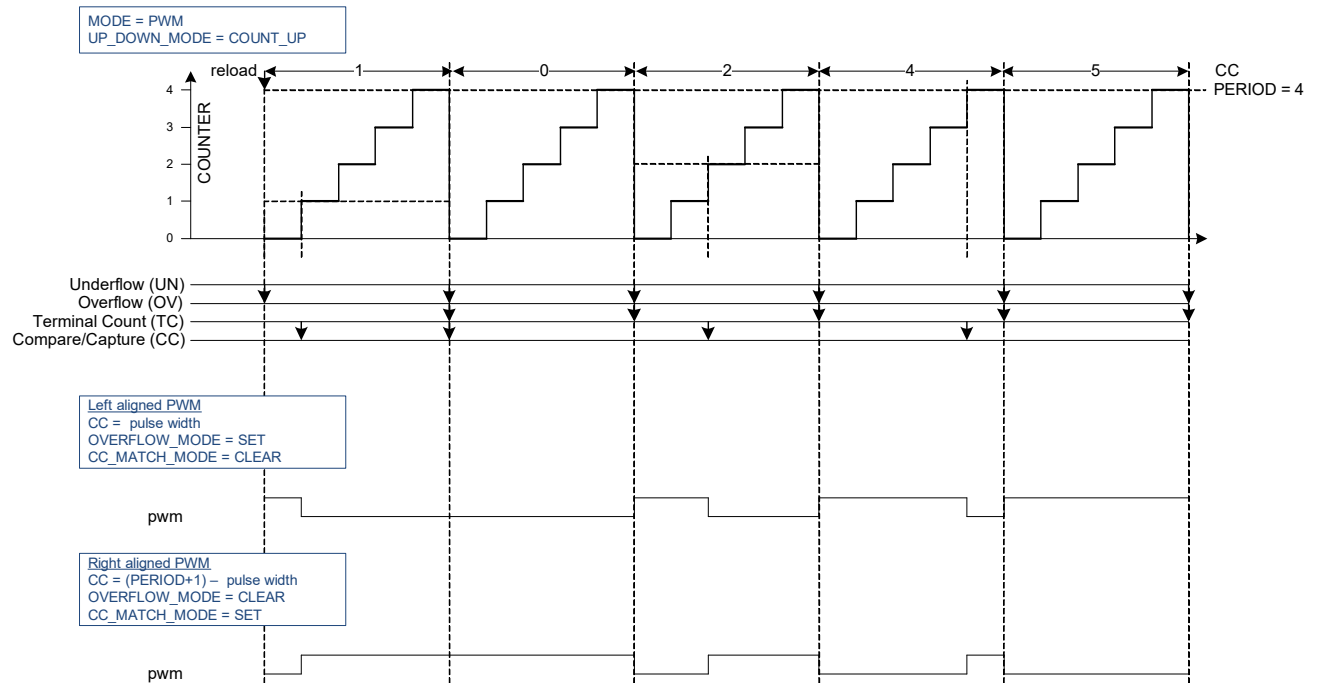


Figure 29-25 illustrates a PWM in down counting mode. The counter is initialized (to PERIOD) and started with a software-based reload event.

Notes:

- When the counter changes from a state in which COUNTER is 0, an underflow and tc event are generated.
- When the counter changes to a state in which COUNTER is 2, a cc_match event is generated.
- PERIOD is 4, resulting in an effective repeating counter pattern of $4+1 = 5$ counter clock periods.

Figure 29-25. PWM in Down Counting Mode

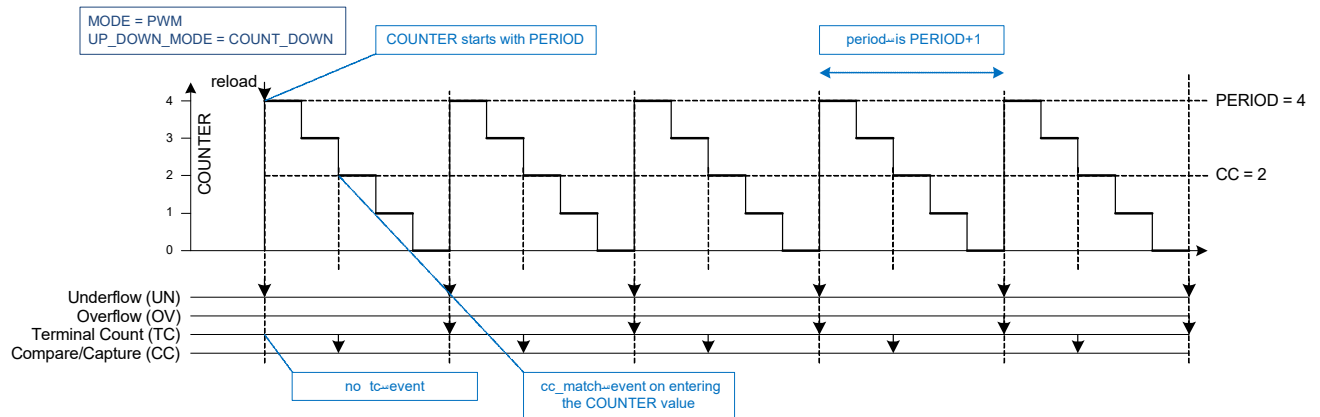


Figure 29-26 illustrates a PWM in down counting mode with different CC values. The figure also illustrates how a right-aligned PWM can be creating using the PWM in down counting mode. Note that the CC is changed (to CC_BUFF, which is not depicted) on a tc event.

Figure 29-26. Right-Aligned Down Counting PWM




Figure 29-27 illustrates a PWM in up/down counting mode. The counter is initialized (to 1) and started with a software-based reload event.

Notes:

- When the counter changes from a state in which COUNTER is 4, an overflow is generated.
- When the counter changes from a state in which COUNTER is 0, an underflow and tc event are generated.
- When the counter changes from a state in which COUNTER is 2, a cc_match event is generated. Note that the actual counter value COUNTER from before the reload event is NOT used, instead the counter value before the reload event is considered to be 0.
- PERIOD is 4, resulting in an effective repeating counter pattern of $2 \times 4 = 8$ counter clock periods.

MODE = PWM
UP_DOWN_MODE = COUNT_UPDN1

COUNTER starts with 4

period is $2 \times \text{PERIOD}$

PERIOD = 4

CC = 2

reload

COUNTER

Underflow (UN)

Overflow (OV)

Terminal Count (TC)

Compare/Capture (CC)

no tc-event

cc_match--event on leaving the COUNTER value

Note:

- The actual counter value COUNTER from before the reload event is NOT used. Instead the counter value before the reload event is considered to be 0. As a result, when the first CC value at the reload event is 0, a cc_match event is generated.
- CC is changed (to CC_BUFF, which is not depicted) on a tc event.

MODE = PWM
UP_DOWN_MODE = COUNT_UPDN1

reload ← 0 1 4 CC PERIOD = 4

COUNTER

Underflow (UN)
Overflow (OV)
Terminal Count (TC)
Compare/Capture (CC)

cc_match—event at the start of the period

overflow—and cc_match—events coincide

Center aligned PWM
CC = PERIOD – pulse width/2
UNDERFLOW_MODE = CLEAR
OVERFLOW_MODE = SET
CC_MATCH_MODE = INVERT

pwm

The following three modes are supported:

- **PWM_STOP_ON_KILL is '1' (PWM_SYNC_KILL is don't care):** Stop on Kill mode. This mode stops the counter on a stop/kill event. Reload or start event is required to run the counter. Both software and external trigger input can be selected as stop kill. Edge detection mode is required.
- **PWM_STOP_ON_KILL is '0' and PWM_SYNC_KILL is '0':** Asynchronous Kill mode. This mode keeps the counter running, but suppresses the PWM output signals synchronously with the next count clock ("active count" pre-scaled `clk_counter`) and continues to do so for the duration of the stop/kill event. Only the external trigger input can be selected as asynchronous kill. Pass through detection mode is required.

- PWM_STOP_ON_KILL is '0' and PWM_SYNC_KILL is '1': Synchronous Kill mode. This mode keeps the counter running, but suppresses the PWM output signals synchronously with the next count clock ("active count" pre-scaled clk_counter) and continues to do so until the next tc event without a stop/kill event. Only the external trigger input can be selected as synchronous kill. Rising edge detection mode is required.

Figure 29-29, Figure 29-30, and Figure 29-31 illustrate the above three modes.

Figure 29-29. PWM Stop on Kill




Figure 29-30. PWM Async Kill




Figure 29-31. PWM Sync Kill


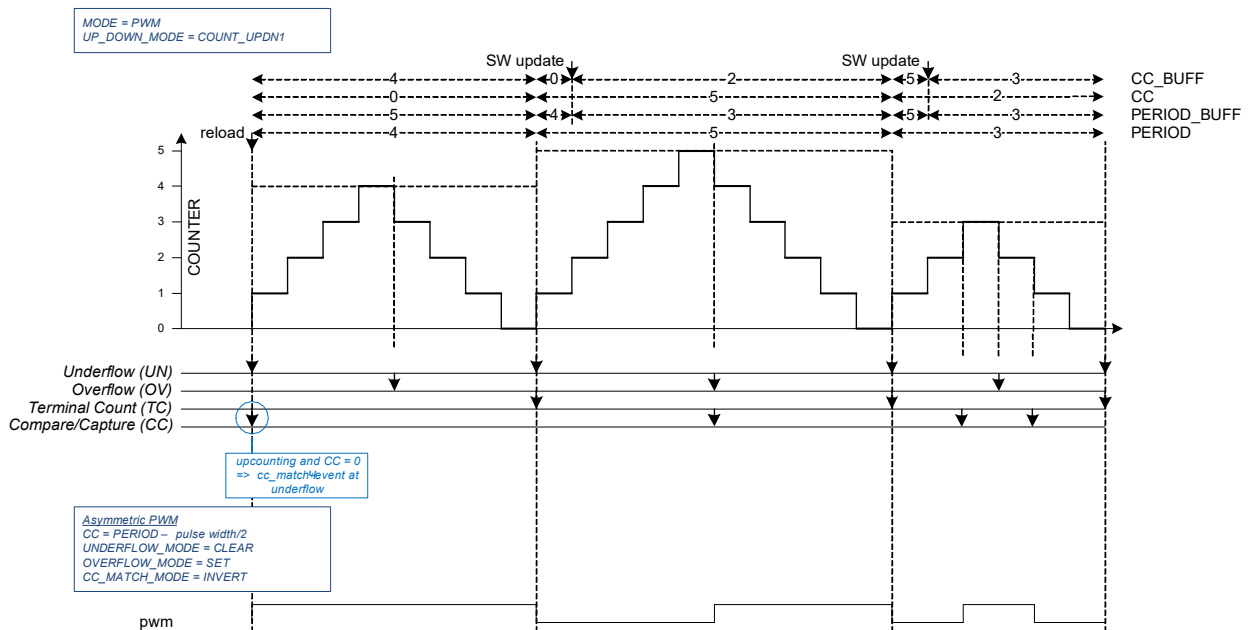


Figure 29-32 illustrates center-aligned PWM with PERIOD/PERIOD_BUFF and CC/CC_BUFF registers (up/down counting mode 1). At the TC condition, the PERIOD and CC registers are automatically exchanged with the PERIOD_BUFF and CC_BUFF registers. The swap event is generated by hardware trigger 1, which is a constant '1' and therefore always active at the TC condition. After the hardware exchange, the software handler on the tc interrupt cause updates PERIOD_BUFF and CC_BUFF.

Figure 29-32. PWM Mode CC Swap Event




A potential problem arises when software updates are not completed before the next tc event with an active pending swap event. For example, if software updates PERIOD_BUFF before the tc event and CC_BUFF after the tc event, swapping does not reflect the CC_BUFF register update. To prevent this from happening, the swap event should be generated by software through a register write after both the PERIOD_BUFF and CC_BUFF registers are updated. The swap event is kept pending by the hardware until the next tc event occurs.

The previous section addressed synchronized updates of the CC/CC_BUFF and PERIOD/PERIOD_BUFF registers of a single PWM using a software-generated swap event. During motor control, three PWMs work in unison and updates to all period and compare register pairs should be synchronized. All three PWMs have synchronized periods and as a result have synchronized tc events. The swap event for all three PWMs is generated by software through a single register write. The software should generate the swap events after the PERIOD_BUFF and CC_BUFF registers of all three PWMs are updated.

Figure 29-22 uses CTRL.QUADRATURE_MODE[0] for “pwm” polarity and CTRL.QUADRATURE_MODE[1] for “pwm_n” polarity. Figure 29-33 illustrates how the polarity settings control the PWM output signals “pwm” and “pwm_n”.

Note: When the counter is not enabled or not running ((temporarily) stopped or killed), the PWM output signals values are determined by their respective polarity settings.

Figure 29-33. PWM Outputs When Killed



29.3.4.1 Asymmetric PWM

The PWM mode supports the generation of an asymmetric PWM. For an asymmetric PWM, the `pwm_dt_input` pulse is not necessarily centered in the middle of the period. This functionality is realized by having a different CC value when counting up and when counting down. The CC and CC_BUFF values are exchanged on an overflow event. Note that this restricts the asymmetry of the generated `pwm_dt_input` pulse.

The COUNT_UPDN2 mode should use the same period value when counting up and counting down. When PERIOD and PERIOD_BUFF are swapped on a tc event (overflow or underflow event), care should be taken to ensure that:

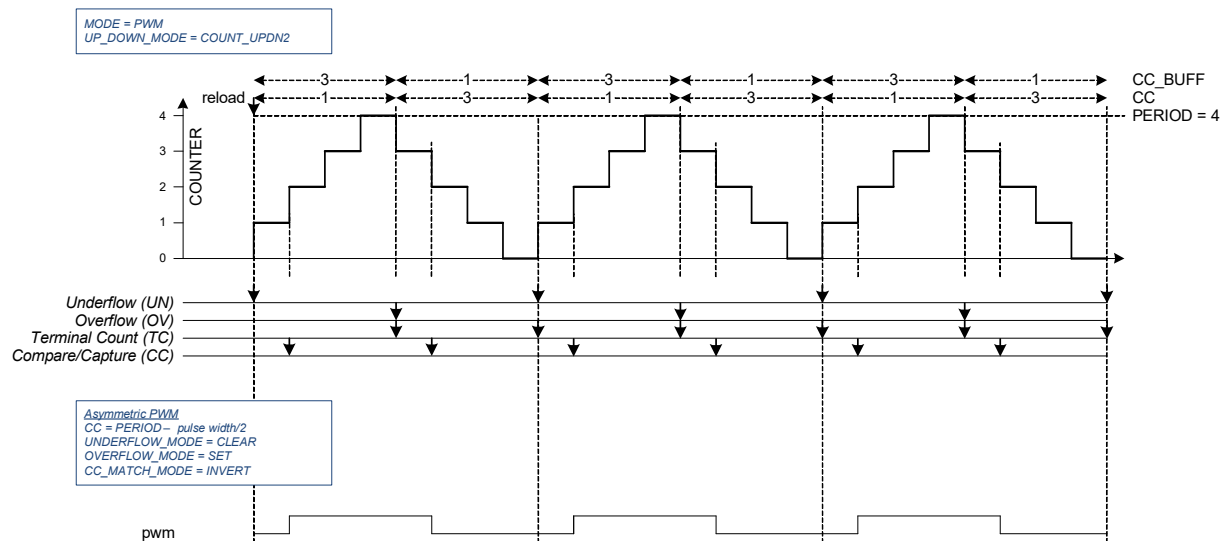
- Within a PWM period (tc event coincides with an overflow event), the period values are the same (an overflow swap of PERIOD and PERIOD_BUFF should not change the period value; that is, PERIOD_BUFF should be PERIOD)
- Between PWM periods (tc event coincides with an underflow event), the period value can change (an underflow swap of PERIOD and PERIOD_BUFF may change the period value; that is, PERIOD_BUFF may be different from PERIOD).

Figure 29-34 illustrates how the COUNT_UPDN2 mode is used to generate an asymmetric PWM.

Notes:

- When up counting, when CC value at the underflow event is 0, a `cc_match` event is generated.
- When down counting, when CC value at the overflow event is PERIOD, a `cc_match` event is generated.
- A tc event is generated for both an underflow and overflow event. The tc event is used to exchange the CC and CC_BUFF values.

Figure 29-34. Asymmetric PWM



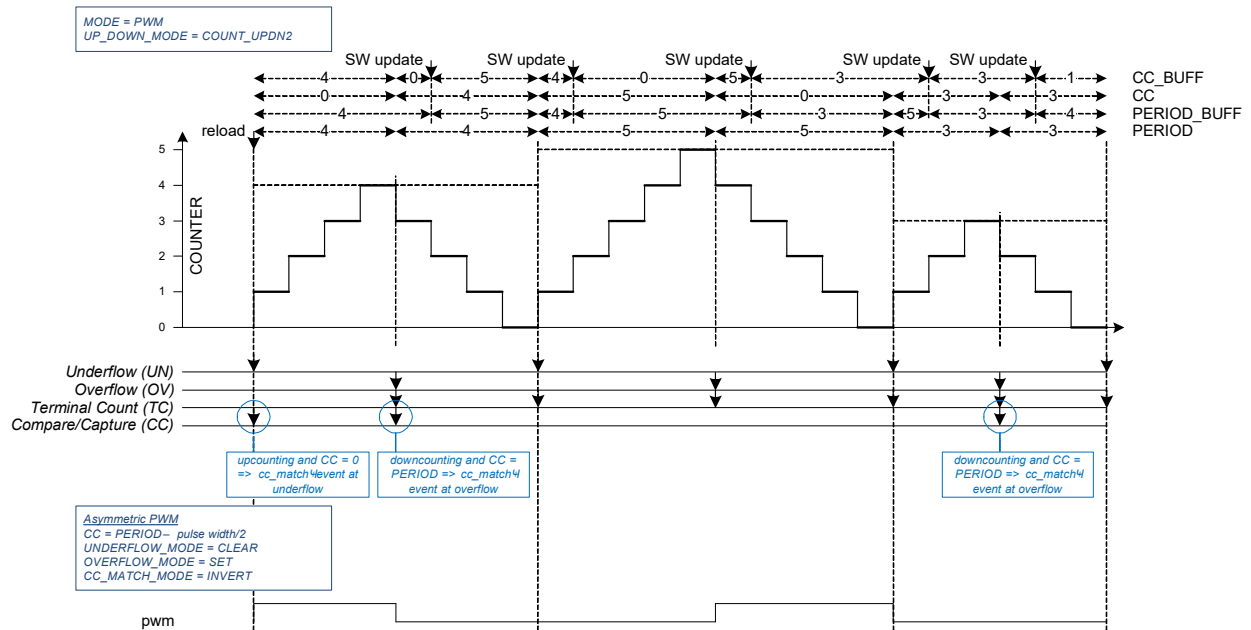
The previous waveform illustrated functionality when the CC values are neither "0" nor PERIOD. Corner case conditions in which the CC values equal "0" or PERIOD are illustrated as follows.

Figure 29-35 illustrates how the COUNT_UPDN2 mode is used to generate an asymmetric PWM.

Notes:

- When up counting, when CC value at the underflow event is 0, a `cc_match` event is generated.
- When down counting, when CC value at the overflow event is PERIOD, a `cc_match` event is generated.
- A tc event is generated for both an underflow and overflow event. The tc event is used to exchange the CC and CC_BUFF values.
- Software updates CC_BUFF and PERIOD_BUFF in an interrupt handler on the tc event (and overwrites the hardware updated values from the CC/CC_BUFF and PERIOD/PERIOD_BUFF exchanges).

Figure 29-35. Asymmetric PWM when Compare = 0 or Period



29.3.4.2 Configuring Counter for PWM Mode

The steps to configure the counter for the PWM mode of operation and the affected register bits are as follows.

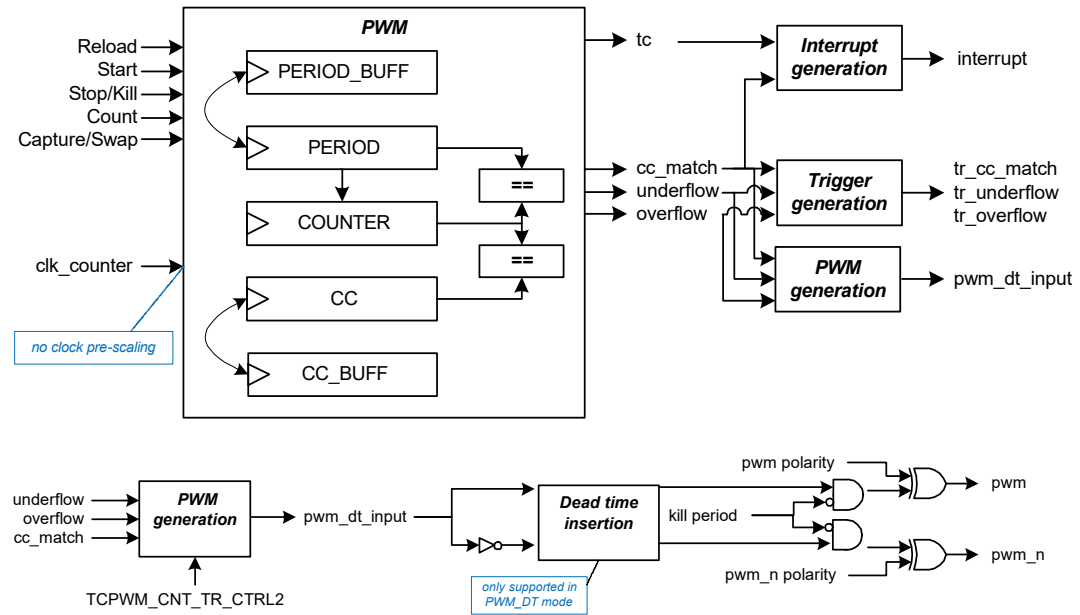
1. Disable the counter by writing '1' to the TCPWM_CTRL_CLR register.
2. Select PWM mode by writing '100b' to the MODE[26:24] field of the TCPWM_CNT_CTRL register.
3. Set clock prescaling by writing to the GENERIC[15:8] field of the TCPWM_CNT_CTRL register.
4. Set the required 16-bit period in the TCPWM_CNT_PERIOD register and the buffer period value in the TCPWM_CNT_PERIOD_BUFF register to swap values, if required.
5. Set the 16-bit compare value in the TCPWM_CNT_CC register and buffer compare value in the TCPWM_CNT_CC_BUFF register to swap values, if required.
6. Set the direction of counting by writing to the UP_DOWN_MODE[17:16] field of the TCPWM_CNT_CTRL register to configure left-aligned, right-aligned, or center-aligned PWM.
7. Set the PWM_STOP_ON_KILL and PWM_SYNC_KILL fields of the TCPWM_CNT_CTRL register as required.
8. Set the TCPWM_CNT_TR_CTRL0 register to select the trigger that causes the event (reload, start, kill, swap, and count).
9. Set the TCPWM_CNT_TR_CTRL1 register to select the edge that causes the event (reload, start, kill, swap, and count).
10. pwm and pwm_n can be controlled by the TCPWM_CNT_TR_CTRL2 register to set, reset, or invert upon CC, OV, and UN conditions.
11. If required, set the interrupt upon TC or CC condition.
12. Enable the counter by writing '1' to the TCPWM_CTRL_SET register. A start trigger must be provided through firmware (TCPWM_CMD_START register) to start the counter if the hardware start signal is not enabled.

29.3.5 Pulse Width Modulation with Dead Time Mode

The PWM-DT functionality is the same as PWM functionality, except for the following differences:

- PWM_DT supports dead time insertion; PWM does not support dead time insertion.
- PWM_DT does not support clock pre-scaling; PWM supports clock pre-scaling.

Figure 29-36. PWM with Dead Time Functionality



Dead time insertion is a step that operates on a preliminary PWM output signal **pwm_dt_input**, as illustrated in Figure 29-36.

Figure 29-37 illustrates dead time insertion for different dead times and different output signal polarity settings.

Figure 29-37. Dead-time Timing





Figure 29-38 illustrates how the polarity settings and stop/kill functionality combined control the PWM output signals “pwm” and “pwm_n”.

Figure 29-38. Dead Time and Kill



29.3.5.1 *Configuring Counter for PWM with Dead Time Mode*

The steps to configure the counter for PWM with Dead Time mode of operation and the affected register bits are as follows:

1. Disable the counter by writing '1' to the TCPWM_CTRL_CLR register.
2. Select PWM with Dead Time mode by writing '101' to the MODE[26:24] field of the TCPWM_CNT_CTRL register.
3. Set the required dead time by writing to the GENERIC[15:8] field of the TCPWM_CNT_CTRL register.
4. Set the required 16-bit period in the TCPWM_CNT_PERIOD register and the buffer period value in the TCPWM_CNT_PERIOD_BUFF register to swap values, if required.
5. Set the 16-bit compare value in the TCPWM_CNT_CC register and the buffer compare value in the TCPWM_CNT_CC_BUFF register to swap values, if required.
6. Set the direction of counting by writing to the UP_DOWN_MODE[17:16] field of the TCPWM_CNT_CTRL register to configure left-aligned, right-aligned, or center-aligned PWM.
7. Set the PWM_STOP_ON_KILL and PWM_SYNC_KILL fields of the TCPWM_CNT_CTRL register as required.
8. Set the TCPWM_CNT_TR_CTRL0 register to select the trigger that causes the event (reload, start, kill, swap, and count).
9. Set the TCPWM_CNT_TR_CTRL1 register to select the edge that causes the event (reload, start, kill, swap, and count).
10. pwm and pwm_n can be controlled by the TCPWM_CNT_TR_CTRL2 register to set, reset, or invert upon CC, OV, and UN conditions.
11. If required, set the interrupt upon TC or CC condition.
12. Enable the counter by writing '1' to the TCPWM_CTRL_SET register. A start trigger must be provided through firmware (TCPWM_CMD_START register) to start the counter if the hardware start signal is not enabled.

29.3.6 Pulse Width Modulation Pseudo-Random Mode (PWM_PR)

The PWM_PR functionality changes the counter value using the linear feedback shift register (LFSR). This results in a pseudo random number sequence. A signal similar to PWM signal is created by comparing the counter value COUNTER with CC. The generated signal has different frequency/noise characteristics than a regular PWM signal.

Table 29-26. PWM_PR Mode Trigger Inputs

Trigger Inputs	Usage
reload	Same behavior as start event. Can be used only when the counter is not running.
start	Starts the counter. The counter is not initialized by hardware. The current counter value is used. Behavior is independent on UP_DOWN_MODE. Can be used only when the counter is not running.
stop/kill	Stops the counter. Different stop/kill modes exist.
count	Not used.
capture	This event acts as a swap event. When this event is active, the CC/CC_BUFF and PERIOD/PERIOD_BUFF registers are exchanged on a tc event (when specified by CTRL.AUTO_RELOAD_CC and CTRL.AUTO_RELOAD_PERIOD). A swap event requires rising, falling, or rising/falling edge event detection mode. Pass-through mode is not supported, unless the selected event is a constant '0' or '1'. When a swap event is detected and the counter is running, the event is kept pending until the next tc event. When a swap event is detected and the counter is not running, the event is cleared by hardware.

Note: Event detection is on the peripheral clock, clk_peri.

Table 29-27. PWM_PR Supported Features

Supported Features	Description
Clock pre-scaling	Pre-scales the counter clock, clk_counter.
One-shot	Counter is stopped by hardware, after a single period of the counter (counter value equals period value PERIOD).
Auto reload CC	CC and CC_BUFF are exchanged on a swap event AND tc event (when specified by CTRL.AUTO_RELOAD_CC).
Auto reload PERIOD	PERIOD and PERIOD_BUFF are exchanged on a swap event and tc event (when specified by CTRL.AUTO_RELOAD_PERIOD).
Kill modes	Specified by PWM_STOP_ON_KILL. See memory map for further details.

Note: The count event is not used. As a result, the PWM_PR functionality operates on the pre-scaled counter clock (clk_counter), rather than on an “active count” pre-scaled counter clock.

Table 29-28. PWM_PR Trigger Outputs

Trigger Outputs	Description
cc_match (CC)	Counter changes from a state in which COUNTER equals CC.
Underflow (UN)	Not used.
Overflow (OV)	Not used.

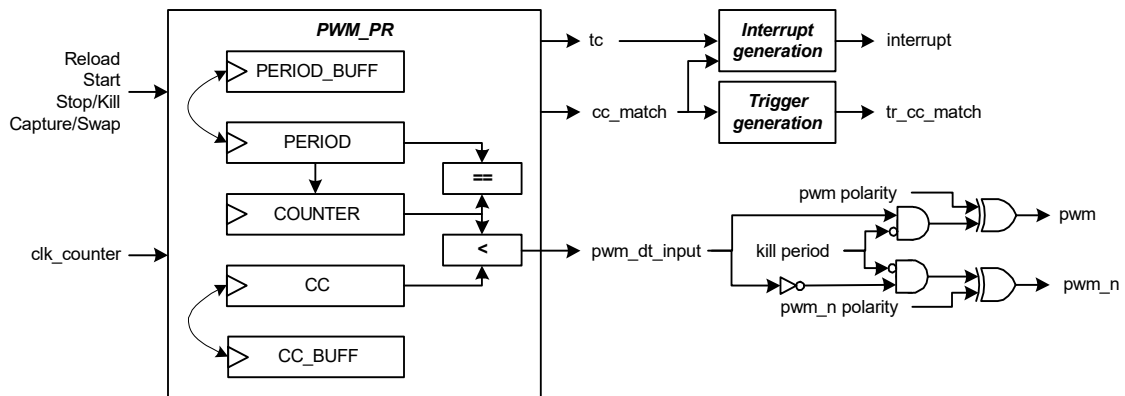
Table 29-29. PWM_PR Interrupt Outputs

Interrupt Outputs	Description
cc_match (CC)	Counter changes from a state in which COUNTER equals CC.
tc	Counter changes from a state in which COUNTER equals PERIOD.

Table 29-30. PWM_PR PWM Outputs

PWM Outputs	Description
pwm	PWM output.
pwm_n	Complementary PWM output.

Figure 29-39. PWM_PR Functionality



The PWM_PR functionality is described as follows:

- The counter value COUNTER is initialized by software (to a value different from 0).
- A reload or start event starts PWM_PR operation.
- During PWM_PR operation:
 - The counter value COUNTER is changed based on the LFSR polynomial: $x^{16} + x^{14} + x^{13} + x^{11} + 1$ (en.wikipedia.org/wiki/Linear_feedback_shift_register).

$$\text{temp} = \text{COUNTER}[16:16] \wedge \text{COUNTER}[16:14] \wedge \text{COUNTER}[16:13] \wedge \text{COUNTER}[16:11] \text{ or}$$

$$\text{temp} = \text{COUNTER}[0] \wedge \text{COUNTER}[2] \wedge \text{COUNTER}[3] \wedge \text{COUNTER}[5];$$

$$\text{COUNTER} = (\text{temp} \ll 15) | (\text{COUNTER} \gg 1)$$
 This will result in a pseudo random number sequence for COUNTER. For example, when COUNTER is initialized to 0xACE1, the number sequence is: 0xACE1, 0x5670, 0xAB38, 0x559C, 0x2ACE, 0x1567, 0x8AB3... This sequence will repeat itself after $2^{16} - 1$ or 65535 counter clock cycles.
 - A 32-bit counter uses the LFSR polynomial: $x^{32} + x^{30} + x^{26} + x^{25} + 1$
 - When the counter value COUNTER equals CC, a cc_match event is generated.
 - When the counter value COUNTER equals PERIOD, a tc event is generated.
 - On a tc event, the CC/CC_BUFF and PERIOD/PERIOD_BUFF can be conditionally exchanged under control of the capture/swap event and the CTRL.AUTO_RELOAD_CC and CTRL.AUTO_RELOAD_PERIOD field (see PWM functionality).
 - The output pwm_dt_input reflects: COUNTER[14:0] < CC[15:0]. Note that only the lower 15 bits of COUNTER are used for comparison, while the COUNTER itself can run up to 16- or 32-bit values. As a result, for CC greater or equal to 0x8000, pwm_dt_input is always 1. The pwm polarity can be inverted (as specified by CTRL.QUADRATURE_MODE[0]).

