

# TriCore<sup>®</sup> 1

## 32-bit Unified Processor Core

### Volume 1

### Core Architecture

### V1.3 & V1.3.1 Architecture

# 32bit

## Microcontrollers



Never stop thinking

**Edition 2008-01**

**Published by  
Infineon Technologies AG  
81726 Munich, Germany**

**© 2008 Infineon Technologies AG  
All Rights Reserved.**

#### **Legal Disclaimer**

The information given in this document shall in no event be regarded as a guarantee of conditions or characteristics. With respect to any examples or hints given herein, any typical values stated herein and/or any information regarding the application of the device, Infineon Technologies hereby disclaims any and all warranties and liabilities of any kind, including without limitation, warranties of non-infringement of intellectual property rights of any third party.

#### **Information**

For further information on technology, delivery terms and conditions and prices, please contact your nearest Infineon Technologies Office ([www.infineon.com](http://www.infineon.com)).

#### **Warnings**

Due to technical requirements components may contain dangerous substances. For information on the types in question please contact the nearest Infineon Technologies Office.

Infineon Technologies components may be used in life-support devices or systems only with the express written approval of Infineon Technologies, if a failure of such components can reasonably be expected to cause the failure of that life-support device or system or to affect the safety or effectiveness of that device or system. Life support devices or systems are intended to be implanted in the human body, or to support and/or maintain and sustain and/or protect human life. If they fail, it is reasonable to assume that the health of the user or other persons may be endangered.

# TriCore<sup>®</sup> 1

## 32-bit Unified Processor Core

### Volume 1

### Core Architecture

### V1.3 & V1.3.1 Architecture

**Microcontrollers**



Never stop thinking

---

Previous Version: V1.3.6

---

Version	Subjects (major changes since last revision)
	<p>The Instruction Set Overview chapter was missing from Volume 2 of the initial release of this document (v1.3.8), dated 2007-11.</p> <p>This version (dated 2008-01) supercedes that release.</p> <p>There are no other changes to the document (vol1 or vol2) aside from the inclusion of the Instruction Set Overview chapter.</p>

---

TriCore® is a registered trademark of Infineon Technologies AG.

### We Listen to Your Comments

Is there any information in this document that you feel is wrong, unclear or missing?  
Your feedback will help us to continuously improve the quality of our documentation.  
Please send feedback (including a reference to this document) to:

**ipdoc@infineon.com**



<b>Table of Contents</b>		<b>Page</b>
<b>1</b>	<b>Architecture Overview</b>	<b>1-1</b>
1.1	Introduction	1-1
1.1.1	Feature Summary	1-2
1.2	Programming Model	1-2
1.2.1	Architectural Registers	1-3
1.2.2	Data Types	1-4
1.2.3	Memory Model	1-4
1.2.4	Addressing Modes	1-6
1.3	Tasks and Contexts	1-7
1.4	Interrupt System	1-8
1.4.1	Interrupt Priority	1-8
1.5	Trap System	1-9
1.6	Protection System	1-9
1.7	Memory Management Unit	1-10
1.8	Core Debug Controller	1-11
1.9	Coprocessor Interface	1-11
<b>2</b>	<b>Programming Model</b>	<b>2-1</b>
2.1	Data Types	2-1
2.1.1	Boolean	2-1
2.1.2	Bit String	2-1
2.1.3	Byte	2-1
2.1.4	Signed Fraction	2-2
2.1.5	Address	2-2
2.1.6	Signed and Unsigned Integers	2-2
2.1.7	IEEE-754 Single-Precision Floating-Point Number	2-2
2.2	Data Formats	2-2
2.2.1	Alignment Requirements	2-4
2.2.2	Byte Ordering	2-5
2.3	Memory Model	2-6
2.4	Semaphores and Atomic Operations	2-7
2.5	Addressing Modes	2-7
2.5.1	Absolute Addressing	2-8
2.5.2	Base + Offset Addressing	2-8
2.5.3	Pre-Increment and Pre-Decrement Addressing	2-9
2.5.4	Post-Increment and Post-Decrement Addressing	2-9
2.5.5	Circular Addressing	2-9
2.5.6	Bit-Reverse Addressing	2-12
2.5.7	Synthesized Addressing Modes	2-13

<b>Table of Contents</b>	<b>Page</b>
<b>3 General Purpose and System Registers</b>	<b>3-1</b>
3.1 General Purpose Registers (GPRs)	3-2
3.1.1 Data General Purpose Registers	3-3
3.1.2 Address General Purpose Registers	3-3
3.2 Program State Information Registers	3-5
3.2.1 Program Counter (PC)	3-5
3.2.2 Program Status Word Register (PSW)	3-6
3.2.3 Previous Context Information and Pointer Register (PCXI)	3-12
3.3 Stack Management Registers	3-13
3.3.1 Address Register A[10] (SP)	3-14
3.3.2 Interrupt Stack Pointer (ISP)	3-15
3.4 System Control Register (SYSCON)	3-16
3.5 CPU Identification Register (CPU_ID)	3-17
3.6 Compatibility Mode Register (COMPAT)	3-18
3.7 Access Control Registers	3-19
3.7.1 BIST Mode Access Control Register (BMACON)	3-19
3.7.2 SIST Mode Access Control Register (SMACON)	3-20
3.8 Interrupt Registers	3-21
3.9 Memory Protection Registers	3-21
3.10 Trap Registers	3-21
3.11 Memory Management Unit Registers	3-21
3.12 Core Debug Controller Registers	3-21
3.13 Floating Point Registers (TriCore 1.3.1)	3-21
3.14 Updating Core Special Function Registers (CSFRs)	3-22
<b>4 Tasks and Functions</b>	<b>4-1</b>
4.1 Context Types	4-1
4.1.1 Context Save Area	4-3
4.2 Task Switching Operation	4-4
4.2.1 Save and Restore Context Operations	4-5
4.3 Context Save Areas (CSAs) and Context Lists	4-5
4.4 Context Switching with Interrupts and Traps	4-7
4.5 Context Switching for Function Calls	4-8
4.6 Context Save and Restore Examples	4-9
4.6.1 Context Save	4-9
4.6.2 Context Restore	4-11
4.7 Context Management Registers	4-13
4.7.1 Free CSA List Head Pointer Register (FCX)	4-14
4.7.2 Previous Context Pointer Register (PCX)	4-15
4.7.3 Free CSA List Limit Pointer Register (LCX)	4-16
4.8 Accessing CSA Memory Locations	4-17

<b>Table of Contents</b>	<b>Page</b>
<b>5 Interrupt System</b>	<b>5-1</b>
5.1 Service Request Node (SRN)	5-1
5.1.1 Service Request Control Register (SRC)	5-3
5.2 Interrupt Control Unit (ICU)	5-6
5.2.1 ICU Interrupt Control Register (ICR)	5-7
5.2.2 Interrupt Control Unit Operation	5-7
5.2.3 Arbitration Scheme	5-8
5.3 Entering an Interrupt Service Routine (ISR)	5-8
5.4 Exiting an Interrupt Service Routine (ISR)	5-9
5.5 Interrupt Vector Table	5-10
5.6 Using the TriCore Interrupt System	5-12
5.6.1 Spanning Interrupt Service Routines across Vector Entries	5-12
5.6.2 Interrupt Priority Groups	5-12
5.6.3 Dividing ISRs into Different Priorities	5-14
5.6.4 Using Different Priorities for the Same Interrupt Source	5-14
5.6.5 Software-Posted Interrupts	5-15
5.6.6 Interrupt Priority Level One	5-15
<b>6 Trap System</b>	<b>6-1</b>
6.1 Trap Types	6-1
6.1.1 Synchronous Traps	6-3
6.1.2 Asynchronous Traps	6-3
6.1.3 Hardware Traps	6-3
6.1.4 Software Traps	6-3
6.1.5 Unrecoverable Traps	6-4
6.2 Trap Handling	6-4
6.2.1 Trap Vector Format	6-4
6.2.2 Accessing the Trap Vector Table	6-4
6.2.3 Return Address (RA)	6-4
6.2.4 Trap Vector Table	6-5
6.2.5 Initial State upon a Trap	6-6
6.3 Trap Descriptions	6-7
6.3.1 MMU Traps (Trap Class 0)	6-7
6.3.2 Internal Protection Traps (Trap Class 1)	6-7
6.3.3 Instruction Errors (Trap Class 2)	6-8
6.3.4 Context Management (Trap Class 3)	6-10
6.3.5 System Bus and Peripheral Errors (Trap Class 4)	6-12
6.3.6 Assertion Traps (Trap Class 5)	6-14
6.3.7 System Call (Trap Class 6)	6-14
6.3.8 Non-Maskable Interrupt (Trap Class 7)	6-14
6.3.9 Debug Traps	6-14
6.4 Exception Priorities	6-15
6.5 Interrupt and Trap Control Registers	6-17

<b>Table of Contents</b>	<b>Page</b>
6.5.1 ICU Interrupt Control Register (ICR) .....	6-17
6.5.2 Base Interrupt Vector Table Pointer (BIV) .....	6-19
6.5.3 Base Trap Vector Table Pointer (BTv) .....	6-20
<b>7 Memory Integrity Error Mitigation (TriCore 1.3.1) .....</b>	<b>7-1</b>
7.1 Memory Integrity Error Classification .....	7-1
7.2 Memory Integrity Error Traps .....	7-2
7.2.1 Program Memory Integrity Error (PIE) .....	7-2
7.2.2 Data Memory Integrity Error (DIE) .....	7-2
7.3 Corrected Error Counts .....	7-3
7.3.1 Count of Corrected Program Memory Integrity Errors Register .....	7-3
7.3.2 Count of Corrected Data Integrity Errors Register .....	7-4
7.4 Error Information Registers .....	7-5
7.4.1 Program Integrity Error Trap Register (PIETR) .....	7-5
7.4.2 Program Integrity Error Address Register (PIEAR) .....	7-6
7.4.3 Data Integrity Error Trap Register (DIETR) .....	7-7
7.4.4 Data Integrity Error Address Register (DIEAR) .....	7-8
7.4.5 Memory Integrity Error Control Register .....	7-9
7.5 Summary .....	7-10
<b>8 Physical Memory Attributes (PMA) .....</b>	<b>8-1</b>
8.1 Physical Memory Properties (PMP) .....	8-1
8.2 Physical Memory Attributes (PMA) .....	8-3
8.2.1 Physical Memory Attributes of the Address Map .....	8-3
8.3 Scratchpad RAM .....	8-4
8.3.1 Scratchpad RAM (TriCore 1.3) .....	8-4
8.3.2 Scratchpad RAM (TriCore 1.3.1) .....	8-4
8.4 Permitted versus Valid Accesses .....	8-5
<b>9 Memory Protection System .....</b>	<b>9-1</b>
9.1 Memory Protection Subsystems .....	9-1
9.2 Range Based Memory Protection .....	9-2
9.2.1 Memory Protection Register Sets .....	9-2
9.3 Memory Protection Registers .....	9-6
9.3.1 Data Segment Protection Register - Upper .....	9-6
9.3.2 Data Segment Protection Register - Lower .....	9-7
9.3.3 Code Segment Protection Register - Upper .....	9-8
9.3.4 Code Segment Protection Register - Lower .....	9-9
9.3.5 Data Protection Mode Register .....	9-10
9.3.6 Code Protection Mode Register .....	9-12
9.4 Access Permissions for Intersecting Memory Ranges .....	9-14
9.4.1 Example Data Protection Register Set .....	9-14
9.5 Using the Memory Protection System .....	9-16
9.5.1 Protection Enable bit .....	9-16



<b>Table of Contents</b>	<b>Page</b>
9.5.2 Set Selection .....	9-16
9.5.3 Address Range .....	9-16
9.5.4 Traps .....	9-17
9.6 Crossing Protection Boundaries .....	9-17
<b>10 Memory Management Unit (MMU) .....</b>	<b>10-1</b>
10.1 Address Spaces .....	10-2
10.2 Address Translation .....	10-3
10.2.1 Address Translation for CSFR Pointers .....	10-3
10.3 Translation Lookaside Buffers (TLBs) .....	10-4
10.3.1 TLB Table Entry (TTE) Contents .....	10-5
10.4 Multiple Address Spaces .....	10-5
10.5 MMU Traps .....	10-5
10.6 Virtual Mode Protection .....	10-7
10.6.1 Direct Translation .....	10-7
10.6.2 Page Table Entry (PTE) Based Translation .....	10-7
10.7 Cacheability .....	10-7
10.7.1 Direct Translation Virtual Address Cacheability .....	10-7
10.7.2 PTE Translation Cacheability .....	10-7
10.7.3 Cacheability of a Virtual Address Flow .....	10-8
10.8 MMU Instructions .....	10-8
10.8.1 TLBMAP (TLB Map) .....	10-9
10.8.2 TLBDEMAP (TLB Demap) .....	10-10
10.8.3 TLBFLUSH (TLB Flush) .....	10-10
10.8.4 TLBPROBE (TLB Probe) .....	10-11
10.9 TLB Usage .....	10-12
10.10 MMU Core Special Function Registers .....	10-13
10.10.1 MMU Configuration Register (MMU_CON) .....	10-13
10.10.2 Address Space Identifier Register (MMU_ASI) .....	10-15
10.10.3 Translation Virtual Address Register (MMU_TVA) .....	10-16
10.10.4 Translation Physical Address Register (MMU_TPA) .....	10-17
10.10.5 Translation Page Index Register (MMU_TPX) .....	10-19
10.10.6 Translation Fault Page Address Register (MMU_TFA) .....	10-20
<b>11 Floating Point Unit (FPU) .....</b>	<b>11-1</b>
11.1 Functional Overview .....	11-1
11.2 IEEE-754 Compliance .....	11-2
11.2.1 IEEE-754 Single Precision Data Format .....	11-2
11.2.2 Denormal Numbers .....	11-3
11.2.3 NaNs (Not a Number) .....	11-3
11.2.4 Underflow .....	11-4
11.2.5 Fused MACs .....	11-4
11.2.6 Traps (TriCore 1.3.1) .....	11-4

<b>Table of Contents</b>	<b>Page</b>
11.2.7 Software Routines .....	11-5
11.3 Rounding .....	11-6
11.3.1 Round to Nearest: Even .....	11-7
11.3.2 Round to Nearest: Denormals and Zero Substitution .....	11-7
11.3.3 Round Towards $\pm \infty$ : Denormals and Zero Substitution .....	11-8
11.4 Exceptions .....	11-8
11.5 Asynchronous Traps (TriCore 1.3.1) .....	11-13
11.6 FPU CSFR Registers (TriCore 1.3.1) .....	11-14
11.6.1 FPU Trap Control Register .....	11-14
11.6.2 FPU Trapping Instruction Program Counter Register .....	11-17
11.6.3 FPU Trapping Instruction Opcode Register .....	11-18
11.6.4 FPU Trapping Instruction Operand SRC1 Register .....	11-19
11.6.5 FPU Trapping Instruction Operand SRC2 Register .....	11-20
11.6.6 FPU Trapping Instruction Operand SRC3 Register .....	11-21
11.6.7 FPU Identification Register .....	11-22
<b>12 Core Debug Controller (CDC) .....</b>	<b>12-1</b>
12.1 Run Control Features .....	12-1
12.2 Debug Events .....	12-3
12.2.1 External Debug Event .....	12-3
12.2.2 Debug Instruction .....	12-3
12.2.3 MTCR and MFCR Instructions .....	12-3
12.2.4 Trigger Event Unit .....	12-4
12.3 Debug Triggers .....	12-5
12.3.1 Combining Debug Triggers .....	12-6
12.4 Debug Actions .....	12-7
12.4.1 Update Debug Status Register (DBGSR) .....	12-7
12.4.2 Indicate on Core Break-Out Signal .....	12-8
12.4.3 Indicate on Core Suspend-Out Signal .....	12-8
12.4.4 Halt .....	12-8
12.4.5 Breakpoint Trap .....	12-8
12.4.6 Breakpoint Interrupt .....	12-10
12.4.7 Suspend Out .....	12-11
12.4.8 Performance Counter Start/Stop (TriCore 1.3.1) .....	12-11
12.4.9 None (TriCore 1.3.1) .....	12-11
12.4.10 Disabled .....	12-12
12.4.11 Suspend In Halt (TriCore 1.3.1) .....	12-12
12.5 Priority of Debug Events .....	12-12
12.6 Call Tracing .....	12-13
12.7 The CDC Control Registers .....	12-14
12.8 CDC Control Registers (TriCore 1.3) .....	12-15
12.8.1 DBGSR Debug Status Register .....	12-15
12.8.2 External Event Register .....	12-17

<b>Table of Contents</b>	<b>Page</b>
12.8.3 Core Register Access Event Register .....	12-18
12.8.4 Software Debug Event Register .....	12-19
12.8.5 Trigger Event Registers .....	12-20
12.8.6 Debug Monitor Start Address Register .....	12-23
12.8.7 Debug Context Save Area Pointer Register .....	12-24
12.9 CDC Control Registers (TriCore 1.3.1) .....	12-25
12.9.1 Debug Status Register .....	12-25
12.9.2 External Event Register .....	12-27
12.9.3 Core Register Access Event Register .....	12-29
12.9.4 Software Debug Event Register .....	12-31
12.9.5 Trigger Event Registers .....	12-33
12.9.6 Debug Monitor Start Address Register .....	12-37
12.9.7 Debug Context Save Area Pointer Register .....	12-38
12.9.8 Debug Trap Control Register .....	12-39
12.9.9 Software Breakpoint Service Request Control Register .....	12-40
12.10 Core Performance Measurement and Analysis (TriCore 1.3.1) .....	12-42
12.11 Performance Counter Registers (TriCore 1.3.1) .....	12-44
12.11.1 Counter Control Register .....	12-45
12.11.2 CPU Clock Cycle Count Register .....	12-47
12.11.3 Instruction Count Register .....	12-48
12.11.4 Multi-Count Register 1 .....	12-49
12.11.5 Multi-Count Register 2 .....	12-50
12.11.6 Multi-Count Register 3 .....	12-51
<b>13 TriCore 1.3.1 Architectural Extensions</b> .....	<b>13-1</b>
13.1 TriCore 1.3.1 Architectural Extensions - Trap System .....	13-1
13.2 TriCore 1.3.1 Architectural Extensions - Core Registers .....	13-3
13.3 TriCore 1.3.1 Architectural Extensions - Instruction Set .....	13-4
13.4 TriCore 1.3.1 - Documentation References .....	13-5
<b>14 Core Register Table</b> .....	<b>14-1</b>
<b>15 List of Registers</b> .....	<b>15-1</b>
<b>16 Index</b> .....	<b>L-1</b>



## Preface

TriCore is a unified, 32-bit microcontroller-DSP, single-core architecture optimized for real-time embedded systems.

This document has been written for system developers and programmers, and hardware and software engineers.

- Volume 1 (this volume) provides a detailed description of the Core Architecture and system interaction.
- Volume 2 gives a complete description of the TriCore Instruction Set including optional extensions for the Memory Management Unit (MMU) and Floating Point Unit (FPU).

It is important to note that this document describes the TriCore architecture, not an implementation. An implementation may have features and resources which are not part of the Core Architecture. The product documentation for that implementation will describe all implementation specific features.

When working with a specific TriCore based product always refer to the appropriate supporting documentation.

## TriCore versions

There have been several versions of the TriCore Architecture implemented in production devices. This manual documents the following architectures: TriCore 1.3, TriCore 1.3.1.

- Unless defined otherwise in the text, or in the margin, all descriptions are common to both the TriCore 1.3 and the TriCore 1.3.1 architecture.
- Information specific to the TriCore 1.3 or TriCore 1.3.1 architecture only is always labelled.

The chapter **TriCore 1.3.1 Architectural Extensions, page 13-1** summarises the new features of the TriCore 1.3.1 architecture.

## Additional Documentation

For information and links to documentation for Infineon products that use TriCore, visit: <http://www.infineon.com/32-bit-microcontrollers>

## Text Conventions

This document uses the following text conventions:

- The default radix is decimal.
  - Hexadecimal constants are suffixed with a subscript letter 'H', as in:  $FFC_H$ .
  - Binary constants are suffixed with a subscript letter 'B', as in:  $111_B$ .
- Register reset values are not generally architecturally defined, but require setting on startup in a given implementation of the architecture. Only those reset values that are architecturally defined are shown in this document. Where no value is shown, the reset value is not defined. Refer to the documentation for a specific TriCore implementation.
- Bit field and bits in registers are in general referenced as 'Register name.Bit field', for example PSW.IS. The Interrupt Stack Control bit of the PSW register.
- Units are abbreviated as follows:
  - MHz = Megahertz.
  - kBaud, kBit = 1000 characters/bits per second.
  - MBaud, MBit = 1,000,000 characters per second.
  - KByte = 1024 bytes.
  - MByte = 1048576 bytes of memory.
  - GByte = 1,024 megabytes.
- Data format quantities referenced are as follows:
  - Byte = 8-bit quantity.
  - Half-word = 16-bit quantity.
  - Word = 32-bit quantity.
  - Double-word = 64-bit quantity.
- Pins using negative logic are indicated by an overbar:  $\overline{BRKOUT}$ .

In tables where register bit fields are defined, the conventions shown in [Table 1](#) are used in this document.

**Table 1 Bit Type Abbreviations**

Abbreviation	Description
r	Read-only. The bit or bit field can only be read.
w	Write-only. The bit or bit field can only be written.
rw	The bit or bit field can be read and written.
h	The bit or bit field can be modified by hardware (such as a status bit). 'h' can be combined with 'rw' or 'r' bits to form 'rwh' or 'rh' bits.
-	Reserved Field. Read value is undefined, must be written with 0.

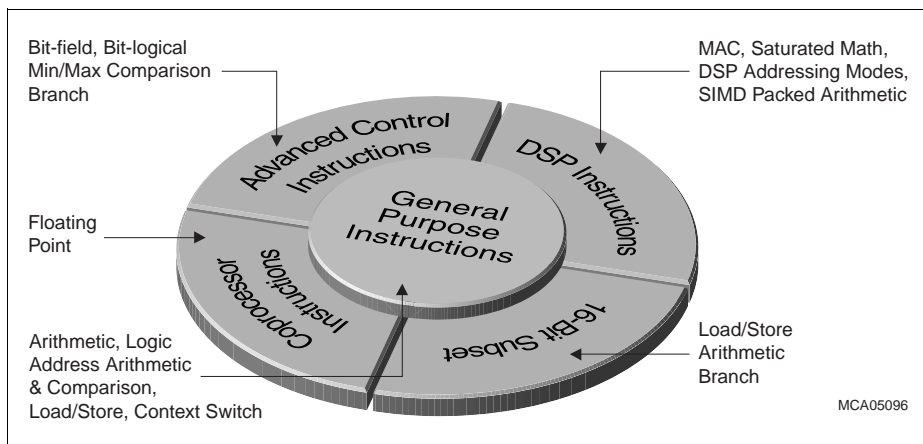
*Note: In register layout tables, a 'Reserved Field' is indicated with '-' in the Field and Type column.*

# 1 Architecture Overview

This chapter gives an overview of the TriCore® architecture.

## 1.1 Introduction

TriCore is the first unified, single-core, 32-bit microcontroller-DSP architecture optimized for real-time embedded systems. The TriCore Instruction Set Architecture (ISA) combines the real-time capability of a microcontroller, the computational power of a DSP, and the high performance/price features of a RISC load/store architecture, in a compact re-programmable core.



**Figure 1 TriCore Architecture Overview**

The ISA supports a uniform, 32-bit address space, with optional virtual addressing and memory-mapped I/O. The architecture allows for a wide range of implementations, ranging from scalar through to superscalar, and is capable of interacting with different system architectures, including multiprocessing. This flexibility at the implementation and system levels allows for different trade-offs between performance and cost at any point in time.

The architecture supports both 16-bit and 32-bit instruction formats. All instructions have a 32-bit format. The 16-bit instructions are a subset of the 32-bit instructions, chosen because of their frequency of use. These instructions significantly reduce code space, lowering memory requirements, system and power consumption.

Real-time responsiveness is largely determined by interrupt latency and context-switch time. The high-performance architecture minimizes interrupt latency by avoiding long multi-cycle instructions and by providing a flexible hardware-supported interrupt scheme. The architecture also supports fast-context switching.

### **1.1.1 Feature Summary**

The key features of the TriCore Instruction Set Architecture (ISA) are:

- 32-bit architecture.
- 4 GBytes of address space.
- 16-bit and 32-bit instructions for reduced code size.
- Most instructions executed in one cycle.
- Branch instructions (using branch prediction).
- Low interrupt latency with fast automatic context switch using wide pathway to on-chip memory.
- Dedicated interface to application-specific coprocessors to allow the addition of customised instructions.
- Zero overhead loop capabilities.
- Dual single-clock-cycle 16×16-bit multiply-accumulate unit (with optional saturation).
- Optional Floating-Point Unit (FPU) and Memory Management Unit (MMU).
- Extensive bit handling capabilities.
- Single Instruction Multiple Data (SIMD) packed data operations (2×16-bit or 4×8-bit operands).
- Flexible interrupt prioritization scheme.
- Byte and bit addressing.
- Little-endian byte ordering for data memory and CPU registers.
- Memory protection.
- Coprocessor support.
- Debug support.

### **1.2 Programming Model**

This section covers aspects of the architecture that are visible to software:

- Architectural Registers [page 1-3](#).
- Data Types [page 1-4](#).
- Memory Model [page 1-4](#).
- Addressing Modes [page 1-6](#).

The Programming Model is described in detail in the chapter [Programming Model, page 2-1](#).



### 1.2.1 Architectural Registers

The architectural registers consist of:

- 32 General Purpose Registers (GPRs).
- Program Counter (PC).
- Two 32-bit registers containing status flags, previous execution information and protection information (PCXI - Previous Context Information register, and PSW - Program Status Word).

Address		Data		System	
31	0	31	0	31	0
A[15] (Implicit Base Address)		D[15] (Implicit Data)		PCXI	
A[14]		D[14]		PSW	
A[13]		D[13]		PC	
A[12]		D[12]			
A[11] (Return Address)		D[11]			
A[10] (Stack Return)		D[10]			
A[9] (Global Address Register)		D[9]			
A[8] (Global Address Register)		D[8]			
A[7]		D[7]			
A[6]		D[6]			
A[5]		D[5]			
A[4]		D[4]			
A[3]		D[3]			
A[2]		D[2]			
A[1] (Global Address Register)		D[1]			
A[0] (Global Address Register)		D[0]			

MCA05246

**Figure 2 Architectural Registers**

The PCXI, PSW and PC registers are crucial to the procedure for storing and restoring a task's context.

The 32 General Purpose Registers (GPRs) are divided into sixteen 32-bit data registers (D[0] through D[15]) and sixteen 32-bit address registers (A[0] through A[15]).

Four of the General Purpose Registers (GPRs) also have special functions:

- D[15] is used as an Implicit Data register.
- A[10] is the Stack Pointer (SP) register.
- A[11] is the Return Address (RA) register.
- A[15] is the Implicit Address register.

Registers [0<sub>H</sub> - 7<sub>H</sub>] are referred to as the 'lower registers' and registers [8<sub>H</sub> - F<sub>H</sub>] are called the 'upper registers'.

Registers A[0], A[1], A[8], and A[9] are defined as system global registers. These are not included in either the upper or lower context (see [Tasks and Functions, page 4-1](#)) and are not saved and restored across calls or interrupts. They are normally used by the operating system to reduce system overhead.

In addition to the General Purpose Registers (GPRs), the core registers are composed of a certain number of Core Special Function Registers (CSFRs). See [General Purpose and System Registers, page 3-1](#).

### 1.2.2 Data Types

The instruction set supports operations on:

- Boolean.
- Bit String.
- Byte.
- Signed Fraction.
- Address.
- Signed / Unsigned Integer.
- IEEE-754 Single-Precision Floating-Point.