As mentioned, different stop/kill modes exist. The mode is specified by PWM_STOP_ON_KILL (PWM_SYNC_KILL should be '0' – asynchronous kill mode). The memory map describes the modes and the desired settings for the stop/kill event. The following two modes are supported:

- PWM_STOP_ON_KILL is '1'. This mode stops the counter on a stop/kill event.
- PWM_STOP_ON_KILL is '0'. This mode keeps the counter running, but suppresses the PWM output signals immediately and continues to do so for the duration of the stop/kill event.


Note that the LFSR produces a deterministic number sequence (given a specific counter initialization value). Therefore, it is possible to calculate the COUNTER value after a certain number of LFSR iterations, n . This calculated COUNTER value can be used as PERIOD value, and the tc event will be generated after precisely n counter clocks.

Figure 29-40 illustrates PWM_PR functionality.

Notes:

- The grey shaded areas represent the counter region in which the pwm_dt_input value is '1', for a CC value of 0x4000. There are two areas, because only the lower 15 bits of the counter value are used.
- When CC is set to 0x4000, roughly one-half of the counter clocks will result in a pwm_dt_input value of '1'.

Figure 29-40. PWM_PR Output



29.3.6.1 Configuring Counter for Pseudo-Random PWM Mode

The steps to configure the counter for pseudo-random PWM mode of operation and the affected register bits are as follows.

1. Disable the counter by writing '1' to the TCPWM_CTRL_CLR register.
2. Select pseudo-random PWM mode by writing '110' to the MODE[26:24] field of the TCPWM_CNT_CTRL register.
3. Set the required period (16 bit) in the TCPWM_CNT_PERIOD register and buffer period value in the TCPWM_CNT_PERIOD_BUFF register to swap values, if required.
4. Set the 16-bit compare value in the TCPWM_CNT_CC register and the buffer compare value in the TCPWM_CNT_CC_BUFF register to swap values.
5. Set the PWM_STOP_ON_KILL and PWM_SYNC_KILL fields of the TCPWM_CNT_CTRL register as required.
6. Set the TCPWM_CNT_TR_CTRL0 register to select the trigger that causes the event (reload, start, kill, and swap).
7. Set the TCPWM_CNT_TR_CTRL1 register to select the edge that causes the event (reload, start, kill, and swap).
8. pwm and pwm_n can be controlled by the TCPWM_CNT_TR_CTRL2 register to set, reset, or invert upon CC, OV, and UN conditions.
9. If required, set the interrupt upon TC or CC condition.
10. Enable the counter by writing '1' to the TCPWM_CTRL_SET register. A start trigger must be provided through firmware (TCPWM_CMD_START register) to start the counter if the hardware start signal is not enabled.

29.4 TCPWM Registers

Table 29-31. List of TCPWM Registers

Register	Comment	Features
TCPWM_CTRL	TCPWM control register	Enables the counter block
TCPWM_CTRL_CLR	TCPWM control clear register	Used to avoid race-conditions on read-modify-write attempt to the CTRL register
TCPWM_CTRL_SET	TCPWM control set register	Used to avoid race-conditions on read-modify-write attempt to the CTRL register
TCPWM_CMD_CAPTURE	TCPWM capture command register	Generate a capture trigger input from software
TCPWM_CMD_RELOAD	TCPWM reload command register	Generate a reload trigger input from software
TCPWM_CMD_STOP	TCPWM stop command register	Generate a stop trigger input from software
TCPWM_CMD_START	TCPWM start command register	Generate a start trigger input from software
TCPWM_INTR_CAUSE	TCPWM counter interrupt cause register	Determines the source of the combined interrupt signal
TCPWM_CNT_CTRL	Counter control register	Configures counter mode, encoding modes, one-shot mode, swap mode, kill mode, dead time, clock pre-scaling, and counting direction
TCPWM_CNT_STATUS	Counter status register	Reads the direction of counting, dead time duration, and clock pre-scaling; checks whether the counter is running
TCPWM_CNT_COUNTER	Count register	Contains the 16- or 32-bit counter value
TCPWM_CNT_CC	Counter compare/capture register	Captures the counter value or compares the value with counter value
TCPWM_CNT_CC_BUFF	Counter buffered compare/capture register	Buffer register for counter CC register; swaps period value
TCPWM_CNT_PERIOD	Counter period register	Contains upper value of the counter
TCPWM_CNT_PERIOD_BUFF	Counter buffered period register	Buffer register for counter period register; swaps compare value
TCPWM_CNT_TR_CTRL0	Counter trigger control register 0	Selects trigger for specific counter events
TCPWM_CNT_TR_CTRL1	Counter trigger control register 1	Determine edge detection for specific counter input signals
TCPWM_CNT_TR_CTRL2	Counter trigger control register 2	Controls counter output lines upon CC, OV, and UN conditions
TCPWM_CNT_INTR	Interrupt request register	Sets the register bit when TC or CC condition is detected
TCPWM_CNT_INTR_SET	Interrupt set request register	Sets the corresponding bits in interrupt request register
TCPWM_CNT_INTR_MASK	Interrupt mask register	Mask for interrupt request register
TCPWM_CNT_INTR_MASKED	Interrupt masked request register	Bitwise AND of interrupt request and mask registers

30. Inter-IC Sound Bus



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation
- Application notes
- Code examples

The Inter-IC Sound Bus (I²S) is a serial bus interface standard used to connect digital audio devices together. The specification is from Philips[®] Semiconductor (I²S bus specification: February 1986, revised June 5, 1996). In addition to the standard I²S format, the I²S block also supports the Left Justified (LJ) format and the Time Division Multiplexed (TDM) format.

30.1 Features

- Supports standard Philips I²S, LJ, and eight-channel TDM digital audio interface formats
- Supports both master and slave mode operation in all the digital audio formats
- Supports independent operation of Receive (Rx) and Transmit (Tx) directions
- Supports operating from an external master clock provided through an external IC such as audio codec
- Provides configurable clock divider registers to generate the required sample rates
- Supports data word length of 8-bit, 16-bit, 18-bit, 20-bit, 24-bit, and 32-bit per channel
- Supports channel length of 8-bit, 16-bit, 18-bit, 20-bit, 24-bit, and 32-bit per channel (channel length fixed at 32-bit in TDM format)
- Provides two hardware FIFO buffers, one each for the Tx block and Rx block, respectively
- Supports both DMA- and CPU-based data transfers

30.2 Architecture


 Figure 30-1. I²S Block Diagram


Figure 30-1 shows the high-level block diagram of the I²S block, which consists of two sub-blocks – I²S Transmitter (Tx) and I²S Receiver (Rx). The digital audio interface format and master/slave mode configuration can be done independently for the Tx and Rx blocks. In the master mode, the word select (ws) and serial data clock (sck) are generated by the I²S block in the PSoC 6 MCU. In the slave mode, the ws and sck signals are input signals to the PSoC 6 MCU, and are generated by the external master device. The I²S block configuration, control, and status registers, along with the FIFO data buffers are accessible through the AHB bus. AHB bus masters such as CPU and DMA can access the I²S registers through the AHB interface. Refer to the [device datasheet](#) for information on port pin assignments of the I²S block signals and AC/DC electrical specifications.

30.3 Digital Audio Interface Formats

The I²S block supports the following digital audio interface formats.

- Standard I²S format
- Left Justified format
- Time Division Multiplexed (TDM) format

The Tx and Rx sub-blocks can be independently configured to support one of the above formats in either master or slave mode. The I2S_MODE bits in the I2S_TX_CTL and I2S_RX_CTL registers are used to configure the digital audio interface format for the Tx and Rx blocks respectively. The MS (Master/Slave) bit in the I2S_TX_CTL and I2S_RX_CTL registers is used to configure the blocks in master or slave mode.

30.3.1 Standard I²S Format

Figure 30-2 shows the timing diagrams for the different word length and channel length combinations in the standard I²S digital audio format. In the standard I²S format, the word select signal (ws) is low for left channel data, and high for right channel data. The ws signal transitions one bit-clock (sck) early relative to the start of the left/right channel data. All the serial data (sd), ws signal transitions on the falling edge of the sck signal, and the read operations on the ws and sd lines are usually done on the rising edge of sck. Therefore, the I²S Tx block writes to the serial data (tx_sdo) line on the falling edge of tx_sck,

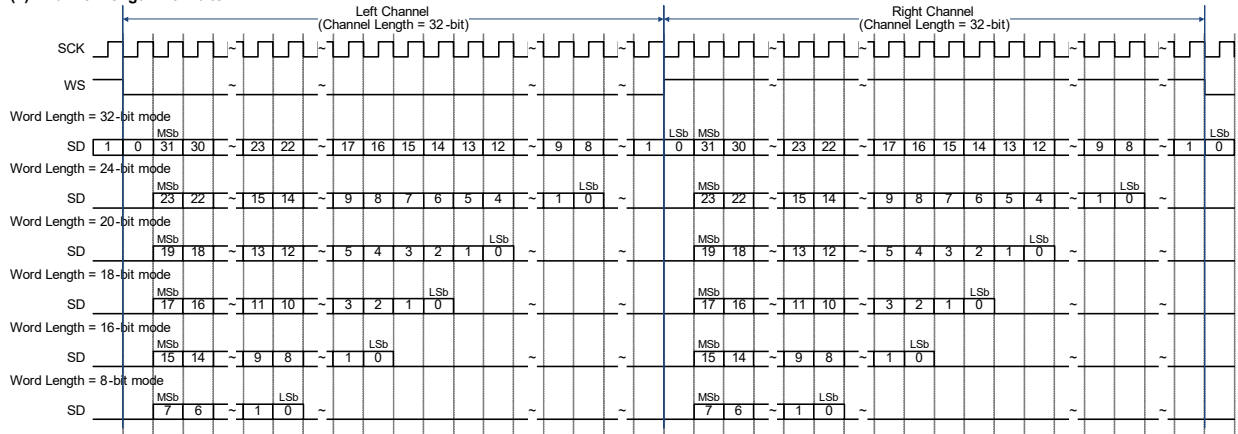
and the I²S Rx block reads the data (rx_sdi) on the rising edge of rx_sck. The serial data is transmitted most significant bit (MSb) first. Depending on whether the block is in master or slave mode, the ws/sck signals are either generated by the block (master mode) or input signals to the block (slave mode).

The I²S block supports configurable word length and channel length selection options. The word length for the Tx and Rx blocks can be configured using the WORD_LEN bits in the I2S_TX_CTL and I2S_RX_CTL registers, respectively. The channel length for the Tx and Rx blocks can be configured using the CH_LEN bits in the I2S_TX_CTL and I2S_RX_CTL registers respectively. The channel length configuration should always be greater than or equal to the word length configuration. Ensure that when the I²S Rx block is operated in slave mode, the master Tx device ensures that its channel length configuration aligns with the I²S Rx block channel length setting. If there is channel length mismatch, the PSoC I²S Rx block in slave mode will not operate correctly.

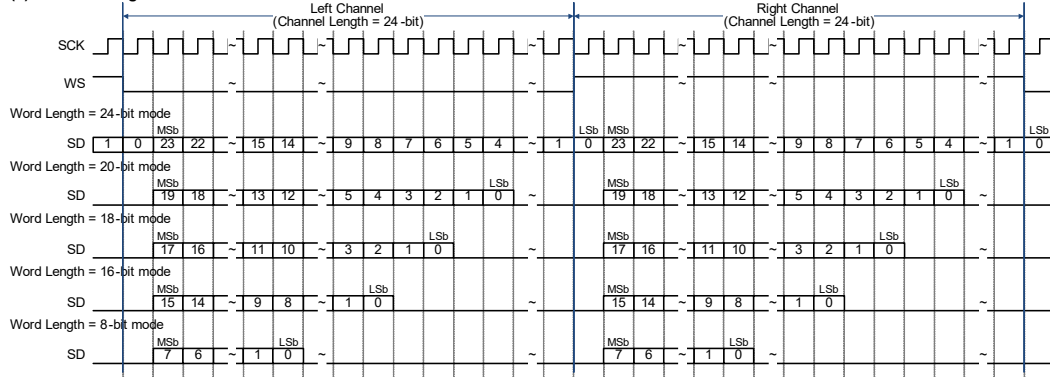
In the Tx block, when the channel length is greater than the word length, the unused bits can be transmitted either as '0' or '1'. This selection is made using the OVHDATA bit in the I2S_TX_CTL register. In the Rx block, when the word length is less than 32 bits, the unused most significant bits written to the 32-bit Rx FIFO register can either be set to '0' or sign bit extended. This selection is made using the BIT_EXTENSION bit in the I2S_RX_CTL register.

Figure 30-2. Standard I²S Format (Word Length and Channel Length Combination Timing Diagrams)

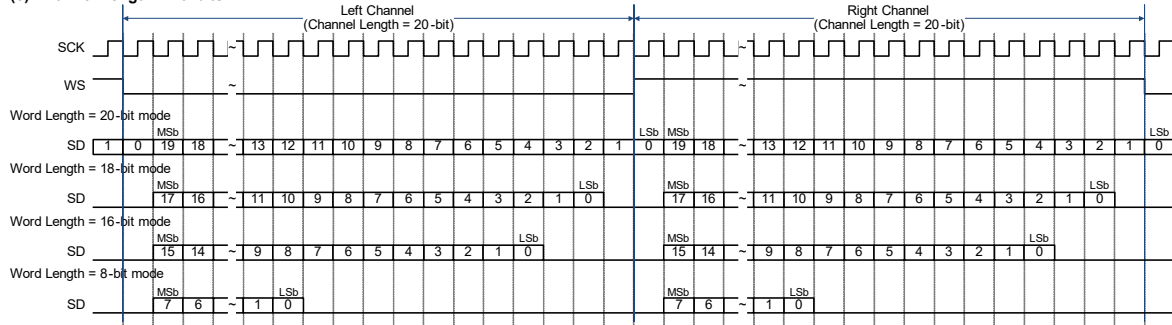
(1) Channel Length = 32-bits



(2) Channel Length = 24-bits



(3) Channel Length = 20-bits



(4) Channel Length = 8-bits




Table 30-1 lists the supported word length and channel length combinations.

Table 30-1. Word Length and Channel Length Combinations

		Word Length					
		8-bit	16-bit	18-bit	20-bit	24-bit	32-bit
Channel Length	32-bit	Valid	Valid	Valid	Valid	Valid	Valid
	24-bit	Valid	Valid	Valid	Valid	Valid	Invalid
	30-bit	Valid	Valid	Valid	Valid	Invalid	Invalid
	18-bit	Valid	Valid	Valid	Invalid	Invalid	Invalid
	16-bit	Valid	Valid	Invalid	Invalid	Invalid	Invalid
	8-bit	Valid	Invalid	Invalid	Invalid	Invalid	Invalid


30.3.2 Left Justified (LJ) Format

Figure 30-3 shows the timing diagrams for the Left Justified interface format using the 32-bit channel length and 32-bit word length configuration as an example. The only differences between the standard I²S and LJ formats are:

- In the standard I²S format, WS signal is low for left channel data and high for right channel data. In the LJ format, WS signal is high for left channel data and low for right channel data.
- In the standard I²S format, WS signal transitions one bit-clock (sck) early relative to the start of the channel data (coincides with LSb of the previous channel). In the LJ format, there is no early transition, and the WS signal transitions coincide with the start of the channel data.

Apart from these differences, all the features explained in the standard I²S format section apply to the LJ format as well.

Figure 30-3. Left Justified Digital Audio Format



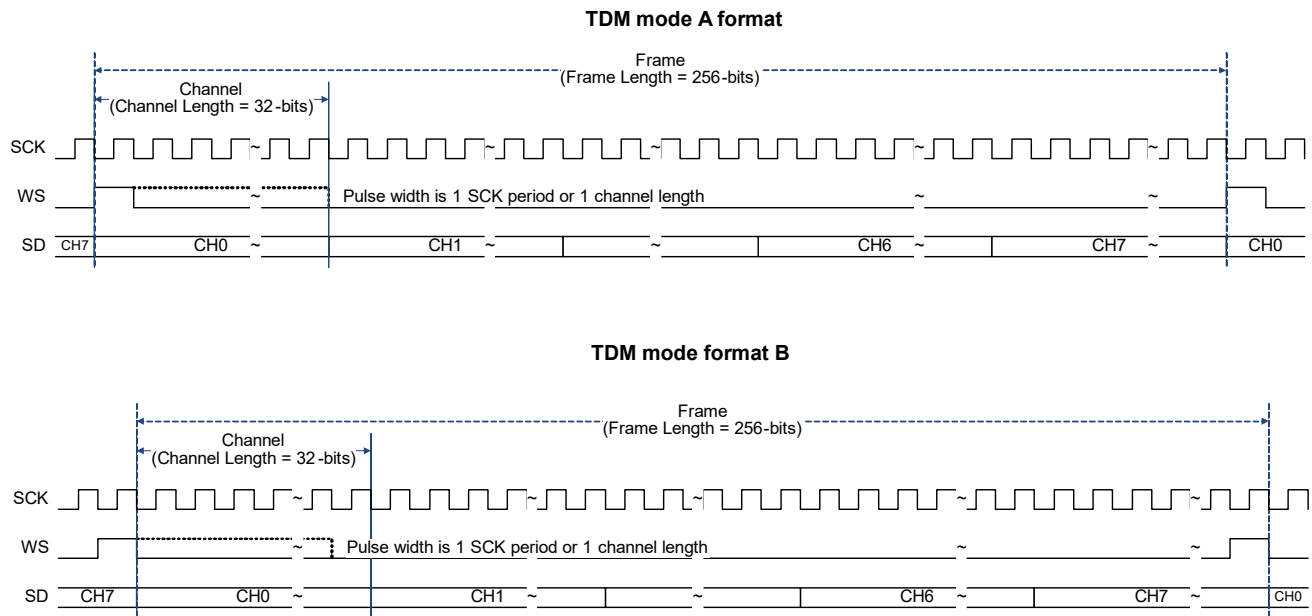
30.3.3 Time Division Multiplexed (TDM) Format

Figure 30-4 shows the timing diagrams for the two types of Time Division Multiplexed (TDM) formats supported by the I²S block. The differences between the standard I²S/LJ formats and the TDM format are as follows:

- Standard I²S/LJ formats support only two channels (left/right) per frame, while TDM format supports up to eight channels per frame.
- In the TDM format, channel length for all eight channels is fixed at 32 bits. In the standard I²S/LJ formats, the channel length is configurable. The word length per channel is configurable similar to the standard I²S and the data is also transmitted most significant bit first. Similar to I²S, when the word length per channel is less than the 32-bit channel length for Tx block, the OVHDATA bit in the I2S_TX_CTL register is used to fill the unused least significant channel data bits with either all zeros or all ones
- In the TDM format, all eight channels of data are always present in a frame, and thus the frame width is fixed at 256 bits. You have the option to configure the number of active channels in a frame by configuring the CH_NR bits in the I2S_TX_CTL and I2S_RX_CTL registers. In the standard I²S/LJ format, the CH_NR should always be configured for two channels. The number of active channels in the TDM format can be less than or equal to eight channels. The unused (inactive) channels always follow the active channels in a frame. As an example, if CH_NR is set for four channels, CH0 to CH3 are the active channels and CH4 to CH7 are the unused channels. The OVHDATA bit in the I2S_TX_CTL register is used to fill the unused channels with either all zeros or all ones.

- The pulse width of the word select (WS) signal in the TDM format can be configured to be either one bit clock (sck) wide or one channel wide. The selection is made using the WS_PULSE bit in the I2S_TX_CTL and I2S_RX_CTL registers. The pulse width is fixed to one channel width in the I²S/LJ format.
- Two types of TDM formats are supported. In TDM mode A, the WS rising edge signal to signify the start of frame coincides with the start of CH0 data. In TDM mode B, the WS rising edge signal to signify the start of frame is one bit clock (sck) early, relative to the start of CH0 data (coincides with the last bit of the previous frame). The selection between the two TDM formats is made using the I2S_MODE bits in the I2S_TX_CTL and I2S_RX_CTL registers.

Figure 30-4. TDM Digital Audio Interface Format



30.4 Clocking Polarity and Delay Options

The I²S block supports configurable clock polarity and delay options to alleviate any timing issues in the system involving PCB signal propagation delays, and delays associated with internal device signal routing.

When the I²S Tx block operates in the slave mode, the tx_sck and tx_ws signals are input signals to the PSoC 6 MCU, and the tx_sdo output signal is transmitted off the tx_sck falling edge. The tx_sdo signal is sampled by the external master device Rx block on the subsequent tx_sck rising edge. Timing issues arise if the tx_sdo signal reaching the master side Rx block does not meet the setup and hold time requirements for input data on the master side. The I²S Tx block in the PSoC 6 MCU has an option to advance the serial data transmission by 0.5 SCK cycles when the B_CLOCK_INV bit in the I2S_TX_CTL register is set. This feature can be used if there are timing issues while operating the I²S Tx block in slave mode.

Similarly, when the I²S Rx block operates in the master mode, the rx_sck and rx_ws signals are output signals from the PSoC 6 MCU, and the rx_sdi signal is transmitted by the external master device on the falling edge of rx_sck. The PSoC I2S Rx block samples the rx_sdi signal on the subsequent rx_sck rising edge. Timing issues arise if the rx_sdi signal reaching the PSoC block does not meet the setup and hold time requirements for input data. The I²S Rx block has an option to delay the serial data capture by 0.5 SCK cycles when the B_CLOCK_INV bit in the I2S_RX_CTL register is set. This feature can be used if there are timing issues while operating the I²S Rx block in master mode.

In addition to these clock delay options, there is also an option to invert the outgoing bit clock (sck) in master mode by setting the SCKO_POL bit in the I2S_TX_CTL and I2S_RX_CTL registers. Similarly, in the slave mode, there is an option to invert the incoming bit clock (sck) by setting the SCKI_POL bit in the I2S_TX_CTL and I2S_RX_CTL registers.

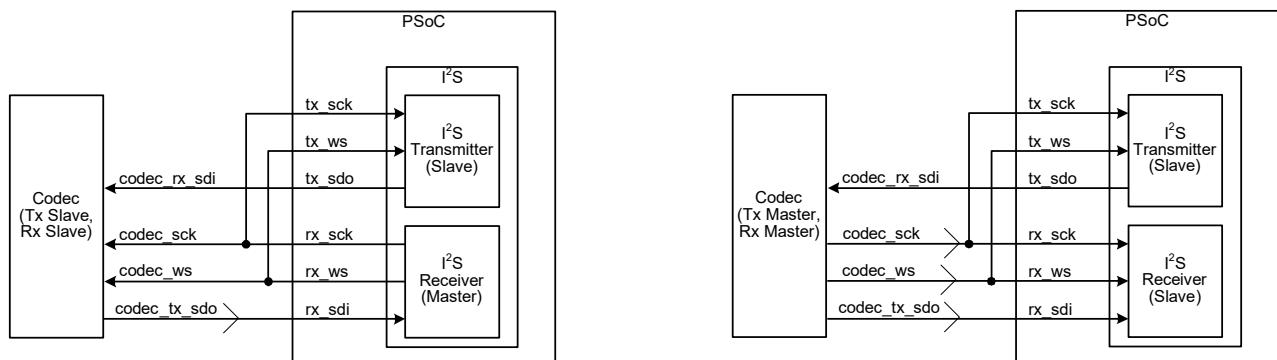
Refer to the [registers TRM](#) for detailed description of the B_CLOCK_INV, SCKI_POL, and SCKO_POL register configurations.

30.5 Interfacing with Audio Codecs

The I²S block in the PSoC 6 MCU interfaces with an audio codec device based on the choice of codec device and the end application requirements. Some scenarios and the connection diagrams are as follows:

- Codecs with separate ws and sck signals for the Rx and Tx directions: To interface with these codecs, the connections between the PSoC I²S block and the codec device will be as shown in [Figure 30-1](#) where the PSoC I²S Tx signals (tx_sck, tx_ws, tx_sdo) connect to the codec Rx signals, and the PSoC I²S Rx signals (rx_sck, rx_ws, rx_sdi) connect to the codec Tx signals. The direction of sck (tx_sck, rx_sck) and ws (tx_ws, rx_ws) signals depends on which device is the master and which device is the slave.
- Codecs with common ws and sck signals for both Rx and Tx directions: There are two possible configurations to interface these codecs with the PSoC 6 MCU as shown in [Figure 30-5](#). In both configurations, the sck signals (tx_sck, rx_sck, codec_sck) are shorted externally. The same goes for the ws signal connections as well (tx_ws, rx_ws, codec_ws). Ensure that only one block is driving the sck and ws lines. So when the codec acts as the slave device, the PSoC I²S Rx block should be in the master mode, and the PSoC I²S Tx block should be in the slave mode (or PSoC I²S Rx as slave and PSoC I²S Tx as master). When the codec acts as the master device, both the PSoC I²S Rx and PSoC I²S Tx blocks should be in slave mode.

Figure 30-5. Interfacing with Codecs having Common ws and sck Signals



30.6 Clocking Features


The I²S unit has three clock inputs.

Table 30-2. Clock Inputs

Signal	DESCRIPTION
clk_sys_i2s	System clock. This clock is used for the AHB slave Interface, control, status, and interrupt registers, and also clocks the DMA trigger control logic.
clk_audio_i2s	I ² S internal clock. This clock is used for I ² S transmitter (Tx)/receiver (Rx) blocks; it is asynchronous with the clk_sys_i2s. This clock is connected to the CLK_HF[1] high-frequency clock in the device. Refer to the Clocking System chapter on page 207 for more details on high frequency clocks.
clk_i2s_if	I ² S external clock. This clock is provided from an external I ² S bus host through a port pin. It is used in place of the clk_audio_i2s clock to synchronize I ² S data to the clock used by the external I ² S bus host.

[Figure 30-6](#) shows the clocking divider structure in the I²S block. In the master mode, the sck and ws signals are generated either using the clk_audio_i2s internal clock or the clk_i2s_if external clock. Refer to the [device datasheet](#) for the port pin assignment of clk_i2s_if clock. The CLOCK_SEL bit in the I2S_CLOCK_CTL register controls the selection between internal and external clocks.

Figure 30-6. Clocking Divider Structure



There are two stages of clock dividers in the I²S block as follows.

- The first stage clock divider is used to generate the internal I²S master clock (MCLK_SOC). The input clock to the first stage divider is either clk_audio_i2s or clk_i2s_if. The first stage clock divider is configured using the CLOCK_DIV bits in I2S_CLOCK_CTL register. Divider values from 1 to 64 are supported.
- The second stage clock divider is used to generate the sck signals. The input clock is the output from the first stage clock divider. This divider value is fixed at '8' (FTX_SCK = FRX_SCK = FMCLK_SOC/8). The word select (ws) signal frequency depends on the sck frequency, and the configured channel length value.

When in slave mode, the internal clock (MCLK_SOC) frequency should still be eight times the frequency of the input serial clock. You must choose the appropriate clock source and the CLOCK_DIV divider value to guarantee this condition is met in the slave mode of operation. Usually, when the PSoC I²S block operates in the slave mode, the host sends a master clock which is an integral multiple of the sampling rate. This master clock can be routed to the clk_i2s_if port pin. The CLOCK_DIV divider value can then be adjusted to ensure that the MCLK_SOC is eight times the input SCK frequency.

Table 30-3 gives an example of the clock divider settings for operating the I²S block at the standard sampling rates in the standard I²S format. Note that the first stage divider values in the table are the register field values – the actual divider values are one more than the configured register values as explained in the clock divider section. Refer to the [device datasheet](#) for details on maximum values of SCK frequency, and the output sampling rates.

Table 30-3. I²S Divider Values for Standard Audio Sampling Rates in Standard I²S Format

Sampling Rate (SR) (kHz)	WORD_LEN (bits)	SCK (2*WORD_LEN*SR) (MHz)	CLK_HF1 (or clk_i2s_if) (MHz)	(CLK_HF[1])/SCK (Total Divider Ratio)	CLK_CLOCK_DIV (First Divider)	Second Stage Divider (Fixed at 8)
8	32	0.512	49.152	96	11	8
16	32	1.024		48	5	
32	32	2.048		24	2	
48	32	3.072		16	1	
44.1	32	2.8224	45.1584	32	3	

30.7 FIFO Buffer and DMA Support

The I²S block has two FIFO buffers - one each for the Tx block and Rx block, respectively. The ordering format of the channel data in both Tx and Rx FIFOs depends on the configured digital audio format. This ordering format should be considered when writing to the Tx FIFO or reading from the Rx FIFO. In the standard I²S and LJ digital audio formats, the ordering of the data is (L, R, L, R, L, ...) where L refers to the left channel data and R refers to the right channel data. In the TDM format with the number of active channels set to four, the data order will be (CH0, CH1, CH2, CH3, CH0, CH1, CH2, CH3, CH0,). If the number of active channels is set to eight, the cycle will repeat after CH0-CH7 data.

I²S Tx FIFO: The I²S Tx block has a hardware FIFO of depth 256 elements where each element is 32-bit wide. In addition to this 256-element FIFO, the I²S block has an internal transmit buffer that can store four 32-bit data to be transmitted. This four-element buffer is used as an intermediary to hold data to be transferred on the I²S bus, and is not exposed to the AHB BUS interface.

The TX FIFO can be paused by setting the TX_PAUSE bit in I2S_CMD. When the TX_PAUSE bit is set, the data sent over I²S is "0", instead of TX FIFO data. To resume normal operation, the TX_PAUSE bit must be cleared.

The I2S_TX_FIFO_CTL register is used for FIFO control operations. The TRIGGER_LEVEL bits in the I2S_TX_FIFO_CTL register can be used to generate a transmit trigger event when the Tx FIFO has less entries than the value configured in the TRIGGER_LEVEL bits.

The FIFO freeze operation can be enabled by setting the FREEZE bit in the I2S_TX_FIFO_CTL register. When the FREEZE bit is set and the Tx block is operational (TX_START bit in I2S_CMD is set), hardware reads from the Tx FIFO do not remove the FIFO entries. Also, the Tx FIFO read pointer will not be advanced. Any writes to the I2S_TX_FIFO register will increment the Tx FIFO write pointer; when the Tx FIFO becomes full, the internal write pointer stops incrementing. The freeze operation may be used for firmware debug purposes. This operation is not intended for normal operation. To return to normal operation after using the freeze operation, the I²S must be reset by clearing the TX_ENABLED bit in the I2S_CTL register, and then setting the bit again.

The CLEAR bit in the I2S_TX_FIFO_CTL register is used to clear the Tx FIFO by resetting the read/write pointers associated with the FIFO. Write accesses to the Tx FIFO using the I2S_TX_FIFO_WR or I2S_TX_FIFO_WR_SILENT registers are not allowed while the CLEAR bit is set.

The I2S_TX_FIFO_STATUS register provides FIFO status information. This includes number of used entries in the Tx FIFO and the current values of the Tx FIFO read/write

pointers. This register can be used for debug purposes. The I²S Tx FIFO read pointer is updated whenever the data is transferred from the Tx FIFO to the internal transmit buffer. Tx FIFO write pointer is updated whenever the data is written to the I2S_TX_FIFO_WR register, either through the CPU or the DMA controller.

For Tx FIFO data writes using the CPU, the hardware can be used to trigger an interrupt event for any of the FIFO conditions such as TX_TRIGGER, TX_NOT_FULL, and TX_EMPTY. As part of the interrupt handler, the CPU can write to the I2S_TX_FIFO_WR register. The recommended method is to write (256 - TRIGGER_LEVEL) words to the I2S_TX_FIFO_WR register every time the TX_TRIGGER interrupt event is triggered. In addition, interrupt events can be generated for FIFO overflow/underflow conditions.

For DMA-based Tx data transfers, the I²S Tx DMA trigger signal (tr_i2s_tx_req) can be enabled by writing '1' to the TX_REQ_EN bit in I2S_TR_CTL register. The trigger signal output will become high whenever the Tx FIFO has less entries than that configured in the TRIGGER_LEVEL field. The DMA channel can be configured to transfer up to (256 - TRIGGER_LEVEL) words from the applicable source address (such as Flash and SRAM regions). The destination address of the DMA should always be the I2S_TX_FIFO_WR register address, with the destination address increment feature disabled in the DMA channel configuration. This FIFO address increment logic is handled internally to adjust the write pointer, and the DMA should not increment the destination address. For more details on DMA channel configuration, refer to the [DMA Controller \(DW\) chapter on page 93](#).

The data in the I2S_TX_FIFO is always right-aligned. The I2S_TX_FIFO_WR format for different word length configurations is provided in [Figure 30-7](#).

Figure 30-7. I2S_TX_FIFO_WR Register Format for Different Word Lengths

	write data format of I2S_TX_FIFO																															
	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 24-bit mode									MSb																LSb							
	fixed "0"								23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 20-bit mode									MSb																LSb							
	fixed "0"												19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 18-bit mode									MSb																LSb							
	fixed "0"														17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 16-bit mode									MSb																LSb							
	fixed "0"																15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

I²S Rx FIFO: The I²S Rx block has a hardware FIFO of depth 256 elements where each element is 32-bit wide. In addition to this 256-element FIFO, the I²S block has an internal receive buffer that can store four 32-bit data to be received. This four-element buffer is used as an intermediary to hold data received on the I²S bus, and is not exposed to the AHB BUS interface.

The I2S_RX_FIFO_CTL register is used for FIFO control operations. The TRIGGER_LEVEL bits in the I2S_RX_FIFO_CTL register is used to generate a receive trigger event when the Rx FIFO has more entries than the value configured in the TRIGGER_LEVEL bits. In the standard I²S/LJ format, the TRIGGER_LEVEL bits can be configured up to the allowed maximum value of 253. In the TDM format, the maximum value of TRIGGER_LEVEL is (254-CH_NR) where CH_NR is the number of active channels in the TDM frame.

The FIFO freeze operation can be enabled by setting the FREEZE bit in the I2S_RX_FIFO_CTL register. When the FREEZE bit is set and the Rx block is operational (RX_START bit in the I2S_CMD register is set), hardware will not write to the Rx FIFO. Also, the Rx FIFO write pointer will not be advanced. Any reads from the I2S_RX_FIFO register will increment the Rx FIFO read pointer; when the Rx FIFO becomes empty, the internal read pointer stops incrementing. The freeze operation may be used for firmware debug purposes. This operation is not intended for normal operation. To return to normal operation after using the freeze operation, the I²S must be reset by clearing the RX_ENABLED bit in the I2S_CTL register and then setting the bit again.

The CLEAR bit in I2S_RX_FIFO_CTL register is used to clear the Rx FIFO by resetting the read/write pointers associated with the FIFO. Read accesses from the Rx FIFO using the I2S_RX_FIFO_RD or I2S_RX_FIFO_RD_SILENT registers are not allowed while the CLEAR bit is set.

The I2S_RX_FIFO_STATUS register provides FIFO status information. This includes number of used entries in the Rx FIFO and the current values of the Rx FIFO read/write pointers. This register can be used for debug purposes. The I²S Rx FIFO write pointer is updated whenever the data is transferred to the Rx FIFO from the internal receive buffer. Rx FIFO read pointer is updated whenever the data is read

from the I2S_RX_FIFO_RD register, either through the CPU or the DMA controller. For debug purposes, the I2S_RX_FIFO_RD_SILENT register is available, which always returns the top element of the Rx FIFO without updating the read pointer.

For Rx FIFO data reads using the CPU, the hardware can be used to trigger an interrupt event for any of the FIFO conditions such as RX_TRIGGER, RX_NOT_EMPTY, and RX_FULL. As part of the interrupt handler, the CPU can read from the I2S_RX_FIFO_RD register. The recommended method is to read (TRIGGER_LEVEL + 1) words from the I2S_RX_FIFO_RD register every time the RX_TRIGGER interrupt event is triggered. In addition, interrupt events can be generated for FIFO overflow/underflow conditions.

For DMA-based Rx data transfers, the I²S Rx DMA trigger signal (tr_i2s_rx_req) can be enabled by writing '1' to the RX_REQ_EN bit in the I2S_TR_CTL register. The trigger signal output will become high whenever the Rx FIFO has more entries than that configured in the TRIGGER_LEVEL field. The DMA channel can be configured to transfer up to (TRIGGER_LEVEL + 1) words to the applicable destination address (such as SRAM regions). The source address of the DMA should always be the I2S_RX_FIFO_RD register address, with the source address increment feature disabled in the DMA channel configuration. This FIFO address increment logic is handled internally to adjust the read pointer, and the DMA should not increment the source address. For more details on DMA channel configuration, refer to the [DMA Controller \(DW\) chapter on page 93](#).

The data in the I2S_RX_FIFO is always right aligned. The I2S_RX_FIFO_RD format for different word length configurations is provided in [Figure 30-8](#). Note that the unused most significant bits are either set as '0' or sign-bit extended depending on the BIT_EXTENSION bit in the I2S_RX_CTL register.


Figure 30-8. I2S_RX_FIFO_RD Register Format for Different Word Lengths

		read data format of I2S_RX_FIFO																															
		31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 24-bit mode not Bit extension										MSb																LSb							
		fixed "0"								23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
	Bit extension	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 20-bit mode not Bit extension										MSb																LSb							
		fixed "0"												19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
	Bit extension	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 18-bit mode not Bit extension										MSb																LSb							
		fixed "0"														17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
	Bit extension	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Word Length = 16-bit mode not Bit extension																		MSb								LSb							
		fixed "0"																15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
	Bit extension	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

30.8 Interrupt Support

The I²S block has one interrupt output signal that goes to the interrupt controller in the CPU. Refer to the [Interrupts chapter on page 54](#) for details on the vector number of the I²S interrupt and the procedure to configure the interrupt priority, vector address, and enabling/disabling. The I²S interrupt can be triggered under any of the following events - TX_TRIGGER, TX_NOT_FULL, TX_EMPTY, TX_OVERFLOW, TX_UNDERFLOW, TX_WD, RX_TRIGGER, RX_NOT_EMPTY, RX_FULL, RX_OVERFLOW, RX_UNDERFLOW, or RX_WD. Each of the interrupt events can be individually enabled/disabled to generate an interrupt condition. The I2S_INTR_MASK register is used to enable the required events by writing '1' to the corresponding bit. Irrespective of the INTR_MASK settings, if any of the events occur, the corresponding event status bit will be set by the hardware in the I2S_INTR register. The I2S_INTR_MASKED register is the bitwise AND of the I2S_INTR_MASK and I2S_INTR registers. The final I²S interrupt signal is the logical OR of all the bits in the I2S_INTR_MASKED register. So only those events that are enabled in the I2S_INTR_MASK register are propagated as interrupt events to the interrupt controller. Interrupts can also be triggered in software by writing to the corresponding bits in I2S_INTR_SET register. [Figure 30-9](#) illustrates the interrupt signal generation.

Figure 30-9. Interrupt Signal Generation




In the interrupt service routine (ISR), the I2S_INTR_MASKED register should be read to know the events that triggered the interrupt event. Multiple events can trigger the interrupt because the final interrupt signal is the logical OR output of the events. The ISR should do the tasks corresponding to each interrupt event that was triggered. At the end of the ISR, the value read in the I2S_INTR_MASKED register earlier should be written to the I2S_INTR register to clear the bits whose interrupt events were processed in the ISR. Unless the bits are not cleared by writing '1' to the I2S_INTR register, the interrupt signal will always be high. A dummy read of the I2S_INTR register should be done for the earlier register write to take effect.

30.9 Watchdog Timer

The Tx and Rx blocks have independent watchdog timers, which can be used to generate an interrupt event if the Word Select (WS) input is idle for more than the configured time period. This feature is available only in the slave mode of operation where the external master drives the WS input lines (tx_ws, rx_ws). This feature can be used to detect any signal transmission issues, master device issues, or if the master has halted communication. If the master drives the same word select signal to both the tx_ws and rx_ws lines, then only one of the watchdog timers can be enabled to cause the interrupt event. Although the following explanation covers Tx watchdog, the same explanation applies to Rx watchdog as well.

To enable the Tx watchdog timer feature, WD_EN bit in the I2S_TX_CTL register should be set. The watchdog timer reload value (32-bit timer) is configured by writing to the I2S_TX_WATCHDOG register. A value of zero written to the I2S_TX_WATCHDOG register will also disable the watchdog timer. Figure 30-10 illustrates the watchdog behavior when the timer is enabled. The timer runs off the CLK_PERI system clock. Refer to the [Clocking System chapter on page 207](#) for details on generation of CLK_PERI. The timer starts running when WD_EN and TX_START bits are set. The timer reload happens either on a rising edge event on tx_ws input signal, or when the timer values reaches zero. When the timer value reaches zero, the TX_WD interrupt event is generated. The TX_WD bit in the I2S_INTR_MASK register should be set to enable interrupt generation by the watchdog timer interrupt event. The interrupt event can be cleared by writing '1' to the TX_WD bit in the I2S_INTR register.

Figure 30-10. Watchdog Timer Working



31. PDM-PCM Converter



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation
- Application notes
- Code examples

The PDM-PCM unit accepts a stereo or mono serial data stream (pulse modulated 1-bit stream) coming from external digital PDM microphones. The PDM-PCM converter consists of a fifth order cascaded integrator comb (CIC) filter followed by a decimator, and a final stage high-pass filter. This block simplifies the conversion process by exposing the different configuration settings as registers, which you can program to meet the application needs. The entire PDM-PCM conversion process is handled in hardware; the PCM output data streaming can be done using the DMA controller thus freeing up the CPU bandwidth from performing periodic audio streaming activities.

31.1 Features

- Supports Mono/Stereo mode pulse density modulation (PDM) to pulse code modulation (PCM) conversion
- Accepts 1-bit PDM input and can generate 16-, 18-, 20-, or 24-bit PCM digital data output
- Configurable PDM microphone clock frequency
- Ability to generate standard audio sampling rates by adjusting the decimation rate and clock divider values
- Digital volume control: Programmable gain amplifier (PGA) control from –12 dB to +10.5 dB in 1.5-dB steps
- Smooth PGA and soft-mute control
- Hardware receive buffer: 24-bit wide, 255-element FIFO with support for DMA controller-based data transfer
- Optional high-pass filter to remove DC and low-frequency noise

31.2 Architecture

Figure 31-1. Block Diagram

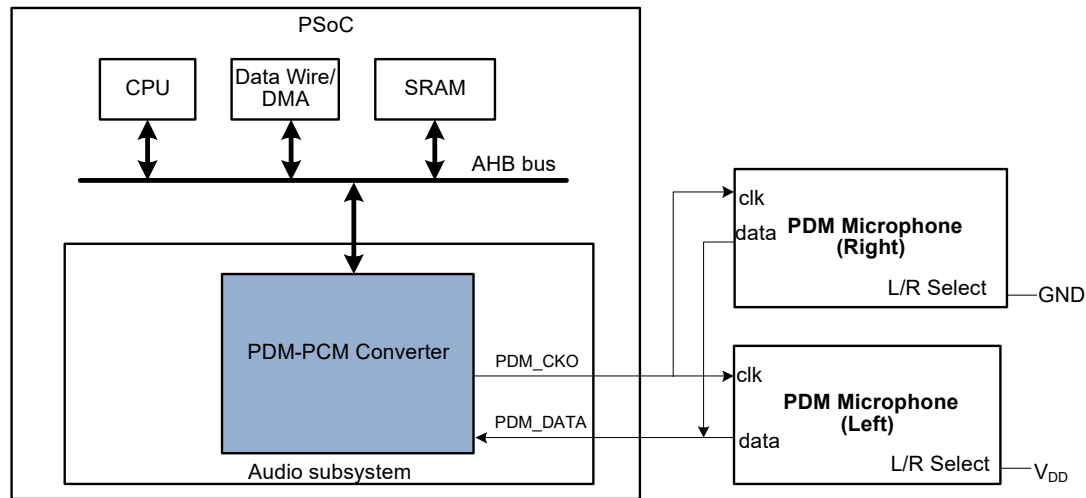


Figure 31-1 shows the block diagram of the PDM-PCM converter. Refer to the [device datasheet](#) for information on port pin assignments of the PDM block signals, electrical specifications, and the interrupt vector number.

31.2.1 Enable/Disable Converter

The PDM-PCM converter can be powered ON or OFF by using the ENABLED bit in the PDM_CTL register. The block can be turned OFF when not used to save power. When the block is powered off by writing '0' to the ENABLED bit, the non-retention registers lose their current values. Refer to the [registers TRM](#) to know which registers are retention and non-retention type. When the block is enabled again, the non-retention registers are reset to their default values. User

firmware should ensure that all non-retention registers are configured as required after the ENABLED bit is set, and before the PDM-PCM conversion is started by setting the STREAM_EN bit in the PDM_CMD register. See [Operating Procedure on page 413](#) for the recommended procedure to configure this PDM-PCM converter before starting the PDM-PCM conversion.

31.2.2 Clocking Features

The block uses the CLK_HF[1] root high-frequency clock for performing the PDM-PCM conversion process. Refer to the [Clocking System chapter on page 207](#) for details on selecting the clock source for CLK_HF[1], and configuring the divider registers to generate CLK_HF[1].

Figure 31-2. PDM-PCM Clocking Dividers




Figure 31-2 shows the clock divider structure in the block. The block has three stages of clock dividers to generate the clock (PDM_CKO), which goes to the external PDM microphone clock input.

1. The first stage clock divider (CLK_CLOCK_DIV field in the PDM_CLOCK_CTL register) is used to generate the actual clock signal (PDM_CLK) that goes to the PDM-PCM converter. The input is CLK_HF[1]; the CLK_CLOCK_DIV can be a value between 0 and 3 (divider value between 1 and 4).

$$f_{\text{PDM_CLK}} = f_{\text{CLK_HF[1]}} / (\text{CLK_CLOCK_DIV} + 1)$$

2. The second stage clock divider (MCLKQ_CLOCK_DIV field in the PDM_CLOCK_CTL register) is used to generate the internal master clock and can take a value between 0 and 3 (divider value between 1 and 4). The input clock is pdm_clk and the output of the divider is MCLKQ.

$$f_{\text{MCLKQ}} = f_{\text{PDM_CLK}} / (\text{MCLKQ_CLOCK_DIV} + 1)$$

- The third stage clock divider (CKO_CLOCK_DIV field in the PDM_CLOCK_CTL register) is used to generate the clock that goes to the PDM microphone clock (PDM_CKO) through the output pin of the PSoC 6 MCU. The input clock is MCLKQ. The divider register can be between 1 and 15 (divider value between 2 and 16).

$$f_{\text{PDM_CKO}} = f_{\text{MCLKQ}} / (\text{CKO_CLOCK_DIV} + 1), \text{ if CKO_CLOCK_DIV} \geq 1$$

$$f_{\text{PDM_CKO}} = f_{\text{MCLKQ}} / 2, \text{ if CKO_CLOCK_DIV} = 0$$

31.2.3 Over-Sampling Ratio

Over-sampling ratio (OSR) is the ratio between the frequency of the PDM microphone clock ($f_{\text{PDM_CKO}}$) and the output PCM sampling rate frequency (f_s). The OSR is determined by the SINC_RATE bits in the PDM_CLOCK_CTL register. The relation is as follows.

$$\text{OSR} = f_{\text{PDM_CKO}} / f_s = 2 \times \text{SINC_RATE}$$

Table 31-1 gives an example of the PDM clock divider and SINC_RATE register configurations to generate the PCM output at standard sampling rates. Note that the PDM divider values in the table are the register field values – the actual divider values are one more than the configured register values as explained in the clock divider section. Refer to the [device datasheet](#) for details on maximum values of PDM_CKO frequency, PDM_CLK frequency, and the output sampling rates.

Table 31-1. PDM Clock Divider Values for Standard Audio Sampling Rates

Sampling Rate (SR) (kHz)	SINC_RATE (= OSR/2)	PDM_CKO (=2*SINC_RATE*SR) (MHz)	CLK_HF1 (MHz)	(CLK_HF1)/ (PDM_CKO) (Total Divider Ratio)	CLK_CLOCK_DIV (First Divider)	MCLKQ_CLOCK_DIV (Second Divider)	CKO_CLOCK_DIV (Third Divider)
8	32	0.512	49.152	96	1	3	11
16	32	1.024		48	1	3	5
32	32	2.048		24	0	3	5
48	32	3.072		16	0	1	7
44.1	32	2.8224	45.1584	32	0	1	7

There are various methods to generate the CLK_HF[1] frequencies listed in the table depending on the clocking options available on the device. Refer to the [Clocking System chapter on page 207](#) for details on the clocking options available in the device, including the clock sources and PLL/FLL circuitry. For example, an external crystal oscillator (ECO) can be used in conjunction with a phase-locked loop (PLL) to generate the CLK_HF[1] at the desired frequency of 49.152 MHz or 45.1584 MHz.

One possible combination of PLL divider values to generate the 49.152 MHz frequency from a 17.2032 MHz ECO are: REFERENCE_DIV = 7, FEEDBACK_DIV = 100, OUTPUT_DIV = 5. One possible combination of PLL divider values to generate the 45.1584 MHz frequency from a 17.2032 MHz ECO are: REFERENCE_DIV = 8, FEEDBACK_DIV = 105, OUTPUT_DIV = 5.

in Table 31-2. The table also lists the invalid register settings, which you must not use in the firmware.

31.2.4 Mono/Stereo Microphone Support

The PDM-PCM converter supports mono-left, mono-right, stereo, and swapped stereo modes of operation. The operation mode is controlled by the PDM_CH_SET and SWAP_LR bits in the PDM_MODE_CTL register. The register settings for the different operation modes are given


Table 31-2. Operation Mode Register Settings

Register Setting	Operation Mode
PCM_CH_SET = 0, SWAP_LR = 0 or 1	Recording OFF
PCM_CH_SET = 1, SWAP_LR = 0	Mono-Left recording mode. Only the left microphone channel is sampled on the rising edge of PDM_CKO. FIFO buffer contains only left channel data.
PCM_CH_SET = 2, SWAP_LR = 1	Mono-Right recording mode. Only the right microphone channel is sampled on the falling edge of PDM_CKO. FIFO buffer contains only right channel data.
PCM_CH_SET = 3, SWAP_LR = 0	Stereo recording mode. The right microphone channel is sampled on the falling edge of PDM_CKO and left channel on rising edge. FIFO buffer contains data in L/R format (left channel followed by right channel)
PCM_CH_SET = 3, SWAP_LR = 1	Swapped Stereo recording mode. The right microphone channel is sampled on the rising edge of PDM_CKO and left channel on falling edge. FIFO buffer contains data in L/R format (left channel followed by right channel)
PCM_CH_SET = 1, SWAP_LR = 1 or PCM_CH_SET = 2, SWAP_LR = 0	Invalid setting (not supported). Do not operate the PDM-PCM converter in these configuration settings.


Figure 31-3 shows the timing diagrams for the different operating modes.

Figure 31-3. PDM Mono/Stereo Timing


Mono left mode (PCM_CH_SET = 1, SWAP_LR = 0)




Mono right mode (PCM_CH_SET = 2, SWAP_LR = 1)



Stereo mode (PCM_CH_SET = 3, SWAP_LR = 0)



Swapped Stereo mode (PCM_CH_SET = 3, SWAP_LR = 1)



To alleviate uncertain board delay impact on PDM_IN setup and hold timing constraints, the PDM-PCM provides CKO_DELAY bits in the PDM_MODE_CTL register to add extra delay for the PDM_CKO path to internal sampler. CKO_DELAY can be a value between 0 and 7. A value of '0' implies that internal sampling of PDM data is advanced by three PDM_CLK clock cycles for the PDM_CKO transition. A value of '7' implies that internal sampling of PDM data is delayed by four PDM_CLK clock cycles for the PDM_CKO transition. Refer to the [registers TRM](#) for details on the meaning of different CKO_DELAY values.

Note: Variations have been observed in the recommendation for left/right sampling logic among the

different PDM microphone manufacturers. The SWAP_LR bit in the PDM-PCM converter ensures that you can adjust the sampling logic according to the microphone datasheet recommendations. Refer to the PDM microphone manufacturer datasheet for the exact timing details. Also, in stereo mode, use the same manufacturer for both the left/right PDM microphones to ensure the timing behavior is uniform for both channels.

31.2.5 Hardware FIFO Buffers and DMA Controller Support

The PDM-PCM converter has a hardware FIFO depth of 255 elements where each element is 24-bit wide.

The PDM_RX_FIFO_CTL register is used for FIFO control operations. Refer to the register description in the [registers TRM](#) for more details. The TRIGGER_LEVEL field in the PDM_RX_FIFO_CTL register is used to generate a receive trigger event (interrupt event, DMA trigger signal) when the Rx FIFO has more entries than the value configured in the TRIGGER_LEVEL field. The TRIGGER_LEVEL field can be configured up to 254 in the mono microphone recording mode and up to 253 in the stereo microphone recording mode.

The FIFO freeze operation can be enabled by setting the FREEZE bit in the PDM_RX_FIFO_CTL register. When the FREEZE bit is set and the Rx block is operational (STREAM_EN bit in the PDM_CMD register is set), hardware will not write to the Rx FIFO. Also, the Rx FIFO write pointer will not be advanced. Any reads from the PDM_RX_FIFO_RD register will increment the Rx FIFO read pointer; when the Rx FIFO becomes empty, the internal read pointer stops incrementing. The freeze operation may be used for firmware debug purposes. This operation is not intended for normal operation. To return to normal operation after using the freeze operation, the PDM-PCM must be reset by clearing the ENABLED bit in PDM_CTL register, and then setting the bit again.

The CLEAR bit in the PDM_RX_FIFO_CTL register is used to clear the Rx FIFO by resetting the read/write pointers associated with the FIFO. Read accesses from the Rx FIFO using PDM_RX_FIFO_RD or PDM_RX_FIFO_RD_SILENT registers are not allowed while the CLEAR bit is set.

The PDM_RX_FIFO_STATUS register provides FIFO status information. This includes number of used entries in the Rx FIFO and the current values of the Rx FIFO read/write pointers. This register can be used for debug purposes. The Rx FIFO write pointer is updated whenever the data is transferred to the Rx FIFO from the internal receive buffer. Rx FIFO read pointer is updated whenever the data is read from the PDM_RX_FIFO_RD register, either through the CPU or the DMA controller. For debug purposes, the PDM_RX_FIFO_RD_SILENT register is available, which always returns the top element of the Rx FIFO without updating the read pointer.

For Rx FIFO data reads using the CPU, the hardware can be used to trigger an interrupt event for any of the FIFO conditions such as RX_TRIGGER and RX_NOT_EMPTY. As part of the interrupt handler, the CPU can read from the PDM_RX_FIFO_RD register. The recommended method is to read (TRIGGER_LEVEL + 1) words from the PDM_RX_FIFO_RD register every time the RX_TRIGGER interrupt event is triggered. In addition, interrupt events can be generated for FIFO overflow and underflow conditions.

For DMA-based data transfers, the DMA trigger signal (tr_pdm_rx_req) can be enabled by writing '1' to the RX_REQ_EN bit in the PDM_TR_CTL register. The trigger signal output will become high whenever the Rx FIFO has more entries than that configured in the TRIGGER_LEVEL field. Refer to the [Trigger Multiplexer Block chapter on page 260](#) for details on how to connect the DMA trigger signal to a particular DMA channel. The DMA channel can be configured to transfer up to (TRIGGER_LEVEL + 1) words to the applicable destination address (such as SRAM regions). The source address of the DMA should always be the PDM_RX_FIFO_RD register address, with the source address increment feature disabled in the DMA channel configuration. This FIFO address increment logic is handled internally to adjust the read pointer, and the DMA should not increment the source address. For more details on DMA channel configuration, refer to the [DMA Controller \(DW\) chapter on page 93](#).

The successive data read from the PDM_RX_FIFO_RD follows the Left 1/Right 1/Left 2/Right 2/... format in stereo and swapped stereo modes of operation. For mono left and mono right recording modes, the data read from FIFO contains either the left channel data (mono left mode) or the right channel data (mono right mode).

The data in the PDM_RX_FIFO_RD is always right-aligned. The PDM_TX_FIFO_RD format for different word length configurations is provided in [Figure 31-4](#). Note that the unused most significant bits are either set as '0' or sign-bit extended depending on the BIT_EXTENSION field setting in the PDM_DATA_CTL register.

Figure 31-4. FIFO Register Structure

		read data format of PDM_RX_FIFO																																							
		31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0								
WORD_LEN = 24-bit mode												MSb																				LSb									
	BIT_EXTENSION = 0	fixed "0"										23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0						
	BIT_EXTENSION = 1 (Sign Bit Extension)	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0							
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0							
WORD_LEN = 20-bit mode												MSb																				LSb									
	BIT_EXTENSION = 0	fixed "0"																				19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
	BIT_EXTENSION = 1 (Sign Bit Extension)	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	"1"	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0							
		"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	"0"	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0						

31.2.6 Interrupt Support

The block has one output signal (interrupt_pdm) that goes to the interrupt controller in the CPU. Refer to the [device datasheet](#) for details on the vector number of the PDM-PCM interrupt. Refer to the [Interrupts chapter on page 54](#) for the procedure to configure the interrupt priority, vector address, and enabling/disabling.

The PDM interrupt can be triggered under any of the following events – RX_TRIGGER, RX_NOT_EMPTY, RX_OVERFLOW, or RX_UNDERFLOW. [Table 31-3](#) lists the trigger conditions and details of these events.

Table 31-3. Interrupt Event Trigger Condition

Interrupt Event	Trigger Condition
RX_TRIGGER	The PDM Rx FIFO has more entries than the value specified in the TRIGGER_LEVEL field in the PDM_RX_FIFO_CTL register. At least (TRIGGER_LEVEL + 1) words can be read from the PDM_RX_FIFO_RD register in the interrupt service routine.
RX_NOT_EMPTY	The PDM Rx FIFO has at least one word that can be read from the PDM_RX_FIFO_RD register.
RX_OVERFLOW	The PDM Rx FIFO content is overwritten by the PDM-PCM converter due to the number of unread words exceeding the 256 word buffer capacity. The output PCM data after the overflow condition is discarded and not transferred to the FIFO. This event can be used to detect bandwidth constraints in the end application due to DMA or the interrupt used for data transfer not getting the required priority for executing the data transfers.
RX_UNDERFLOW	Attempt to read from an empty PDM Rx FIFO. This can happen due to incorrect configuration of the DMA or the interrupt service routine code used to do the data transfer.


Each of the interrupt events can be individually enabled or disabled to generate the interrupt condition. The PDM_INTR_MASK register is used to enable the required events by writing '1' to the corresponding bit.

Irrespective of the INTR_MASK settings, if any event occurs, the corresponding event status bit will be set by the hardware in the PDM_INTR register. The PDM_INTR_MASKED register is the bitwise AND of the PDM_INTR_MASK and PDM_INTR registers. The final PDM interrupt signal is the logical OR of all the bits in the PDM_INTR_MASKED register. So only those events that are enabled in the PDM_INTR_MASK register are propagated as interrupt events to the interrupt controller.

Interrupt events can also be triggered in software by writing to the corresponding bits in PDM_INTR_SET register.

[Figure 31-5](#) illustrates the interrupt signal generation from the PDM-PCM block as explained above. Only the RX_TRIGGER interrupt generation is highlighted in the figure; the remaining interrupt events also follow the same generation logic.

Figure 31-5. PDM Interrupt Signal Generation



In the interrupt service routine (ISR) corresponding to the interrupt vector number of `interrupt_pdm`, the `PDM_INTR_MASKED` register should be read to know the events that triggered the interrupt event. Multiple events can trigger the interrupt because the final interrupt signal is the logical OR output of the events. The ISR should do the tasks corresponding to each interrupt event that was triggered. At the end of the ISR, the value read in the `PDM_INTR_MASKED` register earlier should be written to the `PDM_INTR` register to clear the bits whose interrupt events were processed in the ISR. A dummy read of the `PDM_INTR` register should be done for the earlier register write to `PDM_INTR` to take effect.

All of the interrupt event bits in `PDM_INTR` register will continue to indicate the event condition regardless of the true state until that bit is cleared (for example, when set, the `RX_OVERFLOW` bit will continue to indicate an `OVERFLOW` state until the `RX_OVERFLOW` bit is cleared regardless of the true state of the FIFO. A mere FIFO read of the FIFO will not clear the `RX_OVERFLOW` bit.). Unless the `PDM_INTR` bits that are used to generate the interrupt are not cleared by writing '1' to the `PDM_INTR` bits, the interrupt signal will always be high.

31.2.7 Digital Volume Gain

The PDM-PCM converter supports independent digital volume control on the left/right channels with a range from -12.5 dB to +10.5 dB in steps of 1.5 dB. It is programmed by configuring the `PGA_R` and `PGA_L` bits in the `PDM_CTL` register. PGA gain may be changed on the fly during normal operation, or as a one-time setting before starting the PDM-PCM conversion process.

31.2.8 Smooth Gain Transition

To reduce zipper or clip noise during on-the-fly gain transition or during soft mute operation, a built-in volume smoother is implemented with fine gain steps and fine time steps that enable soft ramp up or ramp down of the volume levels. Two fine gain options of 0.13 dB and 0.26 dB step sizes are available. The fine gain is set by the `STEP_SEL` bit in the `PDM_CTL` register. In addition to the fine gain steps, a time step is available for the fine gain change in terms of the number of sample cycles. The time step can be

configured to a value between 64 sample periods and 512 sample periods using the `S_CYCLES` bits in the `PDM_MODE_CTL` register. So the `STEP_SEL` and `S_CYCLES` bit settings together determine the rate at which PGA gain or the soft mute transitions take effect.

31.2.9 Soft Mute

The PDM-PCM contains a built-in software-controlled mute function that digitally attenuates signals to imperceptible levels or zero. When mute function is enabled by setting the `SOFT_MUTE` bit in the `PDM_CTL` register, the corresponding PCM output is decreased from current level to mute state through predefined granular gain step per time constant transition. The `STEP_SEL` bit setting determines the gain step and the `S_CYCLES` bits determine the time constant. During soft-mute, the block is still ON and the PCM data streaming is operational; the DMA or CPU-based data transfer also happens as usual. Only the PCM output level is muted. When mute function is disabled by setting `SOFT_MUTE` = 0, the mute function is OFF and the PDM-PCM returns to normal operation where output signal level goes up to normal with current PGA gain.

31.2.10 Word Length and Sign Bit Extension

The PCM output word length can be configured for either 16-bits, 18-bits, 20-bits, or 24-bits using `WORD_LEN` bits in the `PDM_DATA_CTL` register. Irrespective of the word length setting, the PCM output is always read from the FIFO data buffer register (`PDM_RX_FIFO_RD`) as a 32-bit value. The unused most significant bits in the 32-bit value can either be sign extended or extended by '0' by using the `BIT_EXTENSION` bit in the `PDM_DATA_CTL` register.

31.2.11 High-Pass Filter

The PDM-PCM converter has a final stage high-pass filter (HPF) that blocks DC offset and low-frequency noise in signal band. The HPF is enabled when the `HPF_EN_N` bit in the `PDM_MODE_CTL` register is zero, and disabled when the `HPF_EN_N` bit is 1.