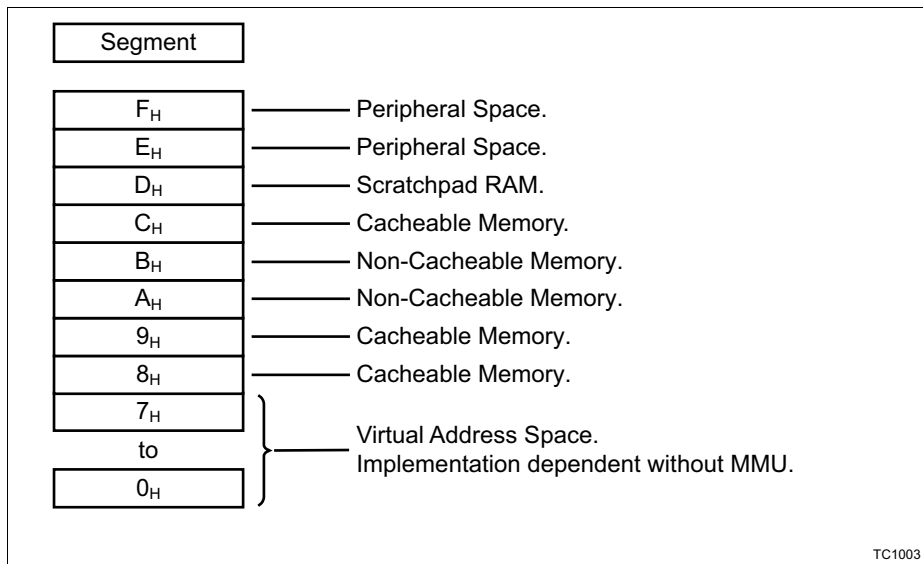
Most instructions work on a specific data type, while others are useful for manipulating several data types.

### 1.2.3 Memory Model

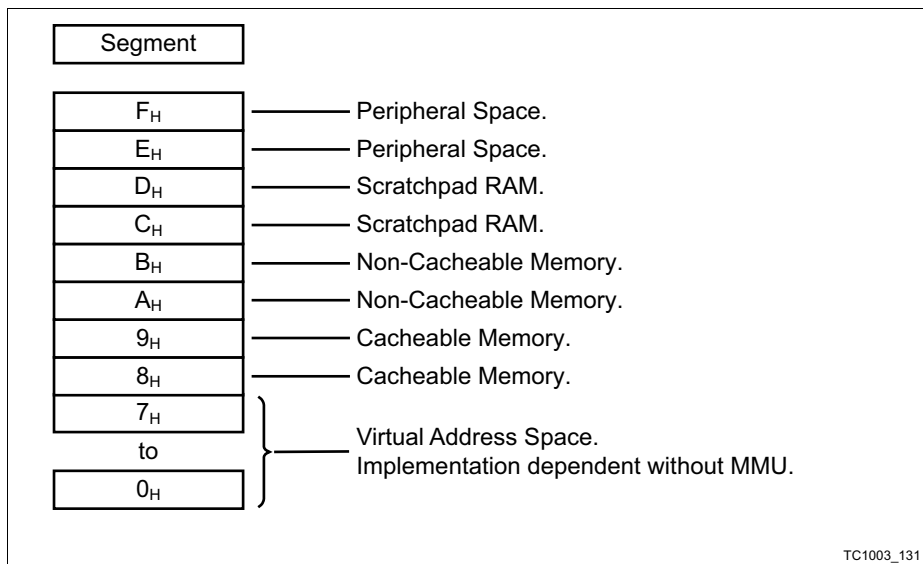
The architecture can access up to 4 GBytes (address width is 32-bits) of unified program and I/O memory.

The address space is divided into 16 regions or segments [0<sub>H</sub> - F<sub>H</sub>], each of 256 MBytes. The upper four bits of an address select the specific segment. The first 16 KBytes of each segment can be accessed directly using absolute addressing.

The diagram which follows shows the TriCore architecture address space mapping.



**Figure 3 Address Map and Memory Model (TriCore 1.3)**



**Figure 4 Address Map and Memory Model (TriCore 1.3.1)**

### **1.2.4 Addressing Modes**

Addressing modes allow load and store instructions to efficiently access simple data elements within data structures such as records, randomly and sequentially accessed arrays, stacks and circular buffers.

The architecture supports seven addressing modes. The simple data elements are 8-bits, 16-bits, 32-bits and 64-bits wide.

These addressing modes support efficient compilation of C/C++ programs, easy access to peripheral registers and efficient implementation of typical DSP data structures (circular buffers for filters and bit-reversed indexing for Fast Fourier Transformations).

Addressing modes which are not directly supported in the hardware can be synthesized through short instruction sequences.

For more information see [Synthesized Addressing Modes, page 2-13](#).

### 1.3 Tasks and Contexts

A task is an independent thread of control. There are two types: Software Managed Tasks (SMTs) and Interrupt Service Routines (ISRs).

SMTs are created through the services of a real-time kernel or Operating System, and are dispatched under the control of scheduling software. ISRs are dispatched by hardware in response to an interrupt. An ISR is the code that is invoked directly by the processor on receipt of an interrupt. SMTs are sometimes referred to as user tasks, assuming that they execute in User Mode.

Each task is allocated its own mode, depending on the task's function:

- **User-0 Mode:** Used for tasks that do not access peripheral devices. This mode cannot enable or disable interrupts.
- **User-1 Mode:** Used for tasks that access common, unprotected peripherals. Typically this would be a read or write access to serial port, a read access to timer, and most I/O status registers. Tasks in this mode may disable interrupts for a short period.
- **Supervisor Mode:** Permits read/write access to system registers and all peripheral devices. Tasks in this mode may disable interrupts.

Individual modes are enabled or disabled primarily through the I/O mode bits in the Processor Status Word (PSW).

A set of state elements are associated with any task, and these are known collectively as the task's context. The context is everything the processor needs to define the state of the associated task and enable its continued execution. This includes the CPU General Registers that the task uses, the task's Program Counter (PC), and its Program Status Information (PCXI and PSW). The architecture efficiently manages and maintains the context of the task through hardware. The context is subdivided into the upper context and the lower context.

#### Context Save Areas

The architecture uses linked lists of fixed-size Context Save Areas (CSAs). A CSA consists of 16 words of memory storage, aligned on a 16-word boundary. Each CSA can hold exactly one upper or one lower context. CSAs are linked together through a Link Word.

The architecture saves and restores context more quickly than conventional microprocessors and microcontrollers. The unique memory subsystem design with a wide data path allows the architecture to perform rapid data transfers between processor registers and on-chip memory.

Context switching occurs when an event or instruction causes a break in program execution. The CPU then needs to resolve this event before continuing with the program.

The events and instructions which cause a break in program execution are:

- Interrupt or service requests.
- Traps.
- Function calls.

See [Tasks and Functions, page 4-1](#).

## **1.4 Interrupt System**

A key feature of the architecture is its powerful and flexible interrupt system. The interrupt system is built around programmable Service Request Nodes (SRNs).

A Service Request is defined as an interrupt request or a DMA (Direct Memory Access) request. A service request may come from an on-chip peripheral, external hardware, or software.

Conventional architectures generally take a long time to service interrupt requests, and they are normally handled by loading a new Program Status (PS) from a vector table in data memory. In the TriCore architecture, service requests jump to vectors in code memory to reduce response time. The entry code for the ISR is a block within a vector of code blocks. Each code block provides an entry for one interrupt source.

### **1.4.1 Interrupt Priority**

Service requests are prioritized, and prioritization allows for nested interrupts. The rules for prioritization are:

- A service request can interrupt the servicing of a lower priority interrupt.
- Interrupt sources with the same priority cannot interrupt each other.
- The Interrupt Control Unit (ICU) determines which source will win arbitration based on the priority number.

All Service Requests are assigned Priority Numbers (SRPNs). Even the ISR has its own priority number. Different service requests must be assigned different priority numbers.

The maximum number of interrupt sources is 255. Programmable options range from one priority level with 255 sources, up to 255 priority levels with one source each.

Interrupt numbers are assumed to be assigned in linear order of interrupt priority. This is feasible because interrupt numbers are not hardwired to individual sources, but are assigned by software executed during the power-on boot sequence.

See [Interrupt System, page 5-1](#).

## 1.5 Trap System

A trap occurs as a result of an event such as a Non-Maskable Interrupt (NMI), an instruction exception or illegal access. The TriCore architecture contains eight trap classes and these traps are further classified as synchronous or asynchronous, hardware or software. Each trap is assigned a Trap Identification Number (TIN) that identifies the cause of the trap within its class. The entry code for the trap handler is comprised of a vector of code blocks. Each code block provides an entry for one trap. When a trap is taken, the TIN is placed in data register D[15].

The trap classes are:

- MMU (Memory Management Unit).
- Internal Protection.
- Instruction Error.
- Context Management.
- System Bus and Peripherals.
- Assertion Trap.
- System Call.
- Non-Maskable Interrupt (NMI).

See [Trap System, page 6-1](#).

## 1.6 Protection System

One of the domains that TriCore supports is safety-critical embedded applications. The architecture features a protection system designed to protect core system functionality from the effects of software errors in less critical application tasks, and to prevent unauthorised tasks from accessing critical system peripherals. The protection system also facilitates debugging. It detects and traps errors that might otherwise go unnoticed until it was too late to identify the cause of the error.

The overall protection system is composed of three main subsystems:

1. **The Trap System:** Described briefly in [Section 1.5](#), but covered in detail in [Trap System, page 6-1](#).
2. **The I/O Privilege Level:** TriCore supports three I/O modes: User-0 mode, User-1 mode and Supervisor mode. The User-1 mode allows application tasks to directly access non-critical system peripherals. This allows embedded systems to be implemented efficiently, without the loss of security inherent in the common practice of running everything in Supervisor mode.
3. **The Memory Protection System:** This protection system provides control over which regions of memory a task is allowed to access, and what types of access it is permitted.

For TriCore v1.3 and later architecture revisions, there are actually two independent memory protection systems. For applications that require virtual memory, the optional Memory Management Unit (MMU) supports a familiar page-based model for memory protection. That model gives each memory page its own access permissions. The relatively conventional MMU design and the page-based memory protection model facilitate porting of standard operating systems that expect this model. The MMU is detailed in [Memory Management Unit \(MMU\), page 10-1](#).

For the smaller and lower cost applications there is a range-based memory protection system. This is designed to provide coarse-grained memory protection for systems that do not require virtual memory. The range-based memory protection system and its interaction with I/O privilege level for access to peripherals, is detailed in [Memory Protection System, page 9-1](#).

## **1.7 Memory Management Unit**

TriCore can make use of an optional Memory Management Unit (MMU). When configured with an MMU, the memory space has two addressing regions; physical and virtual. The physical and virtual address space is 4 GBytes in each instance, with those 4 GBytes each divided into sixteen, 256 MByte segments.

Segments [ $8_H$ - $F_H$ ] bypass virtual mapping and are directly, physically used. Segments [ $0_H$ - $7_H$ ] are virtually mapped by the MMU when it is present and enabled, or physically mapped when the MMU is not present or disabled.

Virtual addresses are always translated into physical addresses before accessing memory. This translation to a physical address is either a Direct Translation or a Page Table Entry (PTE) Translation, depending on MMU mode and virtual address region:

- **Direct Translation:** If the virtual address belongs to the upper half of the virtual address space, then the virtual address is directly used as the physical address. If the virtual address belongs to the lower half of the address space and the processor is operating in Physical mode, then the virtual address is used indirectly as the physical address.
- **Page Table Entry (PTE) Translation:** If the processor is operating in Virtual mode and the virtual address belongs to the lower half of the address space, then the virtual address is translated using PTE. PTE translation is performed by replacing the Virtual Page Number (VPN) of the virtual address by a Physical Page Number (PPN) to obtain a physical address.

See [Memory Management Unit \(MMU\), page 10-1](#).



## **1.8 Core Debug Controller**

The Core Debug Controller (CDC) is designed to support real-time systems that require non-intrusive debugging. Most of the architectural state in the CPU Core and Core on-chip memories can be accessed through the system Address Map. The debug functionality is an interface of architecture, implementation and software tools.

Access to the CDC is typically provided via the On-Chip Debug Support (OCDS) of the system containing the CPU.

A general description of the mechanism and registers is detailed in [Core Debug Controller \(CDC\), page 12-1](#)

## **1.9 Coprocessor Interface**

The TriCore architecture may be extended with implementation defined application specific instructions. These instructions are executed on dedicated coprocessor hardware attached to the coprocessor interface.



## **2 Programming Model**

This chapter discusses the following aspects of the TriCore® architecture that are visible to software:

- Supported data types [page 2-1](#).
- Data formats in registers and memory [page 2-2](#).
- The Memory model [page 2-6](#).
- Addressing modes [page 2-7](#).

### **2.1 Data Types**

The instruction set supports operations on the following Data Types:

- Boolean [page 2-1](#).
- Bit String [page 2-1](#).
- Byte [page 2-1](#).
- Signed Fraction [page 2-2](#).
- Address [page 2-2](#).
- Signed and Unsigned Integers [page 2-2](#).
- IEEE-754 Single-precision Floating-point Number [page 2-2](#).

Most instructions operate on a specific Data Type, while others are useful for manipulating several Data Types.

#### **2.1.1 Boolean**

A Boolean is either TRUE or FALSE:

- TRUE is the value one (1) when generated and non-zero when tested.
- FALSE is the value zero (0).

Booleans are produced as the result in comparison and logic instructions, and are used as source operands in logical and conditional jump instructions.

#### **2.1.2 Bit String**

A bit string is a packed field of bits.

Bit strings are produced and used by logical, shift, and bit field instructions.

#### **2.1.3 Byte**

A byte is an 8-bit value that can be used for a character or a very short integer. No specific coding is assumed.

### **2.1.4 Signed Fraction**

The architecture supports 16-bit, 32-bit and 64-bit signed fractional data for DSP arithmetic. Data values in this format have a single high-order sign bit, where 0 represents positive (+) and 1 represents negative (-), followed by an implied binary point and fraction. Their values are therefore in the range [-1,1).

### **2.1.5 Address**

An address is a 32-bit unsigned value.

### **2.1.6 Signed and Unsigned Integers**

Signed and unsigned integers are normally 32 bits. Shorter signed or unsigned integers are sign-extended or zero-extended to 32 bits when loaded from memory into a register.

### **Multi-precision**

Multi-precision integers are supported with addition and subtraction using carry. Integers are considered to be bit strings for shifting and masking operations. Multi-precision shifts can be made using a combination of single-precision shifts and bit field extracts.

### **2.1.7 IEEE-754 Single-Precision Floating-Point Number**

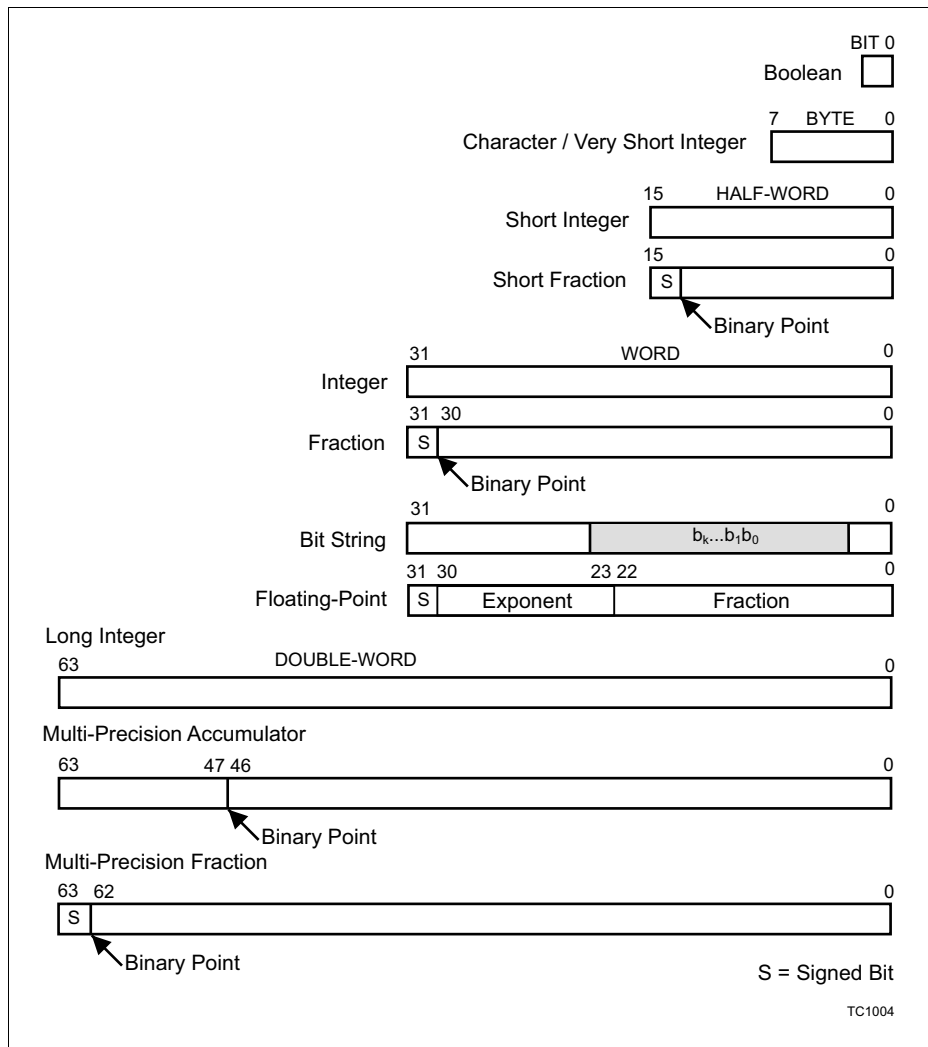
Depending on the particular implementation of the core architecture, IEEE-754 floating-point numbers are supported by coprocessor hardware instructions or by software calls to a library.

## **2.2 Data Formats**

All General Purpose Registers (GPRs) are 32 bits wide, and most instructions operate on word (32-bit) values. When byte or half-word data elements are loaded from memory, they are automatically sign-extended or zero-extended to fill the register. The type of filling is implicit in the load instruction. For example, LD.B to load a byte with sign extension, or LD.BU to load a byte with zero extension.

The supported Data Formats are:

- Bit.
- Byte: signed, unsigned.
- Half-word: signed, unsigned, fraction.
- Word: signed, unsigned, fraction, floating-point.
- 48-bit: signed, unsigned, fraction.
- Double-word: signed, unsigned, fraction.



**Figure 5 Supported Data Formats**

## 2.2.1 Alignment Requirements

Alignment requirements differ for addresses and data (see [Table 1](#)). Address variables loaded into or stored from address registers, must always be word-aligned.

Data can be aligned on any half-word boundary, regardless of size, except where noted below. This facilitates the use of packed arithmetic operations in DSP applications, by allowing two or four packed 16-bit data elements to be loaded or stored together on any half-word boundary.

There are some restrictions of which programmers must be aware, specifically:

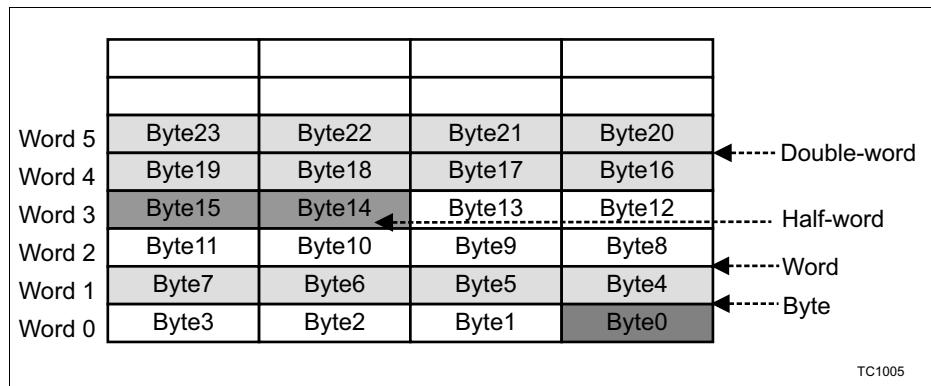
- The LDMST and SWAP instructions require their operands to be word-aligned.
- Half-word alignment for LD.D and ST.D is only allowed when the source or destination address is targeted at cached memory or data scratchpad RAM (see [Scratchpad RAM, page 8-4](#)). For all other addresses double-word accesses must be word-aligned.
- The byte operations LD.B, ST.B, LD.BU, ST.T may be byte aligned.

**Table 1 Alignment Rules**

Access type	Access size	Alignment of address in memory	
Load, Store Data Register	Byte	Byte (1 <sub>H</sub> )	
	Half-word	2 bytes (2 <sub>H</sub> )	
	Word	2 bytes (2 <sub>H</sub> )	
	Double-word	2 bytes (2 <sub>H</sub> )	
Load, Store Address Register	Word	4 bytes (4 <sub>H</sub> )	
	Double-word	4 bytes (4 <sub>H</sub> )	
SWAP.W, LDMST	Word	4 bytes (4 <sub>H</sub> )	
ST.T	Byte	Byte (1 <sub>H</sub> )	
Context Load / Store / Restore / Save	16 x 32-bit registers	64 bytes (40 <sub>H</sub> )	

### 2.2.2 Byte Ordering

The data memory and CPU registers store data in little-endian byte order (the least-significant bytes are at lower addresses). The following figure illustrates byte ordering. Little-endian memory referencing is used consistently for data and instructions.



**Figure 6 Byte Ordering**

## 2.3 Memory Model

The architecture has an address width of 32 bits and can access up to 4 GBytes of memory. The address space is divided into 16 regions or segments,  $[0_H - F_H]$ . Each segment is 256 MBytes. The upper 4 bits of an address select the specific segment. The first 16 KBytes of each segment can be accessed using absolute addressing.

Many data accesses use addresses computed by adding a displacement to the value of a base address register. Using a displacement to cross one of the segment boundaries is not allowed and if attempted causes a MEM trap. This restriction allows direct determination of the accessed segment from the base address.

See [Trap System, page 6-1](#) for more information on Traps.

### Physical Memory Attributes

The physical memory attributes of segments zero to seven are implementation dependent. If an MMU is present and enabled, segments  $[0_H - 7_H]$  are considered virtual addresses that must be translated. If an MMU is not present the access characteristics are implementation dependent and may cause a trap.

### Physical Memory Addresses

Physical memory addresses in segment  $F_H$  are guaranteed to have the peripheral space attribute and therefore all accesses are non-speculative and are not accessible to User-0 mode. This segment can therefore be used for mapping peripheral registers.

The Core Special Function Registers (CSFRs) are mapped to a 64 KBytes space in the memory map. The base location of this 64 KBytes space is implementation-dependent.

Segments  $8_H$  to  $D_H$  have further limitations placed upon them in some implementations. For example, specific segments for program and data may be defined by device-specific implementations. Other details of the memory mapping are implementation-specific.

For more information see [Physical Memory Attributes \(PMA\), page 8-1](#).

**Table 2 Physical Address Space**

Address	Segments	Description
$FFFF\ FFFF_H : E000\ 0000_H$	$E_H - F_H$	Peripheral space.
$DFFF\ FFFF_H : 8000\ 0000_H$	$8_H - D_H$	Detailed limitations are implementation specific.
$7FFF\ FFFF_H : 0000\ 0000_H$	$0_H - 7_H$	Implementation dependent.



## 2.4 Semaphores and Atomic Operations

The TriCore architecture has two instructions which read and write memory in atomic fashion, which are supported by all versions of the architecture:

- LDMST (Load, Modify, Store)
- SWAP.W (Swap register with memory).

LDMST uses a mask register to write selected bits from a source register into a memory word. However it does not return a value, so it can not be used as an atomic "test and set" type operations for binary semaphores. The SWAP.W is provided for this purpose.

## 2.5 Addressing Modes

Addressing modes allow load and store instructions to access simple data elements such as records, randomly and sequentially accessed arrays, stacks, and circular buffers.

The simple data elements are 8-bits, 16-bits, 32-bits, or 64-bits wide. The architecture supports seven addressing modes.

The addressing modes support efficient compilation of C/C++, give easy access to peripheral registers, and efficient implementation of typical DSP data structures (circular buffers for filters and bit-reversed indexing for FFTs).

**Table 3 Addressing Modes**

Addressing Mode	Address Register Use
Absolute	None
Base + Short Offset	Address Register
Base + Long Offset	Address Register
Pre-increment	Address Register
Post-increment	Address Register
Circular	Address Register Pair
Bit-reverse	Address Register Pair

Addressing modes which are not directly supported in the hardware can be synthesized through short instruction sequences.

For more information see [Synthesized Addressing Modes, page 2-13](#).



### 2.5.3 Pre-Increment and Pre-Decrement Addressing

Pre-increment and pre-decrement addressing (where pre-decrement addressing is obtained by the use of a negative offset), may be used to push onto an upward or downward-growing stack, respectively.

The pre-increment addressing mode uses the sum of the address register and the offset both as the effective address and as the value written back into the address register.


### 2.5.4 Post-Increment and Post-Decrement Addressing

Post-increment and post-decrement addressing (where post-decrement addressing is obtained by the use of a negative offset), may be used for forward or backward sequential access of arrays respectively. Furthermore, the two versions of the mode may be used to pop from a downward-growing or upward-growing stack, respectively.

The post-increment addressing mode uses the value of the address register as the effective address and then updates this register by adding the sign-extended 10-bit offset to its previous value.

### 2.5.5 Circular Addressing

The primary use of circular addressing ([Figure 8](#)) is for accessing data values in circular buffers while performing filter calculations.



**Figure 8 Circular Addressing Mode**

The circular addressing mode uses an address register pair to hold the state it requires:

- The even register is always a base address ( $B$ ).
- The most significant half of the odd register is the buffer size ( $L$ ).
- The least significant half holds the index into the buffer ( $I$ ).
- The effective address is  $(B+I)$ .
- The buffer occupies memory from addresses  $B$  to  $B+L-1$ .

The index is post-incremented using the following algorithm:

```
tmp = I + sign_ext(offset10);  
  
if (tmp < 0)  
    I = tmp + L;  
else if (tmp >= L)  
    I = tmp - L;  
else  
    I = tmp;
```


TC1009

**Figure 9 Circular Addressing Index Algorithm**

The 10-bit offset is specified in the instruction word and is a byte-offset that can be either positive or negative. Note that correct 'wrap around' behaviour is guaranteed as long as the magnitude of the offset is smaller than the size of the buffer.

To illustrate the use of circular addressing, consider a circular buffer consisting of 25, 16-bit values. If the current index is 48, then the next item is obtained using an offset of two (2-bytes per value). The new value of the index 'wraps around' to zero. If we are at an index of 48 and use an offset of four, the new value of the index is two. If the current index is four and we use an offset of -8, then the new index is 46 ( $4-8+50$ ).

In the end case, where a memory access runs off the end of the circular buffer (**Figure 10**), the data access also wraps around to the start of the buffer. For example, consider a circular buffer containing  $n+1$  elements where each element is a 16-bit value. If a load word is performed using the circular addressing mode and the effective address of the operation points to element  $n$ , the 32-bit result contains element  $n$  in the bottom 16 bits and element 0 in the top 16 bits.



**Figure 10 Circular Buffer End Case**

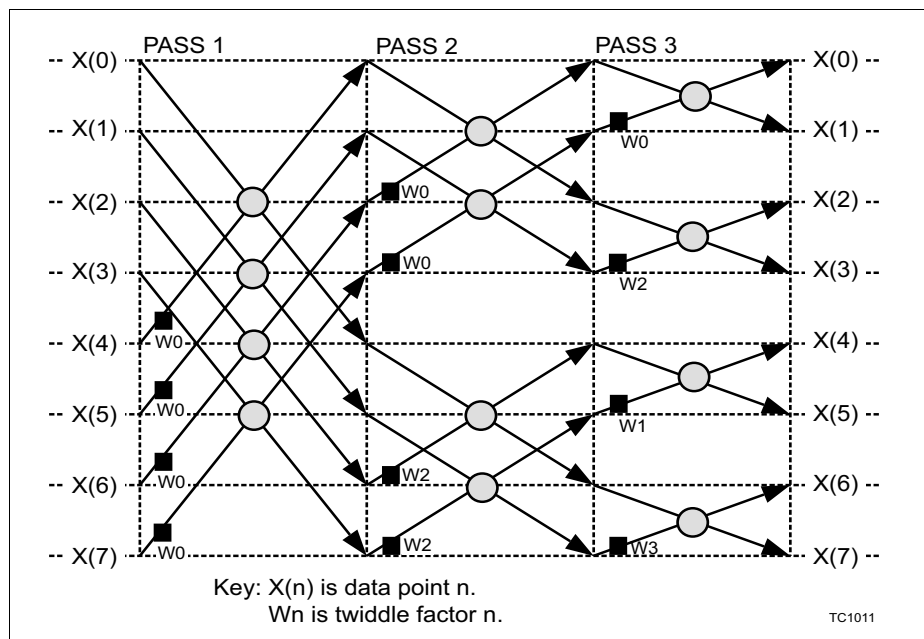
The size and length of a circular buffer has the following restrictions:

- The start of the buffer must be aligned to a 64-bit boundary. An implementation is free to advise the user of optimal alignment of circular buffers etc., but must support alignment to the 64-bit boundary.
- The length of the buffer must be a multiple of the data size, where the data size is determined from the instruction being used to access the buffer. For example, a buffer accessed using a load-word instruction must be a multiple of 4 bytes in length, and a buffer accessed using a load double-word instruction must be a multiple of 8-bytes in length.