The filter response for HPF is characterized as:

$$H(z) = \frac{1 - z^{-1}}{1 - (1 - 2^{-HPF_GAIN})z^{-1}}$$

The HPF operates at the final PCM output sampling rate. HPF_GAIN is a 4-bit gain configuration setting in the PDM_MODE_CTL register. In default mode, HPF_GAIN = 0xB, so the HPF can be formulated by polynomial:

$$H(z) = \frac{1 - z^{-1}}{1 - 0.99951z^{-1}}$$

The HPF_GAIN setting can be tuned to adjust HPF cutoff corner frequency for better system configuration.

31.2.12 Enable/Disable Streaming

The PDM-PCM conversion process can be dynamically enabled/disabled by using the STREAM_EN bit in the PDM_CMD register.

31.2.13 Power Modes

The PDM-PCM can operate in Active and Sleep CPU modes while in LP or ULP system power modes. It is not operational in system Deep Sleep or Hibernate power modes. When the device transitions from Deep Sleep/ Hibernate power modes to the LP/ULP power modes, the non-retention registers lose their previous configuration values. So the non-retention registers must be appropriately configured before enabling the PDM-PCM again for LP/ULP mode operation. One option is to store the non-retention register values in SRAM before entering Deep Sleep/ Hibernate modes. When returning to the LP/ULP modes, the SRAM values can be copied to the registers after enabling the PDM-PCM by setting the ENABLED bit in the PDM_CTL register. Refer to the [registers TRM](#) to identify the non-retention registers for the PDM.

31.3 Operating Procedure

31.3.1 Initial Configuration

The sequence of steps for initial configuration of the PDM-PCM converter before starting the conversion process is as follows:

1. Configure the clock dividers and decimation rate in the PDM_CLOCK_CTL register. This register configuration should be done before enabling the PDM-PCM converter. If the ENABLED bit in the PDM_CTL register is set, it should be cleared before changing the clock configuration.
2. Enable the block; set the PGA gain and fine gain step setting as required by writing to the PDM_CTL register.
3. Configure the PDM_MODE_CTL and PDM_DATA_CTL registers as required.

4. Configure the Rx FIFO trigger level setting by writing to the TRIGGER_LEVEL bits in the PDM_RX_FIFO_CTL register. The CLEAR and FREEZE bits in PDM_RX_FIFO_CTL are not set for normal operation.
5. Configure the events that must generate the interrupt by setting the corresponding bits in the PDM_INTR_MASK register, and clearing the remaining bits.
6. Configure the interrupt PDM interrupt vector and enable the interrupt vector. See the [Interrupts chapter on page 54](#) for details.
7. If a DMA-based data transfer is required, connect the PDM DMA trigger signal (tr_pdm_rx_req) to the trigger input of the required DMA channel. See the [Trigger Multiplexer Block chapter on page 260](#) for details on how to connect to the DMA channel trigger input. Configure the DMA channel as required - the source address of the DMA descriptor is PDM_RX_FIFO_RD register with the source address increment feature disabled and the source data length is word type (32-bits). The DMA channel can be used to transfer (TRIGGER_LEVEL + 1) words from the PDM_RX_FIFO_RD register whenever the trigger signal becomes high. The destination address configuration depends on the application requirements. See the [DMA Controller \(DW\) chapter on page 93](#) for details on DMA channel configuration.
8. Enable the DMA trigger signal generation by setting the RX_REQ_EN bit in the PDM_TR_CTL register.

31.3.2 Interrupt Service Routine (ISR) Configuration

The code for the PDM interrupt service routine should have the following flow:

1. The events that triggered the interrupt can be found by reading the PDM_INTR_MASKED register in the ISR. All the bits that are set causes the interrupt event. The register value should also be in a variable "var".
2. For each of the event bits that are set in PDM_INTR_MASKED, appropriate application level tasks can be executed. For example, the RX_TRIGGER event can be used for CPU-based data transfers if a DMA-based data transfer is not used. DMA transfers should use the tr_pdm_rx_req trigger signal (by setting the RX_REQ_EN bit in the PDM_TR_CTL register). The DMA trigger should not use the RX_TRIGGER interrupt event to reduce CPU usage for data transfer. The RX_OVERFLOW event can be used to take appropriate counter measures such as giving higher priority to PDM-PCM DMA channel. The RX_UNDERFLOW event typically indicates wrong data transfer logic in the application – either in the CPU-based data transfer code or in the DMA channel configuration used to transfer data.
3. After the event conditions have been processed, the "var" value read from PDM_INTR_MASKED should be written to the PDM_INTR register to clear the events that are set in the register. Due to the buffered write logic, the

PDM_INTR register should also be read after the write process to ensure the write process is completed in the slower peripheral clock domain.

31.3.3 Enabling / Disabling Streaming

The PDM-PCM conversion process starts after the STREAM_EN bit is set in the PDM_CMD register. Depending on the application needs, the streaming can be dynamically started and stopped using the STREAM_EN bit. Clear the Rx FIFO before starting the streaming process to reset the read/write pointers and FIFO state. The procedure to clear the FIFO is to write a '1' to the CLEAR bit in PDM_RX_FIFO_CTL followed by writing a '0' to the CLEAR bit. When the CLEAR bit is set, all the data entries in the Rx FIFO are cleared by resetting the internal read/write pointers. Read accesses to the PDM_RX_FIFO_RD and PDM_RX_FIFO_RD_SILENT registers are prohibited when the CLEAR bit is 1. Therefore, the CLEAR bit should be cleared before starting the streaming operation.

32. Universal Serial Bus (USB) Device Mode



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - USB
- Application notes
- Code examples

The PSoC 6 MCU USB block can act as a USB device that communicates with a USB host. The USB block is available as a fixed-function digital block in the PSoC 6 MCU. It supports full-speed communication (12 Mbps) and is designed to be compliant with the USB Specification Revision 2.0. USB devices can be designed for plug-and-play applications with the host and also support hot swapping. This chapter details the PSoC 6 MCU USB block and transfer modes. For details about the USB specification, see the USB Implementers Forum website.

32.1 Features


The USB in the PSoC 6 MCU has the following features:

- Complies with USB Specification 2.0
- Supports full-speed peripheral device operation with a signaling bit rate of 12 Mbps
- Supports eight data endpoints and one control endpoint
- Provides shared 512-byte buffer for data endpoints
- Provides dedicated 8-byte memory for control endpoint (EP0)
- Supports four types of transfers – bulk, interrupt, isochronous, and control
- Supports bus- and self-powered configurations
- Enables USB suspend mode for low power
- Supports three types of logical transfer modes:
 - No DMA mode (Mode 1)
 - Manual DMA mode (Mode 2)
 - Automatic DMA mode (Mode 3)
- Supports maximum packet size of 512 bytes using Mode 1 and Mode 2, and maximum packet size of 1023 bytes for isochronous transfer using Mode 3
- Provides integrated 22-Ω USB termination resistors on D+ and D– lines, and 1.5-kΩ pull-up resistor on the D+ line
- Supports USB 2.0 Link Power Management (LPM)

32.2 Architecture

Figure 32-1 illustrates the device architecture of the USB block in PSoC 6 MCUs. It consists of the USB Physical Layer (USB PHY), Serial Interface Engine (SIE), and the local 512-byte memory buffer.

Figure 32-1. USB Device Block Diagram



32.2.1 USB Physical Layer (USB PHY)

The USB PHY allows physical layer communication with the USB host through the D+, D-, and VBUS pins. It handles the differential mode communication with the host, VBUS detection, and monitoring events such as SE0 on the USB bus.

32.2.2 Serial Interface Engine (SIE)

The SIE handles the decoding and creating of data and control packets during transmit and receive. It decodes the USB bit streams into USB packets during receive, and creates USB bit streams during transmit. The following are the features of the SIE:

- Conforms to USB Specification 2.0
- Supports one device address
- Supports eight data endpoints and one control endpoint
- Supports interrupt trigger events for each endpoint
- Integrates an 8-byte buffer in the control endpoint

The registers for the SIE are mainly used to configure the data endpoint operations and the control endpoint data buffers. This block also controls the interrupt events available for each endpoint.

32.2.3 Arbiter

The Arbiter handles access of the SRAM memory by the endpoints. The SRAM memory can be accessed by the CPU, DMA, or SIE. The arbiter handles the arbitration between the CPU, DMA, and SIE. The arbiter consists of the following blocks:

- SIE Interface Module
- CPU/DMA Interface
- Memory Interface
- Arbiter Logic

The arbiter registers are used to handle the endpoint configurations, read address, and write address for the endpoints. It also configures the logical transfer type required for each endpoint.

32.2.3.1 SIE Interface Module

This module handles all the transactions with the SIE. The SIE reads data from the SRAM memory and transmits to the host. Similarly, it writes the data received from the host to the SRAM memory. These requests are registered in the SIE Interface module and are handled by this module.

32.2.3.2 CPU/DMA Interface Block

This module handles all transactions with the CPU and DMA. The CPU requests for reads and writes to the SRAM memory for each endpoint. These requests are registered in this interface and are handled by the block. When the DMA is configured, this interface is responsible for all transactions between the DMA and USB. The block supports the DMA request line for each data endpoint. The behavior of the DMA depends on the type of logical transfer mode configured in the configuration register.

32.2.3.3 Memory Interface

The memory interface is used to control the interface between the USB and SRAM memory unit. The maximum memory size supported is 512 bytes organized as 256×16 -bit memory unit. This is a dedicated memory for the USB. The memory access can be requested by the SIE or by the CPU/DMA. The SIE Interface block and the CPU/DMA Interface block handle these requests.

32.2.3.4 Arbiter Logic

This is the main block of the arbiter. It is responsible for arbitrations for all the transactions that happen in the arbiter. It arbitrates the CPU, DMA, and SIE access to the memory unit and the registers. This block also handles memory management, which is either 'Manual' or 'Automatic'. In Manual memory management, the read and write address manipulations are done by the firmware. In Automatic management, all the memory handling is done by the block itself. This block takes care of the buffer size allocation. It also handles common memory area. This block also handles the interrupt requests for each endpoint.

32.3 Operation

32.3.1 USB Clocking Scheme

The USB device block should be clocked at 48 MHz with an

accuracy of $\pm 0.25\%$. In the PSoC 6 MCU, CLK_HF3 is the clock source. The USB device block also requires a 100-kHz peripheral clock for USB bus reset timing. The required USB clock can be generated using one of the following clocking schemes:

- IMO (trimmed with USB) -> PLL -> CLK_HF3
- ECO (with the required accuracy) -> FLL -> CLK_HF3
- ECO (with the required accuracy) -> PLL -> CLK_HF3
- Use external clock (EXTCLK) with the required accuracy

32.3.2 USB PHY

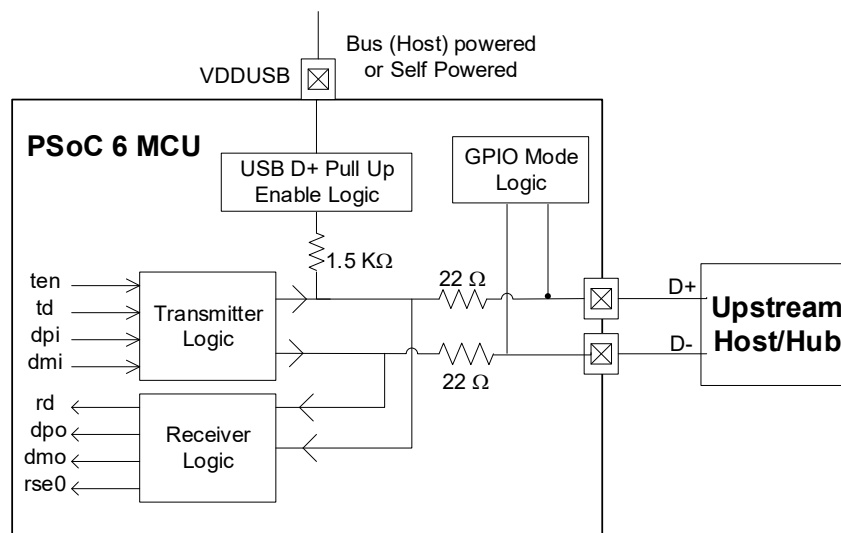
The USB includes the transmitter and receiver (transceiver), which corresponds to the USB PHY. Figure 32-2 shows the PHY architecture. The USB PHY also includes the pull-up resistor on the D+ line to identify the device as full-speed type to the host. The PHY integrates the $22\text{-}\Omega$ series termination resistors on the USB line. The signal between the USB device and the host is a differential signal. The receiver receives the differential signal from the host and converts it to a single-ended signal for processing by the SIE. The transmitter converts the single-ended signal from the SIE to the differential signal, and transmits it to the host. The differential signal is given to the upstream devices at a nominal voltage range of 0 V to 3.3 V.

32.3.2.1 Power Scheme

The USB PHY is powered by the VBUS power pad of the PSoC 6 MCU. The VBUS pad can be driven either by the host VBUS (bus-powered) or an external power supply (self-powered).

The USB PHY needs a nominal voltage of 3.3 V for its communication with the host.

Figure 32-2. USB PHY Architecture



32.3.2.2 VBUS Detection

USB devices can either be bus-powered (power sourced from the host) or self-powered (power sourced from an external power supply). The VDDUSB power pad pin powers the USB PHY and USB I/Os (D+ and D– pins). The presence of VBUS can be detected using the following steps:

1. Enable the interrupt on VDDUSB power pad. For this write '1' to the VDDIO_ACTIVE[5] bit of VDD_IN-TR_MASK register.
2. VDDIO_ACTIVE[5] bit of the supply detection interrupt register (VDD_INTR) is set to '1' whenever a change to supply is detected. Clear the interrupt cause by writing '1' to the bitfield.
3. Check the status of the VDD_ACTIVE[5] bit of the external power supply detection register (VDD_ACTIVE). The bit is set to '1' when there is supply and '0' when there is no supply.

32.3.2.3 USB D+ Pin Pull-up Enable Logic

When a USB device is self-powered, the USB specification warrants that the device enable the pull-up resistor on its D+ pin to identify itself as a full-speed device to the host. When the host VBUS is removed, the device should disable the pull-up resistor on the D+ line to not back power the host. The USB PHY includes an internal 1.5-k Ω pull-up resistor on the D+ line to indicate to the host that the PSoC 6 MCU is a full-speed device. The pull-up resistor can be enabled or disabled by configuring the DP_UP_EN bit in the USBLPM_POWER_CTRL register.

32.3.2.4 Transmitter and Receiver Logic

The transceiver block transmits and receives USB differential signals with an upstream device, and includes the USB D+ pull-up resistor used to maintain an idle state on the bus. Output data is differentially transmitted to upstream devices at a nominal voltage of 3.3 V. The differential inputs received from upstream devices are converted into single-ended data and sent to the core logic at a nominal voltage of 1.8 V. The D+ and D– pins are terminated with 22- Ω resistors to meet the USB impedance specification.

32.3.2.5 GPIO Mode Logic

The D+ and D– pins can be used either as GPIO pins or USB I/O pins. This is controlled by the IOMODE bit of the USBDEV_USBIO_CR1 register. This bit should be set HIGH for GPIO functionality and LOW for USB operation.

32.3.2.6 Link Power Management (LPM)

The USB PHY supports link power management (LPM), which is similar to the suspend mode, but has transitional latencies in tens of microseconds between power states, compared to the greater than 20 ms latency associated with suspend/resume modes. For more details on LPM, refer to the USB 2.0 specification. The following features are supported for LPM.

- The LPM_CTL register should be configured to enable/disable LPM, type of response when LPM is enabled,

and the response when a sub PID other than the LPM token is received from the host.

- The LPM_STAT register stores the values of the Best Effort Service Latency (BESL) and the remote wakeup feature as sent by the host. The firmware should read this register on the LPM interrupt event and enter the appropriate low-power mode (Deep Sleep or Sleep) based on the BESL value from the host.

32.3.3 Endpoints

The SIE and arbiter support eight data endpoints (EP1 to EP8) and one control endpoint (EP0). The data endpoints share the SRAM memory area of 512 bytes. The endpoint memory management can be either manual or automatic. The endpoints are configured for direction and other configuration using the SIE and arbiter registers. The endpoint read address and write address registers are accessed through the arbiter.

The endpoints can be individually made active. In the Auto Management mode, the register EP_ACTIVE is written to control the active state of the endpoint. The endpoint activation cannot be dynamically changed during runtime. In Manual Memory Management mode, the firmware decides the memory allocation, so it is not required to specify the active endpoints. The EP_ACTIVE register is ignored during the manual memory management mode. The EP_TYPE register is used to control the transfer direction (IN, OUT) for the endpoints. The control endpoint has separate eight bytes for its data (EP0_DR registers).

32.3.4 Transfer Types

The PSoC 6 MCU USB supports full-speed transfers and is compliant with the USB 2.0 specification. It supports four types of transfers:

- Interrupt Transfer
- Bulk Transfer
- Isochronous Transfer
- Control Transfer

For further details about these transfers, refer to the USB 2.0 specification.

32.3.5 Interrupt Sources

The USB device block generates 14 interrupts to the CPU. These interrupts are mapped to three general-purpose interrupt lines – INTR_LO, INTR_MED, and INTR_HI. Each of these three interrupt lines has an associated status register, which identifies the cause of the interrupt event. These are the USBLPM_INTR_CAUSE_LO, USBLPM_INTR_CAUSE_MED, and USBLPM_INTR_CAUSE_HI registers. The routing of these interrupts is controlled by the USBLPM_INTR_LVL_SEL register fields.

The following events generate an interrupt on one of the three interrupt lines:

- USB start of frame (SOF) event

- USB bus reset event
- Eight data endpoint (EP1 – EP8) interrupt events
- Control endpoint (EP0) interrupt event
- Link power management (LPM) event
- Resume event
- Arbiter Interrupt event

32.3.5.1 USB Start of Frame (SOF) Event

The SOF interrupt is generated upon receiving an SOF packet from the USB host. The SOF interrupt is enabled using the SOF_INTR_MASK bitfield in the USBLPM_INTR_SIE_MASK register.

- The SOF interrupt status is reflected in the SOF_INTR status bit in the USBLPM_INTR_SIE status register.
- The SOF interrupt status is also available in the SOF_INTR_MASKED bit of the USBLPM_INTR_SIE_MASKED register – this bit is the logical AND of the corresponding SOF bits in the USBLPM_INTR_SIE_MASK register and the USBLPM_INTR_SIE register.
- If there is no SOF interrupt for 3 ms, the USB device goes into SUSPEND state.

32.3.5.2 USB Bus Reset Event

The USB bus reset interrupt is generated when a USB bus reset condition occurs. The bus reset interrupt is enabled by setting the BUS_RESET_INTR_MASK bit in the USBLPM_INTR_SIE_MASK register.

- The bus reset interrupt status is reflected in the BUS_RESET_INTR bit in the USBLPM_INTR_SIE status register.
- The bus reset interrupt status is also available in the BUS_RESET_INTR_MASKED bit of the USBLPM_INTR_SIE_MASKED register – this bit is the logical AND of the corresponding bus reset bits in the USBLPM_INTR_SIE_MASK register and the USBLPM_INTR_SIE register.
- The SIE logic triggers the counter to start running on the divided version of CLK_PERI when an SE0 condition is detected on the USB bus. When the counter reaches the count value configured in the USBDEV_BUS_RST_CNT register, the bus reset interrupt is triggered. Typically, divided CLK_PERI is set to 100 kHz and USBDEV_BUS_RST_CNT is set to '10'.

32.3.5.3 Data Endpoint Interrupt Events

These are eight interrupt events corresponding to each data endpoint (EP1-EP8). Each of the endpoint interrupt events can be enabled/disabled by using the corresponding bit in the USBDEV_SIE_EP_INT_EN register. The interrupt status of each endpoint can be known by reading the USBDEV_SIE_EP_INT_SR status register. An endpoint whose interrupt is enabled can trigger the interrupt on the following events:

- Successful completion of an IN/OUT transfer
- NAK-ed IN/OUT transaction if the corresponding NAK_INT_EN bit in the SIE_EPx_CR0 register is set
- When there is an error in the transaction, the ERR_IN_TXN bit in the SIE_EPx_CR0 register is set and interrupt is generated.
- If the STALL bit in SIE_EPx_CR0 is set, then stall events can generate interrupts. This stall event can occur if an OUT packet is received for an endpoint whose mode bits in SIE_EPx_CR0 are set to ACK_OUT or if an IN packet is received with mode bits set to ACK_IN.

32.3.5.4 Control Endpoint Interrupt Event

The interrupt event corresponding to the control endpoint (EP0) is generated under the following events:

- Successful completion of an IN/OUT transfer
- When a SETUP packet is received on the control endpoint

The EP0 interrupt is setup using the EP0_INTR_SET bit of the USBLPM_INTR_SIE_SET register.

32.3.5.5 Link Power Management (LPM) Event

Generated whenever the LPM token packet is received. The LPM interrupt is enabled by setting the LPM_INTR_MASK bit in the USBLPM_INTR_SIE_MASK register. The LPM interrupt status is reflected in the LPM_INTR status bit in the USBLPM_INTR_SIE status register.

The LPM interrupt status is also available in the LPM_INTR_MASKED bit of the USBLPM_INTR_SIE_MASKED register; this bit is the logical AND of the corresponding LPM bits in the USBLPM_INTR_SIE_MASK and USBLPM_INTR_SIE registers.

The firmware needs to read the USBLPM_LPM_STAT register to read the BESL remote wakeup values and appropriately enter the desired low-power mode. Enter the low-power mode in the main code. The exit from LPM is identical to the resume event wakeup in the case of suspend mode.

32.3.5.6 RESUME Interrupt

The RESUME interrupt is asserted by the USB block when it detects '0' on the DP pad. The RESUME interrupt is enabled by setting the RESUME_INTR_MASK bitfield of the USBLPM_INTR_SIE_MASK register.

- The RESUME interrupt status is reflected in the RESUME_INTR status bit in the USBLPM_INTR_SIE status register.
- The RESUME interrupt status is also available in the RESUME_INTR_MASKED bit of the USBLPM_INTR_SIE_MASKED register – this bit is the logical AND of the corresponding RESUME bits in the USBLPM_INTR_SIE_MASK register and the USBLPM_INTR_SIE register.

The RESUME interrupt is an Active mode interrupt and not available in Deep Sleep or Hibernate mode.

32.3.5.7 Arbiter Interrupt Event

The arbiter interrupt can arise from five possible sources. Each interrupt source is logically ANDed with its corresponding ENABLE bit and the results are logically ORed to result in a single arbiter interrupt event.

The arbiter interrupt event can arise under any of the following five scenarios:

- DMA Grant
- IN Buffer Full
- Buffer Overflow
- Buffer Underflow
- DMA Termin

DMA Grant

This event is applicable in Mode 2 or Mode 3. (See [Logical Transfer Modes on page 421](#) for details on DMA modes). This event is triggered when the DMA controller pulses the Burstend signal corresponding to that endpoint, for which a DMA request had been raised to the DMA controller earlier. The request may have been either a manual DMA request or an automatic arbiter DMA request. A common grant status exists for both modes of requests. This grant status indicates completion of the DMA transaction. This status indication can be used by firmware to determine when the next manual DMA request can be raised. Multiple requests raised for the same endpoint before the DMA grant status is set will be dropped by the block. Only the first of multiple requests will be transmitted to DMA controller.

IN Buffer Full

This event status can occur in any of the DMA modes (Mode 1, 2, or 3) and is applicable only for IN endpoints.

- Store and Forward Mode (Modes 1 and 2): This status is set when the entire packet data is transferred to the local memory. The check is that data written for the particular endpoint is equal to the programmed byte count for that endpoint in the USBDEV_SIE_EPx_CNT0 and USBDEV_SIE_EPx_CNT1 registers.
- Cut Through Mode (Mode 3): In this mode, the IN buffer full status is set when the IN endpoint's dedicated buffer is filled with the packet data. The size of this buffer is determined by the value programmed in bits [3:0] of the USBDEV_BUF_SIZE register. This status indication can be used to determine when the mode value in the USBDEV_SIE_EPx_CR0 register can be programmed to acknowledge an IN token for that endpoint.

Buffer Overflow

This event status is active only in the Cut Through Mode (Mode 3). The following conditions can cause this bit to be set:

- Data overflow on the endpoint dedicated buffer space
 - In an IN endpoint, the dedicated buffer can overflow if the DMA transfer writes a larger number of bytes than the space available in the dedicated buffer. Until

an IN token is received for that endpoint, it cannot use the common buffer area, hence resulting in an overflow of data. The possible causes of this buffer overflow can be incorrect programming of either the DMA transfer descriptor transfer size or the USBDEV_BUF_SIZE register.

- In an OUT endpoint, the dedicated buffer can overflow if two OUT transactions occur consecutively. The data from the previous transaction is still present in the common area and the current ongoing transaction fills up the OUT endpoint's dedicated buffer space and overflows. The possible causes of this overflow can be the overall DMA bandwidth constraint due to other DMA transactions or reduced size of the dedicated OUT buffer size.

Common area data overflow

- In an IN endpoint, the common area overflow occurs when the DMA transfer writes a larger number of bytes than the space available in the common area. This situation may arise due to incorrect programming of either the DMA transfer descriptor transfer size or the USBDEV_DMA_THRESH and USBDEV_DMA_THRESH_MSB registers.
- In an OUT endpoint, the common area overflow occurs when the data written to the common area has not yet been read and new data overwrites the existing data.

Buffer Underflow

This event is applicable only in the Cut Through mode (Mode 3). This underflow condition can occur only for an IN endpoint. The underflow condition can occur either in the dedicated buffer space or common buffer space. The underflow condition on the dedicated buffer space can either be due to the reduced dedicated IN buffer size or DMA bandwidth constraint. The underflow condition can occur on the common buffer space due to DMA bandwidth constraint and/or lower DMA channel priority.

DMA Termin

This status is set when USBDEV generates a dma_termin signal to indicate the total programmed/ received bytes that are written/read by the DMA controller. This status indication can be used by the firmware to reprogram the IN/OUT endpoint for the next transfer. For an OUT endpoint, this indicates that the OUT packet data is available in the system memory for further processing by the application.

32.3.6 DMA Support

Each of the eight data endpoints has one DMA channel available to transfer data between the endpoint buffer and the SRAM memory. The USB generates the DMA request signals (usb.dma_req[7:0]) to the respective DMA channels to initiate the data transfer for the endpoint. This goes to the Trigger Group 13 multiplexer as input triggers that can be routed to one of the trigger inputs of the DMA block. The Burstend signals from the DMA channel to the correspond-

ing endpoint is routed using the Trigger Group 9 multiplexer. For more details, see the [DMA Controller \(DW\) chapter on page 93](#) and [Trigger Multiplexer Block chapter on page 260](#).

32.4 Logical Transfer Modes

The USB block in PSoC 6 MCUs supports two types of logical transfers. The logical transfers can be configured using the register setting for each endpoint. Any of the logical transfer methods can be adapted to support the three types of data transfers (Interrupt, Bulk, and Isochronous) mentioned in the USB 2.0 specification. The control transfer is mandatory in any USB device.

The logical transfer mode is a combination of memory management and DMA configurations. The logical transfer modes are related to the data transfer within the USB (to

and from the SRAM memory unit for each endpoint). It does not represent the transfer methods between the device and the host (the transfer types specified in the USB 2.0 specification).

The USB supports two basic types of transfer modes:

- Store and Forward mode
 - Manual Memory Management with No DMA Access (Mode 1)
 - Manual Memory Management with Manual DMA Access (Mode 2)
- Cut Through mode
 - Automatic Memory Management with Automatic DMA Access (Mode 3)

[Table 32-1](#) gives a comparison of the two transfer modes.

Table 32-1. USB Transfer Modes

Feature	Store and Forward Mode	Cut Through Mode
SRAM Memory Usage	Requires more memory	Requires less memory
SRAM Memory Management	Manual	Auto
SRAM Memory Sharing	512 bytes of SRAM shared between endpoints. Sharing is done by firmware.	Each endpoint is allocated a lesser share of memory automatically by the block. The remaining memory is available as "common area." This common area is used during the transfer.
IN Command	Entire packet present in SRAM memory before the IN command is received.	Memory filled with data only when SRAM IN command is received. Data is given to host when enough data is available (based on DMA configuration). Does not wait for the entire data to be filled.
OUT Command	Entire packet is written to SRAM memory on OUT command. After entire data is available, it is copied from SRAM memory to the USB device.	Waits only for enough bytes (depends on DMA configuration) to be written in SRAM memory. When enough bytes are present, it is immediately copied from SRAM memory to the USB device.
Transfer of Data	Data is transferred when all bytes are written to the memory.	Data is transferred when enough bytes are available. It does not wait for the entire data to be filled.
Types Based on DMA	No DMA mode Manual DMA mode	Only Auto DMA mode
Supported Transfer Types	Ideal for interrupt and bulk transfers	Ideal for Isochronous transfer

Every endpoint has a set of registers that need to be handled during the modes of operation, as detailed in [Table 32-2](#).

Table 32-2. Endpoint Registers

Register	Comment	Content	Usage
USBDEV_ARB_RWx_WA	Endpoint Write Address Register	Address of the SRAM	This register indicates the SRAM location to which the data in the data register is to be written.
USBDEV_ARB_RWx_RA	Endpoint Read Address Register	Address of the SRAM	This register indicates the SRAM location from which the data must be read and stored to the data register.
USBDEV_ARB_RWx_DR	Endpoint Data Register	8-Bit Data	<p>Data register is read/written to perform any transaction.</p> <p>IN command: Data written to the data register is copied to the SRAM location specified by the WA register. After write, the WA value is automatically incremented to point to the next memory location.</p> <p>OUT command: Data available in the SRAM location pointed by the USBDEV_ARB_RWx_RA register is read and stored to the DR. When the DR is read, the value of USBDEV_ARB_RWx_RA is automatically incremented to point to the next SRAM memory location that must be read.</p>
USBDEV_SIE_EPx_CNT0 and USBDEV_SIE_EPx_CNT1	Endpoint Byte Count Register	Number of Bytes	<p>Holds the number of bytes that can be transferred.</p> <p>IN command: Holds the number of bytes to be transferred to host.</p> <p>OUT command: Holds the maximum number of bytes that can be received. The firmware programs the maximum number of bytes that can be received for that endpoint. The SIE updates the register with the number of bytes received for the endpoint.</p>
"Mode" bits in USBDEV_SIE_EPx_CR0	Mode Values	Response to the Host	Controls how the USB device responds to the USB traffic and the USB host. Some examples of modes are ACK, NAK, and STALL.

In Manual memory management, the endpoint read and endpoint write address registers are updated by the firmware. So the memory allocation can be done by the user. The memory allocation decides which endpoints are active; that is, you can decide to share the 512 bytes for all the eight endpoints or a lesser number of endpoints.

In Automatic memory management, the endpoint read and endpoint write address registers are updated by the USB block. The block assigns memory to the endpoints that are activated using the USBDEV_EP_ACTIVE register. The size of memory allocated depends on the value in the USBDEV_BUF_SIZE register. The remaining memory, after allocation, is called the common area memory and is used for data transfer.

In all of these modes, either the 8-bit endpoint data register or the 16-bit endpoint data register can be used to read/write to the endpoint buffer. While transferring data to the 16-bit data registers, ensure that the corresponding SRAM memory address locations are also 16-bit aligned.

In the following text, the algorithm for the IN and OUT transaction for each mode is discussed. An IN transaction is when the data is read by the USB host (for example, PC). An OUT transaction is when the data is written by the USB host to the USB device. The choice of using the DMA and memory management can be configured using the USBDEV_ARB_EPx_CFG register.

32.4.1 Manual Memory Management with No DMA Access

All operations in this mode are controlled by the CPU and works in a store-and-forward operation mode. An entire packet is transferred to the memory and a mode bit (such as ACK IN or ACK OUT) is set by the CPU. The SIE responds appropriately to an IN/OUT token received from the host. All memory space management is handled by the CPU.

Figure 32-3. No DMA Access IN Transaction

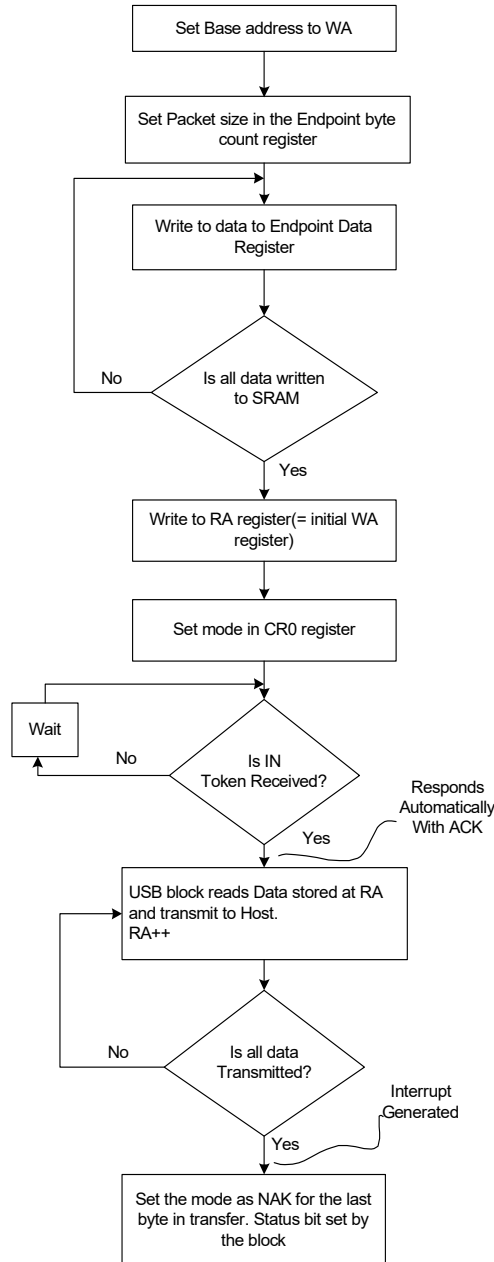
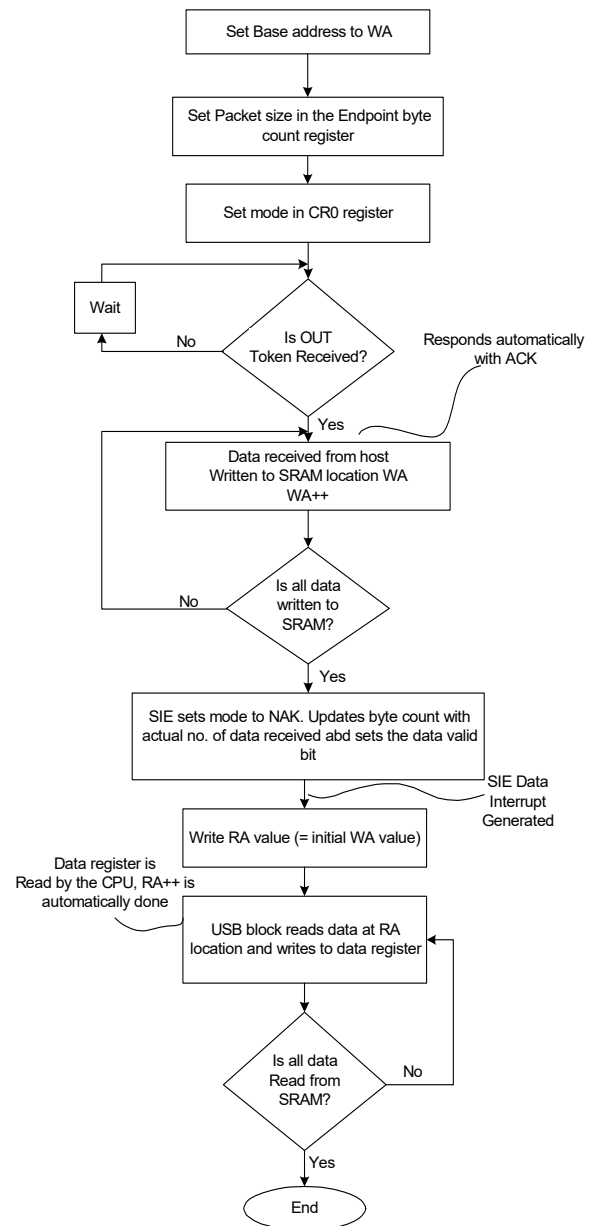


Figure 32-4. No DMA Access OUT Transaction



32.4.2 Manual Memory Management with DMA Access

This mode is similar to the No DMA Access except that write/read of packets is performed by the DMA. A DMA request for an endpoint is generated by setting the DMA_CFG bit in the USBDEV_ARB_EPx_CFG register. When the DMA service is granted and is done (DMA_GNT), an arbiter interrupt can be programmed to occur. The transfer is done using a single DMA cycle or multiple DMA cycles. After completion of every DMA cycle, the arbiter interrupt (DMA_GNT) is generated. Similarly, when all the data bytes (programmed in byte count) are written to the memory, the arbiter interrupt occurs and the IN_BUF_FULL bit is set.

Figure 32-5 and Figure 32-6 show the flow charts for manual DMA IN and OUT transactions respectively.

Figure 32-5. Manual DMA IN Transaction

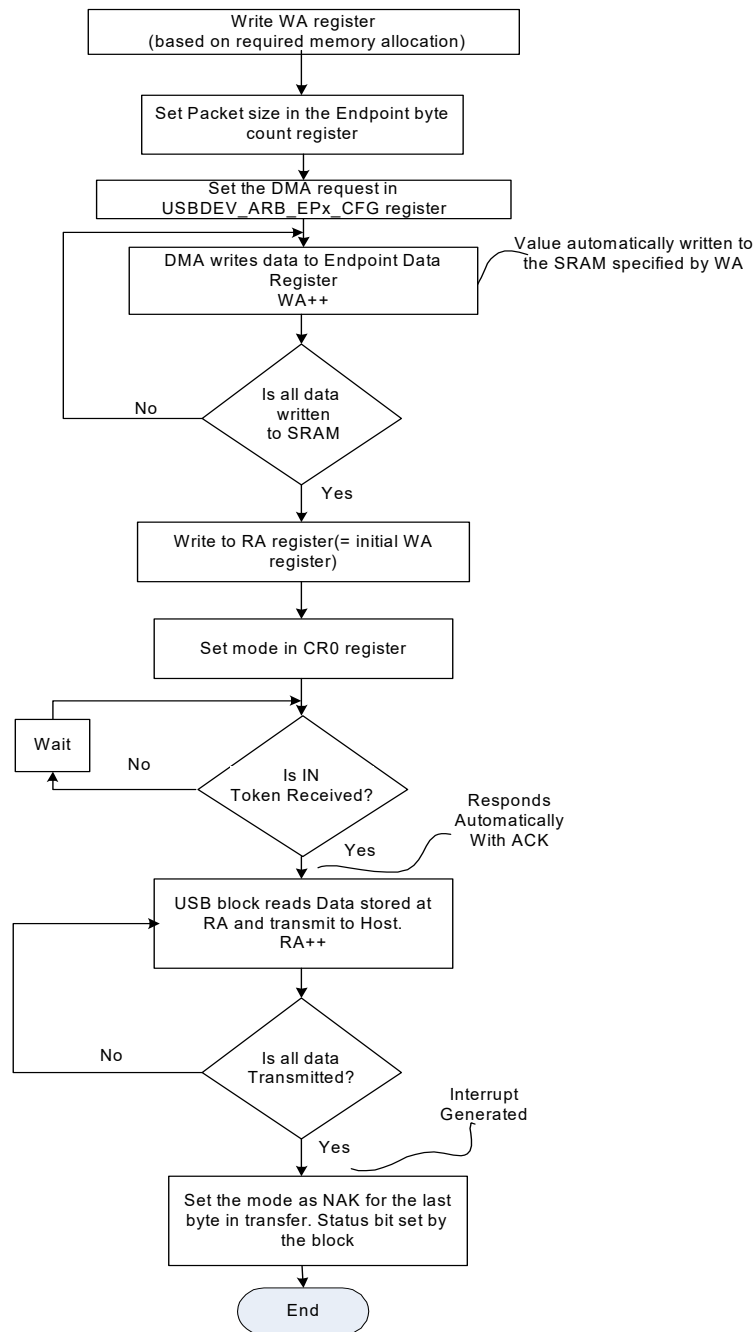
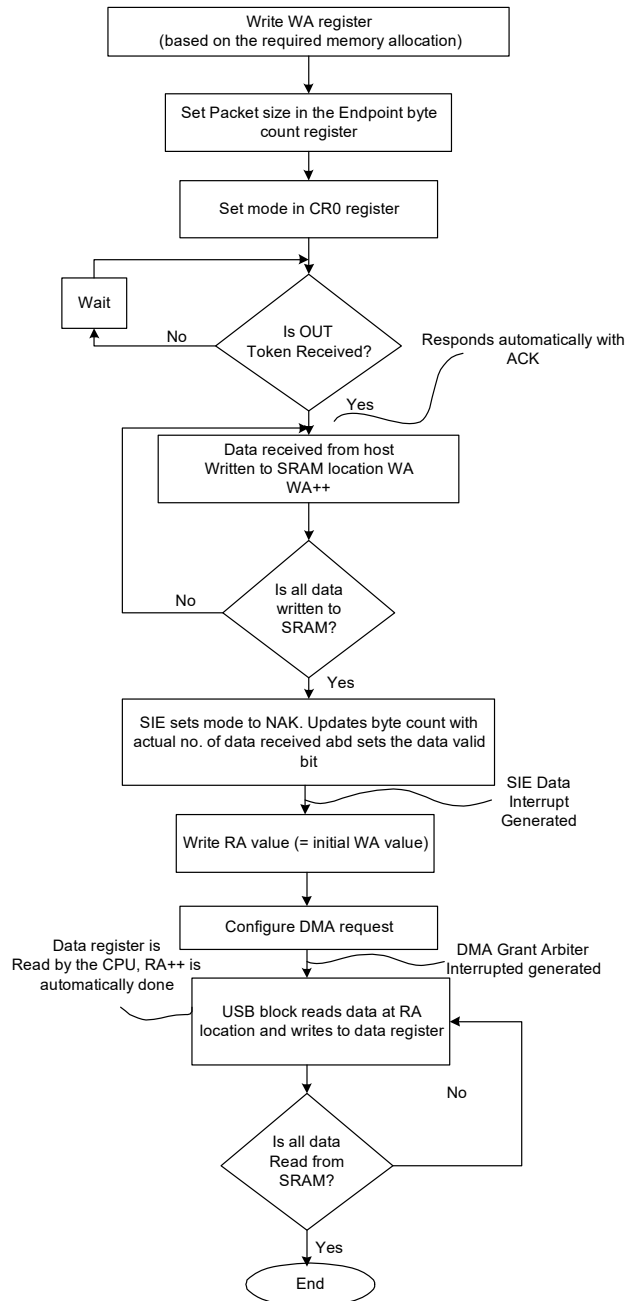


Figure 32-6. Manual DMA OUT Transaction



remaining memory (256 bytes) is left as common area and is common for all endpoints.

In this mode, the memory requirement is less and it is suitable for full-speed isochronous transfers up to 1023 bytes.

When an IN command is sent by the host, the device responds with the data present in the dedicated memory area for that endpoint. It simultaneously issues a DMA request for more data for that EP. This data fills up the common area. The device does not wait for the entire packet of data to be available. It waits only for the (USBDEV_DMA_THRES_MSB, USBDEV_DMA_THRES) number of data available in the SRAM memory and begins the transfer from the common area.

Similarly, when an OUT command is received, the data for the OUT endpoint is written to the common area. When some data (greater than USBDEV_DMA_THRES_MSB, USBDEV_DMA_THRES) is available in the common area, the arbiter block initiates a DMA request and the data is immediately written to the device. The device does not wait for the common area to be filled.

This mode requires configuration of the USBDEV_DMA_THRES and USBDEV_DMA_THRES_MSB registers to hold the number of bytes that can be transferred in one DMA transfer (32 bytes). Similarly, the burst count of the DMA should always be equal to the value set in the USBDEV_DMA_THRES registers. Apart from the DMA configuration, this mode also needs the configuration of the USBDEV_BUF_SIZE for the IN and the OUT buffers and the USBDEV_EP_ACTIVE and the USBDEV_EP_TYPE registers.

Each DMA channel has two descriptors and both of them are used in this mode. Each descriptor is considered as a data chunk of 32 bytes and it executes according to the trigger mechanism. The descriptors are chained and hence 64 bytes can be transferred without firmware interaction. When both descriptors complete the endpoint DMA done interrupt and the DMA error interrupt triggers (due to the lack of data to transfer). The descriptors are updated to advance the source SRAM (IN endpoint) or destination SRAM (OUT endpoint) pointer locations and then enabled again. This sequence continues till all data is transferred.

The steps for IN and OUT transactions using automatic DMA mode are shown in [Figure 32-7](#) and [Figure 32-8](#).

32.4.3 Automatic DMA Mode

This is the Automatic memory management mode with auto DMA access. The CPU programs the initial buffer size requirement for IN/OUT packets and informs the arbiter block of the endpoint configuration details for the particular application. The block then controls memory partitioning and handling of all memory pointers. During memory allocation, each active IN endpoint (set by the USBDEV_EP_ACTIVE and USBDEV_EP_TYPE registers) is allocated a small amount of memory configured using the USBDEV_BUF_SIZE register (32 bytes for each of the eight endpoints). The

Figure 32-7. IN Transaction using Automatic DMA Mode

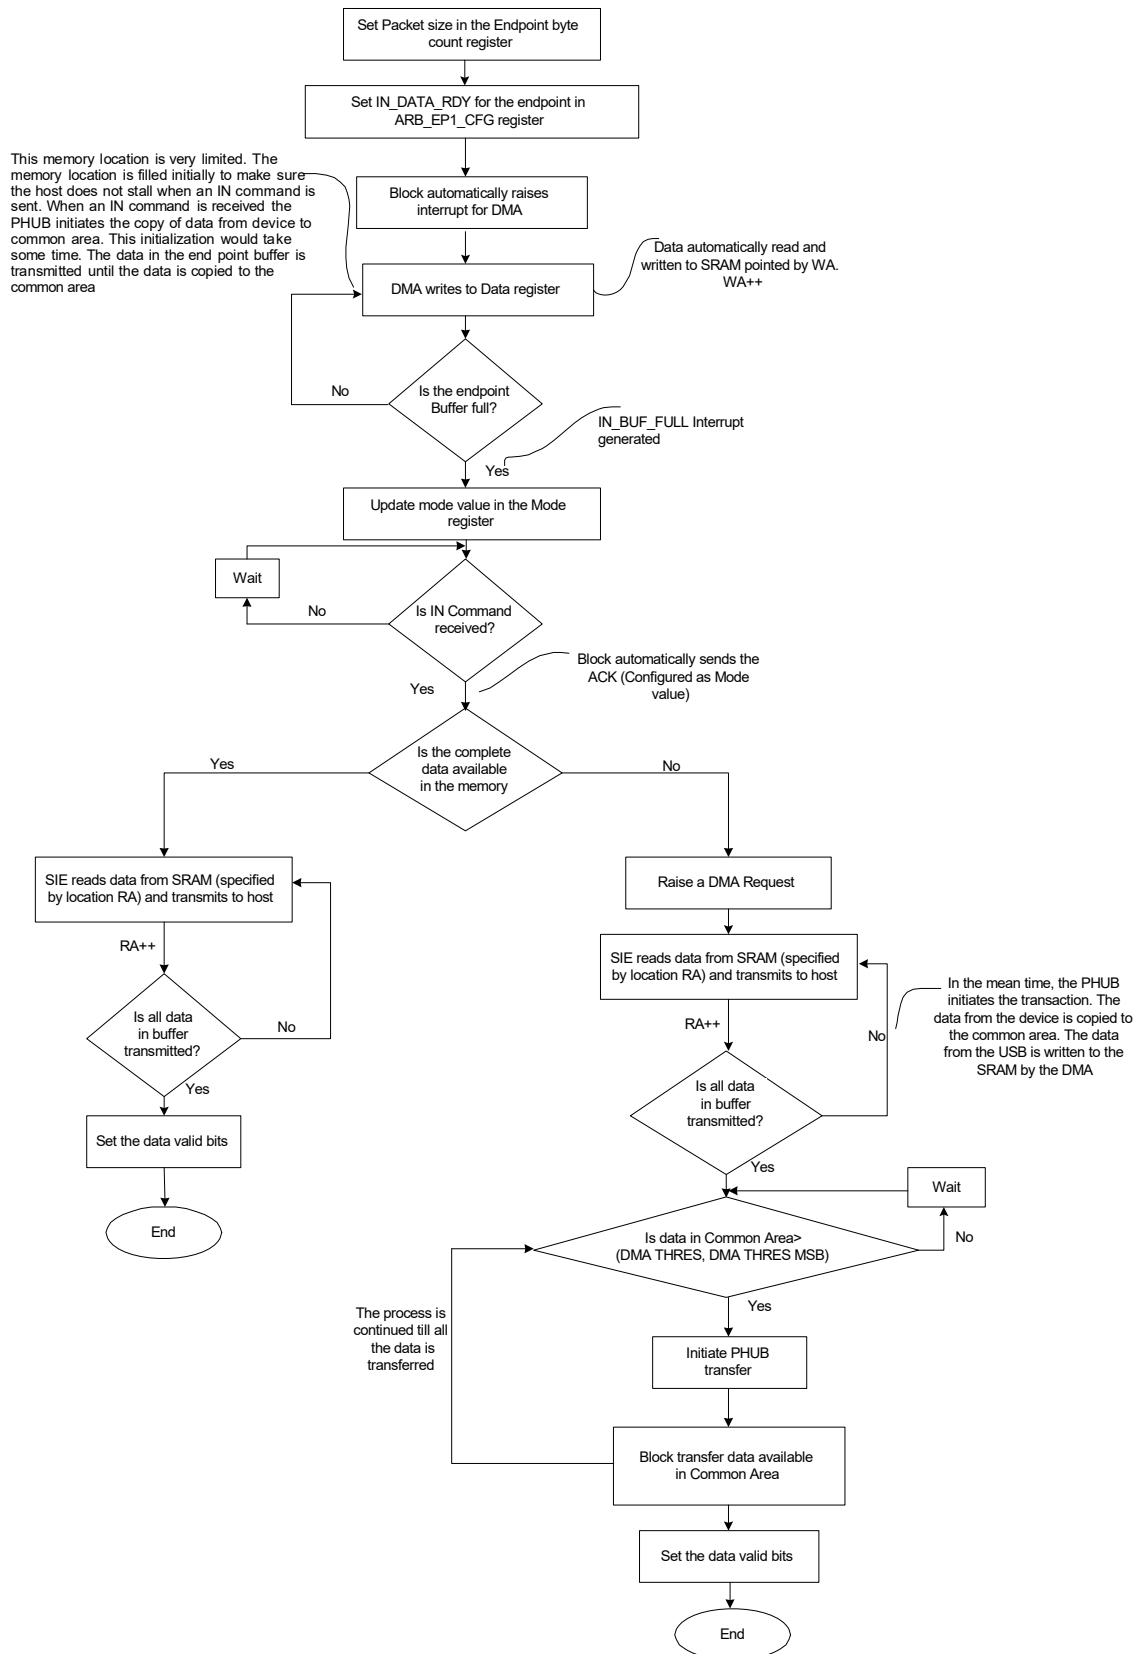



Figure 32-8. OUT Transaction using Automatic DMA Mode



32.4.4 Control Endpoint Logical Transfer

The control endpoint has a special logical transfer mode. It does not share the 512 bytes of memory. Instead, it has a dedicated 8-byte register buffer (USBDEV_EP0_DRx registers). The IN and OUT transaction for the control endpoint is detailed in the following figures.

Figure 32-9. Control Endpoint IN Transaction


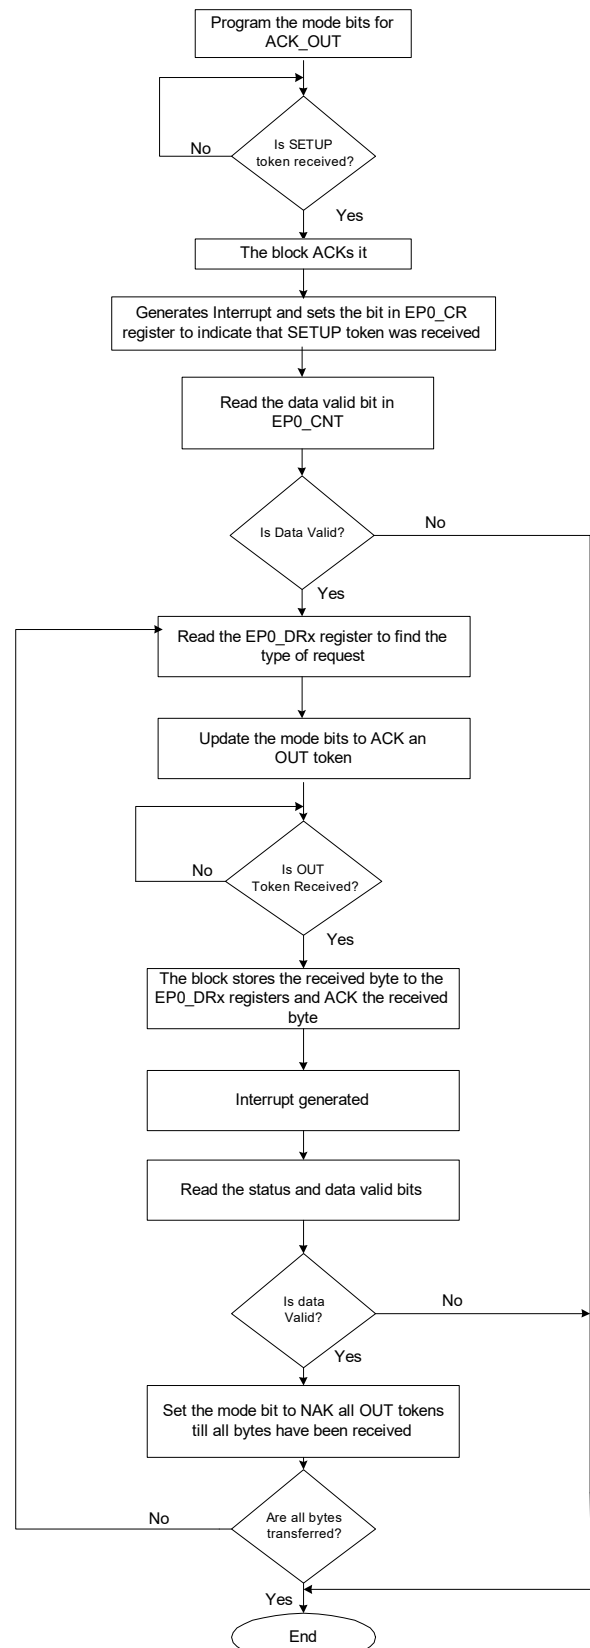


Figure 32-10. Control Endpoint OUT Transaction



32.5 USB Power Modes

The USB supports two modes of operation:

- Active mode: In this mode, the USB is powered up and clocks are turned on.
- Low-power (Deep Sleep) mode: In this mode, all clocks except the low-frequency clock are turned off.

Before entering low-power mode, the firmware should enable the GPIO interrupt (falling edge interrupt) on the D+ pin. The USB suspend mode can be determined by monitoring the SOF interrupt; this means, if there is no SOF interrupt from the USB for more than 3 ms, the USB goes into suspend mode and the block can be put into low-power mode by the firmware.

If there is any activity on the USB bus, D+ will be pulled low, which will cause a CPU interrupt. This interrupt can be used to wake up the USB device.

32.6 USB Device Registers

Name	Description
USBDEV_EP0_DR	Control endpoint EP0 data register
USBDEV_CR0	USB control 0 register
USBDEV_CR1	USB control 1 register
USBDEV_SIE_EP_INT_EN	USB SIE data endpoint interrupt enable register
USBDEV_SIE_EP_INT_SR	USB SIE data endpoint interrupt status
USBDEV_SIE_EPx_CNT0	Non-control endpoint count register
USBDEV_SIE_EPx_CNT1	Non-control endpoint count register
USBDEV_SIE_EPx_CR0	Non-control endpoint's control register
USBDEV_USBIO_CR0	USBIO control 0 register
USBDEV_USBIO_CR2	USBIO control 2 register
USBDEV_USBIO_CR1	USBIO control 1 register
USBDEV_DYN_RECONFIG	USB dynamic reconfiguration register
USBDEV_SOF0	Start of frame register
USBDEV_SOF1	Start of frame register
USBDEV_OSCLK_DR0	Oscillator lock data register 0
USBDEV_OSCLK_DR1	Oscillator lock data register 1
USBDEV_EP0_CR	Endpoint0 control register
USBDEV_EP0_CNT	Endpoint0 count register
USBDEV_ARB_RWx_WA	Endpoint write address value register. Pointer value increments by 1 when accessed by CPU/debugger
USBDEV_ARB_RWx_WA_MSB	Endpoint write address value register
USBDEV_ARB_RWx_RA	Endpoint read address value register. Pointer value increments by 1 when accessed by CPU/debugger
USBDEV_ARB_RWx_RA_MSB	Endpoint read address value register
USBDEV_ARB_RWx_DR	Endpoint data register
USBDEV_BUF_SIZE	Dedicated endpoint buffer size register
USBDEV_EP_ACTIVE	Endpoint active indication register
USBDEV_EP_TYPE	Endpoint type (IN/OUT) indication register
USBDEV_ARB_EPx_CFG	Endpoint configuration register
USBDEV_ARB_EPx_INT_EN	Endpoint interrupt enable register
USBDEV_ARB_EPx_SR	Endpoint interrupt enable register
USBDEV_ARB_CFG	Arbiter configuration register
USBDEV_USB_CLK_EN	USB block clock enable register

Name	Description
USBDEV_ARB_INT_EN	Arbiter interrupt enable register
USBDEV_ARB_INT_SR	Arbiter interrupt status register
USBDEV_CWA	Common area write address register
USBDEV_CWA_MSB	Endpoint read address value register
USBDEV_DMA_THRES	DMA burst / threshold configuration register
USBDEV_DMA_THRES_MSB	DMA burst / threshold configuration register
USBDEV_BUS_RST_CNT	Bus reset count register
USBDEV_MEM_DATA	Data register
USBDEV_SOF16	Start of frame register
USBDEV_OSCLK_DR16	Oscillator lock data register
USBDEV_ARB_RWx_WA16	Endpoint write address value register. Pointer value increments by 2 when accessed by CPU/debugger
USBDEV_ARB_RWx_RA16	Endpoint read address value register. Pointer value increments by 2 when accessed by CPU/debugger
USBDEV_ARB_RWx_DR16	Endpoint data register
USBDEV_DMA_THRES16	DMA burst / threshold configuration register
USBLPM_POWER_CTL	Power control register
USBLPM_USBIO_CTL	USB IO control register
USBLPM_FLOW_CTL	Flow control register
USBLPM_LPM_CTL	LPM control register
USBLPM_LPM_STAT	LPM status register
USBLPM_INTR_SIE	USB SOF, BUS RESET, and EP0 interrupt status register
USBLPM_INTR_SIE_SET	USB SOF, BUS RESET, and EP0 interrupt set register
USBLPM_INTR_SIE_MASK	USB SOF, BUS RESET,, and EP0 interrupt mask register
USBLPM_INTR_SIE_MASKED	USB SOF, BUS RESET, and EP0 interrupt masked register
USBLPM_INTR_LVL_SEL	Select interrupt level for each interrupt source register
USBLPM_INTR_CAUSE_HI	High-priority interrupt cause register
USBLPM_INTR_CAUSE_MED	Medium-priority interrupt cause register
USBLPM_INTR_CAUSE_LO	Low-priority interrupt cause register

33. LCD Direct Drive



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - LCD Drive
- Application notes
- Code examples

The PSoC 6 MCU Liquid Crystal Display (LCD) drive system is a highly configurable peripheral that allows the device to directly drive STN and TN segment LCDs.

33.1 Features

The PSoC 6 MCU LCD segment drive block has the following features:

- Supports up to 56 segments and 8 commons
- Supports Type A (standard) and Type B (low-power) drive waveforms
- Any GPIO can be configured as a common or segment
- Supports five drive methods:
 - Digital correlation
 - PWM at 1/2 bias
 - PWM at 1/3 bias
 - PWM at 1/4 bias
- Ability to drive 3-V displays from 1.8 V V_{DD} in Digital Correlation mode
- Operates in Active, Sleep, and Deep Sleep modes
- Digital contrast control

33.2 Architecture

33.2.1 LCD Segment Drive Overview

A segmented LCD panel has the liquid crystal material between two sets of electrodes and various polarization and reflector layers. The two electrodes of an individual segment are called commons (COM) or backplanes and segment electrodes (SEG). From an electrical perspective, an LCD segment can be considered as a capacitive load; the COM/SEG electrodes can be considered as the rows and columns in a matrix of segments. The opacity of an LCD segment is controlled by varying the root-mean-square (RMS) voltage across the corresponding COM/SEG pair.

The following terms/voltages are used in this chapter to describe LCD drive:

- **V_{RMSOFF}** : The voltage that the LCD driver can realize on segments that are intended to be off.
- **V_{RMSON}** : The voltage that the LCD driver can realize on segments that are intended to be on.

- **Discrimination Ratio (D):** The ratio of V_{RMSON} and V_{RMSOFF} that the LCD driver can realize. This depends on the type of waveforms applied to the LCD panel. Higher discrimination ratio results in higher contrast.

Liquid crystal material does not tolerate long term exposure to DC voltage. Therefore, any waveforms applied to the panel must produce a 0-V DC component on every segment (on or off). Typically, LCD drivers apply waveforms to the COM and SEG electrodes that are generated by switching between multiple voltages. The following terms are used to define these waveforms:

- **Duty:** A driver is said to operate in 1/M duty when it drives 'M' number of COM electrodes. Each COM electrode is effectively driven 1/M of the time.
- **Bias:** A driver is said to use 1/B bias when its waveforms use voltage steps of $(1/B) \times V_{DRV}$. V_{DRV} is the highest drive voltage in the system (equals V_{DD}). The PSoC 6 MCU supports 1/2, 1/3, and 1/4 biases in PWM drive modes.
- **Frame:** A frame is the length of time required to drive all the segments. During a frame, the driver cycles through the commons in sequence. All segments receive 0-V DC (but non-zero RMS voltage) when measured over the entire frame.

The PSoC 6 MCU supports two different types of drive waveforms in all drive modes. These are:

- **Type-A Waveform:** In this type of waveform, the driver structures a frame into M sub-frames. 'M' is the number of COM electrodes. Each COM is addressed only once during a frame. For example, COM[i] is addressed in sub-frame i.
- **Type-B Waveform:** The driver structures a frame into 2M sub-frames. The two sub-frames are inverses of each other. Each COM is addressed twice during a frame. For example, COM[i] is addressed in sub-frames i and M+i. Type-B waveforms are slightly more power efficient because it contains fewer transitions per frame.

33.2.2 Drive Modes


The PSoC 6 MCU supports the following drive modes.

- PWM drive at 1/2 bias
- PWM drive at 1/3 bias
- PWM drive at 1/4 bias with high-frequency clock input
- Digital correlation

33.2.2.1 PWM Drive

In PWM drive mode, multi-voltage drive signals are generated using a PWM output signal together with the intrinsic resistance and capacitance of the LCD. [Figure 33-1](#) illustrates this.