If these restrictions are not met the implementation takes an alignment trap (ALN). An alignment trap is also taken if the index ( $I$ )  $\geq$  length ( $L$ ).


## 2.5.6 Bit-Reverse Addressing

Bit-reverse addressing is used to access arrays used in FFT algorithms. The most common implementation of the FFT ends with results stored in bit-reversed order ([Bit-Reverse Addressing, page 2-12](#)).



**Figure 11 Bit-Reverse Addressing**

Bit-reverse addressing uses an address register pair to hold the required state:



**Figure 12 Register Pair for Bit-Reverse Addressing**

- The even register is the base address of the array (B).
- The least-significant half of the odd register is the index into the array (I).
- The most-significant half is the modifier (M), used to update I after every access.
- The effective address is B+I.
- The index, I, is post-incremented and its new value is *reverse* [*reverse* (I) + *reverse* (M)]. The *reverse*(I) function exchanges bit *n* with bit (15−*n*) for *n* = 0, ... 7.

To illustrate for a 1024 point real FFT using 16-bit values, the buffer size is 2048 bytes. Stepping through this array using a bit-reverse index would give the sequence of byte indices: 0, 1024, 512, 1536, and so on. This sequence can be obtained by initializing I to 0 and M to 0400<sub>H</sub>.

**Table 4 1024-point FFT Using 16-bit Values**

I (decimal)	I (binary)	Reverse(I)	Rev[Rev(I) + Rev(M)]
0	0000000000000000 <sub>B</sub>	0000000000000000 <sub>B</sub>	0000010000000000 <sub>B</sub>
1024	0000010000000000 <sub>B</sub>	0000000000100000 <sub>B</sub>	0000001000000000 <sub>B</sub>
512	0000001000000000 <sub>B</sub>	0000000001000000 <sub>B</sub>	0000011000000000 <sub>B</sub>
1536	0000011000000000 <sub>B</sub>	0000000001100000 <sub>B</sub>	0000010001100000 <sub>B</sub>

The required value of M is given by; buffer size/2, where the buffer size is given in bytes.

### 2.5.7 Synthesized Addressing Modes

This section describes how addressing that is not directly supported in the hardware addressing modes, can be synthesized through short instruction sequences.

#### Indexed Addressing

The Indexed addressing mode can be synthesized using the ADDSC.A instruction (Add Scaled Index to Address), which adds a scaled data register to an address register. The scale factor can be 1, 2, 4 or 8 for addressing indexed arrays of bytes, half-words, words, or double-words.

#### Bit Indexed Addressing

To support addressing of indexed bit arrays, the ADDSC.AT instruction scales the index value by 1/8 (shifts right 3 bits) and adds it to the address register.

The two low-order bits of the resulting byte address are cleared to give the address of the word containing the indexed bit.

To extract the bit, the word in which it is contained, is loaded. The bit index is then used in an EXTR.U instruction.

A bit field, beginning at the indexed bit position, can also be extracted. To store a bit or bit field at an indexed bit position, ADDSC.AT is used in conjunction with the LDMST (Load/Modify/Store) instruction.

## **PC-Relative Addressing**

PC-relative addressing is the normal mode for branches and calls. However the architecture does not support direct PC-relative addressing of data. This is because the separate on-chip instruction and data memories make data access to the program memory expensive.

When PC-relative addressing of data is required, the address of a nearby code label is placed into an address register and used as a base register in base + offset mode to access the data. Once the base register is loaded it can be used to address other PC-relative data items nearby.

A code address can be loaded into an address register in various ways. If the code is statically linked (as it almost always is for embedded systems), then the absolute address of the code label is known and can be loaded using the LEA instruction (Load Effective Address), or with a sequence to load an extended absolute address. The absolute address of the PC relative data is also known, and there is no need to synthesize PC-relative addressing.

For code that is dynamically loaded, or assembled into a binary image from position-independent pieces without the benefit of a relocating linker, the appropriate way to load a code address for use in PC-relative data addressing is to use the JL (Jump and Link) instruction. A jump and link to the next instruction is executed, placing the address of that instruction into the return address (RA) register A[11]. Before this is done though, it is necessary to copy the actual return address of the current function to another register.



### **3 General Purpose and System Registers**

There are two types of Core Register, the General Purpose Registers (GPRs) and the Core Special Function Registers (CSFRs). The GPRs consist of 16 general purpose data and 16 general purpose address registers. The CSFRs control the operation of the core and provide status information about the core.

- General Purpose Registers.
- System registers (PSW, PC, PCXI).
- Stack Management registers are (A[10] and ISP).
- SYSCON and CPU\_ID registers.
- Trap registers.
- Context Management registers.
- Memory Protection registers.
- Memory Management registers.
- Debug registers.
- Floating Point registers.
- Special Function registers associated with the core.

#### **Reset Values**

It should be noted that because this manual describes the TriCore® architecture, not an implementation of that architecture, some reset values are not given. Where they are not given, the values are implementation specific.

#### **ENDINIT Protection**

The architecture supports the concept of an initialisation state prior to an operational state.

When in the initialisation state, all Core Special Function Registers can be modified, using the MTCR instruction. In the operational state only a subset of CSFRs can be modified in this way. All other functions remain identical between these states.

CSFRs that are only writable in the initialisation state are described as ENDINIT protected.

The transition between the initialisation state and the operational state is controlled by the system implementation. This facility adds an extra level of protection to critical CSFRs by only allowing them to be changed in the initialisation state.

The following registers are ENDINIT protected:

- BTV, BIV and ISP
- SMACON, BMACON, COMPAT, MIECON (TriCore 1.3.1)

### **3.1 General Purpose Registers (GPRs)**

The General Purpose Registers (GPRs) are split evenly into:

- 16 Data registers (DGPRs), D[0] to D[15].
- 16 Address registers (AGPRs), A[0] to A[15].

The separation of data and address registers facilitates efficient implementations in which arithmetic and memory operations are performed in parallel. Several instructions allow the interchange of information between data and address registers (used for example, to create or derive table indexes). Two consecutive even-odd data registers can be concatenated to form eight extended-size registers (E[0], E[2], E[4], E[6], E[8], E[10], E[12], and E[14]), in order to support 64-bit values. The address registers (P[0], P[2], P[4], P[6], P[8], P[10], P[12], and P[14]) can be used in the same way.

Registers A[0], A[1], A[8], and A[9] are defined as system global registers. Their contents are not saved or restored across calls, traps or interrupts.

Register A[10] is used as the Stack Pointer (SP). See [Stack Management Registers, page 3-13](#).

Register A[11] is used to store the Return Address (RA) for calls and linked jumps, and to store the return Program Counter (PC) value for interrupts and traps.

While the 32-bit instructions have unlimited use of the GPRs, many 16-bit instructions implicitly use A[15] as their address register and D[15] as their data register. This implicit use eases the encoding of these instructions into 16 bits.

Support of 64-bit data values is provided with the use of odd/even register pairs. In the assembler syntax these register pairs are either referred to as a pair of 32-bit registers (for example, D[9]/D[8]) or as an extended 64-bit register. For example, E[8] is the concatenation of D[9] and D[8], where D[8] is the least significant word of E[8].

In order to support extended addressing modes, an even/odd address register pair holds the extended address reference as a pair of 32-bit address registers (A[8]/A[9] for example).

There are no separate floating-point registers. The data registers are used to perform floating-point operations. The floating-point data is saved and restored automatically using the fast context switch support.


**Figure 13, page 3-4**, shows the 32-bit wide GPRs.

## General Purpose and System Registers


### 3.1.1 Data General Purpose Registers

Dn (n = 0-15)

Data Register n (FF00<sub>H</sub>+n\*4) Reset Value: Implementation Specific



rw



rw

Field	Bits	Type	Description
DATA	[31:1]	rw	Data Register n Value

### 3.1.2 Address General Purpose Registers

An (n = 0-15)

Address Register n (FF80<sub>H</sub>+n\*4) Reset Value: Implementation Specific



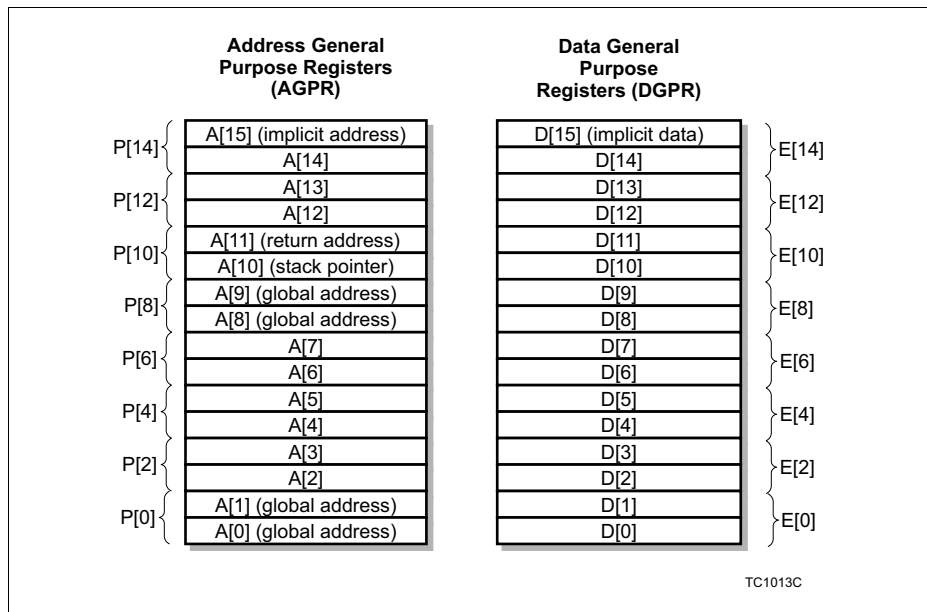
rw



rw

Field	Bits	Type	Description
ADDR	[31:1]	rw	Address Register n Value

**General Purpose and System Registers**



**Figure 13 General Purpose Registers (GPRs)**

The GPRs are an essential part of a task's context. When saving or restoring a task's context to and from memory the context is split into the upper and lower contexts:

- Registers A[2] to A[7] and D[0] to D[7] are part of the lower context.
- Registers A[10] to A[15] and D[8] to D[15] are part of the upper context.

*Note: Upper and lower contexts are described in detail in [Chapter 4 Tasks and Functions](#).*

## 3.2 Program State Information Registers

The PC, PSW, and PCXI registers hold and reflect program state information. These registers are an important part of storing and restoring a task's context, when the contents are stored, restored or modified during this process.


- PC: Program Counter [page 3-5](#).
- PSW: Program Status Word [page 3-6](#).
- PCXI: Previous Context Information [page 3-12](#).

### 3.2.1 Program Counter (PC)

The 32-bit Program Counter (PC) shown below, holds the address of the instruction that is currently running. The Program Counter is part of a task's state information.

#### PC

**Program Counter Register (FE08<sub>H</sub>) Reset Value: Implementation Specific**



Field	Bits	Type	Description
PC	[31:1]	rw	Program Counter
-	0	-	Reserved Field

### 3.2.2 Program Status Word Register (PSW)

The Program Status Word register (PSW) is a 32-bit register that contains a task-specific architectural state not captured in the General Purpose Register values. The lower half holds control values and parameters related to the protection system, including:

- The Protection Register Set (PRS).
- The I/O privilege level (IO).
- The Interrupt Stack flag (IS).
- The Global register Write permission flag (GW).
- The Call Depth Counter (CDC).
- The Call Depth Count Enable field (CDE).

#### PSW

**Program Status Word (FE04<sub>H</sub>)**      **Reset Value: 0000 0B80<sub>H</sub>**

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
USB								-							
rw															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
-		PRS		IO		IS		GW		CDE		CDC			
rw		rw		rw		rw		rw		rw		rw			

Field	Bits	Type	Description
USB	[31:24]	rw	<b>User Status Bits</b> The eight most significant bits of the PSW are designated as User Status Bits. These bits may be set or cleared as execution side effects of user instructions. Refer to the <a href="#">PSW User Status Bits</a> section which follows this table.
-	[23:14]	-	<b>Reserved Field</b>
PRS	[13:12]	rw	<b>Protection Register Set</b> Selects the active Data and Code Memory Protection Register Set. The memory protection register values control load, store and instruction fetches within the current process. 00 <sub>B</sub> : Protection Register Set 0. 01 <sub>B</sub> : Protection Register Set 1. 10 <sub>B</sub> : Protection Register Set 2. 11 <sub>B</sub> : Protection Register Set 3.

**General Purpose and System Registers**

Field	Bits	Type	Description
IO	[11:10]	rw	<p><b>Access Privilege Level Control (I/O Privilege)</b>  Determines the access level to special function registers and peripheral devices.</p> <p><b>00<sub>B</sub> : User-0 Mode</b>  No peripheral access. Access to memory regions with the peripheral space attribute are prohibited and results in a PSE or MPP trap. This access level is given to tasks that need not directly access peripheral devices. Tasks at this level do not have permission to enable or disable interrupts.</p> <p><b>01<sub>B</sub> : User-1 Mode</b>  Regular peripheral access. Enables access to common peripheral devices that are not specially protected, including read/write access to serial I/O ports, read access to timers, and access to most I/O status registers. Tasks at this level may disable interrupts.</p> <p><b>10<sub>B</sub> : Supervisor Mode</b>  Enables access to all peripheral devices. It enables read/write access to core registers and protected peripheral devices. Tasks at this level may disable interrupts.</p> <p><b>11<sub>B</sub> : Reserved Value</b></p>
IS	9	rw	<p><b>Interrupt Stack Control</b>  Determines if the current execution thread is using the shared global (interrupt) stack or a user stack.</p> <p><b>0 : User Stack</b>  If an interrupt is taken when the IS bit is 0, then the stack pointer register is loaded from the ISP register before execution starts at the first instruction of the Interrupt Service Routine (ISR).</p> <p><b>1 : Shared Global Stack</b>  If an interrupt is taken when the PSW.IS bit is 1, then the current value of the stack pointer is used by the Interrupt Service Routine (ISR).</p>

**General Purpose and System Registers**

Field	Bits	Type	Description
GW	8	rw	<p><b>Global Address Register Write Permission</b></p> <p>Determines whether the current execution thread has permission to modify the global address registers. Most tasks and ISRs use the global address registers as 'read only' registers, pointing to the global literal pool and key data structures. However a task or ISR can be designated as the 'owner' of a particular global address register, and is allowed to modify it. The system designer must determine which global address variables are used with sufficient frequency and/or in sufficiently time-critical code to justify allocation to a global address register. By compiler convention, global address register A[0] is reserved as the base register for short form loads and stores. Register A[1] is also reserved for compiler use. Registers A[8] and A[9] are not used by the compiler, and are available for holding critical system address variables.</p> <p>0 : Write permission to global registers A[0], A[1], A[8], A[9] is disabled.</p> <p>1 : Write permission to global registers A[0], A[1], A[8], A[9] is enabled.</p>
CDE	7	rw	<p><b>Call Depth Count Enable</b></p> <p>Enables call-depth counting, provided that the PSW.CDC mask field is not all set to 1.</p> <p>0 : Call depth counting is temporarily disabled. It is automatically re-enabled after execution of the next Call instruction.</p> <p>1 : Call depth counting is enabled.</p> <p>If PSW.CDC = 1111111<sub>B</sub>, call depth counting is disabled regardless of the setting on the PSW.CDE bit.</p>



**General Purpose and System Registers**

Field	Bits	Type	Description
CDC	[6:0]	rw	<p><b>Call Depth Counter</b></p> <p>Consists of two variable width subfields. The first subfield consists of a string of zero or more initial 1 bits, terminated by the first 0 bit.</p> <p>The remaining bits form the second subfield (CDC.COUNT) which constitutes the call depth count value. The count value is incremented on each Call and is decremented on a Return.</p> <p>0cccccc<sub>B</sub> : 6-bit counter; trap on overflow.  10cccccc<sub>B</sub> : 5-bit counter; trap on overflow.  110cccccc<sub>B</sub> : 4-bit counter; trap on overflow.  1110cccc<sub>B</sub> : 3-bit counter; trap on overflow.  11110ccc<sub>B</sub> : 2-bit counter; trap on overflow.  111110cc<sub>B</sub> : 1-bit counter; trap on overflow.  1111110<sub>B</sub> : Trap every call (call trace mode).  1111111<sub>B</sub> : Disable call depth counting.</p> <p>When the call depth count (CDC.COUNT) overflows a trap (CDO) is generated.</p> <p>Setting the CDC to 1111110<sub>B</sub> allows no bits for the counter and causes every call to be trapped. This is used for Call Depth Tracing.</p> <p>Setting the CDC to 1111111<sub>B</sub> disables call depth counting.</p>

### PSW User Status Bits

The eight most significant bits of the PSW are designated as User Status Bits. These bits may be set or cleared as execution side effects of user instructions, typically recording result status. Individual bits can also be used to condition the operation of particular instructions. For example the ADDX (Add Extended) and ADDC (Add with Carry) instructions use bit 31 to record the carry out from the ADD operation, and the pre-execution value of the bit is reflected in the result of the ADDC instruction.

**Table 5 PSW User Status Bits**

Field	Bits	Type	Description
C	31	rw	Carry.
V	30	rw	Overflow.
SV	29	rw	Sticky Overflow.
AV	28	rw	Advance Overflow.
SAV	27	rw	Sticky Advance Overflow.
-	[26:24]	-	Reserved Field.

There are two classes of instructions that employ the user status bits:

- Arithmetic instructions that may produce carry and overflow results.
- Implementation-specific coprocessor instructions which may use any or all of the eight bits, in a manner that is entirely implementation specific.

Bits [23:16] of the PSW are reserved bits with no defined use in current versions of the architecture. They read as zero when the PSW is read via the MFCR (Move From Core Register) instruction after a system reset. Their value after writing to the PSW via the MTCR (Move To Core Register) instruction, is architecturally undefined and should be written as zero.

### Access Privilege Level Control (I/O Privilege)

Software Managed Tasks (SMTs) are created through the services of a real-time kernel or Operating System, and are dispatched under the control of scheduling software. Interrupt Service Routines (ISRs) are dispatched by hardware in response to an interrupt. An ISR is the code that is invoked directly by the processor on receipt of an interrupt. SMTs are sometimes referred to as user tasks, assuming that they execute in User Mode.

Each task is allocated its own mode, depending on the task's function:

- **User-0 Mode:** Used for tasks that do not access peripheral devices. This mode may not enable or disable interrupts.
- **User-1 Mode:** Used for tasks that access common, unprotected peripherals. Typically this would be a read or write access to serial port, a read access to timer, and most I/O status registers. Tasks in this mode may disable interrupts.
- **Supervisor Mode:** Permits read/write access to system registers and all peripheral devices. Tasks in this mode may disable interrupts.

A set of state elements are associated with any task, and these are known collectively as the task's context. The context is everything the processor needs to define the state of the associated task and enable its continued execution. This includes the CPU General Registers that the task uses, the task's Program Counter (PC), and its Program Status Information (PCXI and PSW). The architecture efficiently manages and maintains the context of the task through hardware.

**General Purpose and System Registers**

### 3.2.3 Previous Context Information and Pointer Register (PCXI)

The Previous Context Information Register (PCXI) contains linkage information to the previous execution context, supporting fast interrupts and automatic context switching. The PCXI is part of a task's state information. The Previous Context Pointer (PCX) holds the address of the CSA of the previous task.


#### PCXI, PCX

##### Previous Context Information

(FE00<sub>H</sub>)

##### Previous Context Pointer

**Reset Value: Implementation Specific**



Field	Bits	Type	Description
PCPN	[31:24]	rw	<b>Previous CPU Priority Number</b> Contains the priority level number of the interrupted task.
PIE	23	rw	<b>Previous Interrupt Enable</b> Indicates the state of the interrupt enable bit (ICR.IE) for the interrupted task.
UL	22	rw	<b>Upper or Lower Context Tag</b> Identifies the type of context saved: 0 : Lower Context. 1 : Upper Context. If the type does not match the type expected when a context restore operation is performed, a trap is generated.
-	[21:20]	-	<b>Reserved Field</b>
PCXS	[19:16]	rw	<b>Previous Context Pointer Segment Address</b> Contains the segment address portion of the PCX. This field is used in conjunction with the PCXO field.
PCXO	[15:0]	rw	<b>Previous Context Pointer Offset Field</b> The PCXO and PCXS fields form the pointer PCX, which points to the CSA of the previous context.

### **3.3 Stack Management Registers**

Stack management in the architecture supports a user stack and an interrupt stack. Address register A[10], the Interrupt Stack Pointer (ISP) and a PSW bit are used in the management of the stack.

- A[10](SP): A[10] (Stack Pointer) [page 3-14](#).
- ISP: Interrupt Stack Pointer [page 3-15](#).

A[10] is used as the stack pointer. The initial contents of this register are usually set by an RTOS when a task is created, which allows a private stack area to be assigned to individual tasks.

The ISP helps to prevent Interrupt Service Routines (ISRs) from accessing the private stack areas and possibly interfering with the software managed task's context. An automatic switch to the use of the ISP instead of the private stack pointer is implemented in the architecture. The PSW.IS bit indicates which stack pointer is in effect. When an interrupt is taken and the interrupted task was using its private stack (PSW.IS == 0), the contents are saved with the upper context of the interrupted task and A[10](SP) is loaded with the current contents of the ISP.

When an interrupt or trap is taken and the interrupted task was already using the interrupt stack (PSW.IS == 1), then no pre-loading of A[10](SP) is performed. The Interrupt Service Routine (ISR) continues to use the interrupt stack at the point where the interrupted routine had left it.

Usually it is only necessary to initialize the ISP once during the initialization routine. However, depending on application needs, the ISP can be modified during execution. Note that there is nothing preventing an ISR or system service routine from executing on a private stack.

*Note: Use of A[10](SP) in an ISR is at the discretion of the application programmer.*

## General Purpose and System Registers


### 3.3.1 Address Register A[10] (SP)

The A[10] Stack Pointer (SP) register is defined as follows:

**A[10](SP)**

**Address Register A[10] (Stack Pointer)(FFA8<sub>H</sub>)**

**Reset Value: Implementation Specific**



Field	Bits	Type	Description
A[10](SP)	[31:0]	rw	Address Register A[10] (Stack Pointer)

### 3.3.2 Interrupt Stack Pointer (ISP)


The Interrupt Stack Pointer is defined as follows.

#### ISP

Interrupt Stack Pointer

(FE28<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
ISP	[31:0]	rw	Interrupt Stack Pointer

*Note: This register is ENDINIT protected.*

## General Purpose and System Registers

### 3.4 System Control Register (SYSCON)

The System Configuration Register provides the enable/disable bit for the memory protection system and a status flag for the Free Context List Depletion condition.

#### SYSCON

#### System Configuration Register

**Reset Value: 0000 0000<sub>H</sub>**

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
-														<b>PRO TEN</b>	<b>FCD SF</b>
														rw	rwh

Field	Bits	Type	Description
-	[31:2]	-	<b>Reserved Field</b>
PROTEN	1	rw	<b>Memory Protection Enable</b> Enables the memory protection system. Memory protection is controlled through the memory protection register sets. Note: Initialize the protection register sets prior to setting PROTEN to one. 0 : Memory Protection is disabled. 1 : Memory Protection is enabled.
FCDSF	0	rwh	<b>Free Context List Depleted Sticky Flag</b> This sticky bit indicates that a FCD (Free Context List Depleted) trap occurred since the bit was last cleared by software. 0 : No FCD trap occurred since the last clear. 1 : An FCD trap occurred since the last clear.



**General Purpose and System Registers**

### 3.5 CPU Identification Register (CPU\_ID)


Identification Registers identify the processor type and revision used. Only the CPU core ID register is described here. All other ID registers are described in the product documentation. The CPU Identification Register identifies the CPU type and revision.

#### CPU\_ID

##### CPU Module Identification

(FE18<sub>H</sub>)

**Reset Value: Implementation Specific**



Field	Bits	Type	Description
MOD	[31:16]	r	<b>Module Identification Number</b> Used for module identification.
MOD_32B	[15:8]	r	<b>32-Bit Module Enable</b> A value of C0 <sub>H</sub> in this field indicates a 32-bit module with a 32-bit module ID register.
MOD_REV	[7:0]	r	<b>Module Revision Number</b> Used for revision numbering. The value of the revision starts at 01 <sub>H</sub> (first revision) up to FF <sub>H</sub> .

### 3.6 Compatibility Mode Register (COMPAT)

*Note: TriCore 1.3.1 architecture only*

The COMPAT register is provided to allow implementations to selectively force compatibility of features with previous versions.

The contents of the register are implementation specific.

#### COMPAT

Compatibility Mode Register

(9400<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
-	[31:0]	-	Implementation Specific

*Note: This register is ENDINIT protected.*

## 3.7 Access Control Registers

### 3.7.1 BIST Mode Access Control Register (BMACON)

*Note: TriCore 1.3.1 architecture only.*

Implementations may control the operation of Built in Self Test (BIST) systems using the BMACON register. The contents of this register is implementation specific.

#### BMACON

##### BIST Mode Access Control

(9004<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
-	[31:0]	-	Implementation Specific

*Note: This register is ENDINIT protected*

### 3.7.2 SIST Mode Access Control Register (SMACON)

*Note: TriCore 1.3.1 architecture only.*

Implementations may control the operation of Software in System Test (SIST) systems using the SMACON register. The contents of this register is implementation specific.

#### SMACON

SIST Mode Access Control

(900C<sub>H</sub>)

Reset Value: Implementation Specific

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
Implementation Specific															

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Implementation Specific															

Field	Bits	Type	Description
-	[31:0]	-	Implementation Specific

*Note: This register is ENDINIT protected*

### **3.8 Interrupt Registers**

A typical Service Request Control register in the TriCore architecture holds the individual control bits to enable or disable the request, to assign a priority number, and to direct the request to one of the service providers. The Core Special Function Registers (CSFR) which control the Interrupts are described in [Interrupt System, page 5-1](#).

### **3.9 Memory Protection Registers**

The number of Memory Protection Register Sets is specific to each implementation of the architecture. There can be a maximum number of four sets (one set includes both a data set and a code set). Each register set is made up of several range registers (also called Range Table Entries).

Each Range Table Entry consists of a Segment Protection register pair and a bit field within a common Mode register. The register pair specifies the lower and upper boundary addresses of the memory range.

The Core Special Function Registers (CSFR) which control the Memory Protection Registers are described in [Memory Protection System, page 9-1](#).

### **3.10 Trap Registers**

The Core Special Function Registers (CSFR) which control the Trap Registers are described in [Trap System, page 6-1](#).

### **3.11 Memory Management Unit Registers**

The optional Memory Management Unit (MMU) supports virtual memory and page-based memory access protection. The Core Special Function Registers (CSFR) which control the optional MMU are described in [Memory Management Unit \(MMU\), page 10-1](#).

### **3.12 Core Debug Controller Registers**

TriCore 1 registers that support debugging are described in [Core Debug Controller \(CDC\), page 12-1](#)

### **3.13 Floating Point Registers (TriCore 1.3.1)**

The registers for the optional TriCore Floating Point Unit are described on [FPU\\_TRAP\\_CON, page 11-14](#).

### **3.14 Updating Core Special Function Registers (CSFRs)**

The need for software updates to CSFRs is usually infrequent. Implementations are therefore not required to implement hardware structures to avoid hazard conditions that may result from the update of CSFRs. Such hazard conditions are avoided by the insertion of an ISYNC instruction immediately after the MTCR update of the CSFR. The ISYNC instruction ensures that the effects of the CSFR update are correctly seen by all following instructions.