Figure 33-1. PWM Drive (at 1/3 Bias)



The output waveform of the drive electronics is a PWM waveform. With the Indium Tin Oxide (ITO) panel resistance and the segment capacitance to filter the PWM, the voltage across the LCD segment is an analog voltage, as shown in Figure 33-1. This figure illustrates the generation of a 1/3 bias waveform (four commons and voltage steps of $V_{DD}/3$). See the [Clocking System chapter on page 207](#) for details.

The PWM is derived from either CLK_LF (32 kHz, low-speed operation) or CLK_PERI (high-speed operation). See the [Clocking System chapter on page 207](#) for more details of peripheral and low-frequency clocks. The filtered analog voltage across the LCD segments typically runs at low frequency for segment LCD driving.

Figure 33-2 and Figure 33-3 illustrate the Type A and Type B waveforms for COM and SEG electrodes for 1/2 bias and 1/4 duty. Only COM0/COM1 and SEG0/SEG1 are drawn for demonstration purpose. Similarly, Figure 33-4 and Figure 33-5 illustrate the Type A and Type B waveforms for COM and SEG electrodes for 1/3 bias and 1/4 duty.

Figure 33-2. PWM1/2 Type-A Waveform Example

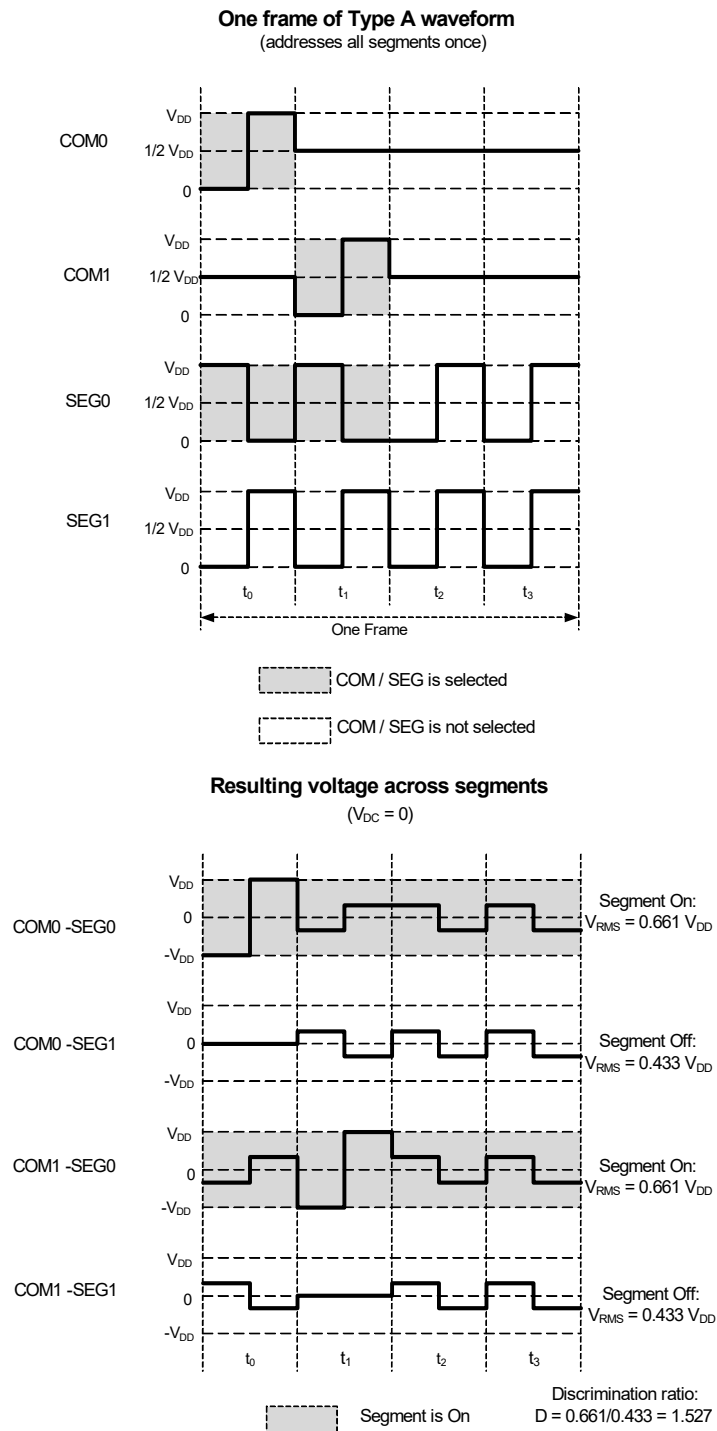


Figure 33-3. PWM1/2 Type-B Waveform Example




Figure 33-4. PWM1/3 Type-A Waveform Example

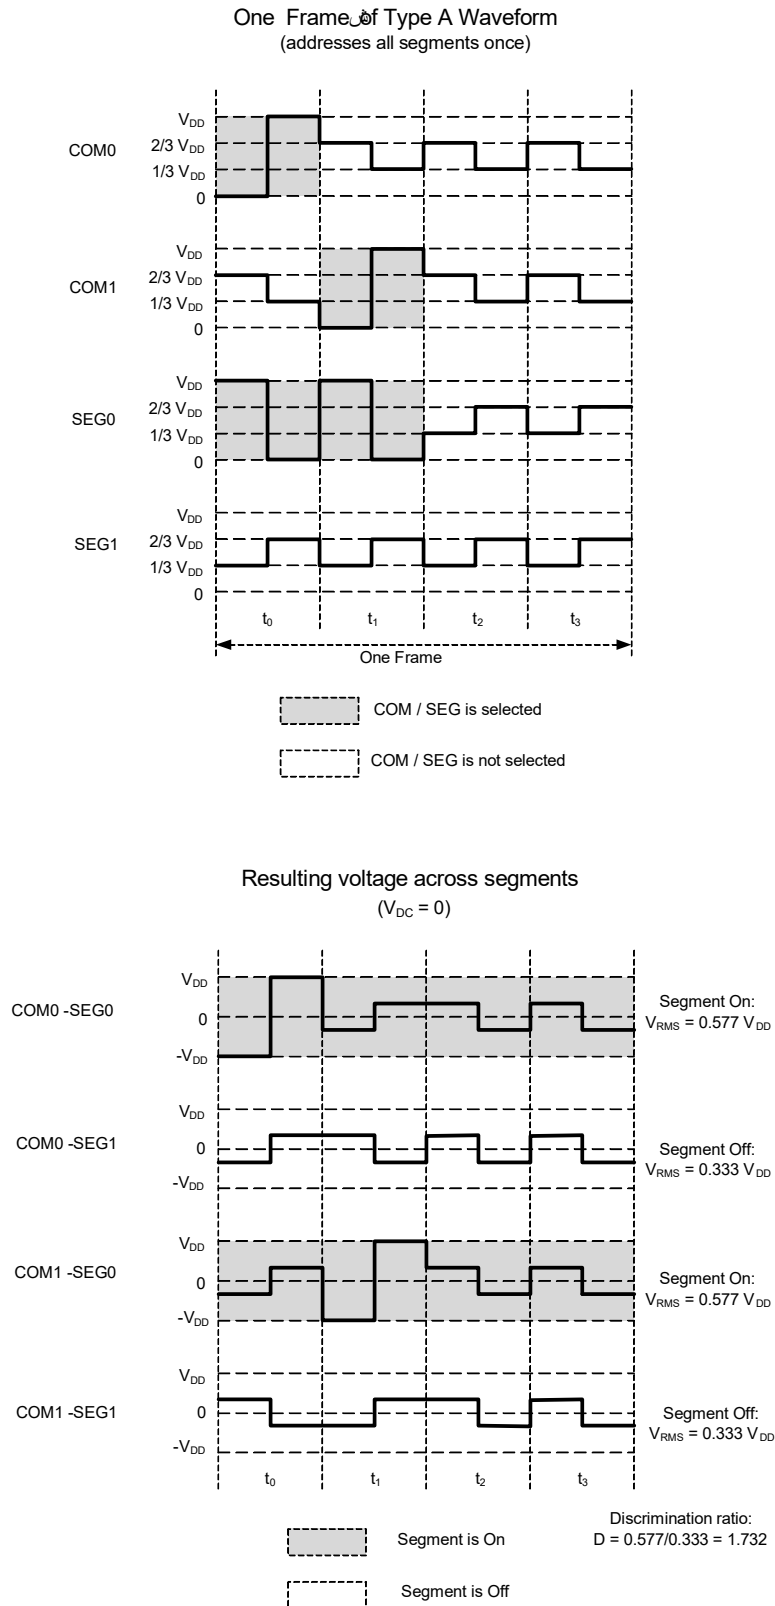
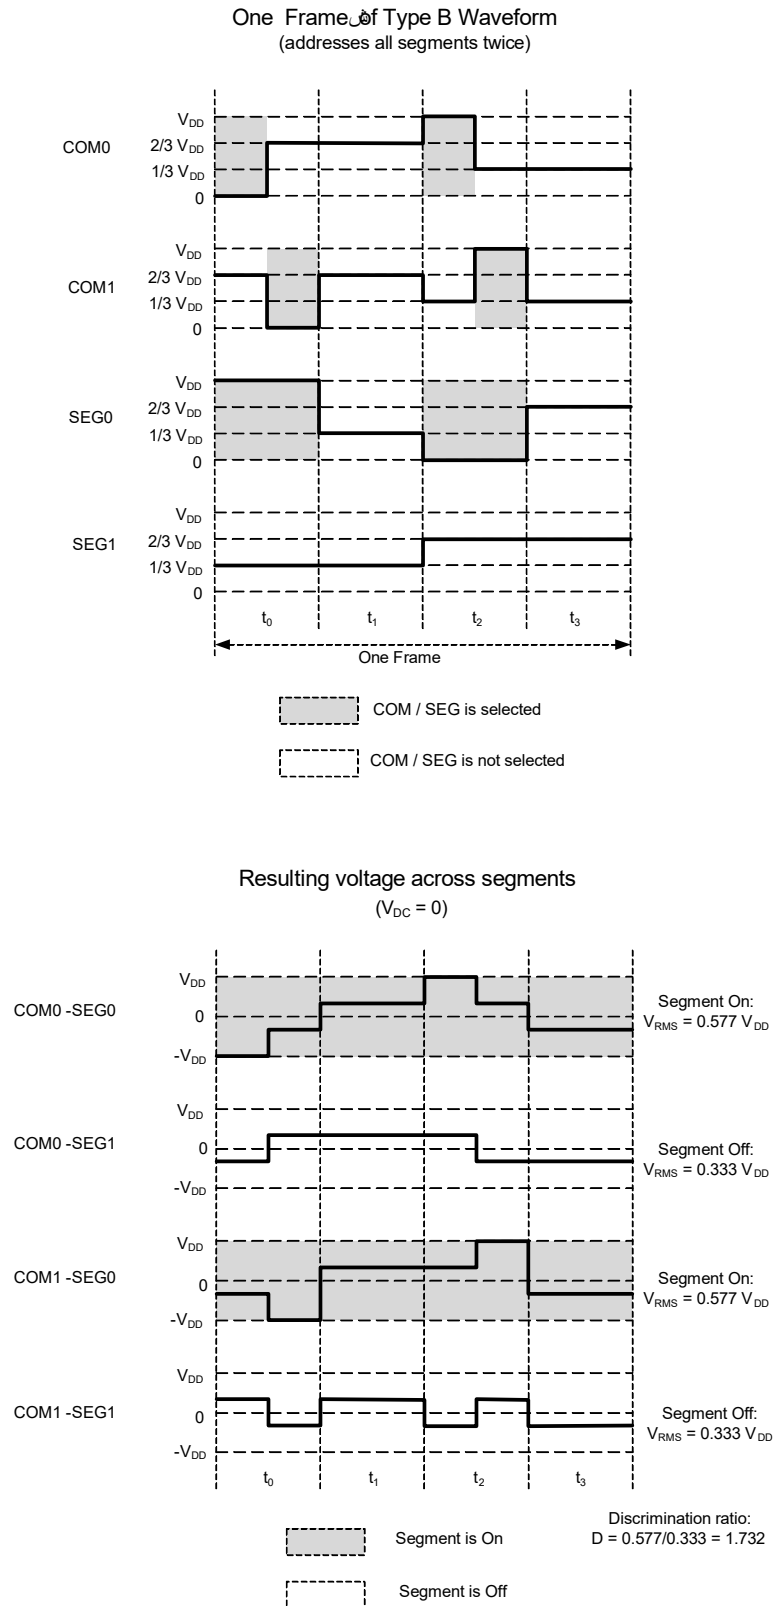


Figure 33-5. PWM1/3 Type-B Waveform Example



The effective RMS voltage for ON and OFF segments can be calculated easily using these equations:

$$V_{RMS(OFF)} = \sqrt{\frac{2(B-2)^2 + 2(M-1)}{2M}} \times \left(\frac{V_{DRV}}{B}\right)$$

Equation 33-1

$$V_{RMS(ON)} = \sqrt{\frac{2B^2 + 2(M-1)}{2M}} \times \left(\frac{V_{DRV}}{B}\right)$$

Equation 33-2

Where B is the bias and M is the duty (number of COMs).

For example, if the number of COMs is four, the resulting discrimination ratios (D) for 1/2 and 1/3 biases are 1.528 and 1.732, respectively. 1/3 bias offers better discrimination ratio in two and three COM drives also. Therefore, 1/3 bias offers better contrast than 1/2 bias and is recommended for most applications. 1/4 bias is available only in high-speed operation of the LCD. They offer better discrimination ratio especially when used with high COM designs (more than four COMs).

When the low-speed operation of LCD is used, the PWM signal is derived from the 32-kHz CLK_LF. To drive a low-capacitance display with acceptable ripple and rise/fall times using a 32-kHz PWM, additional external series resistances of 100 k-1 MΩ should be used. External resistors are not required for PWM frequencies greater than ~1 MHz. The ideal PWM frequency depends on the capacitance of the display and the internal ITO resistance of the ITO routing traces.

The 1/2 bias mode has the advantage that PWM is only required on the COM signals; the SEG signals use only logic levels, as shown in [Figure 33-2](#) and [Figure 33-3](#).

33.2.2.2 Digital Correlation

The digital correlation mode, instead of generating bias voltages between the rails, takes advantage of the characteristic of LCDs that the contrast of LCD segments is determined by the RMS voltage across the segments. In this approach, the correlation coefficient between any given pair of COM and SEG signals determines whether the corresponding LCD segment is on or off. Thus, by doubling the base drive frequency of the COM signals in their inactive sub-frame intervals, the phase relationship of the COM and SEG drive signals can be varied to turn segments on and off. This is different from varying the DC levels of the signals as in the PWM drive approach. [Figure 33-8](#) and [Figure 33-9](#) are example waveforms that illustrate the principles of operation.

Figure 33-6. Digital Correlation Type-A Waveform





Figure 33-7. Digital Correlation Type-B Waveform



The RMS voltage applied to on and off segments can be calculated as follows:

$$V_{RMS(OFF)} = \sqrt{\frac{(M-1)}{2M}} \times (V_{DD})$$

$$V_{RMS(ON)} = \sqrt{\frac{2 + (M-1)}{2M}} \times (V_{DD})$$

Where B is the bias and M is the duty (number of COMs). This leads to a discrimination ratio (D) of 1.291 for four COMs. Digital correlation mode also has the ability to drive 3-V displays from 1.8-V V_{DD} .

33.2.3 Recommended Usage of Drive Modes

The PWM drive mode has higher discrimination ratios compared to the digital correlation mode, as explained in [33.2.2.1 PWM Drive](#) and [33.2.2.2 Digital Correlation](#). Therefore, the contrast in digital correlation method is lower than PWM method but digital correlation has lower power consumption because its waveforms toggle at low frequencies.

The digital correlation mode creates reduced, but acceptable contrast on TN displays, but no noticeable difference in contrast or viewing angle on higher contrast STN displays. Because each mode has strengths and weaknesses, recommended usage is as follows.


Table 33-1. Recommended Usage of Drive Modes

Display Type	Deep Sleep Mode	Sleep/Active Mode	Notes
Four COM TN Glass	Digital correlation	PWM 1/3 bias	Firmware must switch between LCD drive modes before going to deep sleep or waking up.
Four COM STN Glass	Digital correlation		No contrast advantage for PWM drive with STN glass.
Eight COM, STN	Not supported	PWM 1/4 bias	Supported only in the high-speed LCD mode. The low-speed CLK_LF is not fast enough to make the PWM work at high multiplex ratios.

33.2.4 Digital Contrast Control

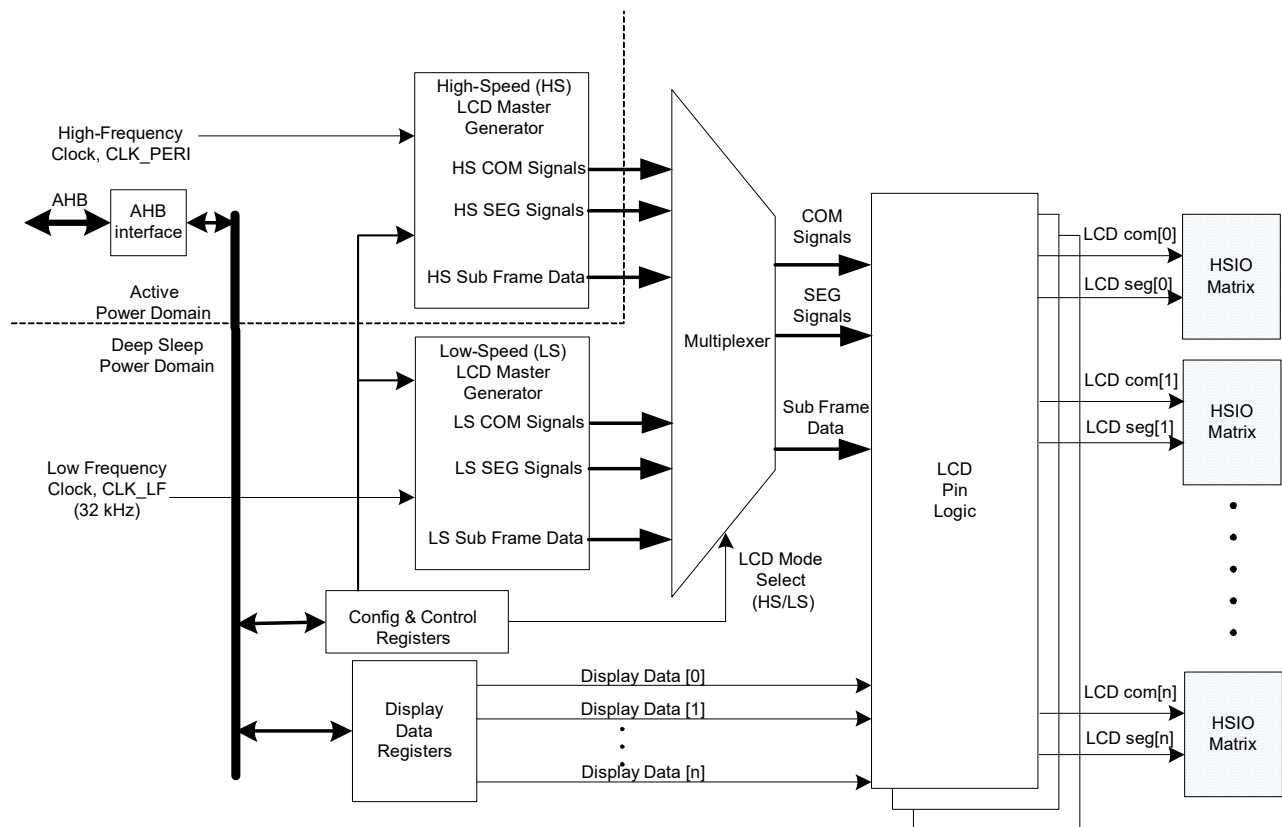
In all drive modes, digital contrast control can be used to change the contrast level of the segments. This method reduces contrast by reducing the driving time of the segments. This is done by inserting a 'Dead-Time' interval after each frame. During dead time, all COM and SEG signals are driven to a logic 1 state. The dead time can be controlled in fine resolution. [Figure 33-8](#) illustrates the dead-time contrast control method for 1/3 bias and 1/4 duty implementation.

Figure 33-8. Dead-Time Contrast Control



33.3 PSoC 6 MCU Segment LCD Direct Drive

Figure 33-9. Block Diagram of LCD Direct Drive System



The LCD controller block contains two generators, one with a high-speed clock source CLK_PERI and the other with a low-speed clock source (32 kHz) derived from the CLK_LF. These are called high-speed LCD master generator and low-speed LCD master generator, respectively. Both the generators support PWM and digital correlation drive modes. PWM drive mode with low-speed generator requires external resistors, as explained in [PWM Drive on page 432](#).

The multiplexer selects one of these two generator outputs to drive LCD, as configured by the firmware. The LCD pin logic block routes the COM and SEG outputs from the generators to the corresponding I/O matrices. Any GPIO can be used as either COM or SEG. This configurable pin assignment for COM or SEG is implemented in GPIO and I/O matrix; see [High-Speed I/O Matrix](#). These two generators share the same configuration registers. These memory mapped I/O registers are connected to the system bus (AHB) using an AHB interface.

The LCD controller works in three device power modes: Active, Sleep, and Deep Sleep. High-speed operation is supported in Active and Sleep modes. Low-speed operation is supported in Active, Sleep, and Deep Sleep modes. The LCD controller is unpowered in Hibernate mode.

33.3.1 High-Speed and Low-Speed Master Generators

The high-speed and low-speed master generators are similar to each other. The only exception is that the high-speed version has larger frequency dividers to generate the frame and sub-frame periods. The high-speed generator is in the active power domain and the low-speed generator is in the Deep Sleep power domain. A single set of configuration registers is provided to control both high-speed and low-speed blocks. Each master generator has the following features and characteristics:

- Register bit configuring the block for either Type A or Type B drive waveforms (LCD_MODE bit in LCD0_CONTROL register).
- Register bits to select the number of COMs (COM_NUM field in LCD0_CONTROL register).
- Operating mode configuration bits enabled to select one of the following:
 - Digital correlation
 - PWM 1/2 bias
 - PWM 1/3 bias

- ❑ PWM 1/4 bias (not supported in low-speed generator)
- ❑ Off/disabled. Typically, one of the two generators will be configured to be Off

OP_MODE and BIAS fields in LCD0_CONTROL bits select the drive mode.

- A counter to generate the sub-frame timing. The SUBFR_DIV field in the LCD0_DIVIDER register determines the duration of each sub-frame. If the divide value written into this counter is C, the sub-frame period is $4 \times (C+1)$. The low-speed generator has an 8-bit counter. This counter generates a maximum half sub-frame period of 8 ms from the 32-kHz CLK_LF. The high-speed generator has a 16-bit counter.
- A counter to generate the dead time period. These counters have the same number of bits as the sub-frame period counters and use the same clocks. DEAD_DIV field in the LCD0_DIVIDER register controls the dead time period.

33.3.2 Multiplexer and LCD Pin Logic

The multiplexer selects the output signals of either high-speed or low-speed master generator blocks and feeds it to the LCD pin logic. This selection is controlled by the configu-

ration and control register. The LCD pin logic uses the sub-frame signal from the multiplexer to choose the display data. This pin logic will be replicated for each LCD pin.

33.3.3 Display Data Registers

Each LCD segment pin is part of an LCD port with its own display data register, LCD0_DATAx. The device has eight such LCD ports. Note that these ports are not real pin ports but the ports/connections available in the LCD hardware for mapping the segments to commons. Each LCD segment configured is considered as a pin in these LCD ports. The LCD0_DATAxx registers are 32-bit wide and store the ON/OFF data for all SEG-COM combination enabled in the design. For example, LCD0_DATA0x holds SEG-COM data for COM0 to COM3 and LCD0_DATA1x holds SEG-COM data for COM4 to COM7. The bits $[4i+3:4i]$ (where 'i' is the pin number) of each LCD0_DATAxx register represent the ON/OFF data for Pin[i], as shown in Table 33-2. The LCD0_DATAxx register should be programmed according to the display data of each frame. The display data registers are Memory Mapped I/O (MMIO) and accessed through the AHB slave interface. See the [device datasheet](#) for the pin connections.

Table 33-2. SEG-COM Mapping Example of LCD0_DATA00 Register (each SEG is a pin of the LCD port)

BITS[31:28] = PIN_7[3:0]				BITS[27:24] = PIN_6[3:0]			
PIN_7-COM3	PIN_7-COM2	PIN_7-COM1	PIN_7-COM0	PIN_6-COM3	PIN_6-COM2	PIN_6-COM1	PIN_6-COM0
BITS[23:20] = PIN_5[3:0]				BITS[19:16] = PIN_4[3:0]			
PIN_5-COM3	PIN_5-COM2	PIN_5-COM1	PIN_5-COM0	PIN_4-COM3	PIN_4-COM2	PIN_4-COM1	PIN_4-COM0
BITS[15:12] = PIN_3[3:0]				BITS[11:8] = PIN_2[3:0]			
PIN_3-COM3	PIN_3-COM2	PIN_3-COM1	PIN_3-COM0	PIN_2-COM3	PIN_2-COM2	PIN_2-COM1	PIN_2-COM0
BITS[7:3] = PIN_1[3:0]				BITS[3:0] = PIN_0[3:0]			
PIN_1-COM3	PIN_1-COM2	PIN_1-COM1	PIN_1-COM0	PIN_0-COM3	PIN_0-COM2	PIN_0-COM1	PIN_0-COM0

33.4 Register List

Table 33-3. LCD Direct Drive Register List

Register Name	Description
LCD0_ID	This register includes the information of LCD controller ID and revision number
LCD0_DIVIDER	This register controls the sub-frame and dead-time period
LCD0_CONTROL	This register is used to configure high-speed and low-speed generators
LCD0_DATA0x	LCD port pin data register for COM0 to COM3
LCD0_DATA1x	LCD port pin data register for COM4 to COM7

Section E: Analog Subsystem

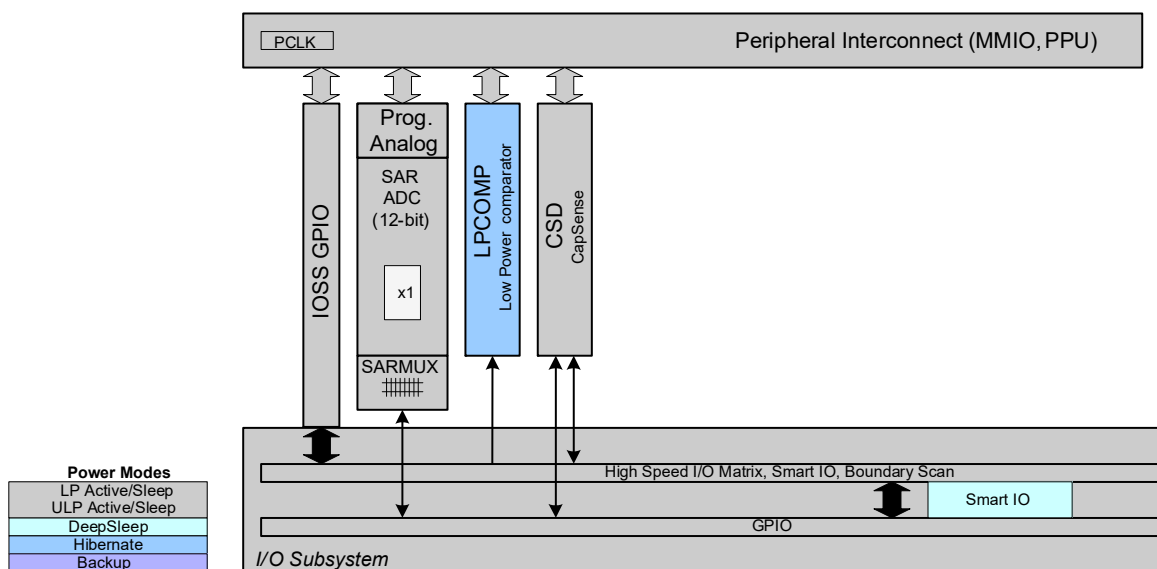


This section encompasses the following chapters:

- [Analog Reference Block chapter on page 445](#)
- [Low-Power Comparator chapter on page 448](#)
- [SAR ADC chapter on page 453](#)
- [Temperature Sensor chapter on page 467](#)
- [CapSense chapter on page 471](#)

Top Level Architecture

Figure E-1. Analog System Block Diagram



34. Analog Reference Block



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SysAnalog
- Application notes
- Code examples

The Analog Reference block (AREF) generates highly accurate reference voltage and currents needed by the programmable analog subsystem (PASS) and CapSense (CSD) blocks.

34.1 Features

- Provides accurate bandgap references for PASS and CSD subsystems
- 1.2-V voltage reference (V_{REF}) generator
- Option to output alternate voltage references routed from SRSS or from an external pin
- Proportional to absolute temperature (IPTAT) current reference generation
- Zero dependency to absolute temperature (IZTAT) flat current reference generation, which is independent of temperature variations
- Option to generate IZTAT from SRSS generated current reference
- Option to enable or disable references in System Deep-Sleep mode

34.2 Architecture

Figure 34-1. AREF Block Diagram




Figure 34-1 shows the architecture of the AREF block. AREF contains a precise bandgap reference, which generates a 1.2-V V_{REF} and a PTAT current that is 1 μ A at room temperature (23 °C). It also generates a current reference, which is stable over temperature (IZTAT).

34.2.1 Bandgap Reference Block

The AREF block contains a local bandgap reference generator, which has tighter accuracy, temperature stability, and lower noise than the SRSS bandgap reference. The bandgap reference block provides a temperature stable voltage and current reference (V_{REF}) and an additional current that tracks temperature. All AREF output currents are sinking.

Table 34-1. Bandgap References in AREF

Bandgap Reference Outputs	Value
VREF	1.2 V
IPTAT	1 μ A

34.2.2 Zero Dependency To Absolute Temperature Current Generator (IZTAT)

The IZTAT current generator uses the output of the bandgap reference block to generate a precise current reference, which has a low variation over temperature.

Table 34-2. IZTAT Reference

IZTAT Reference Outputs	Value
IZTAT	1 μ A

34.2.3 Reference Selection Multiplexers

AREF has two multiplexers to select the sources for the output references.

34.2.3.1 V_{REF} Selection Multiplexer Options

- SRSS – routes the V_{REF} from SRSS (0.8 V)
Note: SRSS references are not available in Deep Sleep power mode. This option should not be used when the device is in Deep Sleep mode.
- Local – routes the locally generated bandgap reference (1.2 V)
- External – routes the reference from the external V_{REF} pin

34.2.3.2 IZTAT Selection Multiplexer Options

- SRSS - uses the SRSS current reference (250 nA) to generate IZTAT.
Note: SRSS references are not available in Deep Sleep power mode. This option should not be used when the device is in Deep Sleep mode.
- Local - internally generates IZTAT

34.2.4 Startup Modes

AREF supports two startup modes, which provide a trade-off between wakeup time and noise performance when transitioning from the Deep Sleep to Active power mode. This is selected using the AREF_MODE bit of the AREF_CTRL register. The FAST_START mode enables faster wakeup, but with higher noise levels. Firmware can switch to NORMAL mode after the FAST_START settling time to achieve better noise performance.

Table 34-3. Startup Modes

AREF_MODE	Description
0 (NORMAL)	Normal startup mode
1 (FAST_START)	Fast startup mode

34.2.5 Low-Power Modes

AREF's voltage and current references can be made available during the Deep Sleep power mode. This is configured using the DEEPSLEEP_ON bit in the AREF_CTRL register. Individual references can be active by configuring DEEPSLEEP_MODE as shown in Table 34-4.