## **4 Tasks and Functions**

Most embedded and real-time control systems are designed according to a model in which interrupt handlers and software-managed tasks are each considered to be executing on their own 'virtual' microcontroller. That model is generally supported by the services of a Real-time Executive or Real-time Operating System (RTOS), layered on top of the features and capabilities of the underlying machine architecture.

In the TriCore® architecture, the RTOS layer can be very 'thin' and the hardware can efficiently handle much of the switching between one task and another. At the same time the architecture allows for considerable flexibility in the tasking model used. System designers can choose the real-time executive and software design approach that best suits the needs of their application, with relatively few constraints imposed by the architecture.

The mechanisms for low-overhead task switching and for function calling within the TriCore architecture are closely related.

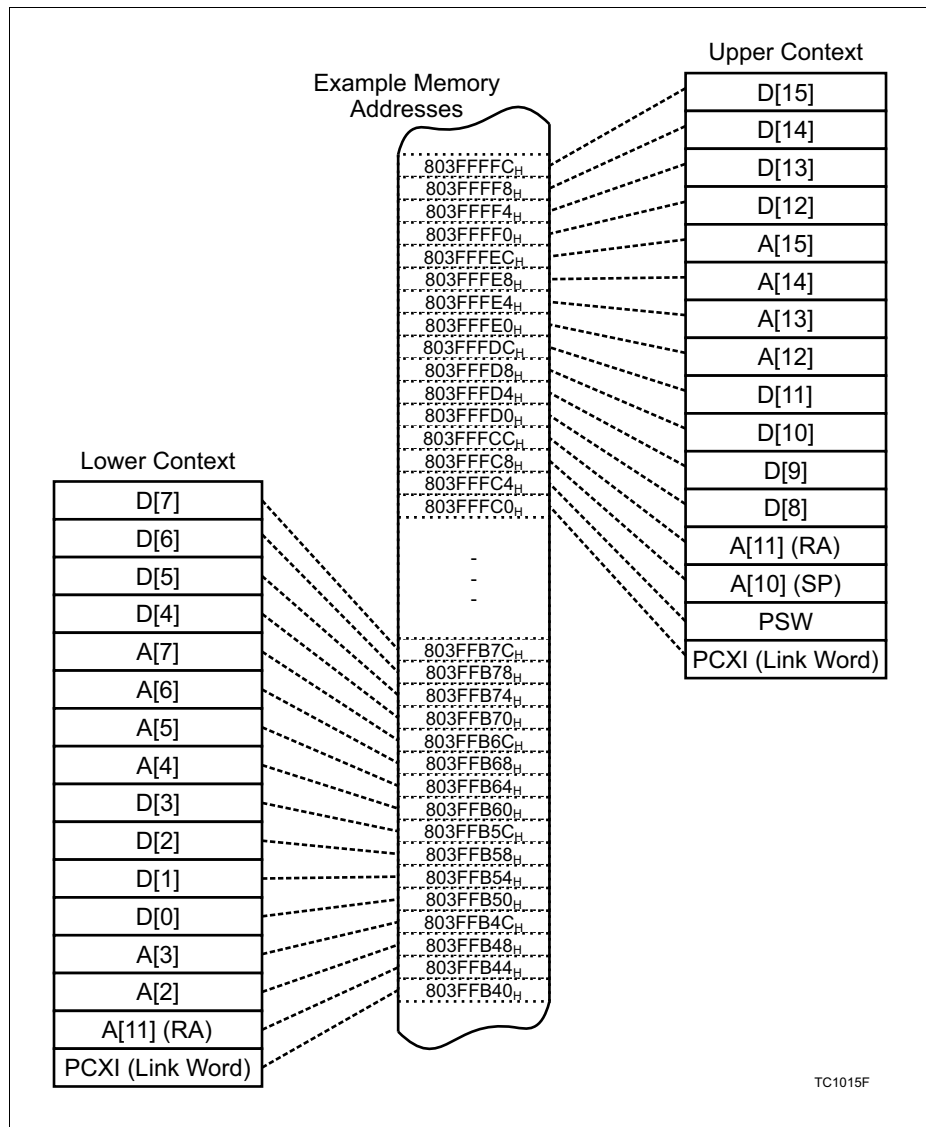
### **4.1 Context Types**

A task is an independent thread of control. The state of a task is defined by its context. When a task is interrupted, the processor uses that task's context to re-enable the continued execution of the task.

The context types are:

- Upper context: Consists of the upper address registers A[10] to A[15] and the upper data registers D[8] to D[15]. The upper context also includes PCXI and PSW. These registers are designated as non-volatile for purposes of function-calling (their contents are preserved across calls).
- Lower context: Consists of the lower address registers A[2] to A[7], the lower data registers D[0] to D[7], A[11] (Return Address) and PCXI.

Contexts, when saved to memory, occupy 16 word blocks of storage, known as Context Save Areas (CSAs).



**Figure 14 Upper and Lower Contexts**

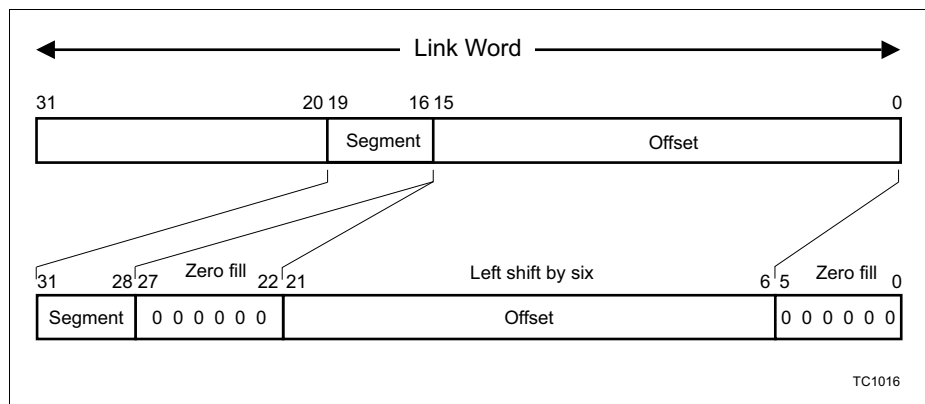


### 4.1.1 Context Save Area

The architecture uses linked lists of fixed-size Context Save Areas. A CSA is 16 words of memory storage, aligned on a 16 word boundary. Each CSA can hold exactly one upper or one lower context. CSAs are linked together through a Link Word.

The Link Word includes two fields that link the given CSA to the next one in a chain. The fields are a 4-bit segment and a 16-bit offset. The segment number and offset are used to generate the Effective Address (EA) of the linked CSA. See [Figure 15](#).

Incrementing the pointer offset value by one always increments the EA to the address that is 16 word locations above the previous one. The total usable range in each address segment for CSAs is 4 MBytes, resulting in storage space for  $2^{16}$  CSAs.



**Figure 15 Generation of the Effective Address of a Context Save Area (CSA)**

If the CSA is in use (for example, it holds an upper or lower context image for a suspended task), then the Link Word also contains other information about the linked context. The entire Link Word is a copy of the PCXI register for the associated task.

For further information on how linked CSAs support context switching, refer to [Context Save Areas \(CSAs\) and Context Lists, page 4-5](#)

## 4.2 Task Switching Operation

The architecture switches tasks when one of the events or instructions listed in [Table 6](#), occurs. When one of these events or instructions is encountered, the upper or lower context of the task is saved or restored. The upper context is saved automatically as a result of an external interrupt, trap or function call. The lower context is saved explicitly through instructions. In [Table 6](#) 'Save' is a store through the Free CSA List Head Pointer register (FCX) after the next value for the FCX is read from the Link Word. 'Store' is a store through the Effective Address of the instruction with no change to the CSA list or the FCX register. 'Restore' is the converse of 'Save'. 'Load' is the converse of 'Store'.

There is an essential difference in the treatment of registers in the upper and lower contexts, in terms of how their contents are maintained. The lower context registers are similar to global registers in the sense that a interrupt handler, trap handler or called function, sees the same values that were present in the registers just before the interrupt, trap or call. Any changes made to those registers that are made in the interrupt, trap handler or called function, remains present after the return from the event, since they are not automatically restored as part of the Return From Call (RET) or Return From Exception (RFE) semantics. That means that the lower context registers can be used to pass arguments to called functions and pass return values from those functions. It also means that interrupt and trap handlers must save the original values they find in these registers before using the registers, and to restore the original values before exiting.

The upper context registers are not guaranteed to be static hardware registers. Conceptually, a function call or interrupt handler always begins execution with its own private set of upper context registers. The upper context registers of the interrupted or calling function are not inherited.

Only the A[10](SP), A[11](RA), PSW, PCXI and (in the case of a trap) D[15] registers start with architecturally defined values in the called function, trap handler or interrupt handler. A function, trap handler or interrupt handler that reads any of the other upper context registers before writing a value into it, is performing an undefined operation.

**Table 6 Context Related Events and Instructions**

Event / Instruction	Context Operation	Complement Instruction	Context Operation
Interrupt	Save Upper	RFE - Return from Exception	Restore Upper
Trap	Save Upper	RFE - Return from Exception	Restore Upper
CALL - Function Call	Save Upper	RET - Return from Call	Restore Upper
BISR - Begin Interrupt Service Routine	Save Lower	RSLCX - Restore Lower Context	Restore Lower
SVLCX - Save Lower Context	Save Lower	RSLCX - Restore Lower Context	Restore Lower

**Table 6 Context Related Events and Instructions (Continued)**

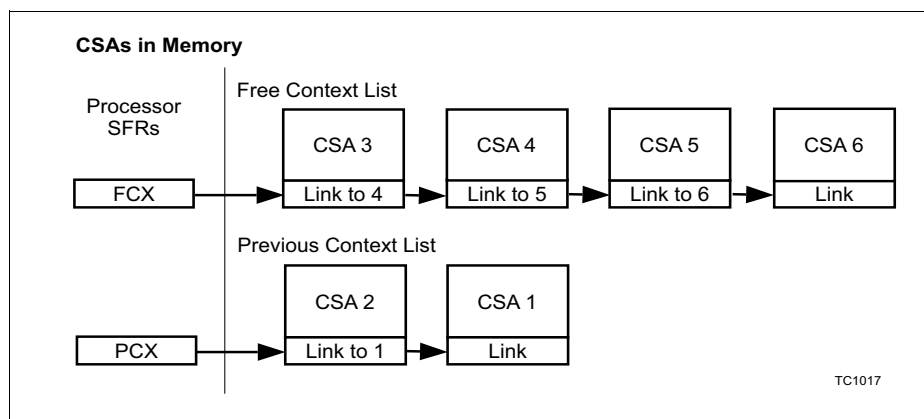
Event / Instruction	Context Operation	Complement Instruction	Context Operation
STLCX - Store Lower Context	Store Lower	LDLCX - Load Lower Context	Load Lower
STUCX - Store Upper Context	Store Upper	LDUCX - Load Upper Context	Load Upper

### 4.2.1 Save and Restore Context Operations

The Effective Address of all context related operations must be a physical memory address which maps to cached memory or data scratchpad RAM (of the processor performing the access). Using address ranges not covered by physical memories will lead to undefined results.

### 4.3 Context Save Areas (CSAs) and Context Lists

The upper and lower contexts are saved in Context Save Areas (CSAs). Unused CSAs are linked together in the free context list. CSAs that contain saved upper or lower contexts are linked together in the previous context list. The following figure (Figure 16) shows a simple configuration of CSAs within both context lists.



**Figure 16 CSAs in Context Lists**

The contents of the FCX register always points to an available CSA in the free context list. That CSAs Link Word points to the next available CSA in the free context list.

**Tasks and Functions**

Before an upper or lower context is saved in the first available CSA, its Link Word is read, supplying a new value for the FCX. To the memory subsystem, context saving is therefore a read/modify/write operation. The new value of FCX, which points to the next available CSA, is available immediately for subsequent upper or lower context saves.

The LCX register points to one of the last CSAs in the free list and is used to recognise impending free CSA list depletion. If the value of FCX matches that of LCX when an operation that performs a context save is attempted, the operation completes and a free CSA list depletion trap (FCD) is taken on the next instruction; i.e., the return address of the FCD trap is the first instruction of the trap/interrupt/called routine or the instruction following an SVLCX or BISR instruction. See [Context Management \(Trap Class 3\), page 6-10](#).

The action taken by the trap handler depends on the software implementation. It might issue a system reset for example, if it is determined that the CSA list depletion resulted from an unrecoverable software error. Normally however it extends the free list, either by allocating additional memory or by terminating one or more tasks and reclaiming their CSA call chains. In those cases the trap handler exits with a RFE instruction.

The link word in the last CSA in a free context list must be set to null before it is first used. This is necessary to support the FCU trap. Before first use of the CSA, the PCX pointer value should be null. This is to support CSU (Call Stack Underflow) traps.

The PCXI.PCX field points to the CSA where the previous context was saved. The PCXI.UL bit identifies whether the saved context is upper (PCXI.UL == 1) or lower (PCXI.UL == 0). If the type does not match the type expected when a context restore operation is performed, a CYTP exception occurs and a context management trap is taken.

After the context save operation has been performed the Return Address A[11](RA) is updated:

- For a call, the A[11](RA) is updated with the function return address.
- For a synchronous trap, the A[11](RA) is updated with the PC of the instruction which raised the trap.
- For a SYSCALL and an asynchronous trap or an interrupt, the A[11](RA) is updated with the PC of the next instruction to be executed.

When a lower context save operation is performed the value of A[11](RA) is included in the saved context and is placed in the second word of the CSA. This A[11](RA) is correspondingly restored by a lower context restore.

The Call Depth Control field (PSW.CDC) consists of two subfields; A call depth counter, and a mask that determines the width of the counter and when it overflows.

The Call Depth Counter is incremented on calls and is restored to its previous value on returns. An exception occurs when the counter overflows. Its purpose is to prevent software errors from causing 'runaway recursion' and depleting the CSA free list.

## **4.4 Context Switching with Interrupts and Traps**

When an interrupt or trap (for example NMI or SYSTRAP) occurs, the processor saves the upper context of the current task in memory, suspends execution of the current task and then starts execution of the interrupt or trap handler.

If, when an interrupt or trap is taken, the processor is not using the interrupt stack (PSW.IS bit == 0), the Stack Pointer is then loaded with the current contents of the ISP (Interrupt Stack Pointer). The PSW.IS bit is then set to one (1) to indicate execution from the interrupt stack.

The Interrupt Control Register (ICR) holds the Current CPU Priority Number (ICR.CCPN), the Interrupt Enable bit (ICR.IE) and Pending Interrupt Priority Number (ICR.PIPN). These fields, together with the Previous CPU Priority Number (PCXI.PCPN) and Previous Interrupt Enable (PCXI.PIE) are all part of the interrupt management system.

ICR.CCPN is typically only non-zero within Interrupt Service Routines (ISRs) where it is used to order interrupt servicing. It is held in a register that is separate from the PSW and is not part of the context that the RTOS handles for switching among Software Managed Tasks (SMTs).

PCXI.PIE is only typically zero within Trap handlers started within ISRs, e.g. an NMI or SYSTRAP occurring during a peripheral service request.

For both interrupts and traps, the existing PCPN and PIE values in the current PCXI are saved in the CSA for the upper context, and the existing IE and CCPN values in the ICR are copied to the PCXI.PIE and PCXI.PCPN fields. Once the interrupt or trap is handled, the saved lower context is reloaded if necessary and execution of the interrupted task is resumed (RFE).

On an interrupt or trap the upper context of the current task context is saved by hardware as an explicit part of the interrupt or trap sequence. For small interrupt and trap handlers that can execute entirely within this set of registers saved on the interrupt, no further context saving is needed. The handler can execute immediately and return. Typically handlers that make calls or require more registers execute the BISR (Begin Interrupt Service Routine) or SVLCX (Save Lower Context) instruction to save the lower context registers that were not saved as part of the interrupt or trap sequence. That instruction must be issued before any of the associated registers are modified, but it need not be the first instruction in the handler.

Interrupt handlers with critical response time requirements can perform their initial, time-critical processing immediately, using upper context registers. After that they can execute a BISR and continue with less time-critical processing. The BISR re-enables interrupts, hence its use dividing time critical from less time critical processing.

Trap handlers typically do not have critical response time requirements, however those that can occur in an ISR or those which might hold off interrupts for too long can also take a similar approach to distinguish between non-interruptible and interruptible execution segments.

## **4.5 Context Switching for Function Calls**

When a function call is made (the CALL instruction is executed), the context of the calling routine must be saved and then restored in order to resume the caller's execution after return from the function.

On a function call the entire set of upper context registers are saved by hardware. Furthermore, the saving of the upper context by the CALL instruction happens in parallel with the call jump. In addition, restoring the upper context is performed by the RET (Return) instruction and takes place in parallel with the return jump. The called function does not need to save and restore the caller's context and is freed of any need to restrict its usage of the upper context registers. The calling and called functions can co-operate on the use of the lower context registers.

## 4.6 Context Save and Restore Examples

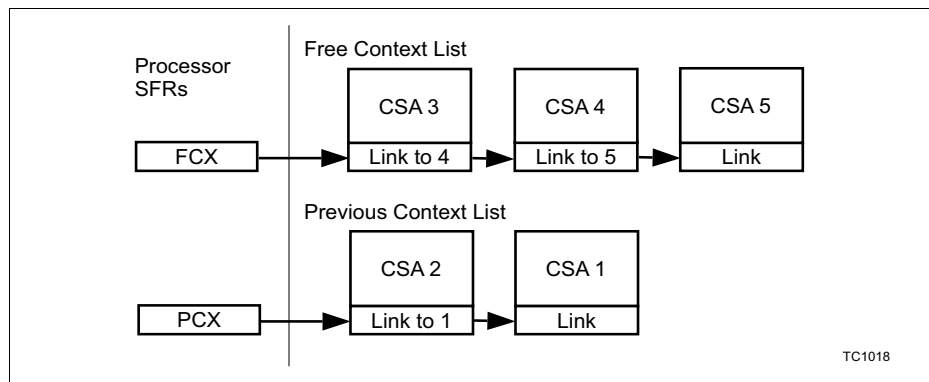
This section provides an example of a context save operation and an example of a context restore operation.

### 4.6.1 Context Save

**Figure 17** shows the free and previous context lists for this example. The free context list (FCX) contains three free CSAs (3, 4, and 5), and the previous context list (PCX) contains two CSAs (2 and 1).


The FCX points to CSA3, the first available CSA. The Link Word of CSA3 points to CSA4; the Link Word of CSA4 points to CSA5. The PCX points to the most recently saved CSA in the previous context list. The Link Word of CSA2 points to CSA1. CSA1 contains the saved context prior to CSA2.

When the context save operation is performed, the first CSA in the free context list (CSA3) is pulled off and is placed on the front of the previous context list.



**Figure 17 CSAs and Processor State Prior to Context Save**

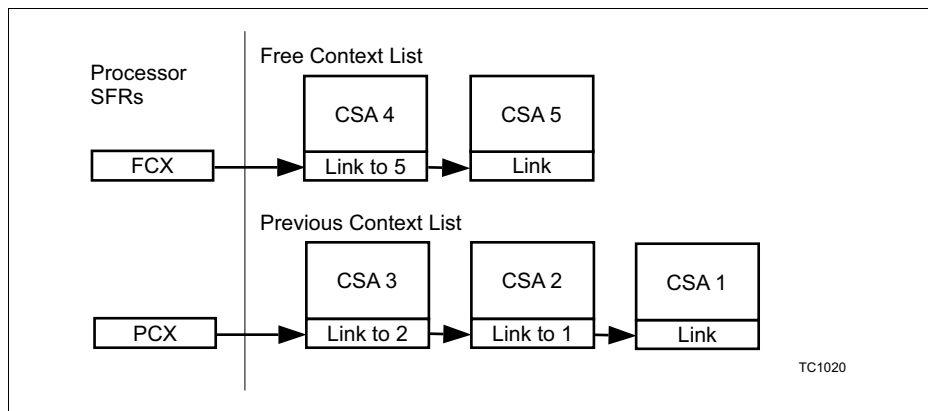
**Figure 18** shows the steps taken during the context save operation. The numbers in the figure correspond to the steps listed after the figure.



**Figure 18 CSA and Processor SFR Updates on a Context Save Process**

1. The contents of the Link Word in CSA3 are loaded into the NEW\_FCX. The NEW\_FCX now points to CSA4. The NEW\_FCX is an internal buffer and is not accessible by the user.
2. The contents of the PCX are written into the Link Word of CSA3. The Link Word of CSA3 now points to CSA2.
3. The contents of FCX are written into the PCX. The PCX now points to CSA3, which is at the front of the Previous Context List.
4. The NEW\_FCX is loaded into the FCX.

The processor SFRs and CSAs look as shown in [Figure 19](#). The processor context to be saved is now written into the rest of CSA3.



**Figure 19 CSAs and Processor State After Context Save**

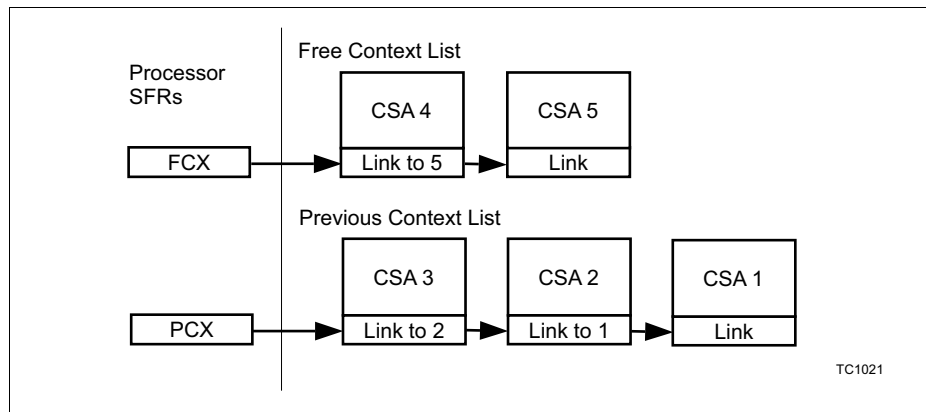


### 4.6.2 Context Restore

The example in [Figure 20](#), shows the previous context list (PCX) with three CSAs (3, 2, and 1) and the free context list (FCX) containing two CSAs (4 and 5).

The FCX points to CSA4, the first available CSA in the free context list. PCX points to CSA3, the most recently saved CSA in the previous context list.

The Link Word of CSA3 points to CSA2; the Link Word of CSA2 points to CSA1; the Link Word of CSA4 points to CSA5.




**Figure 20 CSAs and Processor State Prior to Context Restore**

When the context restore operation is performed, the first CSA in the previous context list (CSA3) is pulled off and is placed on the front of the free context list.

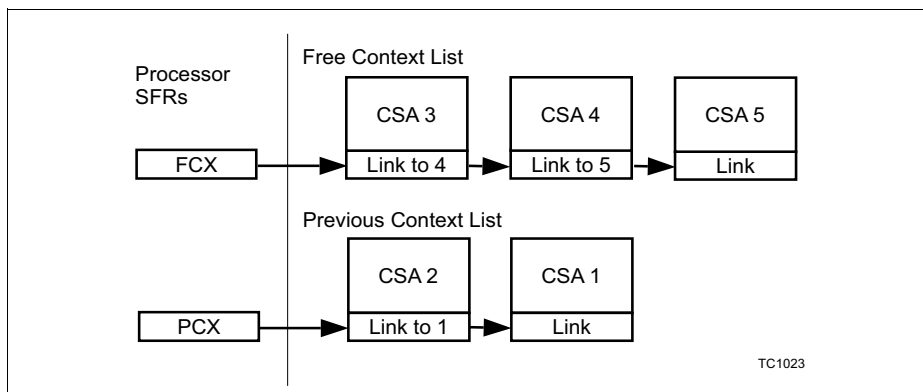
[Figure 21](#) shows the steps taken during the context restore operation. The numbers in the figure correspond to the following steps:

1. The contents of the Link Word in CSA3 are loaded into the NEW\_PCX. The NEW\_PCX now points to CSA2. The NEW\_PCX is an internal buffer and is not accessible by the user.
2. The contents of the FCX are written into the Link Word of CSA3. The Link Word of CSA3 now points to CSA4.
3. The contents of the PCX are written into the FCX. The FCX now points to CSA3, which is at the front of the free context list.
4. The NEW\_PCX is loaded into the PCX.



**Figure 21 CSA and Processor SFR updates on a Context Restore Process**

The processor SFRs and CSAs now look as shown in [Figure 22](#). The restored context is then written into the upper or lower context registers.



**Figure 22 CSAs and Processor State After Context Restore**

## 4.7 Context Management Registers

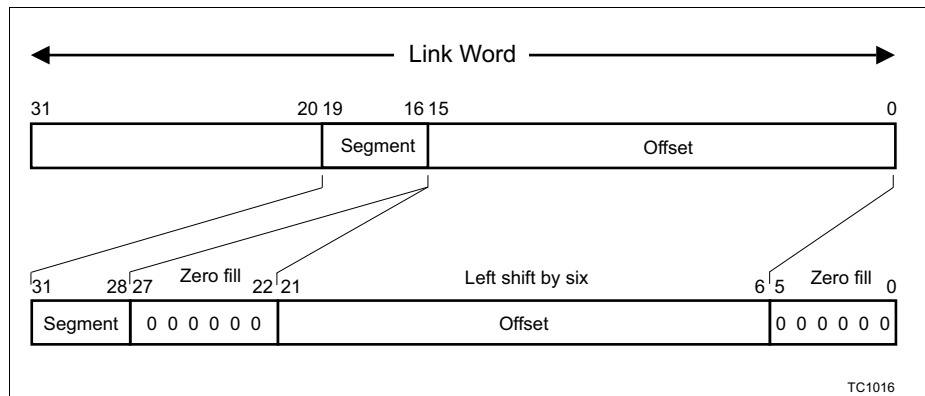
The three context management registers are pointers that are used during context save and restore operations.

- FCX: Free CSA List Head Pointer [page 4-14](#).
- PCX: Previous Context Pointer [page 4-15](#).
- LCX: Free CSA List Limit Pointer [page 4-16](#).

Each pointer consists of two fields:

- A 16-bit offset.
- A 4-bit segment specifier.

**Table 23** shows how the effective address of a Context Save Area (CSA) is generated using these two fields. A Context Save Area is an address range containing 16 word locations (64 bytes), which is the space required to save one upper or one lower context. Incrementing the pointer offset value by one always increments the Effective Address (EA) to the address that is 16 word locations above the previous one. The total usable range in each address segment for CSAs is 4 MBytes, resulting in storage space for 64 KByte CSAs.



**Figure 23 Generation of the Effective Address of a Context Save Area (CSA)**

*Note:* See [Context Save Area, page 4-3](#) for additional constraints on the Effective Address (EA).

### 4.7.1 Free CSA List Head Pointer Register (FCX)


The Free CSA List Head Pointer (FCX) register holds the free CSA list head pointer. This always points to an available CSA.

#### FCX

**Free CSA List Head Pointer**

**(FE38<sub>H</sub>)**

**Reset Value: Implementation Specific**



Field	Bits	Type	Description
-	[31:20]	-	<b>Reserved Field</b>
FCXS	[19:16]	rw	<b>FCX Segment Address Field</b> Used in conjunction with the FCXO field.
FCXO	[15:0]	rw	<b>FCX Offset Address Field</b> The FCXO and FCXS fields together form the FCX pointer, which points to the next available CSA.


#### 4.7.2 Previous Context Pointer Register (PCX)

The Previous Context Pointer (PCX) holds the address of the CSA of the previous task. The PCX is part of the PCXI register.

##### PCX

##### Previous Context Pointer Register (FE00<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
	[31:20]	-	These bits are not relevant to the pointer function and so are not described here. See the PCXI register.
PCXS	[19:16]	rw	<b>Previous Context Pointer Segment Address Field</b> This field is used in conjunction with the PCXO field.
PCXO	[15:0]	rw	<b>Previous Context Pointer Offset Field</b> The PCXO and PCXS fields form the pointer PCX, which points to the CSA of the previous context.