Note: These options are applicable only when DEEPSLEEP_ON = 1.

Table 34-4. Deep Sleep Mode

DEEPSLEEP_MODE[1:0]	Description
00 (OFF)	All references are OFF during Deep Sleep
01 (IPTAT)	IPTAT is ON during Deep Sleep. This mode enables fast wakeup from Deep Sleep.
10 (IPTAT_IPTAT)	IPTAT and IZTAT are ON during Deep Sleep
11 (IPTAT_IPTAT_VREF)	IPTAT, IZTAT, and VREF are ON during Deep Sleep

34.3 Registers

Table 34-5. List of AREF Registers

Register	Comment	Features
AREF_CTRL	AREF control register	Reference selection, startup time, and low-power mode.

35. Low-Power Comparator



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - Low-Power Comparator
- Application notes
- Code examples

PSoC 6 MCUs have two Low-Power comparators, which can perform fast analog signal comparison of internal and external analog signals in all system power modes. Low-Power comparator output can be inspected by the CPU, used as an interrupt/wakeup source to the CPU when in CPU Sleep mode, used as a wakeup source to system resources when in System Deep Sleep or Hibernate mode, or fed to as an asynchronous or synchronous signal (level or pulse).

35.1 Features

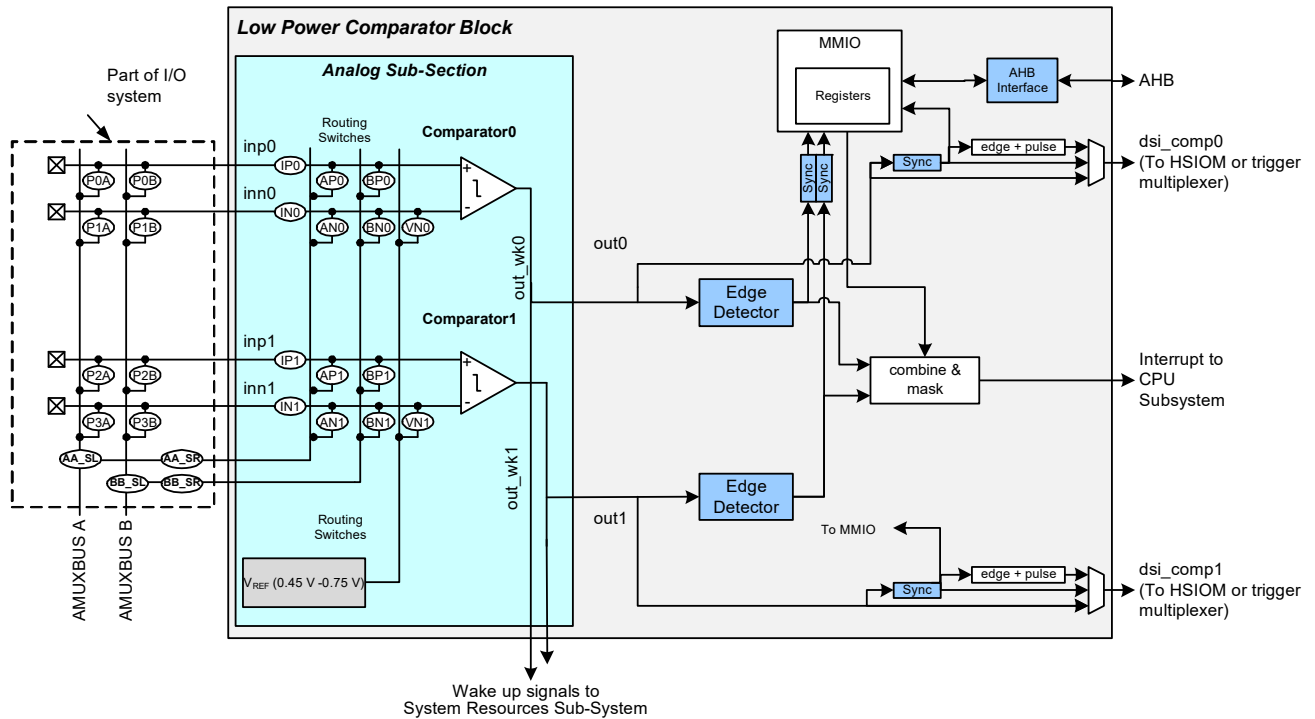
The PSoC 6 MCU comparators have the following features:

- Configurable input pins
- Programmable power and speed
- Ultra low-power mode support
- Each comparator features a one-sided hysteresis option
- Rising edge, falling edge, combined rising and falling edge detection at the comparator output
- Local reference voltage generation
- Wakeup source from low-power modes

35.2 Architecture

Figure 35-1 shows the block diagram for the Low-Power comparator.

Figure 35-1. Low-Power Comparator Block Diagram



The following sections describe the operation of the PSoC 6 MCU Low-Power comparator, including input configuration, power and speed modes, output and interrupt configuration, hysteresis, and wakeup from low-power modes.

35.2.1 Input Configuration

Low-Power comparators can operate with the following input modes:

- Compare two voltages from external pins
- Compare a voltage from an external pin against an internally generated analog-signal (through AMUXBUS).
- Compare two internal voltages through AMUXBUS A/ AMUXBUS B
- Compare internal and external signals with a locally-generated reference voltage. Note that this voltage is not a precision reference and can vary from 0.45 V–0.75 V.

See the [device datasheet](#) for detailed specifications of the Low-Power comparator.

Note that AMUXBUS connections are not available in Deep Sleep and Hibernate modes. If Deep Sleep or Hibernate operation is required, the Low-Power comparator must be connected to the dedicated pins. This restriction also includes routing of any internally-generated signal, which

uses the AMUXBUS for the connection. See the [I/O System chapter on page 226](#) for more details on connecting the GPIO to AMUXBUS A/B or setting up the GPIO for comparator input.

Refer to the LPCOMP_CMP0_SW, LPCOMP_CMP1_SW, LPCOMP_CMP0_SW_CLEAR, and LPCOMP_CMP1_SW_CLEAR registers in the [registers TRM](#) to understand how to control the internal routing switches shown in [Figure 35-1](#).

If the inverting input of a comparator is routed to a local voltage reference, the LPREF_EN bit in the LPCOMP_CONFIG register must be set to enable the voltage reference.

35.2.2 Output and Interrupt Configuration

Both Comparator0 and Comparator1 have hardware outputs available at dedicated pins. See the [device datasheet](#) for the location of comparator output pins.

Firmware readout of Comparator0 and Comparator1 outputs are available at the OUT0 and OUT1 bits of the LPCOMP_STATUS register ([Table 35-1](#)). The output of each comparator is connected to a corresponding edge detector block. This block determines the edge that triggers the interrupt. The edge selection and interrupt enable is

configured using the INTTYPE0 and INTTYPE1 bitfields in the LPCOMP_CMP0_CTRL and LPCOMP_CMP1_CTRL registers for Comparator0 and Comparator1, respectively. Using the INTTYPE0 and INTTYPE1 bits, the interrupt type can be selected to disabled, rising edge, falling edge, or both edges, as described in [Table 35-1](#).

Each comparator's output can also be routed directly to a GPIO pin through the HSIOM. See the [I/O System chapter on page 226](#) for more details.

During an edge event, the comparator will trigger an interrupt. The interrupt request is registered in the COMP0 bit and COMP1 bit of the LPCOMP_INTR register for Comparator0 and Comparator1, respectively. Both Comparator0 and Comparator1 share a common interrupt signal output (see [Figure 35-1](#)), which is a logical OR of the two interrupts and mapped as the Low-Power comparator block's interrupt in the CPU NVIC. Refer to the [Interrupts chapter on page 54](#) for details. If both the comparators are used in a design, the COMP0 and COMP1 bits of the LPCOMP_INTR register must be read in the interrupt service routine to know which one triggered the

interrupt. Alternatively, COMP0_MASK bit and COMP1_MASK bit of the LPCOMP_INTR_MASK register can be used to mask the Comparator0 and Comparator1 interrupts to the CPU. Only the masked interrupts will be serviced by the CPU. After the interrupt is processed, the interrupt should be cleared by writing a '1' to the COMP0 and COMP1 bits of the LPCOMP_INTR register in firmware. If the interrupt to the CPU is not cleared, it stays active regardless of the next compare events. This can interrupt the CPU continuously. Refer to the [Interrupts chapter on page 54](#) for details.

LPCOMP_INTR_SET register bits [1:0] can be used to assert an interrupt for firmware debugging.

In Deep Sleep mode, the wakeup interrupt controller (WIC) can be activated by a comparator edge event, which then wakes up the CPU. Similarly in Hibernate mode, the LPCOMP can wake up the system resources sub-system. Thus, the LPCOMP has the capability to monitor a specified signal in low-power modes. See the [Power Supply and Monitoring chapter on page 183](#) and the [Device Power Modes chapter on page 190](#) for more details.

Table 35-1. Output and Interrupt Configuration

Register[Bit_Pos]	Bit_Name	Description
LPCOMP_STATUS[0]	OUT0	Current/Instantaneous output value of Comparator0
LPCOMP_STATUS[16]	OUT1	Current/Instantaneous output value of Comparator1
LPCOMP_CMP0_CTRL[7:6]	INTTYPE0	Sets on which edge Comparator0 will trigger an IRQ 00: Disabled 01: Rising Edge 10: Falling Edge 11: Both rising and falling edges
LPCOMP_CMP1_CTRL[7:6]	INTTYPE1	Sets on which edge Comparator1 will trigger an IRQ 00: Disabled 01: Rising Edge 10: Falling Edge 11: Both rising and falling edges
LPCOMP_INTR[0]	COMP0	Comparator0 Interrupt: hardware sets this interrupt when Comparator0 triggers. Write a '1' to clear the interrupt
LPCOMP_INTR[1]	COMP1	Comparator1 Interrupt: hardware sets this interrupt when Comparator1 triggers. Write a '1' to clear the interrupt
LPCOMP_INTR_SET[0]	COMP0	Write a '1' to trigger the software interrupt for Comparator0
LPCOMP_INTR_SET[1]	COMP1	Write a 1 to trigger the software interrupt for Comparator1

35.2.3 Power Mode and Speed Configuration

The Low-Power comparators can operate in three power modes:

- Normal
- Low-power
- Ultra low-power

The power or speed setting for Comparator0 is configured using the MODE0 bitfield of the LPCOMP_CMP0_CTRL register. Similarly, the power or speed setting for Comparator1 is configured using the MODE1 bitfield of the LPCOMP_CMP1_CTRL register. The power consumption and response time vary depending on the selected power mode; power consumption is highest in fast mode and lowest in ultra-low-power mode, response time is fastest in fast mode and slowest in ultra-low-power

mode. Refer to the [device datasheet](#) for specifications for the response time and power consumption for various power settings.

The comparators can also be enabled or disabled using these bitfields, as described in [Table 35-2](#).

Note: The output of the comparator may glitch when the power mode is changed while comparator is enabled. To avoid this, disable the comparator before changing the power mode.

Table 35-2. Comparator Power Mode Selection Bits

Register[Bit_Pos]	Bit_Name	Description
LPCOMP_CMP0_CTRL[1:0]	MODE0	Comparator0 power mode selection 00: Off 01: Ultra low-power operating mode. This mode must be used when the device is in Deep Sleep or Hibernate mode. 10: Low-power operating mode 11: Normal, full-speed, full-power operating mode See the datasheet for electrical specifications in each power mode.
LPCOMP_CMP1_CTRL[1:0]	MODE1	Comparator1 power mode selection 00: Off 01: Ultra low-power operating mode. This mode must be used when the device is in Deep Sleep or Hibernate mode. 10: Low-power operating mode 11: Normal, full-speed, full-power operating mode See the datasheet for electrical specifications in each power mode.

Additionally, the entire Low-Power comparator system can be enabled or disabled globally using the LPCOMP_CONFIG[31] bit. See the [registers TRM](#) for details of these bitfields.

35.2.4 Hysteresis

For applications that compare signals close to each other and slow changing signals, hysteresis helps to avoid oscillations at the comparator output when the signals are noisy. For such applications, a fixed hysteresis may be enabled in the comparator block. See the [device datasheet](#) for the hysteresis voltage range.

The hysteresis level is enabled/disabled by using the HYST0 and HYST1 bitfields in the LPCOMP_CMP0_CTRL and LPCOMP_CMP1_CTRL registers for Comparator0 and Comparator1, as described in [Table 35-3](#).

Table 35-3. Hysteresis Control Bits

Register[Bit_Pos]	Bit_Name	Description
LPCOMP_CMP0_CTRL[5]	HYST0	Enable/Disable hysteresis to Comparator0 1: Enable Hysteresis 0: Disable Hysteresis See the datasheet for hysteresis voltage range.
LPCOMP_CMP1_CTRL[5]	HYST1	Enable/Disable hysteresis to Comparator1 1: Enable Hysteresis 0: Disable Hysteresis See the datasheet for hysteresis voltage range.

35.2.5 Wakeup from Low-Power Modes

The comparator is operational in the device's low-power modes, including Sleep, Deep Sleep, and Hibernate modes. The comparator output can wake the device from Sleep and Deep Sleep modes. The comparator output can generate interrupts in Deep Sleep mode when enabled in the LPCOMP_CONFIG register, the INTTYPE_x bits in the LPCOMP_CMP_x_CTRL register should not be set to disabled, and the INTR_MASK_x bit should be set in the LPCOMP_INTR_MASK register for the corresponding comparator to wake the device from low-power modes. Moreover, if the comparator is required in Deep Sleep

or Hibernate modes, then it must be configured into Ultra Low-Power mode before the device is put into Deep Sleep or Hibernate mode.

In the Deep Sleep and Hibernate power modes, a compare event on either Comparator0 or Comparator1 output will generate a wakeup interrupt. The INTTYPEEx bits in the LPCOMP_CONFIG register should be properly configured. The mask bits in the LPCOMP_INTR_MASK register is used to select whether one or both of the comparator's interrupt is serviced by the CPU. See the [Device Power Modes chapter on page 190](#) for more details.

35.2.6 Comparator Clock

The comparator uses the system main clock SYSCLK as the clock for interrupt synchronization. See the [Clocking System chapter on page 207](#) for more details.

35.3 Register List

Table 35-4. Low-Power Comparator Register Summary

Register	Function
LPCOMP_CONFIG	LPCOMP global configuration register
LPCOMP_INTR	LPCOMP interrupt register
LPCOMP_INTR_SET	LPCOMP interrupt set register
LPCOMP_INTR_MASK	LPCOMP interrupt request mask register
LPCOMP_INTR_MASKED	LPCOMP masked interrupt output register
LPCOMP_STATUS	Output status register
LPCOMP_CMP0_CTRL	Comparator0 configuration register
LPCOMP_CMP1_CTRL	Comparator1 configuration
LPCOMP_CMP0_SW	Comparator0 switch control
LPCOMP_CMP1_SW	Comparator1 switch control
LPCOMP_CMP0_SW_CLEAR	Comparator0 switch control clear
LPCOMP_CMP1_SW_CLEAR	Comparator1 switch control clear

36. SAR ADC



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SAR ADC
- Application notes
- Code examples

The PSoC 6 MCU has a 12-bit successive approximation register analog-to-digital converter (SAR ADC). The 12-bit, 1-Msps SAR ADC is designed for applications that require moderate resolution and high data rate.

36.1 Features

- Maximum sample rate of 1 Msps
- Sixteen individually configurable logical channels that can scan eleven unique input channels. Each channel has the following features:
 - Input from eight dedicated pins (eight single-ended mode or four differential inputs) or internal signals (AMUXBUS or temperature sensor)
 - Each channel may choose one of the four programmable acquisition times to compensate for external factors (such as high input impedance sources with long settling times)
 - Single-ended or differential measurement
 - Averaging and accumulation
 - Double-buffered results
- Result may be left- or right-aligned, or may be represented in 16-bit sign extended
- Scan can be triggered by firmware, trigger from other peripherals or pins
 - One-shot – periodic or continuous mode
- Hardware averaging support
 - First order accumulate
 - Supports 2, 4, 8, 16, 32, 64, 128, and 256 samples (powers of 2)
- Selectable voltage references
 - Internal V_{DDA} and $V_{DDA}/2$ references
 - Internal 1.2-V reference with buffer
 - External reference
- Interrupt generation
 - End of scan
 - Saturation detect and over-range (configurable) detect for every channel
 - Scan results overflow
 - Collision detect

36.2 Architecture

Figure 36-1. Block Diagram

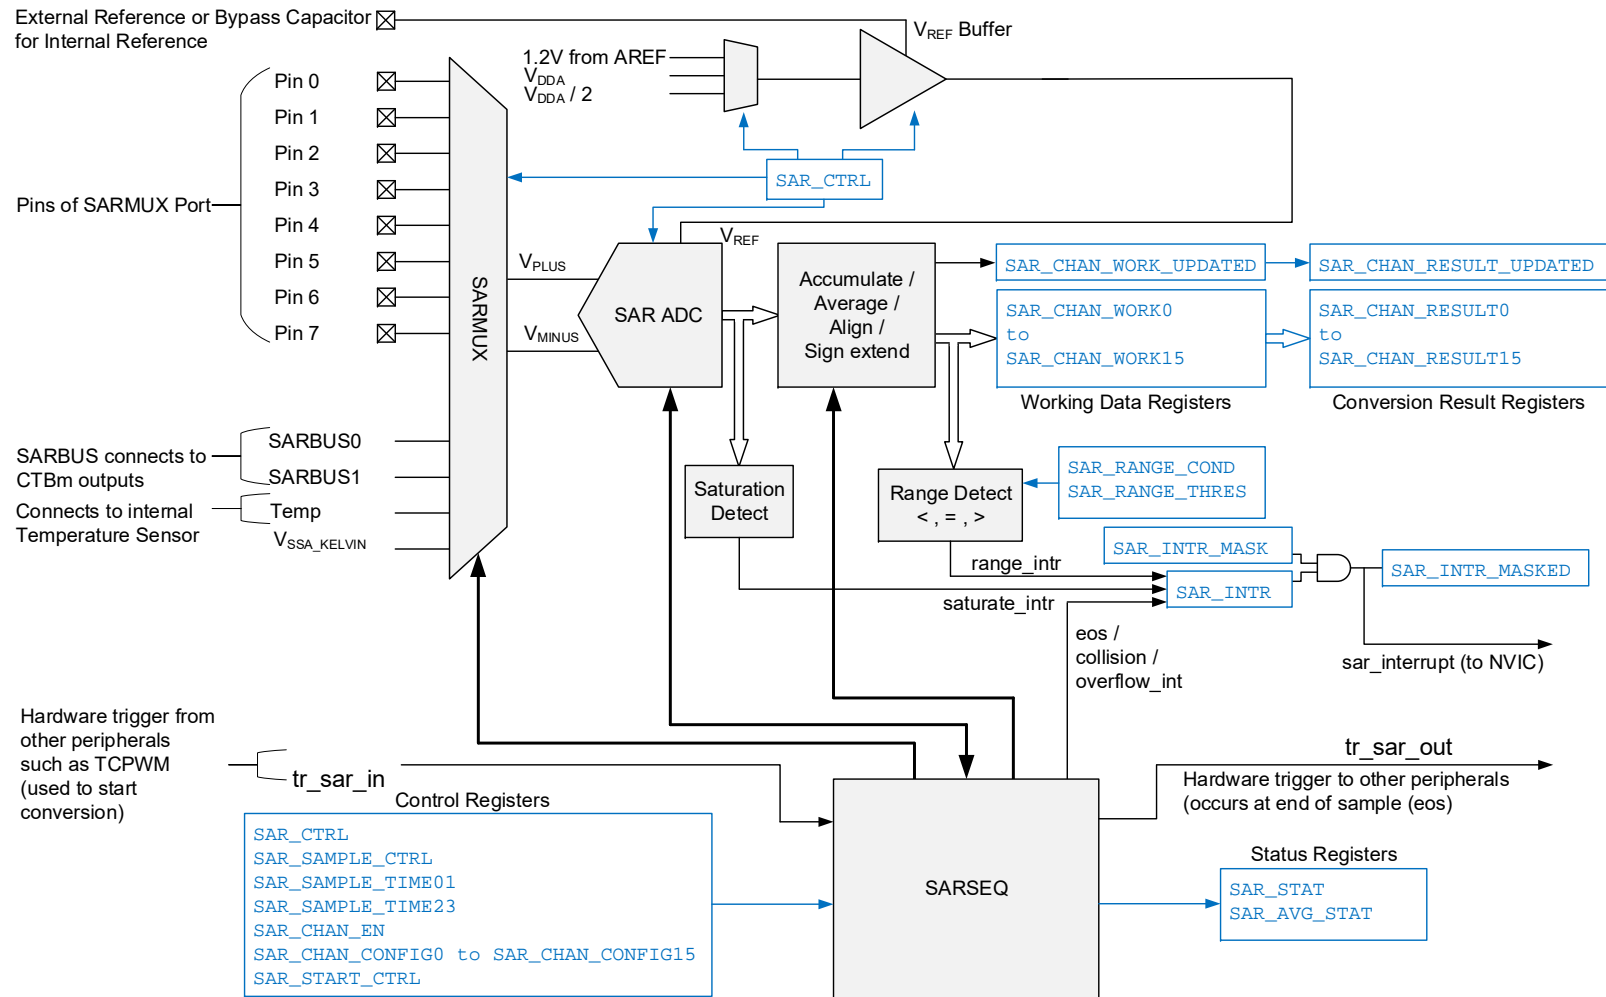


Figure 36-1 shows the simplified block diagram of PSoC 6 MCU SAR ADC system, with important registers shown in blue. Preceding the SAR ADC is the SARMUX, which can route external pins and internal signals (AMUXBUS A/ AMUXBUS B or temperature sensor output) to the internal channels of SAR ADC. The sequencer controller (SARSEQ) is used to control SARMUX and SAR ADC to do an automatic scan on all enabled channels without CPU intervention. SARSEQ also performs pre-processing such as averaging and accumulating the output data.

The result from each channel is double-buffered and a complete scan may be configured to generate an interrupt at the end of the scan. The sequencer may also be configured to flag overflow, collision, and saturation errors that can be configured to assert an interrupt.

36.2.1 SAR ADC Core

The PSoC 6 MCU SAR ADC core is a 12-bit SAR ADC. The maximum sample rate for this ADC is 1 Msps. The SAR ADC core has the following features:

- Fully differential architecture; also supports single-ended mode
- 12-bit resolution
- Four programmable acquisition times
- Seven programmable power levels
- Supports single and continuous conversion mode

SAR_CTRL register contains the bitfields that control the operation of SAR ADC core. See the [registers TRM](#) for more details of this register.

36.2.1.1 Single-ended and Differential Modes

The PSoC 6 MCU SAR ADC can operate in single-ended and differential modes. Differential or single-ended mode can be configured using the DIFFERENTIAL_EN bitfield in the channel configuration register, SAR_CHAN_CONFIGx, where x is the channel number (0–15).

SAR ADC gives full range output (0 to 4095) for differential inputs in the range of $-V_{REF}$ to $+V_{REF}$.

Note: The precise value of the input range in the differential mode is $-V_{REF}$ to $(+V_{REF} - (V_{REF}/2047))$. The positive input range is limited by the resolution of the ADC.

The single-ended mode options of negative input include V_{SSA} , V_{REF} , or an external input from P1, P3, P5, or P7 pins of SARMUX. See the [device datasheet](#) for the exact location of SARMUX pins. This mode is configured by the NEG_SEL bitfield in the global configuration register SAR_CTRL. When V_{MINUS} is connected to these SARMUX pins, the single-ended mode is equivalent to differential mode. However, when the odd pin of each differential pair is connected to the common alternate ground, these conversions are 11-bit because measured signal value cannot go below ground.

To get a single-ended conversion with 12 bits, you must connect V_{REF} to the negative input of the SAR ADC; then, the input range can be from 0 to $2 \times V_{REF}$.

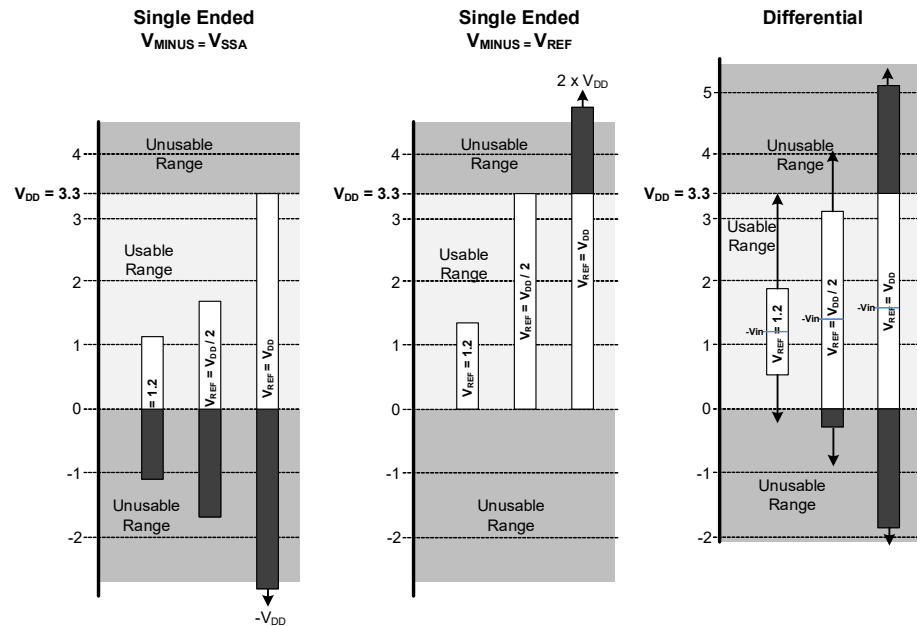
Note that temperature sensor can only be used in single-ended mode.

36.2.1.2 Input Range

All inputs should be in the range of V_{SSA} to V_{DDA} . Input voltage range is also limited by V_{REF} . If voltage on negative input is V_n and the ADC reference is V_{REF} , the range on the positive input is $V_n \pm V_{REF}$. This criterion applies for both single-ended and differential modes. In single-ended mode, V_n is connected to V_{SSA} , V_{REF} or an external input.

Note that $V_n \pm V_{REF}$ should be in the range of V_{SSA} to V_{DDA} . For example, if negative input is connected to V_{SSA} , the range on the positive input is 0 to V_{REF} , not $-V_{REF}$ to V_{REF} . This is because the signal cannot go below V_{SSA} . Only half of the ADC range is usable because the positive input signal cannot swing below V_{SS} , which effectively only generates an 11-bit result.

Figure 36-2. Input Range



36.2.1.3 Result Data Format

Result data format is configurable from two aspects:

- Signed/unsigned
- Left/right alignment

When the result is considered signed, the most significant bit of the conversion is used for sign extension to 16 bits with MSb. For an unsigned conversion, the result is zero extended to 16-bits. The sample value can be either right-aligned or left-aligned within the 16 bits of the result register. By default, data is right-aligned in data[11:0], with sign extension to 16 bits, if required. Left-alignment will cause lower significant bits to be made zero.

Result data format can be controlled by DIFFERENTIAL_SIGNED, SINGLE_ENDED_SIGNED, and LEFT_ALIGN bitfields in the SAR_SAMPLE_CTRL register.

The result data format can be shown as follows.

Table 36-1. Result Data Format

Alignment	Signed/ Unsigned	Result Register															
		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Right	Unsigned	–	–	–	–	11	10	9	8	7	6	5	4	3	2	1	0
Right	Signed	11	11	11	11	11	10	9	8	7	6	5	4	3	2	1	0
Left	–	11	10	9	8	7	6	5	4	3	2	1	0	–	–	–	–

36.2.1.4 Negative Input Selection

The negative input connection choice affects the voltage range and effective resolution (Table 36-2). In single-ended mode, negative input of the SAR ADC can be connected to V_{SSA} , V_{REF} , or P1, P3, P5, or P7 pins of SARMUX. V_x is the common

Table 36-2. Negative Input Selection Comparison

Single-ended/Differential	Signed/Unsigned	V_{MINUS}	V_{PLUS} Range	Result Register
Single-ended	N/A ^a	V_{SSA}	$+V_{REF}$ $V_{SSA} = 0$	0x7FF 0x000
Single-ended	Unsigned	V_{REF}	$+2 \times V_{REF}$ V_{REF} $V_{SSA} = 0$	0xFFFF 0x800 0
Single-ended	Signed	V_{REF}	$+2 \times V_{REF}$ V_{REF} $V_{SSA} = 0$	0x7FF 0x000 0x800
Single-ended	Unsigned	V_x^b	$V_x + V_{REF}$ V_x $V_x - V_{REF}$	0xFFFF 0x800 0
Single-ended	Signed	V_x	$V_x + V_{REF}$ V_x $V_x - V_{REF}$	0x7FF 0x000 0x800
Differential	Unsigned	V_x	$V_x + V_{REF}$ V_x $V_x - V_{REF}$	0xFFFF 0x800 0
Differential	Signed	V_x	$V_x + V_{REF}$ V_x $V_x - V_{REF}$	0x7FF 0x000 0x800

a. For single-ended mode with V_{MINUS} connected to V_{SSA} , conversions are effectively 11-bit because voltages cannot swing below V_{SSA} on any PSoC 6 MCU pin. Because of this, the global configuration bit SINGLE_ENDED_SIGNED (SAR_SAMPLE_CTRL[2]) will be ignored and the result is always (0x000-0x7FF).

b. V_x is the differential input common mode voltage.


mode voltage.

Note that single-ended conversions with V_{MINUS} connected to the pins with SARMUX connectivity are electrically equivalent to differential mode. However, when the odd pin of each differential pair is connected to the common alternate ground, these conversions are 11-bit because measured signal value (SARMUX.vplus) cannot go below ground.

36.2.1.5 Acquisition Time

Acquisition time is the time taken by sample and hold (S/H) circuit inside SAR ADC to settle. After acquisition time, the input signal source is disconnected from the SARADC core, and the output of the S/H circuit will be used for conversion. Each channel can select one from four acquisition time options, from 4 to 1023 SAR clock cycles defined in global configuration registers SAR_SAMPLE_TIME01 and SAR_SAMPLE_TIME23. These two registers contain four programmable sample times ST0, ST1, ST2, and ST3. One of these four sample times can be selected for a particular channel by configuring the SAMPLE_TIME_SEL bitfield in the respective SAR_CHAN_CONFIGx register.

Figure 36-3. Acquisition Time



The acquisition time should be sufficient to charge the internal hold capacitor of the ADC through the resistance of the routing path, as shown in [Figure 36-3](#). The recommended value of acquisition time is:

$$T_{acq} \geq 9 \times (R_{src} + R_{mux} + R_{acqsw}) \times C_{sh}$$

Where:

- R_{src} = Source resistance
- R_{mux} = SARMUX switch resistance
- R_{acqsw} = Sample and hold switch
- $R_{mux} + R_{acqsw} \approx 1000$ ohms
- $C_{sh} \approx 10$ pF

This depends on the routing path (see [Analog Routing on page 460](#) for details).

36.2.1.6 SAR ADC Clock

Note: The maximum SAR ADC clock frequency may be limited to less than 18 MHz by the [device datasheet](#) specification for sample rate maximum.

SAR ADC clock frequency must be between 1.8 MHz and 18 MHz, which comes from the peripheral clock (CLK_PERI) in the system resources subsystem (SRSS). See the [Clocking System chapter on page 207](#) to know how to configure the peripheral clock.


36.2.1.7 SAR ADC Timing

As [Figure 36-4](#) shows, an ADC conversion with the minimum acquisition time of four clocks requires 18 clocks to complete. Note that the minimum acquisition time of four clock cycles at 36 MHz is based on the minimum acquisition time supported by the SAR block (R_{SW1} and C_{SHOLD} in [Figure 36-3 on page 458](#)), which is 97 ns.