#### 4.7.3 Free CSA List Limit Pointer Register (LCX)

The free CSA List Limit Pointer (LCX) register is used to recognize impending free CSA list depletion. If a context save operation occurs and the value of FCX matches LCX then the 'free context depletion' condition is recognized, which triggers an FCD trap immediately after completion of the operation causing the context save; i.e. the return address of the FCD trap is the first instruction of the trap/interrupt/called routine, or the instruction following an SVLCX or BISR instruction.


*Note: Please refer to the FCD trap description for details on the use and setting of LCX. See [FCD - Free Context list Depletion \(TIN 1\)](#), page 6-10.*

#### LCX

Free CSA List Limit Pointer

(FE3C<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
-	[31:20]	-	<b>Reserved Field</b>
LCXS	[19:16]	rw	<b>LCX Segment Address</b> This field is used in conjunction with the LCXO field.
LCXO	[15:0]	rw	<b>LCX Offset Field</b> The LCXO and LCXS fields form the pointer LCX, which points to the last available CSA.

## **4.8 Accessing CSA Memory Locations**

Implementations may internally buffer context information to increase performance. To ensure memory coherency, a DSYNC instruction must be executed prior to any access to an active CSA memory location. The DSYNC instruction forces all internally buffered CSA register state to be written to memory.





## **5 Interrupt System**

This chapter describes the interrupt system, including arbitration, the priority level scheme, and access to the vector table.

In a TriCore® system, multiple sources such as peripherals or external inputs can generate an interrupt signal to the CPU to request for service. The interrupt system also supports the implementation of additional units which are capable of handling interrupt requests, such as a second CPU, a standard DMA (Direct Memory Access) unit, or a PCP (Peripheral Control Processor). In the context of this chapter such units are known as 'service providers'. Interrupt requests are often therefore referred to as 'service requests'.

Besides the main CPU, up to three additional service providers can be handled with an interrupt Service Request Node (SRN). The actual number of additional service providers implemented in a given device is implementation dependent.

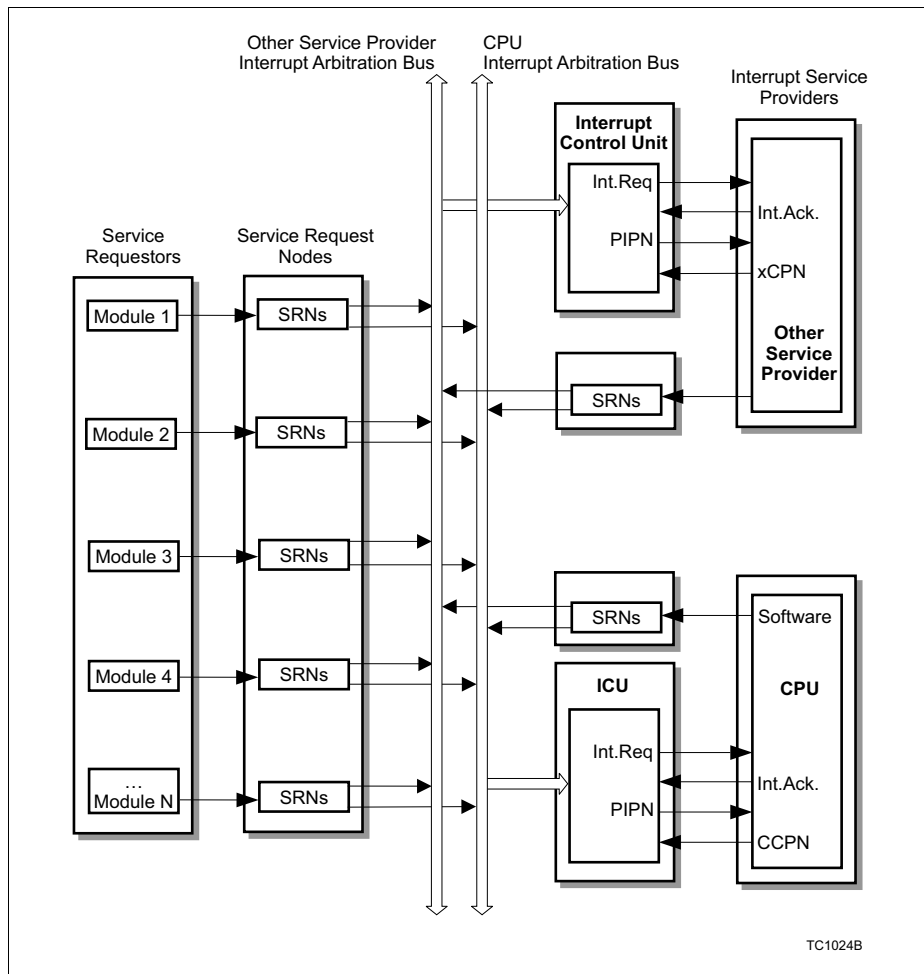
Each interrupt or service request from a module connects to a Service Request Node, containing a Service Request Control Register (SRC). Interrupt arbitration busses connect the SRNs with the interrupt control units of the service providers. These control units handle the interrupt arbitration and communication with the service provider.

**Figure 24, page 5-2** shows an overview of a typical TriCore interrupt system.

### **5.1 Service Request Node (SRN)**

Each Service Request Node contains a Service Request Control Register (SRC) and the necessary logic for communication with the requesting source module and the interrupt arbitration busses. A peripheral or other module can have several service request lines, with each one of them connecting to its own individual SRN.

To support software-posting of interrupts for RTOS code, the TriCore architecture defines four Service Request Nodes (SRNs) which are not attached to a peripheral or any other module on the chip. The interrupt request bit can only be set by software. These SRNs are called the CPU Service Request Nodes. It should be noted however, that the interrupt request can also be set through an external bus master for example.



**Figure 24 Block Diagram of a Typical TriCore Interrupt System**

### 5.1.1 Service Request Control Register (SRC)

A typical Service Request Control register in the TriCore architecture holds the individual control bits to enable or disable the request, to assign a priority number, and to direct the request to one of the service providers. A request status bit shows whether or not the request is active. Besides being activated by the associated module through hardware, each request can also be set or reset through software.

The generic format and description of a Service Request Control register (SRC) is given below.

#### *module\_SRCn*

**Service Request Control (n=0 to 3)**

**Reset Value: Implementation Specific**

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
<b>SET R</b>	<b>CLR R</b>	<b>SRR</b>	<b>SRE</b>	<b>TOS</b>		-		<b>SRPN</b>							
w	w	rh	rw	rw				rw							

Field	Bit	Type	Description
-	[31:16]	-	<b>Reserved Field</b>
SETR	15	w	<b>Service Request Set Bit</b> 0 : No action. 1 : Set SRR (no action if CLRR == 1). Written value is not stored. Read returns 0. No action if CLRR is also set. See <a href="#">Service Request Set and Clear Bits (SETR, CLRR)</a> description.
CLRR	14	w	<b>Service Request Clear Bit</b> 0 : No action. 1 : Clear SRR (no action if SETR == 1). Written value is not stored. Read returns 0. No action if SETR is also set. See <a href="#">Service Request Set and Clear Bits (SETR, CLRR)</a> description.

Field	Bit	Type	Description
SRR	13	rh	<b>Service Request Flag</b> 0 : No Service Request pending. 1 : Service Request is pending. See <a href="#">Service Request Flag (SRR)</a> description.
SRE	12	rw	<b>Service Request Enable Control</b> 0 : Service Request is disabled. 1 : Service Request is enabled. See <a href="#">Service Request Enable Control (SRE)</a> description.
TOS	[11:10]	rw	<b>Type-of-Service Control</b> 00 <sub>B</sub> : Service Provider 0. Typically CPU service is initiated. 01 <sub>B</sub> : Request Service Provider 1. Implementation specific. 10 <sub>B</sub> : Request Service Provider 2. Implementation specific. 11 <sub>B</sub> : Request Service Provider 3. Implementation specific. See <a href="#">Type-of-Service Control (TOS)</a> description.
-	[9:8]	-	<b>Reserved Field</b>
SRPN	[7:0]	rw	<b>Service Request Priority Number</b> 00 <sub>H</sub> : A Service Request on this priority is never serviced. 01 <sub>H</sub> : Service Request, lowest priority. ... FF <sub>H</sub> : Service Request, highest priority. See <a href="#">Service Request Priority Number (SRPN)</a> description.

### Service Request Set and Clear Bits (SETR, CLRR)

These bits enable software to set or clear the actual service request bit SRR.

- Writing 1 to the SETR bit causes the SRR bit to be set to 1.
- Writing 1 to the CLRR bit causes the SRR bit to be cleared to 0.

If hardware attempts to modify SRR during an atomic read-modify-write software operation (such as store) the software operation succeeds and the hardware operation has no effect.

The value written to SETR or CLRR is not stored. Writing zero to these bits has no effect and these bits always return zero when read. If both SETR and CLRR are written to 1 at the same time, the SRR bit is not affected.

### **Service Request Flag (SRR)**

The SRR bit is directly set or reset by the associated hardware. For example, an associated trigger event in a peripheral sets this bit to one and the acknowledgment of the service request by the Service Provider causes this bit to be cleared.

Bit SRR can be set or reset by software via bits SETR or CLRR, respectively. Writing directly to SRR via software has no effect.

SRR can be set or cleared (either by hardware or by software) regardless of the state of the enable bit SRE.

If SRE == 1, a pending service request takes part in the interrupt arbitration of the service provider selected via the TOS bit field. Bit SRR is automatically reset by hardware when the service request is acknowledged and serviced.

If SRE == 0, a pending service request is excluded from interrupt arbitrations. Software can poll SRR to check for a pending service request. SRR must be reset by software in this case (write 1 to CLRR).

### **Service Request Enable Control (SRE)**

The SRE bit controls whether an active interrupt request is passed to the designated interrupt service provider (See the [Type-of-Service Control \(TOS\)](#) description, which follows). If SRE == 1, then the interrupt source associated with this SRN is enabled; i.e. if SRE is set to 1 and the value of the SRR bit moves to 1 (a service request is pending), the Service Request Node (SRN) will participate in interrupt arbitration rounds until the bit is cleared by software or until the interrupt is accepted for presentation to the interrupt service provider indicated by the TOS field. If the SRE bit is set to 0, then the associated interrupt source is disabled.

Disabling an interrupt source by clearing its SRE bit does not affect the setting or clearing of the SRR bit. The SRR bit can still be set by hardware or software (via the SETR bit), and can be read by software, but if the interrupt source is disabled it will not cause a hardware interrupt to be asserted. Users can therefore choose whether to handle the event associated with an individual SRN as an interrupt or through software polling.

### **Type-of-Service Control (TOS)**

The interrupt system is designed to manage up to four Service Providers for service requests from peripherals or other sources. The TOS bit field is used to select the service provider for a request, indicating whether the service request takes part in the interrupt arbitration of the selected service provider. The number of service providers for a given device is implementation specific.

### **Service Request Priority Number (SRPN)**

The 8-bit Service Request Priority Number (SRPN) of a service request, indicates its priority with respect to other sources requesting an interrupt to the same service provider, and to the priority of the service provider itself.

Each SRPN used by active sources requesting the same service provider must be unique at a given time. No active sources can use the same SRPN at the same time, except for the default SRPN of 00<sub>H</sub> which excludes an SRN from taking part in the arbitration. This means that no two or more active sources (requesting CPU service for example) are allowed to use the same SRPN, although they can use the same SRPNs as sources which are requesting another service provider. The term active source in this context means a source which has its request enable bit SRE set to 1, to allow the request to participate in interrupt arbitrations. If a source is not active, meaning its service request enable bit is cleared (SRE == 0), no restrictions are applied to the Service Request Priority Number.

Implementations may look at a subrange of SRPN fields. In such an implementation or configuration the SRPN examined fields must be unique within the examined field.

The SRPN also identifies the entry into the interrupt vector table (or similar structures depending on the nature of the service provider). Unlike other interrupt systems the TriCore vector table provides an entry for each priority number, not for a specific interrupt source. In this way the vector table is de-coupled from the peripherals and a single peripheral can have multiple entry points for different purposes depending on its priority at a given time.

The range for the Service Request Numbers used in a system depends on the number of active service requests and the user-definable organization of the vector table. With the 8-bit SRPN, the interrupt arbitration scheme permits up to 255 sources to be active at one time. More information on the range of SRPNs can be found in [Interrupt Priority Groups, page 5-12](#).

## **5.2 Interrupt Control Unit (ICU)**

The Interrupt Control Unit (ICU) manages the interrupt system and arbitrates incoming interrupt requests to find the one with the highest priority and to determine whether or not to interrupt the service provider. The number of Interrupt Control Units depends on the number of service providers implemented in a TriCore device. Each ICU controls its associated interrupt arbitration bus and manages the communication with its service provider. The ICU is closely coupled with the CPU and its Interrupt Control Register (ICR). This register and the operation of the ICU is described in the sections which follow. In this document, only the CPU Interrupt Control Unit is detailed.

### **5.2.1 ICU Interrupt Control Register (ICR)**

The ICU Interrupt Control Register (ICR) holds the current CPU Priority Number (CCPN), the global Interrupt enable/disable bit (IE) and the Pending Interrupt Priority Number (PIPN), as well as implementation-specific bits to control the interrupt arbitration cycles.

### **5.2.2 Interrupt Control Unit Operation**

When an interrupt service is requested by one or more enabled sources, these requests are serviced depending on their priority ranking. The interrupt system must therefore determine which request has the highest priority each time multiple requests are received. The interrupt system uses a scheme that performs the arbitration in parallel to normal CPU operation. The Interrupt Control Unit (ICU) controls this scheme, which takes place in one or more cycles using the interrupt arbitration bus. The detailed arbitration scheme is implementation specific.

The ICU automatically starts an arbitration when a new interrupt request is detected. At the end of the arbitration the ICU has determined the service request with the highest priority number. This number is stored in the PIPN field of register ICR and generates an interrupt request to the CPU.

The CPU checks the state of the global interrupt enable bit ICR.IE, and compares the current CPU priority number CCPN in register ICR, against the PIPN. The CPU can be interrupted only if ICR.IE == 1 and PIPN is greater than CCPN. If this is true the CPU can enter the service routine; it reads the PIPN to determine the vector entry and acknowledges the ICU, which in turn sends acknowledgement back to the pending interrupt request (the 'winner' of this arbitration round), to inform it that it will be serviced. This node then resets its service request flag (SRR).

After sending the acknowledge, the ICU sets PIPN to 00<sub>H</sub> (no valid pending request) and automatically starts a new arbitration to check whether there is another pending interrupt request. If there is then the priority number of this request is written to PIPN at the end of this arbitration. If there is no pending interrupt request then PIPN remains at 00<sub>H</sub> and the ICU enters an idle state, waiting for the next interrupt request.

*Note: Further CPU interrupt service actions are described in **Entering an Interrupt Service Routine (ISR)**, page 5-8.*

Several conditions could block the CPU from immediately responding to the interrupt request generated by the ICU. These are:

- The interrupt system is globally disabled ( $ICR.IE == 0$ ).
- The current CPU priority CCPN, is equal to or higher than the Pending Interrupt Priority Number (PIP<sub>N</sub>).
- The CPU is in the process of entering an interrupt or trap service routine.
- The CPU is operating on non-interruptible trap services.
- The CPU is executing a multi-cycle instruction.
- The CPU is executing an instruction which modifies the ICR.

The CPU responds to the interrupt request only when these conditions are no longer true.

An arbitration is performed when a new service request is detected, regardless of whether the interrupt system is globally enabled or not, and regardless of whether there are other conditions preventing the CPU from servicing interrupts. In this way the PIP<sub>N</sub> field therefore reflects the pending service request with the highest priority. This can for example, be used for software polling techniques to determine high priority requests while keeping the interrupt system globally disabled.

If a new service request is generated by an SRN while an arbitration is in progress, this request has to wait until at least the end of that arbitration.

### **5.2.3 Arbitration Scheme**

The arbitration scheme is implementation specific and is detailed in the documentation accompanying a specific TriCore product.

### **5.3 Entering an Interrupt Service Routine (ISR)**

When all conditions are clear for the CPU to service an interrupt request, the following actions are performed to enter an Interrupt Service Routine (ISR):

- The upper context of the current task is saved, and A[11] (Return Address) is updated with the current PC.
- If the processor was not previously using the interrupt stack ( $PSW.IS = 0$ ), then the A[10] Stack Pointer is set to the interrupt stack pointer (ISP). The stack pointer bit is then set for using the interrupt stack:  $PSW.IS = 1$ .
- The I/O mode is set to Supervisor mode, which means all permissions are enabled:  $PSW.IO = 10_B$ .
- Memory protection using the interrupt memory protection map is enabled:  $PSW.PRS = 00_B$ .
- The Call Depth Counter ( $PSW.CDC$ ) is cleared, and the call depth limit selector is set for 64:  $PSW.CDC = 0000000_B$ .
- Write permission to global registers A[0], A[1], A[8], A[9] is disabled:  $PSW.GW = 0$ .



- The interrupt system is globally disabled: ICR.IE = 0. The old ICR.IE is saved into PCXI.PIE.
- The Current CPU Priority Number (ICR.CCPN) is saved into the Previous CPU Priority Number (PCXI.PCPN) field.
- The Pending Interrupt Priority Number (ICR.PIPN) is saved into the Current CPU Priority Number (ICR.CCPN) field.
- The interrupt vector table is accessed to fetch the first instruction of the ISR. The effective address is the contents of the BIV register, ORd with the PIPN number left-shifted by 5.

*Note: Global register write permission is disabled (PSW.GW == 0) whenever an Interrupt Service Routine or trap handler is entered. This ensures that all traps and interrupts must assume they do not have write access to the registers controlled by PSW.GW by default.*

An Interrupt Service Routine is entered with the interrupt system globally disabled and the current CPU priority (CCPN) set to the priority (PIPN) of the interrupt being serviced. It is up to the user to enable the interrupt system again and optionally modify the priority number CCPN to implement interrupt priority levels or handle special cases. See [Using the TriCore Interrupt System, page 5-12](#).

The interrupt system can be enabled with the ENABLE instruction. ENABLE sets ICR.IE = 1 (interrupt system enabled). The BISR (Begin Interrupt Service Routine) instruction also enables the interrupt system, sets the ICR.CCPN to a new value, and saves the lower context of the interrupted task. The interrupt enable bit (ICR.IE) and current CPU priority number (ICR.CCPN) can also be modified with the MTCR (Move To Core Register) instruction.

The ENABLE, BISR, and DISABLE (disable interrupts) instructions are all executed such that the CPU is blocked from taking interrupt requests until the instruction is completely finished. This avoids pipeline side effects and eliminates the need for an ISYNC (synchronize instruction stream) following these instructions. MTCR is an exception and must be followed by an ISYNC instruction.

## **5.4 Exiting an Interrupt Service Routine (ISR)**

When an ISR exits with an RFE (Return From Exception) instruction, the hardware automatically restores the upper context. The upper context includes the PCXI register which holds the Previous CPU Priority Number (PCPN) and the Previous Global Interrupt Enable Bit (PIE). The values in these respective bits are used as follows:

- PCXI.PCPN is written to ICR.CCPN to set the CPU priority number to the value before interruption.
- PCXI.PIE is written to ICR.IE to restore the state of this bit.

The interrupted routine then continues.


## 5.5 Interrupt Vector Table

Interrupt Service Routines are associated with interrupts at a particular priority by way of the Interrupt Vector Table. The Interrupt Vector Table is an array of Interrupt Service Routine (ISR) entry points. The Interrupt Vector Table is stored in code memory.

When the CPU takes an interrupt, it calculates an address in the Interrupt Vector Table that corresponds with the priority of the interrupt (the ICR.PIPN bit field). This address is loaded in the program counter. The CPU begins executing instructions at this address in the Interrupt Vector Table. The code at this address is the start of the selected Interrupt Service Routine (ISR). Depending on the code size of the ISR, the Interrupt Vector Table may only store the initial portion of the ISR, such as a jump instruction that vectors the CPU to the rest of the ISR elsewhere in memory.


The Base of Interrupt Vector Table register (BIV) stores the base address of the Interrupt Vector Table. Interrupt vectors are ordered in the table by increasing priority. The BIV register can be modified using the MTCR instruction during the initialization phase of the system (the BIV is ENDINIT protected), before interrupts are enabled. With this arrangement, it is possible to have multiple Interrupt Vector Tables and switch between them by changing the contents of the BIV register.

When interrupted, the CPU calculates the entry point of the appropriate Interrupt Service Routine from the PIPN and the contents of the BIV register. The PIPN is left-shifted by five bits and ORed with the address in the BIV register to generate a pointer into the Interrupt Vector Table. Execution of the ISR begins at this address. Due to this operation, it is recommended that bits [12:5] of register BIV are set to 0. Note that bit 0 of the BIV register is always 0 and cannot be written to (instructions have to be aligned on even byte boundaries).



**Figure 25 Interrupt Vector Table Entry Address Calculation**

Left-shifting the PIPN by 5 bits creates entries in the vector table which are evenly spaced by 8 words. If an interrupt handler is very short it may fit entirely within the 8 words available in the vector code segment. Otherwise the code stored at the entry location can either span several vector entries, or should contain some initial instructions followed by a jump to the rest of the handler. See [Spanning Interrupt Service Routines across Vector Entries, page 5-12](#).



**Figure 26 Interrupt Vector Table**

The BIV register allows the interrupt vector table to be located anywhere in the available code memory. The default on power-up is fixed to 0000 0000<sub>H</sub>, however the BIV register can be written to using the MTCR instruction during the initialization phase of the system, before interrupts are enabled. It is also possible to have multiple interrupt vector tables and switch between them simply by modifying the contents of the BIV register.

## **5.6 Using the TriCore Interrupt System**

The following sections contain examples showing how the TriCore architectures flexible interrupt system can be used to solve both typical and special application requirements.

### **5.6.1 Spanning Interrupt Service Routines across Vector Entries**

Because vector entries are not tied to the interrupt source, it is easy to span Interrupt Service Routines (ISRs) across vector entry locations, as shown previously in [Figure 26, page 5-11](#). Spanning eliminates the need of a jump to the rest of the interrupt handler if it would not fit into the available eight words between entry locations.

Note that priority numbers relating to entries occupied by a spanned service routine must not be used for any of the active Service Request Nodes (SRNs) which request service from the same service provider.

In [Figure 26, page 5-11](#), vector locations three and four are covered through the service routine for entry two. Therefore these numbers must not be assigned to SRNs requesting CPU service, although they can be used to request another service provider. The next available vector entry is now entry five.

Use of this technique increases the range of priority numbers required in a given system, but the size of the vector table must be adjusted accordingly.

### **5.6.2 Interrupt Priority Groups**

Interrupt priority groups describe a set of interrupts which cannot interrupt each others service routine. These groups are easily created with the TriCore interrupt system architecture.

When the CPU starts the service of an interrupt, the interrupt system is globally disabled and the CPU priority CCPN is set to the priority of the interrupt being serviced. This blocks all further interrupts from being serviced until the interrupt system is either enabled again through software, or the service routine is terminated with the RFE (Return From Exception) instruction.

*Note: The RFE instruction automatically re-installs the previous state of the ICR.IE bit. This will be one (ICE.IE = 1), otherwise that interrupt would not have been serviced.*


When Interrupt Service Routine (ISR) software enables the interrupt system again by setting ICR.IE without changing the CCPN, the effect is that all interrupt requests with the same or lower priority than the CCPN are still blocked from being serviced. This includes a re-occurrence of the current interrupt; i.e. it can not interrupt this service.

However this ISR will be interrupted by each request which has a higher priority number than the CCPN. A potential problem (that is easily overcome in the TriCore architecture) is that application requirements often require interrupt requests of similar significance to

## Interrupt System

be grouped together in such a way that no request in that group can interrupt the ISR of another member of the same group.

Creating these Interrupt Priority Groups is easily accomplished in the interrupt system. For a defined group of interrupt requests, the software of their respective service routines sets the CCPN to the number of the highest SRPN used in that group, before enabling the interrupt system again. **Figure 27** shows an example.



**Figure 27 Interrupt Priority Groups**

The interrupt requests with the priority numbers 11 and 12 form one group while the requests with priority numbers 14 to 17 inclusive form another group. Every time one of the interrupts from group one is serviced, the service routine sets the CCPN to 12, the highest number in that group, before re-enabling the interrupt system.

Every time one of the interrupts from group two is serviced, the service routine sets the CCPN to 17 before re-enabling the interrupt system. If interrupt 14 is serviced for example, it can only be interrupted by requests with a priority number higher than 17, but not through a request from its own priority group or requests with lower priority.

One can see the flexibility of this system and its superiority over systems with fixed priority levels. In the example above, the interrupt request with priority number 13 forms its own single member 'group'. Setting the CCPN to the maximum number 255 in each service routine has the same effect as not enabling the interrupt system again; i.e. all interrupt requests can be considered to be in one group.

The flexibility for interrupt priority levels ranges from all interrupts being in one group, to each interrupt request building its own group, and all possible combinations in between.

### **5.6.3 Dividing ISRs into Different Priorities**

Interrupt Service Routines can be easily divided into parts with different priorities. For example, an interrupt is placed on a very high priority because response time and reaction to an event is critical, but further operations in that service routine can run on a lower priority. In this instance the service routine would be divided into two parts, one containing the critical actions, the other part the less critical ones.

The priority of the interrupt node is first set to the high priority, so that when the interrupt occurs the necessary actions are carried out immediately. The priority level of this interrupt is then lowered and the interrupt request bit is set again via software (indicating a pending interrupt) while still in the service routine. Returning to the interrupted program terminates the high priority service routine. The pending interrupt is serviced when the CPU priority is lower than its own priority. After entering the service routine, which is now at a different address in the program memory, the outstanding but low-priority actions of the interrupt are performed.

In other instances the priority of a service request might be low because the response time to an event is not critical, but once it has been granted service it should not be interrupted. To prevent any interruption the TriCore architecture allows the priority level of the service request to be raised within the ISR, and also allows interrupts to be completely disabled.

### **5.6.4 Using Different Priorities for the Same Interrupt Source**

For some applications the priority of an interrupt request in relation to other requests is not fixed, but depends on the current situation in the system. This can be achieved simply by assigning different Service Request Priority Numbers (SRPNs) at different times to an interrupt source depending on the application needs. Usually the ISR for that interrupt executes different code depending on its priority.

In traditional interrupt systems, the ISR would have to check the current priority of that interrupt request and perform a branch to the appropriate code section, causing a delay in the response to the request. In the TriCore system however, the interrupt will automatically have different vector entries for the different priorities. An extra check and branch in the ISR is not necessary, therefore the interrupt latency is reduced.

In case the ISR is independent of the interrupt's priority, branches need to be placed to the common ISR code on each of the vector entries for that interrupt.

*Note: The use of different priority numbers for one interrupt has to be taken into consideration when creating the vector table.*

### **5.6.5 Software-Posted Interrupts**

A software-posted interrupt is a true hardware interrupt, carrying an interrupt priority that is processed through the regular interrupt subsystem when the interrupt is taken. The only difference is that the interrupt request is generated by explicitly setting the service request bit in a Service Request Node (SRN), through a software update of the node's control register.

Once the interrupt request bit in a service request control register is set, there is no way to distinguish between a software-posted interrupt request and a hardware interrupt request. For that reason it is generally advisable to use Service Request Nodes and interrupt priority numbers for software-posted interrupts that are not used for hardware interrupts, such as interrupts which are triggered by a peripheral module. However the number of hardware SRNs available in a given system for such purposes depends on the application requirements. An RTOS can not therefore rely on a certain number of 'free' SRNs for software-posting of interrupts.

To support the use of software-posted interrupts, principally for RTOS code, the architecture provides a number of Service Request Nodes which are intended solely for the purpose of software-posting. They are not connected to any peripheral or any other module on the chip, and the service request flag can only be set by software. This guarantees that there are SRNs available for the RTOS and user code which are not used by hardware modules.

*Note: Current implementations contain four CPU Service Request Nodes.*

### **5.6.6 Interrupt Priority Level One**

Interrupt one is the first and lowest-priority entry in the interrupt vector and is best used for ISRs performing task management.

ISRs whose actions affect the launching of software-managed tasks post a software interrupt request at priority level one to signal the change. This posting is normally from RTOS code in a service function called directly from the ISR. The ISR can then execute a normal return from interrupt, rather than jumping to an ISR exit function in the kernel. There is no need for an exit function to check whether the ISR is returning to the background task level or to a lower priority ISR that it interrupted, in order to determine when to invoke the task dispatch function.

When there is a pending interrupt at a priority higher than the return context for the current interrupt, this interrupt will then be serviced. When a return to the background task level is performed the software-posted interrupt at priority level one will automatically be recognized and serviced.





## 6 Trap System

A trap occurs as a result of an event such as a Non-Maskable Interrupt (NMI), an instruction exception, memory-management exception or an illegal access. Traps are always active; they cannot be disabled by software action. This chapter describes the different traps that can occur and the TriCore® architecture's trap handling mechanism.

### 6.1 Trap Types

The TriCore architecture specifies eight general classes for traps. Each class has its own trap handler, accessed through a trap vector of 32 bytes per entry, indexed by the hardware-defined trap class number. Within each class, specific traps are distinguished by a Trap Identification Number (TIN) that is loaded by hardware into register D[15] before the first instruction of the trap handler is executed. The trap handler must test and branch on the value in D[15] to reach the subhandler for a specific TIN.

Traps can be further classified as synchronous or asynchronous, and as hardware or software generated. These are explained after the following table which lists the trap classes, summarising and classifying the pre-defined set of specific traps within each class.

In the following table: TIN = Trap Identification Number / Synch. = Synchronous / Asynch. = Asynchronous / HW = Hardware / SW = Software.

**Table 7 Supported Traps**

TIN	Name	Synch. / Asynch.	HW / SW	Definition	Page
<b>Class 0 — MMU</b>					
0	VAF	Synch.	HW	Virtual Address Fill.	<a href="#">page 6-7</a>
1	VAP	Synch.	HW	Virtual Address Protection.	<a href="#">page 6-7</a>
<i>Note: For VAF and VAP, see also <a href="#">MMU Traps, page 10-5</a>.</i>					<a href="#">page 10-5</a>
<b>Class 1 — Internal Protection Traps</b>					
1	PRIV	Synch.	HW	Privileged Instruction.	<a href="#">page 6-7</a>
2	MPR	Synch.	HW	Memory Protection Read.	<a href="#">page 6-7</a>
3	MPW	Synch.	HW	Memory Protection Write.	<a href="#">page 6-8</a>
4	MPX	Synch.	HW	Memory Protection Execution.	<a href="#">page 6-8</a>
5	MPP	Synch.	HW	Memory Protection Peripheral Access.	<a href="#">page 6-8</a>
6	MPN	Synch.	HW	Memory Protection Null Address.	<a href="#">page 6-8</a>
7	GRWP	Synch.	HW	Global Register Write Protection.	<a href="#">page 6-8</a>

**Table 7 Supported Traps (Continued)**

TIN	Name	Synch. / Asynch.	HW / SW	Definition	Page
<b>Class 2 — Instruction Errors</b>					
1	IOPC	Synch.	HW	Illegal Opcode.	<a href="#">page 6-8</a>
2	UOPC	Synch.	HW	Unimplemented Opcode.	<a href="#">page 6-8</a>
3	OPD	Synch.	HW	Invalid Operand specification.	<a href="#">page 6-9</a>
4	ALN	Synch.	HW	Data Address Alignment.	<a href="#">page 6-9</a>
5	MEM	Synch.	HW	Invalid Local Memory Address.	<a href="#">page 6-9</a>
<b>Class 3 — Context Management</b>					
1	FCD	Synch.	HW	Free Context List Depletion (FCX = LCX).	<a href="#">page 6-10</a>
2	CDO	Synch.	HW	Call Depth Overflow.	<a href="#">page 6-11</a>
3	CDU	Synch.	HW	Call Depth Underflow.	<a href="#">page 6-11</a>
4	FCU	Synch.	HW	Free Context List Underflow (FCX = 0).	<a href="#">page 6-11</a>
5	CSU	Synch.	HW	Call Stack Underflow (PCX = 0).	<a href="#">page 6-11</a>
6	CTYP	Synch.	HW	Context Type (PCXI.UL wrong).	<a href="#">page 6-11</a>
7	NEST	Synch.	HW	Nesting Error: RFE with non-zero call depth.	<a href="#">page 6-12</a>
<b>Class 4 — System Bus and Peripheral Errors</b>					
1	PSE	Synch.	HW	Program Fetch Synchronous Error.	<a href="#">page 6-12</a>
2	DSE	Synch.	HW	Data Access Synchronous Error.	<a href="#">page 6-12</a>
3	DAE	Asynch.	HW	Data Access Asynchronous Error.	<a href="#">page 6-12</a>
4	CAE	Asynch.	HW	Coprocessor Trap Asynchronous Error. (TriCore 1.3.1)	<a href="#">page 6-13</a>
5	PIE	Synch.	HW	Program Memory Integrity Error. (TriCore 1.3.1)	<a href="#">page 6-13</a>
6	DIE	Asynch/ Synch.	HW	Data Memory Integrity Error. (TriCore 1.3.1)	<a href="#">page 6-13</a>
<b>Class 5— Assertion Traps</b>					
1	OVF	Synch.	SW	Arithmetic Overflow.	<a href="#">page 6-14</a>
2	SOVF	Synch.	SW	Sticky Arithmetic Overflow.	<a href="#">page 6-14</a>
<b>Class 6 — System Call<sup>1)</sup></b>					
	SYS	Synch.	SW	System Call.	<a href="#">page 6-14</a>

**Table 7 Supported Traps (Continued)**

TIN	Name	Synch. / Asynch.	HW / SW	Definition	Page
<b>Class 7 — Non-Maskable Interrupt</b>					
0	NMI	Asynch.	HW	Non-Maskable Interrupt.	<a href="#">page 6-14</a>

<sup>1)</sup> For the system call trap, the TIN is taken from the immediate constant specified in the SYSCALL instruction. The range of values that can be specified is 0 to 255, inclusive.

### 6.1.1 Synchronous Traps

Synchronous traps are associated with the execution or attempted execution of specific instructions, or with an attempt to access a virtual address that requires the intervention of the memory-management system. The instruction causing the trap is known precisely. The trap is taken immediately and serviced before execution can proceed beyond that instruction.

### 6.1.2 Asynchronous Traps

Asynchronous traps are similar to interrupts, in that they are associated with hardware conditions detected externally and signaled back to the core. Some result indirectly from instructions that have been previously executed, but the direct association with those instructions has been lost. Others, such as the Non-Maskable Interrupt (NMI), are external events. The difference between an asynchronous trap and an interrupt is that asynchronous traps are routed via the trap vector instead of the interrupt vector. They can not be masked and they do not change the current CPU interrupt priority number.

### 6.1.3 Hardware Traps

Hardware traps are generated in response to exception conditions detected by the hardware. In most, but not all cases, the exception conditions are associated with the attempted execution of a particular instruction. Examples are the illegal instruction trap, memory protection traps and data memory misalignment traps. In the case of the MMU traps (trap class 0), the exception condition is either the failure to find a TLB (Translation Lookaside Buffer) entry for the virtual page referenced by an instruction (VAF trap), or an access violation for that page (VAP trap). See [MMU Traps, page 10-5](#) for more information.

### 6.1.4 Software Traps

Software traps are generated as an intentional result of executing a system call or an assertion instruction. The supported assertion instructions are TRAPV (Trap on overflow) and TRAPSV (Trap on sticky overflow). System calls are generated by the

SYSCALL instruction. System call traps are described further in [System Call \(Trap Class 6\), page 6-14](#).

### **6.1.5 Unrecoverable Traps**

An unrecoverable trap is one from which software can not recover; i.e. the task that raised the trap can not be simply restarted.

In the TriCore architecture, FCU (a fatal context trap) is an unrecoverable error. See [FCU - Free Context List Underflow \(TIN 4\), page 6-11](#) for more information.

## **6.2 Trap Handling**

The actions taken on traps by the trap handling mechanisms are slightly different from those taken on external or software interrupts. A trap does not change the CPU interrupt priority, so the ICR.CCPN field is not updated. See [Exception Priorities, page 6-15](#).

### **6.2.1 Trap Vector Format**

The trap handler vectors are stored in code memory in the trap vector table. The BTV register specifies the Base address of the Trap Vector table. The vectors are made up of a number of short code segments, evenly spaced by eight words.

If a trap handler is very short it may fit entirely within the eight words available in the vector code segment. If it does not fit the vector code segment then it should contain some initial instructions, followed by a jump to the rest of the handler.

### **6.2.2 Accessing the Trap Vector Table**

When a trap occurs, a trap identifier is generated by hardware. The trap identifier has two components:

- The Trap Class Number (TCN) used to index into the trap vector table.
- The Trap Identification Number (TIN) which is loaded into the data register D[15].

The Trap Class Number is left shifted by five bits and OR'd with the address in the BTV register to generate the entry address of the trap handler.

### **6.2.3 Return Address (RA)**

The return address is saved in the return address register A[11].

For a synchronous trap, the return address is the PC of the instruction that caused the trap. Only the SYS trap and FCD trap are different. On a SYS trap, triggered by the SYSCALL instruction, the return address points to the instruction immediately following SYSCALL. The behaviour for the FCD trap is described in [FCD - Free Context list Depletion \(TIN 1\), page 6-10](#).

For an asynchronous trap, the return address is that of the instruction that would have been executed next, if the asynchronous trap had not been taken. The return address for an interrupt follows the same rule.

### 6.2.4 Trap Vector Table


The entry-points of all Trap Service Routines are stored in memory in the Trap Vector Table. The BTV register specifies the base address of the Trap Vector Table in memory. It can be assigned to any available code memory. The BTV register can be modified using the MTCR instruction during the initialization phase of the system, (the BTV register is ENDINIT protected). This arrangement makes it possible to have multiple Trap Vector Tables and switch between them by changing the contents of the BTV register.

When a trap event occurs, a trap identifier is generated by the hardware detecting the event. The trap identifier is made up of a Trap Class Number (TCN) and a Trap Identification Number (TIN).

The TCN is left-shifted by five bits and OR'd with the address in the BTV register to form the entry address of the TSR. Because of this operation, it is recommended that bits [7:5] of register BTV are set to 0 (see [Figure 28](#)). Note that bit 0 of the BTV register is always 0 and can not be written to (instructions have to be aligned on even byte boundaries).

Left-shifting the TCN by 5 bits creates entries into the Trap Vector Table which are evenly spaced 8 words apart. If a trap handler (TSR) is very short, it may fit entirely within the eight words available in the Trap Vector Table entry. Otherwise, the code at the entry point must ultimately cause a jump to the rest of the TSR residing elsewhere in memory.

Unlike the Interrupt Vector Table, entries in the Trap Vector Table cannot be spanned.



**Figure 28 Trap Vector Table Entry Address Calculation**

### **6.2.5 Initial State upon a Trap**

The initial state when a trap occurs is defined as follows:

- The upper context is saved.
- The return address in A[11] is updated.
- The TIN is loaded into D[15].
- The stack pointer in A[10] is set to the Interrupt Stack Pointer (ISP) when the processor was not previously using the interrupt stack (in case of PSW.IS = 0). The stack pointer bit is set for using the interrupt stack: PSW.IS = 1.
- The I/O mode is set to Supervisor mode, which means all permissions are enabled: PSW.IO = 10<sub>B</sub>.
- The current Protection Register Set is set to 0: PSW.PRS = 00<sub>B</sub>.
- The Call Depth Counter (CDC) is cleared, and the call depth limit is set for 64: PSW.CDC = 0000000<sub>B</sub>.
- Call Depth Counter is enabled, PSW.CDE = 1.
- Write permission to global registers A[0], A[1], A[8], A[9] is disabled: PSW.GW = 0.
- The interrupt system is globally disabled: ICR.IE = 0. The 'old' ICR.IE and ICR.CCPN are saved into PCXI.PIE and PCXI.PCPN respectively. ICR.CCPN remains unchanged.
- The trap vector table is accessed to fetch the first instruction of the trap handler.

Although traps leave the ICR.CCPN unchanged, their handlers still begin execution with interrupts disabled. They can therefore perform critical initial operations without interruptions, until they specifically re-enable interrupts.

For the non-recoverable FCU trap, the initial state is different. The upper context cannot be saved. Only the following states are guaranteed:

- The TIN is loaded into D[15].
- The stack pointer in A[10] is set to the Interrupt Stack Pointer (ISP) when the processor was not previously using the interrupt stack (in case of PSW.IS == 0).
- The I/O mode is set to Supervisor mode (all permissions are enabled: PSW.IO = 10<sub>B</sub>).
- The current Protection Register Set is set to 0: PSW.PRS = 00<sub>B</sub>.
- The interrupt system is globally disabled: ICR.IE = 0. ICR.CCPN remains unchanged.
- The trap vector table is accessed to fetch the first instruction of the FCU trap handler.

## **6.3 Trap Descriptions**

The following sub-sections describe the trap classes and specific traps listed in [Table 7 Supported Traps, page 6-1](#).

### **6.3.1 MMU Traps (Trap Class 0)**

For those implementations that include a Memory Management Unit (MMU), Trap class 0 is reserved for MMU traps. There are two traps within this class, VAF and VAP.

#### **VAF - Virtual Address Fill (TIN 0)**

The VAF trap is generated when the MMU is enabled and the virtual address referenced by an instruction does not have a page entry in the MMU Translation Lookaside Buffer (TLB).

#### **VAP - Virtual Address Protection (TIN 1)**

The VAP trap is generated (when the MMU is enabled) by a memory access undergoing PTE translation that is not permitted by the PTE protection settings, or by a User-0 mode access to an upper segment that does not have the privileged peripheral property.

### **6.3.2 Internal Protection Traps (Trap Class 1)**

Trap class 1 is for traps related to the internal protection system. The memory protection traps in this class, MPR, MPW, and MPX, are for the range-based protection system and are independent of the page-based VAP protection trap of trap class 0. See [Memory Protection Register Sets, page 9-2](#) for more details.

All memory protection traps (MPR, MPW, MPX, MPP, and MPN), are based on the virtual addresses that undergo direct translation.

The following internal Protection Traps are defined:

#### **PRIV - Privilege Violation (TIN 1)**

A program executing in one of the User modes (User-0 or User-1 mode) attempted to execute an instruction not allowed by that mode.

A table of instructions which are restricted to Supervisor mode or User-1 mode, is supplied in the Instruction Set chapter of Volume 2 of this manual.

#### **MPR - Memory Protection Read (TIN 2)**

The MPR trap is generated when the memory protection system is enabled and the effective address of a load, LDMST or SWAP instruction does not lie within any range with read permissions enabled. This trap is not generated when an access violation occurs during a context save/restore operation.

**MPW - Memory Protection Write (TIN 3)**

The MPW trap is generated when the memory protection system is enabled and the effective address of a store, LDMST or SWAP instruction does not lie within any range with write permissions enabled.

This trap is not generated when an access violation occurs during a context save/restore operation.

**MPX - Memory Protection Execute (TIN 4)**

The MPX trap is generated when the memory protection system is enabled and the PC does not lie within any range with execute permissions enabled.

**MPP - Memory Protection Peripheral Access (TIN 5)**

A program executing in User-0 mode attempted a load or store access to a segment that has the privileged peripheral property. See [Physical Memory Attributes \(PMA\)](#), page 8-3.

**MPN - Memory Protection Null address (TIN 6)**

The MPN trap is generated whenever any program attempts a load / store operation to effective address zero.

**GRWP - Global Register Write Protection (TIN 7)**

A program attempted to modify one of the global address registers (A[0], A[1], A[8] or A[9]) when it did not have permission to do so.

**6.3.3 Instruction Errors (Trap Class 2)**

Trap class 2 is for signalling various types of instruction errors. Instruction errors include errors in the instruction opcode, in the instruction operand encodings, or for memory accesses, in the operand address.

**IOPC - Illegal Opcode (TIN 1)**

An invalid instruction opcode was encountered. An invalid opcode is one that does not correspond to any instruction known to the implementation.

**UOPC - Unimplemented Opcode (TIN 2)**

An unimplemented opcode was encountered. An unimplemented opcode corresponds to a known instruction that is not implemented in a given hardware implementation. The instruction may be implemented via software emulation in the trap handler.



Example UOPC conditions are:

- A MMU instruction if the MMU is not present.
- A FPU instruction if the FPU is not present.
- An external coprocessor instruction if the external coprocessor is not present.

### **OPD - Invalid Operand (TIN 3)**

The OPD trap may be raised for instructions that take an even-odd register pair as an operand, if the operand specifier is odd. The OPD trap may also be raised for other cases where operands are invalid.

Implementations are not architecturally required to raise this trap, and may treat invalid operands in an implementation defined manner.

### **ALN - Data Address Alignment (TIN 4)**

An ALN trap is raised when the address for a data memory operation does not conform to the required alignment rules. See [Alignment Requirements, page 2-4](#), for more information on these rules. An ALN trap is also raised when the size, length or index of a circular buffer is incorrect. See [Circular Addressing, page 2-9](#) for more details.

### **MEM - Invalid Memory Address (TIN 5)**

The MEM trap is raised when the address of an access can be determined to either violate an architectural constraint or an implementation constraint.

Defined MEM trap subclasses are *different segment*, *segment crossing*, *CSFR access*, *CSA restriction* and *scratch range*.

An implementation must define which implementation constraint MEM traps it will raise, or the alternative behaviour if the MEM trap is not raised. It must also document any other implementation specific MEM traps it will raise.

Architectural constraints which will raise the MEM trap are:

- An addressing mode that adds an offset to a base address results in an effective address that is in a different segment to the base address (*different segment*).
- A data element is accessed with an address, such that the data object spans the end of one segment and the beginning of another segment (*segment crossing*)

Implementation constraints which can raise the MEM trap are

- A memory address is used to access a Core SFR (CSFR) rather than using a MTCR/MFCR instruction (*CSFR access*)
- A memory address is used for a CSA access and it is not valid for the implementation to place CSA there (*CSA restriction*)
- An access to Scratch memory is attempted using a memory address which lies outside the implemented region of memory (*scratch range error*).

### **6.3.4 Context Management (Trap Class 3)**

Trap class 3 is for exception conditions detected by the context management subsystem, in the course of performing (or attempting to perform) context save and restore operations connected to function calls, interrupts, traps, and returns.

#### **FCD - Free Context list Depletion (TIN 1)**

The FCD trap is generated after a context save operation, when the operation causes the free context list to become 'almost empty'. The 'almost empty' condition is signaled when the CSA used for the save operation is the one pointed to by the context list limit register LCX. The operation responsible for the context save completes normally and then the FCD trap is taken.

If the operation responsible for the context save was the hardware interrupt or trap entry sequence, then the FCD trap handler will be entered before the first instruction of the original interrupt or trap handler is executed. The return address for the FCD trap will point to the first instruction of the interrupt or trap handler.

The FCD trap handler is normally expected to take some form of action to rectify the context list depletion. The nature of that action is OS dependent, but the general choices are to allocate additional memory for CSA storage, or to terminate one or more tasks, and return the CSAs on their call chains to the free list. A third possibility is not to terminate any tasks outright, but to copy the call chains for one or more inactive tasks to uncached external or secondary memory that would not be directly usable for CSA storage, and release the copied CSAs to the free list. In that instance the OS task scheduler would need to recognize that the inactive task's call chain was not resident in CSA storage, and restore it before dispatching the task.

The FCD trap itself uses one additional CSA beyond the one designated by the LCX register, so LCX must not point to the actual last entry on the free context list. In addition, it is possible that an asynchronous trap condition, such as an external bus error, will be reported after the FCD trap has been taken, interrupting the FCD trap handler and using one more CSA. Therefore, to avoid the possibility of a context list underflow, the free context list must include a minimum of two CSAs beyond the one pointed to by the LCX register. If the FCD trap handler makes any calls, then additional CSA reserves are needed.

In order to allow the trap handlers for asynchronous traps to recognize when they have interrupted the FCD trap handler, the FCDSF flag in the SYSCON (system configuration) register is set whenever an FCD trap is generated. The FCDSF bit should be tested by the handler for any asynchronous trap that could be taken while an FCD trap is being handled. If the bit is found to be set, the asynchronous trap handler must avoid making any calls, but should queue itself in some manner that allows the OS to recognize that the trap occurred. It should then carry out an immediate return, back to the interrupted FCD trap handler. See [System Control Register \(SYSCON\), page 3-16](#).

**CDO - Call Depth Overflow (TIN 2)**

A program attempted to execute a CALL instruction with the Call Depth counter enabled and the call depth count value (PSW.CDC.COUNT) at its maximum value. Call Depth Counting guards against context list depletion, by enabling the OS to detect 'runaway recursion' in executing tasks. See [Program Status Word Register \(PSW\), page 3-6](#).

**CDU - Call Depth Underflow (TIN 3)**

A program attempted to execute a RET (return) instruction with the Call Depth counter enabled and the call depth count value (PSW.CDC.COUNT) at zero. A call depth underflow does not necessarily reflect a software error in the currently executing task. An OS can achieve finer granularity in call depth counting by using a deliberately narrow Call Depth Counter, and incrementing or decrementing a separate software counter for the current task on each call depth overflow or underflow trap. A program error would be indicated only if the software counter were already zero when the CDU trap occurred.

**FCU - Free Context List Underflow (TIN 4)**

The FCU trap is taken when a context save operation is attempted but the free context list is found to be empty (i.e. the FCX register contents are null). The FCU trap is also taken if any error is encountered during a context save or restore operation. The context operation cannot be completed. Instead a forced jump is made to the FCU trap handler and D15 updated with the FCU TIN value.

In failing to complete the context save or restore, architectural state is lost, so the occurrence of an FCU trap is a non-recoverable system error. The FCU trap handler should ultimately initiate a system reset.

**CSU - Call Stack Underflow (TIN 5)**

Raised when a context restore operation is attempted and when the contents of the PCX register were null or otherwise invalid. This trap indicates a system software error (kernel or OS) in task setup or context switching among software managed tasks (SMTs). No software error or combination of errors in a user task can generate this condition, unless the task has been allowed write permission to the context save areas which, in itself, can be regarded as a system software error.

**CTYP - Context Type (TIN 6)**

Raised when a context restore operation is attempted but the context type, as indicated by the PCXI.UL bit, is incorrect for the type of restore attempted; i.e. a restore lower context is attempted when PCXI.UL == 1, or a restore upper context is attempted when PCXI.UL == 0. As with the CSU trap, this indicates a system software error in context list management.

### **NEST - Nesting Error (TIN 7)**

A program attempted to execute an RFE (return from exception) instruction with the Call Depth counter enabled and the call depth count value (PSW.CDC.COUNT) non-zero. The return from an interrupt or trap handler should normally occur within the body of the interrupt or trap handler itself, or in code to which the handler has branched, rather than code called from the handler. If this is not the case there will be one or more saved contexts on the residual call chain that must be popped and returned to the free list, before the RFE can be legitimately issued.

### **6.3.5 System Bus and Peripheral Errors (Trap Class 4)**

#### **PSE - Program Fetch Synchronous Error (TIN 1)**

The PSE trap is raised when:

- A bus error<sup>1)</sup> occurred because of an instruction fetch.
- An instruction fetch targets a segment that does not have the code fetch property. See [Physical Memory Attributes \(PMA\), page 8-3](#).
- A Code Fetch operation from Program scratchpad RAM<sup>2)</sup> (PSPR) (See [Scratchpad RAM, page 8-4](#)) where the access is beyond the end of the memory range.

#### **DSE - Data Access Synchronous Error (TIN 2)**

The DSE trap is raised when:

- A data access is attempted to a segment that does not have the data access property. (See [Physical Memory Attributes \(PMA\), page 8-3](#)).
- Whenever a bus error occurred because of a data load operation. It is also raised in the case of a data load operation from Data scratchpad RAM<sup>2)</sup> (DSPR) ([Scratchpad RAM, page 8-4](#)) where the access is beyond the end of the memory range.

*Note: There are implementation-dependent registers for DSE which can be interrogated to determine the source of the error more precisely. Refer to the User's Manual for a specific TriCore implementation for more details.*

#### **DAE - Data Access Asynchronous Error (TIN 3)**

The DAE trap is raised when the memory system reports back an error which cannot immediately be linked to a currently executing instruction. Generally this means an error returned on the system bus from a peripheral or external memory.

<sup>1)</sup> A bus fetch error is also generated for an instruction fetch to the data scratch pad RAM region (D000 0000<sub>H</sub> to D3FF FFFF<sub>H</sub>) when the memory access is outside the range of the actual scratchpad RAMs.

<sup>2)</sup> PSPR is also known as SPRAM (Scratchpad RAM). DSPR is also known as Local Data RAM (LDRAM).

This trap is raised whenever a bus error occurred because of a data store operation, or when:

- There is a data store operation to local scratch memory but the access is beyond the end of the memory range.
- There is an error caused by a cache management instruction.

*Note: There are implementation-dependent registers for DAE which can be interrogated to determine the source of the error more precisely. Refer to the User's Manual for a specific TriCore implementation for more details.*

### **CAE - Coprocessor Trap Asynchronous Error (TIN 4) (TriCore 1.3.1)**

This CAE asynchronous trap is generated by a coprocessor to report an error.

Examples of typical errors that can cause a CAE trap are unimplemented coprocessor instructions and arithmetic errors (as found in the Floating Point Unit for example).

CAE is shared amongst all coprocessors in a given system. A trap handler must therefore inspect all coprocessors to determine the cause of a trap.

### **PIE - Program Memory Integrity Error (TIN 5) (TriCore 1.3.1)**

The PIE trap is raised whenever an uncorrectable memory integrity error is detected in an instruction fetch. The trap is synchronous to the erroneous instruction. A PIE trap is raised if any element within the fetch group contains an unrecoverable error. Hardware is not required to localise the error to a particular instruction.

An implementation may provide additional registers that can be interrogated to determine the source of the error more precisely. Refer to the User manual for a specific Tricore implementation for more details.

### **DIE - Data Memory Integrity Error (TIN 6) (TriCore 1.3.1)**

The DIE trap is raised whenever an uncorrectable memory integrity error is detected in a data access.

Implementations may choose to implement the DIE trap as either an asynchronous or synchronous trap.

A DIE trap is raised if any element accessed by a load or store contains an uncorrectable error. Hardware is not required to localise the error to the access width of the operation.

An implementation may provide additional registers that can be interrogated to determine the source of the error more precisely. Refer to the User manual for a specific Tricore implementation for more details.

### **6.3.6 Assertion Traps (Trap Class 5)**

#### **OVF - Arithmetic Overflow (TIN 1)**

Raised by the TRAPV instruction, if the overflow bit in the PSW is set (PSW.V == 1).

#### **SOVF - Sticky Arithmetic Overflow (TIN 2)**

Raised by the TRAPSV instruction, if the sticky overflow bit in the PSW is set (PSW.SV == 1).

### **6.3.7 System Call (Trap Class 6)**

#### **SYS - System Call (TIN = 8-bit unsigned immediate constant in SYSCALL)**

The SYS trap is raised immediately after the execution of the SYSCALL instruction, to initiate a system call. The TIN that is loaded into D[15] when the trap is taken is not fixed, but is specified as an 8-bit unsigned immediate constant in the SYSCALL instruction. The return address points to the instruction immediately following the SYSCALL.

### **6.3.8 Non-Maskable Interrupt (Trap Class 7)**

#### **NMI - Non-Maskable Interrupt (TIN 0)**

The causes for raising a Non-Maskable Interrupt are implementation dependent. Typically there is an external pin that can be used to signal the NMI, but it may also be raised in response to such things as a watchdog timer interrupt, or an impending power failure. Refer to the User's Manual for a specific TriCore implementation for more details.

### **6.3.9 Debug Traps**

#### **BBM - Break Before Make / BAM - Break After Make**

Please refer to the Core Debug Controller chapter for information on debug traps. See [Chapter 12 Core Debug Controller \(CDC\), page 12-1](#).

## 6.4 Exception Priorities

The priority order between an asynchronous trap, a synchronous trap, and an interrupt from the software architecture model, is as follows:

1. Asynchronous trap (highest priority).
2. Synchronous trap.
3. Interrupt (lowest priority).

The following trap rules must also be considered:

1. The older the instruction in the instruction sequence which caused the trap, the higher the priority of the trap. All potential traps with lower priorities are void.
2. Attempting to save a context with an empty free context list (FCX = 0) results in a FCU (Free Context List Underflow) trap. This trap takes priority over all other exceptions.
3. When the same instruction causes several synchronous traps anywhere in the pipeline, priorities follow those shown in the table below.