Total clock cycles for valid output are equal to:

- 4 clock cycles for sampling input (minimum acquisition time set by SAR_SAMPLE_TIME01 or SAR_SAMPLE_TIME23)
- + 12 clock cycle conversions (with 12-bit resolution)
- + 1 clock for EOF output signal
- + 1 clock for continuous conversion mode and Auto-zero
- = 18 clock cycles.

Figure 36-4. SAR ADC Timing



36.2.2 SARMUX

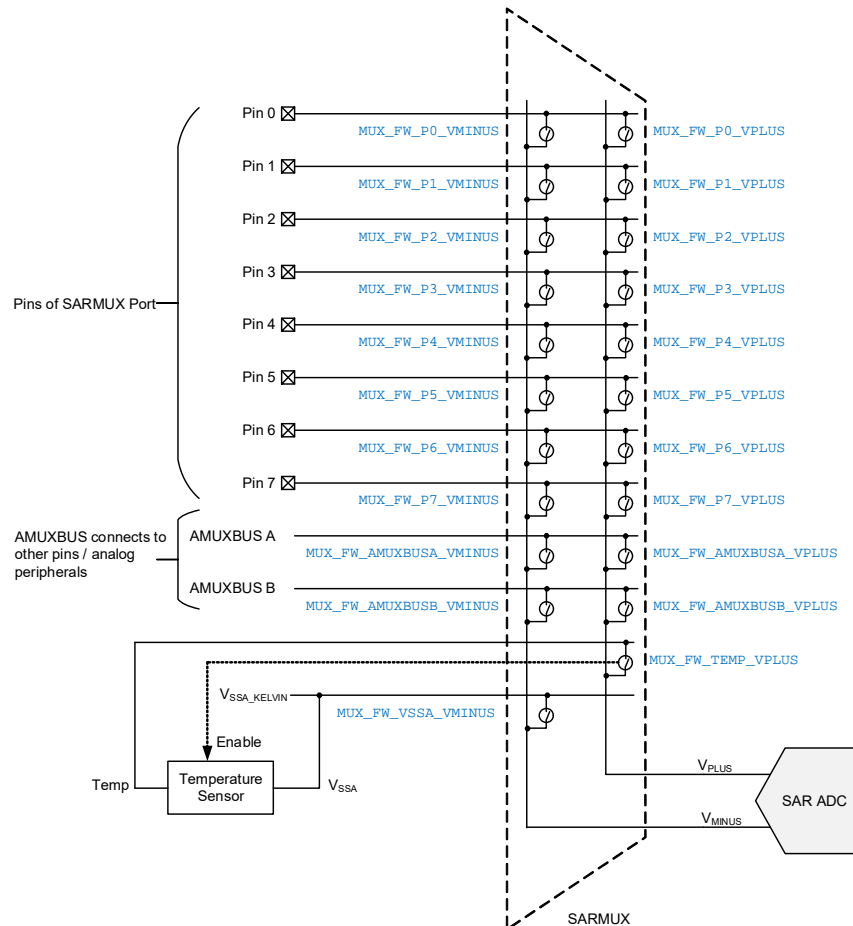
SARMUX is an analog dedicated programmable multiplexer. The main features of SARMUX are:

- Controlled by sequencer controller block (SARSEQ) or firmware
- Internal temperature sensor
- Multiple inputs:
 - Analog signals from pins (port 2)
 - Temperature sensor output (settling time for the temperature sensor is about 1 μ s)
 - AMUXBUS A/AMUXBUS B

36.2.2.1 Analog Routing

SARMUX has many switches that may be controlled by SARSEQ block (sequencer controller) or firmware. Different control methods have different control capability on the switches. [Figure 36-5](#) shows the SARMUX switches. See the [device data-sheet](#) for the exact location of SARMUX pins.

Figure 36-5. SARMUX Switches



Sequencer control: In the sequencer control mode, the SARMUX switches are controlled by the SARSEQ block. After configuring each channel's analog routing, it enables multi-channel, automatic scan in a round-robin fashion, without CPU intervention. Not every switch in analog routing can be controlled by the sequencer, as [Figure 36-5](#) shows.

SAR_MUX_SWITCH_SQ_CTRL register can be used to enable/disable SARSEQ control of SARMUX switches. See section 19.3.4 SARSEQ for more details of sequencer control.

Firmware control: In firmware control, registers are written by the firmware to connect required signals to V_{PLUS}/V_{MINUS} terminals before starting the scan. Firmware can control every switch in SARMUX, as Figure 36-5 shows. However, firmware control needs continuous CPU intervention for multi-channel scans. The SAR_MUX_SWITCH0 register can be used by the firmware to control SARMUX switches. Note that additional register writes may be required to connect blocks outside the SARMUX.

The PSoC 6 MCU analog interconnection is very flexible. SAR ADC can be connected to multiple inputs via SARMUX, including both external pins and internal signals. It can also connect to non-SARMUX ports through AMUXBUS A/AMUXBUS B, at the expense of scanning performance (more parasitic coupling, longer RC time to settle).

Input from SARMUX Port

In this mode, sequencer and firmware control are possible. In addition to SARMUX switch configuration, the GPIOs must be configured properly to connect to SARMUX. See the [I/O System chapter on page 226](#) for more details.

Input from Other Pins through AMUXBUS

Two pins that do not support SARMUX connectivity can be connected to SAR ADC as a differential pair. Additional switches are required to connect these two pins to AMUXBUS A and AMUXBUS B, and then connect AMUXBUS A and AMUXBUS B to the SAR ADC. See the [I/O System chapter on page 226](#) for details of AMUXBUS connections.

The additional switches reduce the scanning performance (more parasitic coupling, longer RC time to settle). This is not recommended for external signals; the dedicated SARMUX port should be used, if possible. Moreover, sequencer control may not be available if more than one AMUXBUS channel is required.

Input from Temperature Sensor

On-chip temperature sensor can be used to measure the device temperature. See the [Temperature Sensor chapter on page 467](#) for more details of this block. For temperature sensors, differential conversions are not available (conversion result is undefined), thus always use it in singled-ended mode. Temperature sensor can be routed to positive input of SAR ADC using a switch that can be controlled by the sequencer, or firmware. Setting the MUX_FW_TEMP_VPLUS bit in the SAR_MUX_SWITCH0 register can enable the temperature sensor and connect its output to V_{PLUS} of SAR ADC; clearing this bit will disable temperature sensor by cutting off its bias current.

36.2.3 SARREF

The main features of the SAR reference block (SARREF) are:

- Reference options: V_{DDA} , $V_{DDA}/2$, 1.2-V bandgap, and external reference
- Reference buffer and bypass capacitor to enhance internal reference drive capability

36.2.3.1 Reference Options

SARREF generates V_{DDA} and $V_{DDA}/2$ voltages. In addition, it can take in a 1.2-V bandgap reference from AREF (see the [Analog Reference Block chapter on page 445](#) for details) or an external V_{REF} connected to a dedicated pin (see the [device datasheet](#) for details). External V_{REF} value should be between 1-V and V_{DDA} .

The external V_{REF} pin is also used to bypass the internal references. The V_{REF_SEL} bitfield in the SAR_CTRL register can be used to select which one of these references is connected to the SAR ADC.

36.2.3.2 Reference Buffer and Bypass Capacitors

The internal references, 1.2 V from bandgap and $V_{DDA}/2$, are buffered with the reference buffer. This reference may be routed to the external V_{REF} pin where a capacitor can be used to filter noise that may exist on the reference.

The VREF_BYP_CAP_EN bitfield in the SAR_CTRL register can be used to enable the bypass capacitor. REFBUF_EN and PWR_CTRL_VREF bitfields in the same register can be used to enable the buffer, and select one of seven available power levels of the reference buffer respectively.

SAR performance varies with the mode of reference and the V_{DDA} supply.

Table 36-3. Reference Modes

Reference Mode	Maximum SAR ADC Clock Frequency	Maximum Sample Rate
External Reference	18 MHz	1 Msps
Internal reference without bypass capacitor	1.8 MHz	100 ksps
Internal reference with bypass capacitor	18 MHz	1 Msps
V_{DDA} as reference	18 MHz	1 Msps

Reference buffer startup time varies with different bypass capacitor size. Table 36-4 lists two common values for the bypass capacitor and its startup time specification. If reference selection is changed between scans or when scanning after device low-power modes in which the ADC is not active (see the [Device Power Modes chapter on page 190](#)), make sure the reference buffer is settled before the SAR ADC starts sampling.

Table 36-4. Bypass Capacitor Values vs Startup Time

Capacitor Value	Startup Time
Internal reference with bypass capacitor (50 nF)	120 μ s
Internal reference with bypass capacitor (100 nF)	210 μ s
Internal reference without bypass capacitor	10 μ s

36.2.3.3 Input Range versus Reference

All inputs should be between V_{SSA} and V_{DDA} . The ADCs input range is limited by V_{REF} selection. If negative input is V_n and the ADC reference is V_{REF} , the range on the positive input is $V_n \pm V_{REF}$. These criteria applies for both single-ended and differential modes as long as both negative and positive inputs stay within V_{SSA} to V_{DDA} .

36.2.4 SARSEQ

SARSEQ is a dedicated control block that automatically sequences the input mux from one channel to the next while placing the result in an array of registers, one per channel. SARSEQ has the following functions:

- Controls SARMUX analog routing automatically without CPU intervention
- Controls SAR ADC core such as selecting acquisition times
- Receives data from SAR ADC and pre-process (average, range detect)
- Results are double-buffered; therefore, the CPU can safely read the results of the last scan while the next scan is in progress.

The features of SARSEQ are:

- Sixteen channels can be individually enabled as an automatic scan without CPU intervention that can scan eleven unique input channels
- Each channel has the following features:
 - Single-ended or differential mode
 - Input from external pin or internal signal (AMUXBUS/temperature sensor)

- One of four programmable acquisition times
- Result averaging and accumulation
- Scan triggering
 - One-shot, periodic, or continuous mode
 - Triggered by any digital signal or input from GPIO pin
 - Triggered by internal block
 - Software triggered
- Hardware averaging support
 - First order accumulate
 - 2, 4, 8, 16, 32, 64, 128, or 256 samples averaging (powers of 2)
 - Results in 16-bit representation
- Double buffering of output data
 - Left or right adjusted results
 - Results in working register and result register
- Interrupt generation
 - Finished scan conversion
 - Channel saturation detection
 - Range (configurable) detection
 - Scan results overflow
 - Collision detect

36.2.4.1 Channel Configuration

The SAR_CHAN_CONFIGx register contains the following bitfields, which control the behavior of respective channels during a SARSEQ scan:

- POS_PORT_ADDR and POS_PIN_ADDR select the connection to V_{PLUS} terminal of the ADC

- **NEG_ADDR_EN**, **NEG_PORT_ADDR**, and **NEG_PIN_ADDR** select the connection to V_{MINUS} terminal of the ADC
- **SAMPLE_TIME_SEL** selects the acquisition time for the channel
- **AVG_EN** enables the hardware averaging feature
- **DIFFERENTIAL_EN** selects single-ended/differential mode

SAR_CHAN_EN contains the enable bits that can be used to include or exclude a channel for the next SARSEQ scan.

36.2.4.2 Averaging

The SARSEQ block has a 20-bit accumulator and shift register to implement averaging. Averaging is performed after sign extension. The **SAR_SAMPLE_CTRL** register controls the global averaging settings.

Each channel configuration register (**SAR_CHAN_CONFIG**) has an enable bit (**AVG_EN**) to enable averaging.

The **AVG_CNT** bitfield in the **SAR_SAMPLE_CTRL** register specifies the number of accumulated samples (N) according to the formula:

$$N = 2^{(\text{AVG_CNT} + 1)} \quad N \text{ range} = [2..256]$$

For example, if **AVG_CNT** = 3, then **N** = 16.

Because the conversion result is 12-bit and the maximum value of **N** is 256 (left shift 8 bits), the 20-bit accumulator will never overflow.

The **AVG_SHIFT** bitfield in the **SAR_SAMPLE_CTRL** register is used to shift the accumulated result to get the averaged value. If this bit is set, the SAR sequencer performs sign extension and then accumulation. The accumulated result is then shifted right **AVG_CNT** + 1 bits to get the averaged result.

The **AVG_MODE** bitfield in the **SAR_SAMPLE_CTRL** can be used to select between accumulate-and-dump and interleaved averaging modes. If a channel is configured for accumulate-and-dump averaging, the SARSEQ will take **N** consecutive samples of the specified channel before moving to the next channel. In the interleaved mode, one sample is taken per channel and averaged over several scans.

36.2.4.3 Range Detection

The SARSEQ supports range detection to allow automatic detection of result values compared to two programmable thresholds without CPU involvement. Range detection is defined by the **SAR_RANGE_THRES** register. The **RANGE_LOW** field in the **SAR_RANGE_THRES** register defines the lower threshold and **RANGE_HIGH** field defines the upper threshold of the range.

The **RANGE_COND** bitfield in the **SAR_RANGE_COND** register define the condition that triggers a channel mas-

kable range detect interrupt (**RANGE_INTR**). The following conditions can be selected:

0: result < **RANGE_LOW** (below the range)

1: **RANGE_LOW** ≤ result < **RANGE_HIGH** (inside the range)

2: **RANGE_HIGH** ≤ result (above the range)

3: result < **RANGE_LOW** || **RANGE_HIGH** ≤ result (outside range)

See [Range Detection Interrupts on page 464](#) for details.

36.2.4.4 Double Buffer

Double buffering is used so that firmware can read the results of a complete scan while the next scan is in progress. The SAR ADC results are written to a set of working registers until the scan is complete, at which time the data is copied to a second set of registers where the data can be read by your application. This action allows sufficient time for the firmware to read the previous scan before the present scan is completed. All input channels are double buffered with 16 registers.

SAR_CHAN_WORKx registers contain the working data and **SAR_CHAN_RESULTx** contain the buffered results of the channels. The **SAR_CHAN_WORK_UPDATED** and **SAR_CHAN_RESULT_UPDATED** registers can be used to track if the working data and the result value of a channel is updated.

36.2.5 SAR Interrupts

SAR ADC can generate interrupts on these events:

- End of Scan – When scanning is complete for all the enabled channels.
- Overflow – When the result register is updated before the previous result is read.
- Collision – When a new trigger is received while the SAR ADC is still processing the previous trigger.
- Range Detection – When the channel result meets the threshold value.
- Saturation Detection – When the channel result is equal to the minimum or maximum value.

This section describes each interrupt in detail. These interrupts have an interrupt mask in the **SAR_INTR_MASK** register. By making the interrupt mask low, the corresponding interrupt source is ignored. The SAR interrupt is generated if the interrupt mask bit is high and the corresponding interrupt source is pending.

When servicing an interrupt, the interrupt service routine (ISR) can clear the interrupt source by writing a '1' to the corresponding interrupt bit in the **SAR_RANGE_INTR** register, after reading the data.

The SAR_INTR_MASKED register is the logical AND between the interrupts sources and the interrupt mask. This register provides a convenient way for the firmware to determine the source of the interrupt.

For verification and debug purposes, a set bit (such as EOS_SET in the SAR_INTR_SET register) is used to trigger each interrupt. This action allows the firmware to generate an interrupt without the actual event occurring.

36.2.5.1 End-of-Scan Interrupt (EOS_INTR)

After completing a scan, the end-of-scan interrupt (EOS_INTR) is raised. Firmware should clear this interrupt after picking up the data from the RESULT registers.

EOS_INTR can be masked by making the EOS_MASK bit 0 in the SAR_INTR_MASK register. EOS_MASKED bit of the SAR_INTR_MASKED register is the logic AND of the interrupt flags and the interrupt masks. Writing a '1' to EOS_SET bit in SAR_INTR_SET register can set the EOS_INTR, which is intended for debug and verification.

36.2.5.2 Overflow Interrupt

If a new scan completes and the hardware tries to set the EOS_INTR and EOS_INTR is still high (firmware does not clear it fast enough), then an overflow interrupt (OVERFLOW_INTR) is generated by the hardware. This usually means that the firmware is unable to read the previous results before the current scan completes. In this case, the old data is overwritten.

OVERFLOW_INTR can be masked by making the OVERFLOW_MASK bit 0 in SAR_INTR_MASK register. OVERFLOW_MASKED bit of SAR_INTR_MASKED register is the logic AND of the interrupt flags and the interrupt masks, which are for firmware convenience. Writing a '1' to the OVERFLOW_SET bit in SAR_INTR_SET register can set OVERFLOW_INTR, which is intended for debug and verification.

36.2.5.3 Collision Interrupt

It is possible that a new trigger is generated while the SARSEQ is still busy with the scan started by the previous trigger. Therefore, the scan for the new trigger is delayed until after the ongoing scan is completed. It is important to notify the firmware that the new sample is invalid. This is done through the collision interrupt, which is raised any time a new trigger, other than the continuous trigger, is received.

The collision interrupt for the firmware trigger (FW_COLLISION_INTR) allows the firmware to identify which trigger collided with an ongoing scan.

The collision interrupts can be masked by making the corresponding bit '0' in the SAR_INTR_MASK register. The corresponding bit in the SAR_INTR_MASKED register is the logic AND of the interrupt flags and the interrupt masks. Writing a '1' to the corresponding bit in SAR_INTR_SET register can

set the collision interrupt, which is intended for debug and verification.

36.2.5.4 Range Detection Interrupts

Range detection interrupt flag can be set after averaging, alignment, and sign extension (if applicable). This means it is not required to wait for the entire scan to complete to determine whether a channel conversion is over-range. The threshold values need to have the same data format as the result data.

Range detection interrupt for a specified channel can be masked by setting the SAR_RANGE_INTR_MASK register specified bit to '0'. Register SAR_RANGE_INTR_MASKED reflects a bitwise AND between the interrupt request and mask registers. If the value is not zero, then the SAR interrupt signal to the NVIC is high.

SAR_RANGE_INTR_SET can be used for debug/verification. Write a '1' to set the corresponding bit in the interrupt request register; when read, this register reflects the interrupt request register.

There is a range detect interrupt for each channel (RANGE_INTR and INJ_RANGE_INTR).

36.2.5.5 Saturate Detection Interrupts

The saturation detection is always applied to every conversion. This feature detects if a sample value is equal to the minimum or maximum value and sets a maskable interrupt flag for the corresponding channel. This action allows the firmware to take action, such as discarding the result, when the SAR ADC saturates. The sample value is tested right after conversion, before averaging. This means that the interrupt is set while the averaged result in the data register is not equal to the minimum or maximum.

Saturation interrupt flag is set immediately to enable a fast response to saturation, before the full scan and averaging. Saturation detection interrupt for specified channel can be masked by setting the SAR_SATURATE_INTR_MASK register specified bit to '0'. SAR_SATURATE_INTR_MASKED register reflects a bit-wise AND between the interrupt request and mask registers. If the value is not zero, then the SAR interrupt signal to the NVIC is high.

SAR_SATURATE_INTR_SET can be used for debug/verification. Write a '1' to set the corresponding bit in the interrupt request register; when read, this register reflects the interrupt request register.

36.2.5.6 Interrupt Cause Overview

INTR_CAUSE register contains an overview of all the pending SAR interrupts. It allows the ISR to determine the cause of the interrupt. The register consists of a mirror copy of SAR_INTR_MASKED. In addition, it has two bits that aggregate the range and saturate detection interrupts of all channels. It includes a logical OR of all the bits in

RANGE_INTR_MASKED and SATURATE_INTR_MASKED registers (does not include INJ_RANGE_INTR and INJ_SATURATE_INTR).

36.2.6 Trigger

A scan can be triggered in the following ways:

- A firmware or one-shot trigger is generated when the firmware writes to the FW_TRIGGER bit of the SAR_START_CTRL register. After the scan is completed, the SARSEQ clears the FW_TRIGGER bit and goes back to idle mode waiting for the next trigger. The FW_TRIGGER bit is cleared immediately after the SAR is disabled.
- Trigger from other peripherals through the trigger multiplexer (see the [Trigger Multiplexer Block chapter on page 260](#) for more details).
- A continuous trigger is activated by setting the CONTINUOUS bit in SAR_SAMPLE_CTRL register. In this mode, after completing a scan the SARSEQ starts the next scan immediately; therefore, the SARSEQ is always BUSY. As a result, all other triggers are essentially ignored. Note that FW_TRIGGER will still get cleared by hardware on the next completion.

For firmware continuous trigger, it takes only one SAR ADC clock cycle before the sequencer tells the SAR ADC to start sampling (provided the sequencer is idle).

36.2.7 SAR ADC Status

The current SAR status can be observed through the BUSY and CUR_CHAN fields in the SAR_STATUS register. The BUSY bit is high whenever the SAR is busy sampling or converting a channel; the CUR_CHAN bits indicate the number of the current channel being sampled. SW_VREF_NEG bit indicates the current switch status, including register controls, of the switch in the SAR ADC that shorts NEG with V_{REF} input.

The CUR_AVG_ACCU and CUR_AVG_CNT fields in the SAR_AVG_STAT register indicate the current averaging accumulator contents and the current sample counter value for averaging (counts down).

The SAR_MUX_SWITCH_STATUS register gives the current switch status of MUX_SWITCH0 register. These status registers help to debug SAR behavior.

36.3 Registers

Name	Description
SAR_CTRL	Global configuration register. Analog control register
SAR_SAMPLE_CTRL	Global configuration register. Sample control register
SAR_SAMPLE_TIME01	Global configuration register. Sample time specification ST0 and ST1
SAR_SAMPLE_TIME23	Global configuration register. Sample time specification ST2 and ST3
SAR_RANGE_THRES	Global range detect threshold register
SAR_RANGE_COND	Global range detect mode register
SAR_CHAN_EN	Enable bits for the channels
SAR_START_CTRL	Start control register (firmware trigger)
SAR_CHAN_CONFIGx	Channel configuration register. There are 16 such registers with x = 0 to 15
SAR_CHAN_WORKx	Channel working data register. There are 16 such registers with x = 0 to 15
SAR_CHAN_RESULTx	Channel result data register. There are 16 such registers with x = 0 to 15
SAR_CHAN_WORK_UPDATED	Channel working data register: updated bits
SAR_CHAN_RESULT_UPDATED	Channel result data register: updated bits
SAR_STATUS	Current status of internal SAR registers (for debug)
SAR_AVG_STAT	Current averaging status (for debug)
SAR_INTR	Interrupt request register
SAR_INTR_SET	Interrupt set request register
SAR_INTR_MASK	Interrupt mask register
SAR_INTR_MASKED	Interrupt masked request register
SAR_SATURATE_INTR	Saturate interrupt request register
SAR_SATURATE_INTR_SET	Saturate interrupt set request register
SAR_SATURATE_INTR_MASK	Saturate interrupt mask register
SAR_SATURATE_INTR_MASKED	Saturate interrupt masked request register
SAR_RANGE_INTR	Range detect interrupt request register
SAR_RANGE_INTR_SET	Range detect interrupt set request register
SAR_RANGE_INTR_MASK	Range detect interrupt mask register
SAR_RANGE_INTR_MASKED	Range interrupt masked request register
SAR_INTR_CAUSE	Interrupt cause register
SAR_MUX_SWITCH0	SARMUX firmware switch controls
SAR_MUX_SWITCH_CLEAR0	SARMUX firmware switch control clear
SAR_MUX_SWITCH_SQ_CTRL	SARMUX switch sequencer control
SAR_MUX_SWITCH_STATUS	SARMUX switch status

37. Temperature Sensor



This PSoC 6 MCU technical reference manual (TRM) provides comprehensive and detailed information about the functions of the PSoC 6 MCU device hardware. It is divided into two books: architecture TRM and registers TRM. The TRM is not recommended for those new to the PSoC 6 MCU, nor as a guide for developing PSoC 6 MCU applications. Use these documents instead:

- [Device datasheet](#)
- Peripheral Driver Library (PDL) documentation - SAR ADC, SysAnalog
- Application notes
- Code examples

PSoC 6 MCUs have an on-chip temperature sensor that is used to measure the internal die temperature. The sensor consists of a transistor connected in diode configuration.


37.1 Features

- Measures device temperature
- Voltage output can be internally connected to SAR ADC for digital readout
- Factory calibrated parameters

37.2 Architecture

The temperature sensor consists of a single bipolar junction transistor (BJT) in the form of a diode. The transistor is biased using a reference current I_{REF} from the analog reference block (see the [Analog Reference Block chapter on page 445](#) for more details). Its base-to-emitter voltage (V_{BE}) has a strong dependence on temperature at a constant collector current and zero collector-base voltage. This property is used to calculate the die temperature by measuring the V_{BE} of the transistor using the SAR ADC, as shown in [Figure 37-1](#).

Figure 37-1. Temperature Sensing Mechanism



The analog output from the sensor (V_{BE}) is measured using the SAR ADC. Die temperature in $^{\circ}\text{C}$ can be calculated from the ADC results as given in the following equation:

$$\text{Temp} = (A \times \text{SAR}_{\text{out}} + 2^{10} \times B) + T_{\text{adjust}} \quad \text{Equation 37-1}$$

- Temp is the slope compensated temperature in $^{\circ}\text{C}$ represented as Q16.16 fixed point number format.
- 'A' is the 16-bit multiplier constant. The value of A is determined using the PSoC 6 MCU characterization data of two point slope calculation. It is calculated as given in the following equation.

$$A = (\text{signed int}) \left(2^{16} \left(\frac{100^{\circ}\text{C} - (-40^{\circ}\text{C})}{\text{SAR}_{100^{\circ}\text{C}} - \text{SAR}_{-40^{\circ}\text{C}}} \right) \right) \quad \text{Equation 37-2}$$

Where,

$\text{SAR}_{100^{\circ}\text{C}}$ = ADC counts at 100°C

$\text{SAR}_{-40^{\circ}\text{C}}$ = ADC counts at -40°C

Constant 'A' is stored in a register SFLASH_SAR_TEMP_MULTIPLIER. See the [registers TRM](#) for more details.

- 'B' is the 16-bit offset value. The value of B is determined on a per die basis by taking care of all the process variations and the actual bias current (I_{REF}) present in the chip. It is calculated as given in the following equation.

$$B = (\text{unsigned int}) \left(2^6 \times 100^{\circ}\text{C} - \left(\frac{A \times \text{SAR}_{100^{\circ}\text{C}}}{2^{10}} \right) \right) \quad \text{Equation 37-3}$$

Where,

$\text{SAR}_{100^{\circ}\text{C}}$ = ADC counts at 100°C

Constant 'B' is stored in a register SFLASH_SAR_TEMP_OFFSET. See the [registers TRM](#) for more details.

- T_{adjust} is the slope correction factor in $^{\circ}\text{C}$. The temperature sensor is corrected for dual slopes using the slope correction factor. It is evaluated based on the result obtained without slope correction, which is given by the following equation:

$$T_{\text{initial}} = (A \times \text{SAR}_{\text{OUT}} + 2^{10} \times B) \quad \text{Equation 37-4}$$

If T_{initial} is greater than the center value (15°C), then T_{adjust} is given by the following equation.

$$T_{\text{adjust}} = \left(\frac{0.5^{\circ}\text{C}}{100^{\circ}\text{C} - 15^{\circ}\text{C}} \times (100^{\circ}\text{C} \times 2^{16} - T_{\text{initial}}) \right)$$


Equation 37-5

If T_{initial} is less than center value, then T_{adjust} is given by the following equation.

$$T_{\text{adjust}} = \left(\frac{0.5^{\circ}\text{C}}{40^{\circ}\text{C} + 15^{\circ}\text{C}} \times (40^{\circ}\text{C} \times 2^{16} - T_{\text{initial}}) \right)$$

Equation 37-6

Figure 37-2. Temperature Error Compensation



Note: A and B are 16-bit constants stored in flash during factory calibration. These constants are valid only when the SAR ADC is running at 12-bit resolution with a 1.2-V reference.

37.3 Temperature Sensor Configuration

The temperature sensor output is routed to the positive input of SAR ADC via dedicated switches, which can be controlled by the sequencer firmware. See the [SAR ADC chapter on page 453](#) for details on how to read the temperature sensor output using the ADC. The temperature sensor is switched ON only when it is connected to the positive input of the SAR ADC.

37.4 Algorithm

1. Configure SAR ADC in single-ended mode with $V_{\text{NEG}} = V_{\text{SS}}$, $V_{\text{REF}} = 1.2 \text{ V}$, 12-bit resolution, and right-aligned result. See the [SAR ADC chapter on page 453](#) for more details.
2. Enable the SARMUX and SAR ADC. See the [SAR ADC chapter on page 453](#) for more details.
3. Enable the temperature sensor by setting the SAR_MUX_FW_TEMP_VPLUS bit in the SAR_MUX_SWITCH0 register.
4. Trigger a SAR ADC conversion and get the digital output from the SAR ADC. See the [SAR ADC chapter on page 453](#) for more details.
5. Fetch 'A' from SFLASH_SAR_TEMP_MULTIPLIER and 'B' from SFLASH_SAR_TEMP_OFFSET.
6. Calculate the die temperature using the following equation:

$$\text{Temp} = (A \times \text{SAR}_{\text{OUT}} + 2^{10} \times B) + T_{\text{ADJUST}}$$

For example, let $A = 0x\text{BC4B}$ and $B = 0x\text{65B4}$. Assume that the output of SAR ADC (V_{BE}) is $0x\text{595}$ at a given temperature. Firmware does the following calculations:

- a. Multiply A and V_{BE} : $0x\text{BC4B} \times 0x\text{595} = (-17333)_{10} \times (1429)_{10} = (-24768857)_{10}$
- b. Multiply B and 1024: $0x\text{65B4} \times 0x\text{400} = (26036)_{10} \times (1024)_{10} = (26660864)_{10}$
- c. Add the result of steps 1 and 2 to get T_{initial} : $(-24768857)_{10} + (26660864)_{10} = (1892007)_{10} = 0x\text{1CDEA7}$
- d. Calculate T_{adjust} using T_{initial} value: T_{initial} is the upper 16 bits multiplied by 2^{16} , that is, $0x\text{1C00} = (1835008)_{10}$. It is greater than 15°C ($0x\text{1C}$ - upper 16 bits). Use Equation 4 to calculate T_{adjust} . It comes to $0x\text{6C6C} = (27756)_{10}$
- e. Add T_{adjust} to T_{initial} : $(1892007)_{10} + (27756)_{10} = (1919763)_{10} = 0x\text{1D4B13}$
- f. The integer part of temperature is the upper 16 bits = $0x\text{001D} = (29)_{10}$

- g. The decimal part of temperature is the lower 16 bits = $0x4B13 = (0.19219)_{10}$
- h. Combining the result of steps f and g, Temp = $29.19219\text{ }^{\circ}\text{C} \sim 29.2^{\circ}\text{C}$

37.5 Registers

Name	Description
SFLASH_SAR_TEMP_MULTIPLIER	Multiplier constant 'A' as defined in Equation 37-1 .
SFLASH_SAR_TEMP_OFFSET	Constant 'B' as defined in Equation 37-1 .

38. CapSense



The CapSense system can measure the self-capacitance of an electrode or the mutual capacitance between a pair of electrodes. In addition to capacitive sensing, the CapSense system can function as an ADC to measure voltage on any GPIO pin that supports the CapSense functionality.

The CapSense touch sensing method in PSoC 6 MCUs, which senses self-capacitance, is known as CapSense Sigma Delta (CSD). Similarly, the mutual-capacitance sensing method is known as CapSense Cross-point (CSX). The CSD and CSX touch sensing methods provide the industry's best-in-class signal-to-noise ratio (SNR), high touch sensitivity, low-power operation, and superior EMI performance.

CapSense touch sensing is a combination of hardware and firmware techniques. Therefore, use the CapSense component provided by the ModusToolbox IDE to implement CapSense designs. See the [PSoC 4 and PSoC 6 MCU CapSense Design Guide](#) for more details.