**Table 8 Synchronous Trap Priorities**

Priority	Type of Trap
<b>Instruction Fetch Traps</b>	
1	Breakpoint (Virtual address, BBM)
2	VA-F-P
3	VAP-P
4	MPX
5	PSE
6	PIE (TriCore 1.3.1)
<b>Instruction Format Traps</b>	
7	IOPC
8	OPD
9	UOPC
<b>Instruction Traps</b>	
10	Breakpoint trap (Instruction, BBM)
11	PRIV
12	GRWP
13	SYS
<b>Context Traps</b>	
14	FCD

**Table 8 Synchronous Trap Priorities (Continued)**

15	FCU (Synchronous)
16	CSU
17	CDO
18	CDU
19	NEST
20	CTYP

**Data Memory Access Traps**

21	MEM (Data address)
22	ALN
23	MPN
24	VAf-D
25	VAP-D
26	MPR
27	MPW
28	MPP
29	DSE
30	DIE (TriCore 1.3.1)

**General Data Traps**

31	SOVF
32	OVF
33	Breakpoint trap (BAM)

**Table 9 Asynchronous Trap Priorities**

Priority	Asynchronous Traps
1	NMI
2	DAE <sup>1)</sup>
3	DIE (TriCore 1.3.1)
4	CAE (TriCore 1.3.1)

<sup>1)</sup> DAE is used for store errors.



## 6.5 Interrupt and Trap Control Registers

Three CSFRs support interrupt and trap handling:

- ICR: Interrupt Control Register [page 6-17](#).
- BIV: Base Interrupt Vector Table Pointer [page 6-19](#).
- BTV: Base Trap Vector Table Pointer [page 6-20](#).

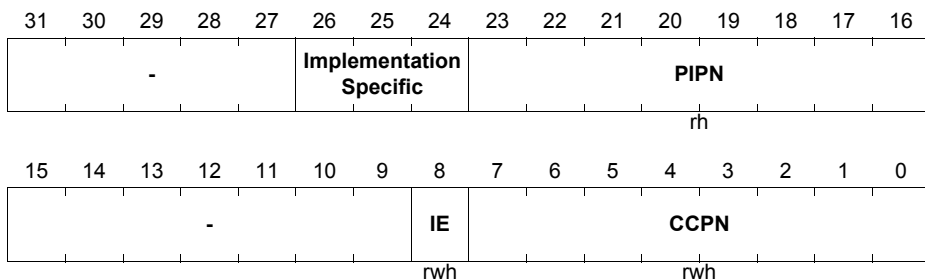
The ICR holds the Current CPU Priority Number (CCPN), the enable/disable bit for the Interrupt System (IE), the Pending Interrupt Priority Number (PIPN), and an implementation specific control for the interrupt arbitration scheme. The other two registers hold the base addresses for the interrupt (BIV) and trap vector tables (BTV). Special instructions control the enabling and disabling of the interrupt system. For more information see [Interrupt System, page 5-1](#).

### 6.5.1 ICU Interrupt Control Register (ICR)

The ICU Interrupt Control register is defined as follows:

#### ICR

**ICU Interrupt Control (FE2C<sub>H</sub>)** **Reset Value: 0000 0000<sub>H</sub>**



Field	Bits	Type	Function
-	[31:27]	-	Reserved Field
	[26:24]	-	Implementation Specific Control of the arbitration. See the relevant documentation for a specific TriCore product implementation.

Field	Bits	Type	Function
PIPN	[23:16]	rh	<b>Pending Interrupt Priority Number</b> A read-only bit field that is updated by the ICU at the end of each interrupt arbitration process. It indicates the priority number of the pending service request. ICR.PIPN is set to 0 when no request is pending, and at the beginning of each new arbitration process. 00 <sub>H</sub> : No valid pending request. 01 <sub>H</sub> : Request pending, lowest priority. ... FF <sub>H</sub> : Request pending, highest priority.
-	[15:9]	-	<b>Reserved Field</b>
IE	8	rwh	<b>Global Interrupt Enable Bit</b> The interrupt enable bit globally enables the CPU service request system. Whether a service request is delivered to the CPU depends on the individual Service Request Enable Bits (SRE) in the SRNs, and the current state of the CPU. ICR.IE is automatically updated by hardware on entry and exit of an Interrupt Service Routine (ISR). ICR.IE is cleared to 0 when an interrupt is taken, and is restored to the previous value when the ISR executes an RFE instruction to terminate itself. ICR.IE can also be updated through the execution of the ENABLE, DISABLE, MTCR, and BISR instructions. 0 : Interrupt system is globally disabled. 1 : Interrupt system is globally enabled.
CCPN	[7:0]	rwh	<b>Current CPU Priority Number</b> The Current CPU Priority Number (CCPN) bit field indicates the current priority level of the CPU. It is automatically updated by hardware on entry or exit of Interrupt Service Routines (ISRs) and through the execution of a BISR instruction. CCPN can also be updated through an MTCR instruction.


### 6.5.2 Base Interrupt Vector Table Pointer (BIV)

The BIV register contains the base address of the interrupt vector table. When an interrupt is accepted, the entry address into the interrupt vector table is generated from the priority number (taken from the PIPN) of that interrupt, left shifted by five bits, and then OR'd with the contents of the BIV register. The left-shift of the interrupt priority number results in a spacing of 8 words (32 bytes) between the individual entries in the vector table.

#### BIV

#### Base Interrupt Vector Table Pointer (FE20<sub>H</sub>)

**Reset Value: Implementation Specific**



Field	Bits	Type	Description
BIV	[31:1]	rw	<b>Base Address of Interrupt Vector Table</b> The address in the BIV register must be aligned to an even byte address (halfword address). Because of the simple ORing of the left-shifted priority number and the contents of the BIV register, the alignment of the base address of the vector table must be to a power of two boundary, dependent on the number of interrupt entries used. For the full range of 256 interrupt entries an alignment to an 8 KByte boundary is required. If fewer sources are used, the alignment requirements are correspondingly relaxed.
-	0	-	<b>Reserved Field</b>

*Note: This register is ENDINIT protected.*


### 6.5.3 Base Trap Vector Table Pointer (BTV)

The BTV contains the base address of the trap vector table. When a trap occurs, the entry address into the trap vector table is generated from the Trap Class of that trap, left-shifted by 5 bits and then OR'd with the contents of the BTV register. The left-shift of the Trap Class results in a spacing of 8 words (32 bytes) between the individual entries in the vector table.

#### **BTV**

**Base Trap Vector Table Pointer (FE24<sub>H</sub>)**

**Reset Value: Implementation Specific**



Field	Bits	Type	Description
BTV	[31:1]	rw	<b>Base Address of Trap Vector Table</b> The address in the BTV register must be aligned to an even byte address (halfword address). Also, due to the simple ORing of the left-shifted trap identification number and the contents of the BTV register, the alignment of the base address of the vector table must be to a power of two boundary. There are eight different trap classes, resulting in Trap Classes from 0 to 7. The contents of BTV should therefore be set to at least a 256 byte boundary (8 Trap Classes * 8 word spacing).
-	0	-	<b>Reserved Field</b>

*Note: This register is ENDINIT protected.*

## 7 Memory Integrity Error Mitigation (TriCore 1.3.1)

*Note: This chapter only applies to the TriCore 1.3.1 architecture.*

This chapter describes the architectural features used to support the mitigation of memory integrity errors within the local memories of TriCore® processors.

### 7.1 Memory Integrity Error Classification

Memory integrity errors are classified as being either Correctable or Uncorrectable.

#### Uncorrectable Memory Integrity Error

If on accessing a memory element containing a memory integrity error, hardware is not able to provide the expected data to the core, the memory integrity error is defined as being uncorrectable.

#### Correctable Memory Integrity Error

If on accessing a memory element containing a memory integrity error, hardware is able to provide the expected data to the core, the memory integrity error is defined as being correctable.

Correctable memory integrity errors are further categorised as either **Resolved** or **Unresolved**. Correctable memory integrity errors always provide the correct data to the core. As part of the correction process hardware may also update the erroneous source data in memory with the corrected data. Such a memory integrity error is defined as being Resolved. If the erroneous source data in memory is not updated the memory integrity error is defined as being Unresolved.

## **7.2 Memory Integrity Error Traps**

When an uncorrectable memory integrity error is encountered one of the following traps is raised.

### **7.2.1 Program Memory Integrity Error (PIE)**

The PIE trap is raised whenever an uncorrectable memory integrity error is detected in an instruction fetch from a local memory. The trap is synchronous to the erroneous instruction. The trap is of Class 4 and TIN 5.

A PIE trap is raised if any element within the fetch group contains an unrecoverable error. Hardware is not required to localise the error to a particular instruction.

*Note: There are implementation specific registers that can be interrogated to determine the source of the error more precisely. Refer to the User manual for a specific Tricore implementation for more details.*

### **7.2.2 Data Memory Integrity Error (DIE)**

The DIE trap is raised whenever an uncorrectable memory integrity error is detected in a data access to a local memory. The trap is of Class 4 and TIN 6.

A TriCore implementation may choose to implement the DIE trap as either an asynchronous or synchronous trap.

A DIE trap is raised if any element accessed by a load/store contains an uncorrectable error. Hardware is not required to localise the error to the access width of the operation.

*Note: There are implementation specific registers that can be interrogated to determine the source of the error more precisely. Refer to the User manual for a specific Tricore implementation for more details.*

## 7.3 Corrected Error Counts

Two architecturally visible registers (CCPIER, CCDIER) are used to maintain a running count of corrected memory integrity errors in the local memory systems.

Each register contains two count fields, one for resolved corrected errors and one for unresolved corrected errors.

### 7.3.1 Count of Corrected Program Memory Integrity Errors Register

- The CCPIE-R counter is incremented on each detection of a corrected-resolved memory integrity error in the local instruction memories. The counter saturates at the value  $FF_H$ .
- The CCPIE-U counter is incremented on each detection of a corrected-unresolved memory integrity error in the local instruction memories. The counter saturates at the value  $FF_H$ .


*Note: TriCore 1.3.1 Architecture Only.*

#### CCPIER

##### Count of Corrected Program Memory Integrity Errors Register

(9218<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>



Field	Bits	Type	Description
-	[31:16]	-	<b>Reserved</b>
CCPIE-R	[15:8]	rwh	<b>Count of Corrected-Resolved Program Integrity Errors.</b> In local instruction memory.
CCPIE-U	[7:0]	rwh	<b>Count of Corrected-Unresolved Program Integrity Errors.</b> In local instruction memory.

**Memory Integrity Error Mitigation (TriCore 1.3.1)**

### 7.3.2 Count of Corrected Data Integrity Errors Register


- The CCDIE-R counter is incremented on each detection of a corrected-resolved memory integrity error in the local data memories. The counter saturates at the value FF<sub>H</sub>.
- The CCDIE-U counter is incremented on each detection of a corrected-unresolved memory integrity error in the local data memories. The counter saturates at the value FF<sub>H</sub>.

*Note: TriCore 1.3.1 Architecture Only.*

#### CCDIER

##### Count of Corrected Data Integrity Errors Register (9028<sub>H</sub>)

**Reset Value: 0000 0000<sub>H</sub>**



Field	Bits	Type	Description
-	[31:16]	-	<b>Reserved</b>
CCDIE-R	[15:8]	rwh	<b>Count of Corrected-Resolved Data Integrity Errors.</b> In local data memory.
CCDIE-U	[7:0]	rwh	<b>Count of Corrected-Unresolved Data Integrity Errors.</b> In local data memory.



## 7.4 Error Information Registers

To provide information for memory integrity error handling and debug, a number of implementation specific registers are provided. The contents of these registers are implementation specific.

### 7.4.1 Program Integrity Error Trap Register (PIETR)

This register contains information allowing software to localise the source of the last detected program memory integrity error.

*Note: TriCore 1.3.1 Architecture Only.*

#### PIETR

#### Program Integrity Error Trap Register

(9214<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>



Field	Bits	Type	Description
-	[31:0]	-	Implementation Specific

## Memory Integrity Error Mitigation (TriCore 1.3.1)

### 7.4.2 Program Integrity Error Address Register (PIEAR)

This register contains the address accessed by the last operation that caused a program memory integrity error.



*Note: TriCore 1.3.1 Architecture Only.*

#### PIEAR

#### Program Integrity Error Address Register

(9210<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>



Field	Bits	Type	Description
-	[31:0]	-	Implementation Specific

## Memory Integrity Error Mitigation (TriCore 1.3.1)

### 7.4.3 Data Integrity Error Trap Register (DIETR)

This register contains information allowing software to localise the source of the last detected data memory integrity error.

*Note: TriCore 1.3.1 Architecture Only.*

#### DIETR

#### Data Integrity Error Trap Register

(9024<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>



Field	Bits	Type	Description
-	[31:0]	-	Implementation Specific

## Memory Integrity Error Mitigation (TriCore 1.3.1)

### 7.4.4 Data Integrity Error Address Register (DIEAR)

This register contains the address accessed by the last operation that caused a data memory integrity error.

*Note: TriCore 1.3.1 Architecture Only.*

#### DIEAR

##### Data Integrity Error Address Register

(9020<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>



Field	Bits	Type	Description
-	[31:0]	-	Implementation Specific

**Memory Integrity Error Mitigation (TriCore 1.3.1)**

### 7.4.5 Memory Integrity Error Control Register

The MIECON register is provided to allow software to control the memory integrity error detection and correction mechanisms. The register is architecturally defined, however the register contents are implementation specific.

*Note: TriCore 1.3.1 Architecture Only.*

#### MIECON

##### Memory Integrity Error Control Register

(9044<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>



Field	Bits	Type	Description
-	[31:0]	-	Implementation Specific

*Note: This register is ENDINIT protected.*

## **7.5 Summary**

A detected memory integrity error in local instruction memory will lead to either:

- a correctable error and an increment of one of the CCPIE counters or
- an uncorrectable error triggering a PIE trap.

A detected memory integrity error in local data memory will lead to either:

- a correctable error and an increment of one of the CCDIE counters or
- an uncorrectable error triggering a DIE trap.

The actual method used for the detection of memory integrity errors is implementation dependent.

## **8 Physical Memory Attributes (PMA)**

This chapter describes the Physical Memory Attributes (PMA) that regions of the TriCore® physical address map may or may not have. These attributes are defined by groups of physical memory properties.

### **8.1 Physical Memory Properties (PMP)**

The TriCore architecture defines properties which physical memory addresses may or may not possess. These properties are:

- Privileged Peripheral (P).
- Cacheable (C).
- Speculative (S).
- Code Fetch (F).
- Data Access (D).

Each property defines a characteristic of the accesses that are possible to a physical memory region. For example, an address that does not have the cacheable property C, would be described as Non-cacheable  $\bar{C}$ .

In the following definitions the concept of necessary and speculative accesses is introduced. Necessary accesses are those required to correctly compute the program and any implementation or simulation of the program execution must perform these accesses. Speculative accesses are those that an implementation may make in order to improve performance either in correct or incorrect anticipation of a necessary access.

#### **Privileged Peripheral (P)**

Only Supervisor and User-1 mode data accesses are possible. No User-0 mode data access is possible. User-0 mode data accesses result in an MPP (Memory Protection Peripheral access) trap. All accesses are exempt from the protection system settings. PTE translation where the physical address targets a region with this property results in undefined behaviour.

#### **Cacheable (C)**

It is possible for data and code fetch accesses to the region to be cached by the CPU if a data cache or code cache is respectively present and enabled.

#### **Speculative (S)**

It is possible to perform speculative data accesses to the memory. A speculative data access is a read access to memory addresses that are not strictly necessary for correct program execution. The processor never performs speculative write accesses which are visible in a memory region.

### Code Fetch (F)

Fetch accesses are possible to this region. The fetch property allows full speculation on all fetch accesses to the region. The cacheable property has no affect on the amount or range of speculation of code fetches. If a necessary fetch access is directed by program flow to a physical memory region that does not have the fetch property then a PSE (Program fetch Synchronous Error) trap occurs.

### Data Access (D)

Data accesses are possible to this region. If a data access is directed by necessary program flow to a physical memory region that does not have the Data Access property, then a DSE (Data access Synchronous Error) trap occurs.

For data accesses, the interpretation of the combinations of the Privileged Peripheral, Cacheable and Speculative properties for a memory region are defined in [Table 10](#). All other combinations of these three properties not present in this table, are reserved.

**Table 10 Data Access - Cacheable and Speculative Properties**

Name	Privileged Peripheral Property	Cacheable Property	Speculative Property	Behaviour of Physical Memory Region
Precise data access	$P$ or $\bar{P}$	$\bar{C}$	$\bar{S}$	The processor only performs necessary accesses, in order, to the region.
Non-Cached access	$\bar{P}$	$\bar{C}$	$S$	The processor may read an entire cache line <sup>1)</sup> containing the address of a necessary access and place it in a buffer for subsequent accesses. The order of accesses is not guaranteed <sup>2)</sup> .
Full Speculation	$\bar{P}$	$C$	$S$	The processor may perform speculative read accesses to entire cache lines in physical memory and place them in the cache. The order of accesses is not guaranteed.

<sup>1)</sup> The size of a cache line is implementation dependant. Examples of implemented cache lines are 16-bytes and 32-bytes, but may be smaller or larger.

<sup>2)</sup> The order of non-cached data accesses can be guaranteed by inserting a DSYNC instruction after each load or store instruction.



## 8.2 Physical Memory Attributes (PMA)

A physical memory attribute is a defined set of physical memory properties. The architecture defines four attributes:

Peripheral Space =  $\overline{\text{PCSFD}}$ .

Emulator Space =  $\overline{\text{PCSFD}}$ .

Cacheable Memory =  $\overline{\text{PCSFD}}$ .

Non-Cacheable Memory =  $\overline{\text{PCSFD}}$ .

All accesses to physical memory that have the Emulator Space attribute are directly translated (See [Memory Management Unit \(MMU\), page 10-1](#)) and are not subject to the protection constraints imposed by the protection system (See [Memory Protection System, page 9-1](#)); i.e. It is not possible to generate an MPX, MPY or MPW trap with a memory access to Emulator Space.

### 8.2.1 Physical Memory Attributes of the Address Map

The 4 GBytes (32-bit) of physical address space is divided into 16 equally sized segments. Each segment has its own physical memory attribute.

Segment  $F_H$  is constrained to be Peripheral Space and the lower 15 segments have defined physical memory attributes, although Segment  $D_H$  is constrained to be either Cacheable or Non-Cacheable Memory. The lower 15 segments have implementation defined physical memory attributes.

The default defined attributes are shown in the following table:

**Table 11 TriCore Default Physical Memory Attributes for all Segments**

Segment	Attributes
$F_H$ <sup>1)</sup>	Peripheral Space.
$E_H$	Peripheral Space.
$D_H$	Non-cacheable Memory.
$C_H$ <sup>2)</sup>	Cacheable Memory.
$B_H$	Non-cacheable Memory.
$A_H$	Non-cacheable Memory.
$9_H$	Cacheable Memory.
$8_H$	Cacheable Memory.
$7_H - 0_H$	Cacheable Memory.

<sup>1)</sup>  $F_H$  is constrained to be Peripheral Space.

<sup>2)</sup> See [Section 8.3.2](#) for Segment C constraints.

**Physical Memory Attributes (PMA)**

The Emulator Space attribute is assigned to an implementation defined region of memory when the Debug Mode is enabled. All physical memory accesses are subject to the constraints imposed by the PMA attributes before being permitted to execute.

## 8.3 Scratchpad RAM

### 8.3.1 Scratchpad RAM (TriCore 1.3)

Segment D contains the scratchpad RAM. There are two different scratchpad RAMs:

- DSPR - Data scratchpad RAM.
- PSPR - Program scratchpad RAM.

**Table 12 Scratchpad RAM (TriCore 1.3)**

Segment D Regions	Properties
DFFFFFFF <sub>H</sub> – D8000000 <sub>H</sub>	Implementation Dependent
D7FFFFFFF <sub>H</sub> – D4000000 <sub>H</sub>	PSPR
D3FFFFFFF <sub>H</sub> – D0000000 <sub>H</sub>	DSPR

### 8.3.2 Scratchpad RAM (TriCore 1.3.1)

Segment C and D contain the scratchpad RAM. There are two different scratchpad RAMs:

- DSPR - Data scratchpad RAM.
- PSPR - Program scratchpad RAM.

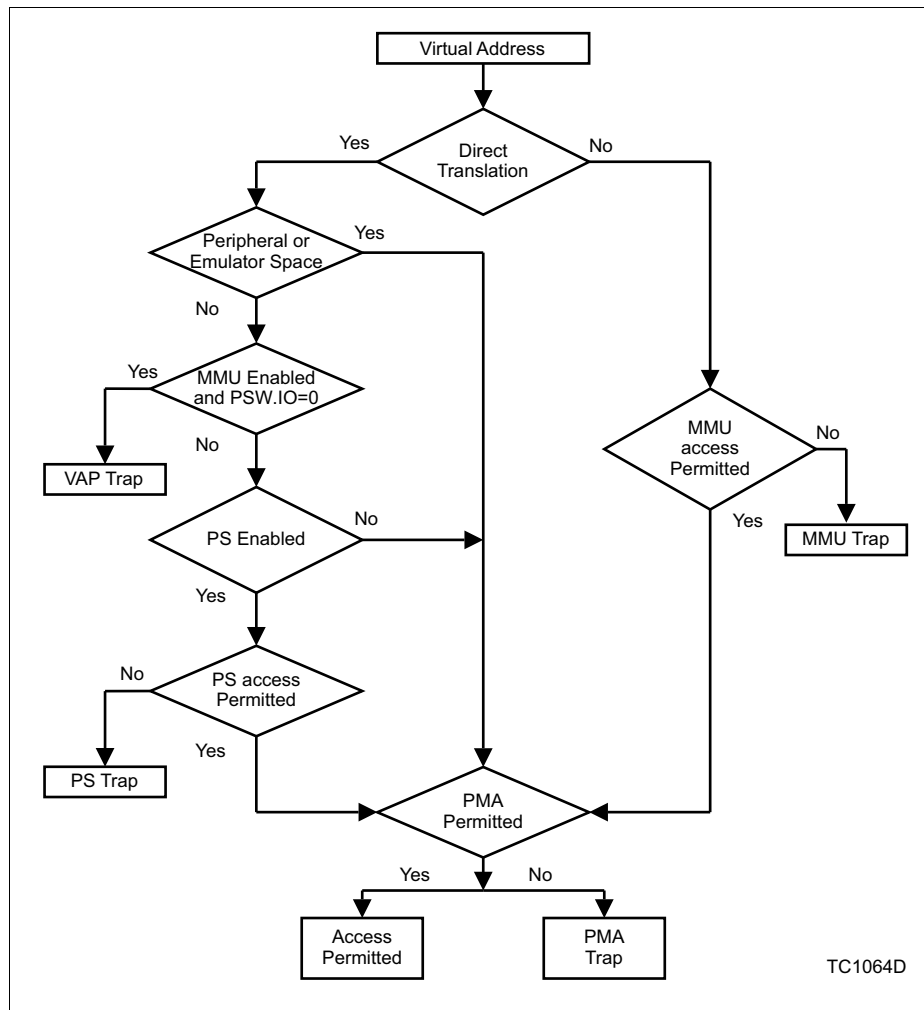
**Table 13 Scratchpad RAM (TriCore 1.3.1)**

Segment C and D Regions	Properties
DFFFFFFF <sub>H</sub> – D8000000 <sub>H</sub>	Implementation Dependent
D7FFFFFFF <sub>H</sub> – D4000000 <sub>H</sub>	PSPR Image
D3FFFFFFF <sub>H</sub> – D0000000 <sub>H</sub>	DSPR
CFFFFFFF <sub>H</sub> – C4000000 <sub>H</sub>	Implementation Dependent
C3FFFFFFF <sub>H</sub> – C0000000 <sub>H</sub>	PSPR

In TriCore 1.3.1, segments C and D are constrained to have the attributes cacheable or non-cacheable, although the cacheable property is only relevant for accesses in the range C8000000<sub>H</sub> – CFFFFFFF<sub>H</sub> and D8000000<sub>H</sub> – DFFFFFFF<sub>H</sub>, implementation defined for data accesses to PSPR and implementation defined for code fetch accesses to DSPR.

## 8.4 Permitted versus Valid Accesses

A memory access can be permitted without necessarily being valid. There are three sources of permission for a memory access, the Protection System (PS), the Memory Management Unit (MMU) and the Physical Memory Attributes (PMA).



**Figure 29 Translation Paths**

---

**Physical Memory Attributes (PMA)**

If a memory access is permitted by the MMU or PS, then it must also be permitted by the PMA for the access to proceed, as shown in [Figure 29](#). A memory access is not valid if the address of the access is to an unimplemented region of memory or is misaligned; therefore an access can be permitted but not valid.

The PS and MMU act upon the direct translation and virtual translation paths respectively, therefore the permission for a memory access that undergoes virtual translation lies only with the MMU, not the PS, and vice-versa.

## **9 Memory Protection System**

The TriCore® protection system provides the essential features needed to isolate errors. The system is unobtrusive, imposing little overhead and avoids non-deterministic run-time behaviour.

The protection system incorporates hardware mechanisms that protect user-specified memory ranges from unauthorized read, write, or instruction fetch accesses.

The protection hardware can also facilitate debugging and generate signals sent to the Core Debug Controller (CDC) (See [Core Debug Controller \(CDC\), page 12-1](#)).

### **9.1 Memory Protection Subsystems**

The following subsystems are involved with Memory Protection.

#### **The Trap System**

A trap occurs as a result of an event such as a Non-Maskable Interrupt (NMI), an instruction exception or illegal access. The TriCore architecture contains eight trap classes and these traps are further classified as synchronous or asynchronous, hardware or software. Covered in detail in [Trap System, page 6-1](#).

#### **The I/O Privilege Level**

There are three I/O modes: User-0 mode, User-1 mode and Supervisor mode. The User-1 mode allows application tasks to directly access non-critical system peripherals. This allows systems to be implemented efficiently, without the loss of security inherent in running in Supervisor mode. Covered in more detail in [Access Privilege Level Control \(I/O Privilege\), page 3-11](#).

#### **Memory Protection**

Provides control over which regions of memory a task is allowed to access, and what types of access it is permitted.

- **Range Based**

The range-based memory protection system is designed for small and low cost applications to provide coarse-grained memory protection for systems that do not require virtual memory. This system is detailed in this chapter.

- **Page Based**

For applications that require virtual memory, the optional Memory Management Unit (MMU) supports a familiar model that gives each memory page its own access permissions. The MMU design and the page-based memory protection model facilitate porting of standard operating systems that expect this model. The MMU is detailed in [Memory Management Unit \(MMU\), page 10-1](#).

## **Peripheral or Emulator Space Protection**

The majority of this chapter is concerned with memory protection, which does not apply to memory regions that have the peripheral space or emulator space attribute. See the chapter [Physical Memory Attributes \(PMA\)](#), page 8-1.

## **Effective Addresses**

Effective addresses are translated into physical addresses using one of two translation mechanisms:

- Direct translation.
- Page Table Entry (PTE) based translation.

Memory protection for addresses that undergo direct address translation is enforced using the range-based memory protection system described in this chapter.

Virtual translation mechanisms are defined in the chapter [Memory Management Unit \(MMU\)](#), page 10-1.

## **9.2 Range Based Memory Protection**

The range-based memory protection system is designed for small and low cost applications to provide coarse-grained memory protection for systems that do not require virtual memory.

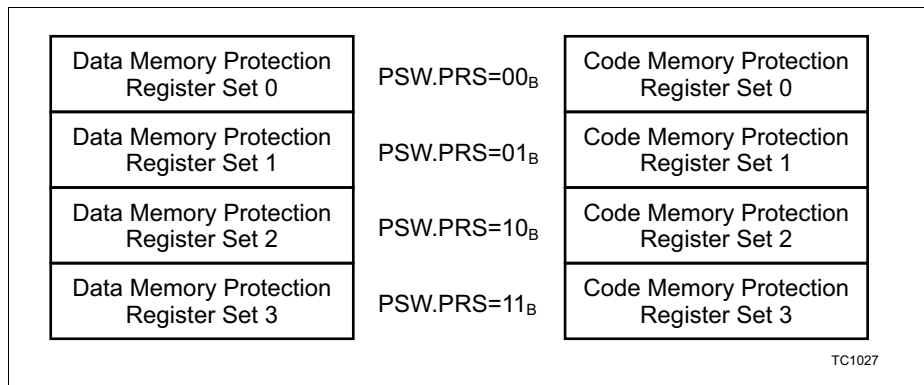
### **9.2.1 Memory Protection Register Sets**

TriCore contains register sets that specify the address range and the access permissions for a number of memory ranges. The PSW.PRS field is used to select which register set is active at a given time. Two register sets are selected simultaneously:

- One Data Memory Protection.
- One Code Memory Protection.

The PSW.PRS field allows selection of up to four such register sets; four for data and four for code. See [Program Status Word Register \(PSW\)](#), page 3-6 for more details on the PSW.

At any given time one of the sets is the current protection register set which determines the legality of memory accesses by the current task or ISR. The PSW.PRS field determines the current protection register set number.



**Figure 30 Memory Protection Register Sets**

The number of register sets provided for memory protection is specific to each TriCore implementation, limited to a minimum of two and a maximum of four. This document only describes the generic format of these register sets. Unimplemented register set addresses are reserved and undefined.

### Range Table Entries

Each register set is made up of several range registers (also called Range Table Entries). The number of range registers in a Data or Code Memory protection set is implementation defined, limited to a minimum of one and a maximum of four for any valid protection set. Implementation may implement differing numbers of code and data range registers in any given protection set. Each Range Table Entry ([Figure 31, page 9-4](#)) consists of a segment Protection register pair and a bit field within a common Mode register.

The register pair specifies the lower and upper boundary addresses of the memory range.

Lower Bound  $\leq$  address  $<$  Upper Bound.

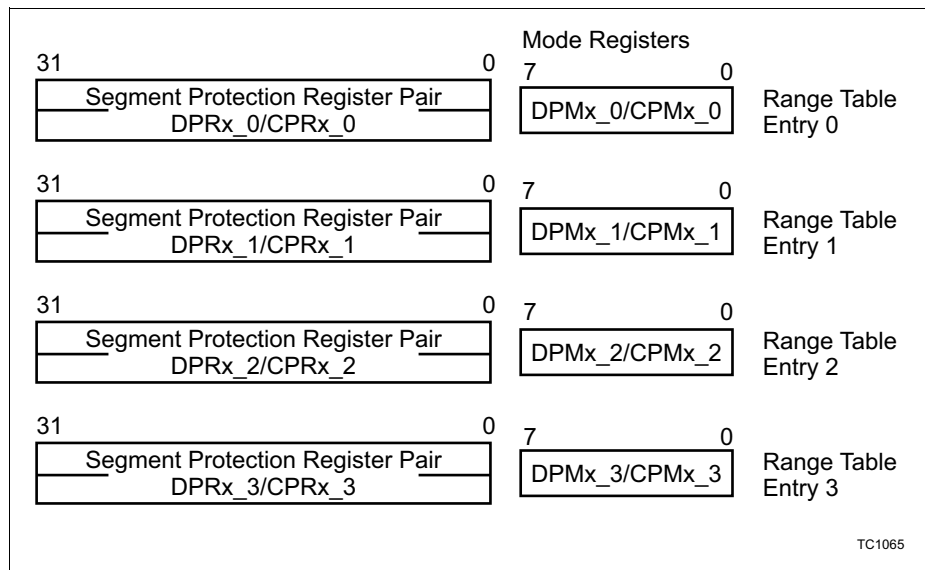
The Mode register contains the access permission. The control options are different for the data and the code memory protection. The Mode register also contains the debug control bits.

For load and store operations, data address values are checked against the entries in the data range table.

On instruction fetches, the PC value for the fetch is checked against the entries in the code range table. When an address is found to fall outside of all ranges defined in the appropriate range table, then permission for the access is denied. When an address is found to fall within a range defined in the appropriate range table, the associated mode table entry is checked for access permissions.

## Memory Protection System

An instruction fetch cannot occur from a byte aligned address, and so the least significant bit of the Code Segment Protection upper and lower bound registers (CPRx\_n) is not writeable and always returns zero.



**Figure 31 Example of a Protection Register Set implementing Four Range Table Entries and Four Data Range Table Entries.**

### Modes of Use for Range Table Entries

Individual range table entries can be used just for memory protection or for debugging. One entry is rarely used for both purposes. If the upper and lower bound values have been set for debug breakpoints they are probably not meaningful for defining protection ranges, and vice-versa. However, it is both possible and reasonable to have some entries used for memory protection and others used for debugging.

To disable an entry for use in memory protection, clear both the RE and WE bits in a data range table entry or clear the XE bit in a code range table entry. The entry can be disabled for use in debugging by clearing any debug signal bits.

When a range entry is being used for debugging, the debug signal bits that are set determine whether it is used as a single range comparator (giving an in-range/not in-range signal) or as a pair of equal comparators. The two uses are not mutually exclusive.



### **Using Protection Register Sets**

Supervisor mode does not automatically disable memory protection. The protection register set that is selected for Supervisor mode tasks will normally be set up to allow write access to regions of memory that are protected from User mode access. In addition Supervisor mode tasks can execute instructions to change the protection maps, or to disable the protection system entirely. But the Supervisor mode does not implicitly override memory protection, and it is possible for a Supervisor mode task to take a memory protection trap.

## 9.3 Memory Protection Registers

### 9.3.1 Data Segment Protection Register - Upper


DPRx\_mU

Data Segment Protection Register x\_m Upper Bound

(x = set number 0 to 3; m = range table entry)

(C004<sub>H</sub>+n\*8<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
UPPBND	[31:0]	rw	DPRx_m Upper Boundary Address

### 9.3.2 Data Segment Protection Register - Lower


DPRx\_mL

Data Segment Protection Register x\_m Lower Bound

(x = set number 0 to 3; m = range table entry)

(C000<sub>H</sub>+n\*8<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
LOWBND	[31:0]	rw	DPRx_m Lower Boundary Address

### 9.3.3 Code Segment Protection Register - Upper


**CPRx\_nU**

**Code Segment Protection Register x\_n Upper Bound**

(x = set number 0 to 3; n = range table entry 0 to 3)

(D004<sub>H</sub>+n\*8<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
UPPND	[31:0]	rw	<b>CPRx_n Upper Boundary Address</b> The least significant bit is not writeable and always returns zero.

### 9.3.4 Code Segment Protection Register - Lower


CPRx\_nL

Code Segment Protection Register x\_n Lower Bound

(x = set number 0 to 3; n = range table entry 0 to 3)

(D000<sub>H</sub>+n\*8<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
LOWBND	[31:0]	rw	<b>CPRx_n Lower Boundary Address</b> The least significant bit is not writeable and always returns zero.

### 9.3.5 Data Protection Mode Register

**DPMx**

**Data Protection Mode Register x**

(x = set number 0 to 3) (E000<sub>H</sub>+n\*80<sub>H</sub>)

**Reset Value: Implementation Specific**

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
<b>WE3</b>	<b>RE3</b>	<b>WS3</b>	<b>RS3</b>	<b>WB L3</b>	<b>RB L3</b>	<b>WB U3</b>	<b>RB U3</b>	<b>WE2</b>	<b>RE2</b>	<b>WS2</b>	<b>RS2</b>	<b>WB L2</b>	<b>RB L2</b>	<b>WB U2</b>	<b>RB U2</b>
rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
<b>WE1</b>	<b>RE1</b>	<b>WS1</b>	<b>RS1</b>	<b>WB L1</b>	<b>RB L1</b>	<b>WB U1</b>	<b>RB U1</b>	<b>WE0</b>	<b>RE0</b>	<b>WS0</b>	<b>RS0</b>	<b>WB L0</b>	<b>RB L0</b>	<b>WB U0</b>	<b>RB U0</b>
rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw	rw

Field	Bits	Type	Description
WE(3-0)	31, 23, 15, 7	rw	<b>Address Field Write Enable</b> 0 : Data write accesses to associated address range not permitted. 1 : Data write accesses to associated address range permitted.
RE(3-0)	30, 22, 14, 6	rw	<b>Address Field Read Enable</b> 0 : Data read accesses to associated address range not permitted. 1 : Data read accesses to associated address range permitted.
WS(3-0)	29, 21, 13, 5	rw	<b>Address Range Data Write Signal</b> 0 : Data write signal disabled. 1 : Signal asserted to Core Debug Controller (CDC) on data write accesses to associated address range.
RS(3-0)	28, 20, 12, 4	rw	<b>Address Range Data Read Signal</b> 0 : Data read signal disabled. 1 : Signal asserted to Core Debug Controller (CDC) on data read accesses to associated address range.

**Memory Protection System**

<b>Field</b>	<b>Bits</b>	<b>Type</b>	<b>Description</b>
WBL(3-0)	27, 19, 11, 3	rw	<b>Data Write Signal on Lower Bound Access</b> 0 : Data write signal disabled. 1 : Signal asserted to Core Debug Controller (CDC) on data write access to an address that matches lower bound address of associated address range.
RBL(3-0)	26, 18, 10, 2	rw	<b>Data Read Signal on Lower Bound Access</b> 0 : Data read signal disabled. 1 : Signal asserted to Core Debug Controller (CDC) on data read access to an address that matches lower bound address of associated address range.
WBU(3-0)	25, 17, 9, 1	rw	<b>Data Write Signal on Upper Bound Access</b> 0 : Write signal disabled. 1 : Signal asserted to Core Debug Controller (CDC) on data write access to an address that matches upper bound address of associated address range.
RBU(3-0)	24, 16, 8, 0	rw	<b>Data Read Signal on Upper Bound Access</b> 0 : Data read signal disabled. 1 : Signal asserted to Core Debug Controller (CDC) on data read access to an address that matches upper bound address of associated address range.

### 9.3.6 Code Protection Mode Register

CPM<sub>x</sub>

Code Protection Mode Register *x*

(*x* = set number 0 to 3) (E200<sub>H</sub>+n\*80<sub>H</sub>)

Reset Value: Implementation Specific

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
<b>XE3</b>	-	<b>XS3</b>	-	<b>BL3</b>	-		<b>BU3</b>	<b>XE2</b>	-	<b>XS2</b>	-	<b>BL2</b>	-		<b>BU2</b>
rw		rw		rw			rw	rw		rw		rw			rw
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
<b>XE1</b>	-	<b>XS1</b>	-	<b>BL1</b>	-		<b>BU1</b>	<b>XE0</b>	-	<b>XS0</b>	-	<b>BL0</b>	-		<b>BU0</b>
rw		rw		rw			rw	rw		rw		rw			rw

Field	Bits	Type	Description
XE(3-0)	31, 23, 15, 7	rw	<b>Address Range Execute Enable</b> 0 : Instruction fetch accesses to associated address range not permitted. 1 : Instruction fetch accesses to associated address range permitted.
-	30, 28, [26:25], 22, 20, [18:17], 14, 12, [10:9], 6, 4, [2:1]	-	<b>Reserved Field</b>
XS(3-0)	29, 21, 13, 5	rw	<b>Address Range Execute Signal</b> 0 : Execute signal disabled. 1 : Signal asserted to Core Debug Controller (CDC) on instruction fetch accesses to associated address range.



**Memory Protection System**

Field	Bits	Type	Description
BL(3-0)	27, 19, 11, 3	rw	<b>Execute Signal on Lower Bound Access</b> 0 : Lower bound execute signal disabled. 1 : Signal asserted to Core Debug Controller (CDC) on instruction fetch access to an address that matches lower bound address of associated address range.
BU(3-0)	24, 16, 8, 0	rw	<b>Execute Signal on Upper Bound Access</b> 0 : Upper bound execute signal disabled. 1 : Signal asserted to Core Debug Controller (CDC) on instruction fetch access to an address that matches upper bound address of associated address range.

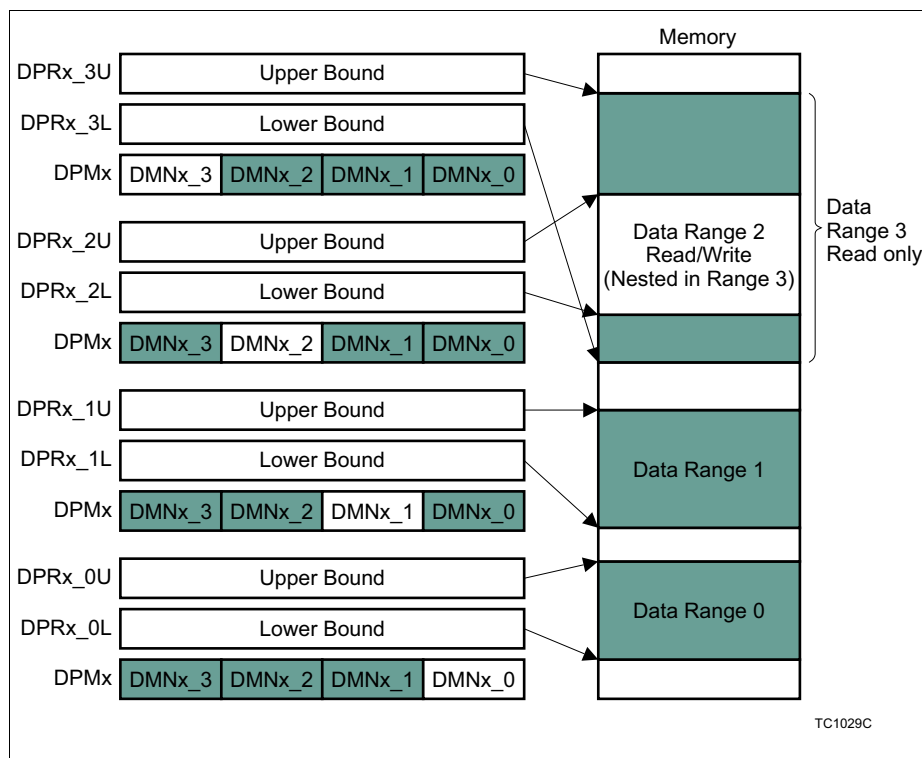
## 9.4 Access Permissions for Intersecting Memory Ranges

The permission to access a memory location is the OR of the memory range permissions. If one of the ranges allows it, the memory access is permitted. This means that when two ranges intersect, the intersecting regions will have the permission of the most permissive range.

For example, if range A is set for read/write permission and range B read-only, the intersecting region of A and B will be read/write. Nesting of ranges can be used for example to allow read/write access to a subrange of a larger range in which the current task is allowed read access.

### 9.4.1 Example Data Protection Register Set

**Figure 32** illustrates the Data Protection Register Set  $x$ , where  $x$  is one of the four sets as selected by the PSW.PRS field. The register set in this example consists of four range table entries.



**Figure 32** Example Configuration of a Data Protection Register Set

In **Figure 32** the four Data Segment and four Data Protection Mode Registers are set up as follows:

- Data Segment Protection Registers DPRx\_3U and DPRx\_3L, define the upper (U) and lower (L) bound for Data Range 3. Data Protection Mode Register 3 (DPMx\_3) defines the read-only permissions for Data Range 3.
- DPRx\_2U and DPRx\_2L define the upper (U) and lower (L) bound for Data Range 2. DPMx\_2 defines the read-write permissions for Data Range 2.

*Note: Data Range 2, which has read/write permission, is nested within Data Range 3, which has read-only permission. Because the intersecting rules state that permission to access a memory location is the OR of the regions permissions, this region therefore has read/write permission.*

- DPRx\_1U and DPRx\_1L define the upper (U) and lower (L) bound for Data Range 1. DPMx\_1 defines the permissions for Data Range 1.
- DPRx\_0U and DPRx\_0L define the upper (U) and lower (L) bound for Data Range 0. DPMx\_0 defines the permissions for Data Range 0.

This same configuration can be used to illustrate Code Protection Register Set x.

## **9.5 Using the Memory Protection System**

When the protection system is enabled every memory access (read, write or execute) is checked for legality before the access is performed. The legality is determined by all of the following:

- The Protection Enable bit in the SYSCON register (SYSCON.PROTEN).
- The currently selected protection register set (PSW.PRS).
- The ranges defined in the protection register set.

### **9.5.1 Protection Enable bit**

For the memory protection system to be active, the Protection Enable bit (SYSCON.PROTEN) must be set to one (SYSCON.PROTEN == 1). If the memory protection system is disabled (SYSCON.PROTEN == 0), then any access to any memory address is permitted.

### **9.5.2 Set Selection**

At any given time, one of the sets is the current protection register set which determines the legality of memory accesses by the current task or Interrupt Service Routine. The PSW.PRS field indicates the current Protection Register Set number.

### **9.5.3 Address Range**

Data addresses (read and write accesses) are checked against the currently selected data address range table, while instruction fetch addresses are checked against the code address range tables. The mode entries for the data range table entries enable only read and write accesses, while the mode entries for the code range table entries enable only execute access.

In order for data to be read from program space, there must be an entry in the data address range table that covers the address being read. Conversely there must be an entry in the code address range table that covers the instruction being read.

The protection system does not differentiate between access permission levels. The data and code protection settings have the same effect, whether the permission level is currently set to Supervisor, User-1 or User-0 mode.

The protection system does not apply to accesses in memory regions with the peripheral space or emulator space attribute. If a memory access is attempted to either of these segments, the access is permitted by the protection system (but not necessarily the Physical Memory Attributes) regardless of the protection system settings. For more on PMA, see [Physical Memory Attributes \(PMA\), page 8-1](#).

Saves or restores of contexts to the context save area (see [Save and Restore Context Operations, page 4-5](#)) do not require the permission of the protection system to proceed.

### 9.5.4 Traps

There are three traps generated by memory protection, each corresponding to the three protection mode register bits:


- MPW (Memory Protection Write) trap = WE bit.
- MPR (Memory Protection Read) trap = RE bit.
- MPX (Memory Protection Execute) trap = XE bit.

Refer to [Chapter 6 Trap System, page 6-1](#) for a complete description of Traps.

### 9.6 Crossing Protection Boundaries

A memory access can straddle two regions defined by the protection system. [Figure 33](#) shows a memory access (code or data) crossing the boundary of a permitted region and a 'not permitted' region of memory. In this situation it is implementation defined as to whether or not a memory protection trap is taken.

To ensure deterministic behaviour in all implementations of TriCore, a region at least twice the size of the largest memory accesses, minus one byte, should be left as a buffer between each memory protection region.



**Figure 33    Protection Boundaries**



## 10 Memory Management Unit (MMU)

This chapter describes the TriCore® Memory Management Unit (MMU) architecture. The MMU is an optional component in TriCore configurations. It need not be present in every system that uses the core, and even when present it can be disabled.

If the MMU is not present and enabled in a system, then virtual memory and page-based memory access protection are not supported for that system. The range-based protection system (a non-optional core component) still provides basic memory protection services. For more details see [Memory Protection System, page 9-1](#).

The memory management features include:

- 4 GBytes of virtual address space divided into 16 segments of 256 MBytes each. The upper half of the virtual address space (segments [ $8_H - F_H$ ]) is global, and mapped directly onto the physical address space. The lower segment (segments [ $0_H - 7_H$ ]) is implicitly qualified by an Address Space Identifier (ASI). The operating system can allocate distinct address spaces to each unique ASI. Two or more processes can also share an address space ID, either serially or in parallel.
- 4 GBytes of physical address space divided into 16, 256 MByte segments.
- Virtual to physical address translation by direct translation or via Page Table Entries (PTE), depending on the segment number of the virtual address and the status of the MMU.
- Cacheability and access permissions based on physical memory attributes for directly translated addresses, or by a combination of physical memory attributes and virtual page attributes for addresses translated via Page Table Entries. Attributes for segments [ $8_H - F_H$ ] are pre-defined in a system memory map.

Virtual addresses are always translated into physical addresses before accessing memory. The virtual address is translated into a physical address using either direct translation or Page Table Entry (PTE) translation.

- Direct translation: If the virtual address belongs to the upper segment of the virtual address space then the virtual address is directly used as the physical address. If the virtual address belongs to the lower segment of the address space, then the virtual address is used directly as the physical address if the processor is operating in Physical mode (i.e. the MMU is disabled or not present).
- PTE translation: If the virtual address belongs to the lower segment of the address space and the processor is operating in Virtual mode (i.e. the MMU is present and enabled), then the virtual address is translated using a Page Table Entry.

PTE translation is performed by replacing the Virtual Page Number (VPN) of the virtual address by a Physical Page Number (PPN) to obtain a physical address.

Six memory-mapped Memory Management Unit (MMU) Core Special Function Registers (CSFRs) control the memory management system.

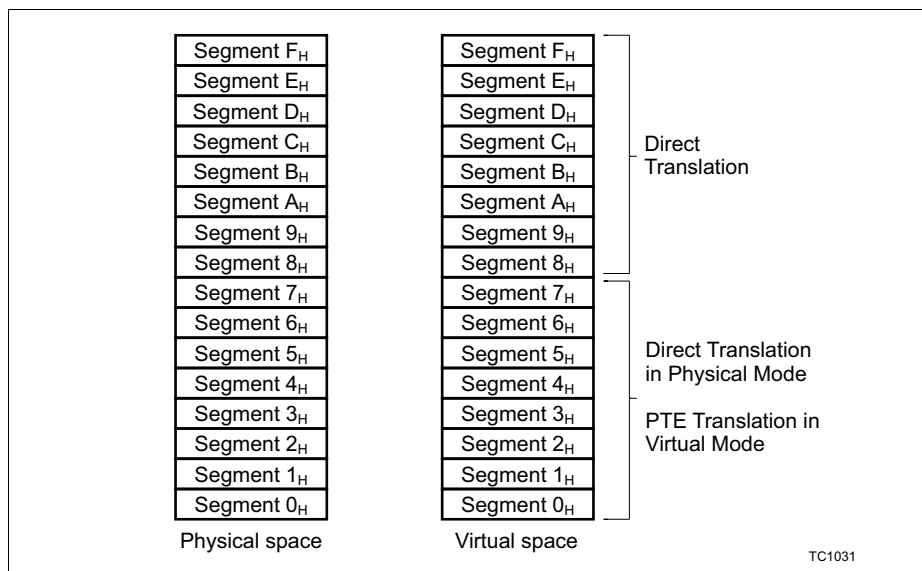
## 10.1 Address Spaces

The TriCore virtual address space is 4 GBytes in size and divided into 16 segments, with each segment consisting of 256 MBytes. The upper 4 bits of the 32-bit virtual address are used to identify the segment. Segments are numbered  $[0_H - F_H]$ .

*Note: A virtual address is always translated into a physical address before accessing memory.*

The physical address space is 4 GBytes in size and is divided into 16 segments of 256 MBytes each.

The physical and virtual address space maps are shown in the following figure:



**Figure 34 Physical and Virtual Address Spaces**

A 32-bit virtual address is comprised of a Virtual Page Number (VPN) concatenated with a Page Offset.

A 32-bit physical address is comprised of a Physical Page Number (PPN) concatenated with a Page Offset.



## 10.2 Address Translation

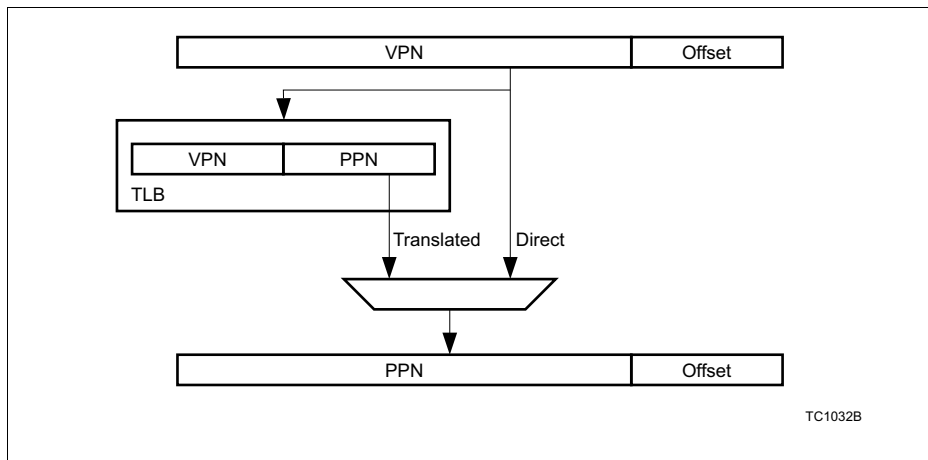
The MMU\_CON.V bit controls the operating mode, Physical or Virtual, of the processor; Physical mode if MMU\_CON.V == 0, or Virtual mode if MMU\_CON.V == 1 (See [MMU Configuration Register \(MMU\\_CON\)](#), page 10-13).

The virtual address is translated into a physical address using either direct translation or Page Table Entry translation, as shown in [Figure 35](#).

Translation using the Page Table Entry (PTE) is performed by replacing the Virtual Page Number (VPN) of the virtual address by a Physical Page Number (PPN), to obtain a physical address.

If the virtual address is from the upper segments of the virtual address space then the virtual address is used directly as the physical address (direct translation).

If the virtual address is from the lower segments of the address space, then the virtual address is used directly as the physical address if the processor is operating in Physical mode, or translated using a Page Table Entry if the processor is operating in Virtual mode.



**Figure 35 Virtual Address Translation**

### 10.2.1 Address Translation for CSFR Pointers

The context pointers (PCX, FCX, and LCX), the Base Trap Vector (BTV) and the Base Interrupt Vector (BIV) are constrained to use segments that undergo direct translation. See [Context Management Registers](#), page 4-13 and [Interrupt and Trap Control Registers](#), page 6-17.

### 10.3 Translation Lookaside Buffers (TLBs)

The MMU provides PTE-based virtual address translation through two Translation Lookaside Buffers (TLBs); TLB-A and TLB-B.

The MMU supports four page sizes; 1 KByte, 4 KBytes, 16 KBytes, and 64 KBytes, although not all of these sizes can be used at once. At any given time each TLB provides translations for only one particular page size. The page size setting of each TLB is determined through the MMU\_CON.SZA and MMU\_CON.SZB fields, as described in [MMU Configuration Register \(MMU\\_CON\), page 10-13](#).


Each TLB contains a number ( $N$ ) of TLB Table Entries (TTEs), where  $N$  is a minimum of 4 and a maximum of 128. The MMU\_CON.TSZ field determines the size of each TLB.

Each TTE has an 8-bit index associated with it:

- Index numbers 0, ..., MMU\_CON.TSZ, are used for the entries in TLB-A.
- Index numbers 128, ..., 128+MMU\_CON.TSZ, are used for the entries in TLB-B.

Each TTE contains a Page Table Entry (PTE).

The organization (associativity) of each TLB is implementation-dependent.



**Figure 36 Translation Lookaside Buffers**

### 10.3.1 TLB Table Entry (TTE) Contents

TLB Table Entries (TTE) contain the following fields:

- **Address Space Identifier (ASI):** Specifies the address space corresponding to the virtual address. ASIs allow mappings of up to 32 virtual address spaces to coexist in a TLB. An ASI is similar to a Process ID.
- **Virtual Page Number (VPN):** Stores  $32 - \log_2 \text{Pagesize}$  bits where *Pagesize* is the size of the page in bytes.
- **Physical Page Number (PPN):** Stores  $32 - \log_2 \text{Pagesize}$  bits where *Pagesize* is the size of the page in bytes.
- **Execute Enable (XE):** Allows instruction fetches from the virtual page.
- **Write Enable (WE):** Allows data writes to the virtual page.
- **Read Enable (RE):** Allows data reads from the virtual page.
- **Cacheability bit (C):** Allows the virtual page to be cached.
- **Global bit (G):** Indicates that the page is globally mapped therefore making it visible in all address spaces.
- **Valid bit (V):** Indicates that the TTE contains a valid mapping.

### 10.4 Multiple Address Spaces

The MMU provides efficient support for multiple and shared virtual address spaces. Each TTE (TLB Table Entry) contains an Address Space Identifier (ASI) which can identify the distinct address space corresponding to the particular mapping. The ASI Register (MMU\_ASI) provides the current address space identifier for all memory accesses.

Virtual address translation is performed by a valid TTE if:

- It is a non-global TTE that matches the incoming VPN of the virtual address and the Address Space Identifier contained in the ASI register.
- It is a global TTE that matches the incoming VPN.

*Note: Global TTEs are indicated by the TTE[ ].G bit. Such mappings are visible to all virtual address spaces.*

### 10.5 MMU Traps

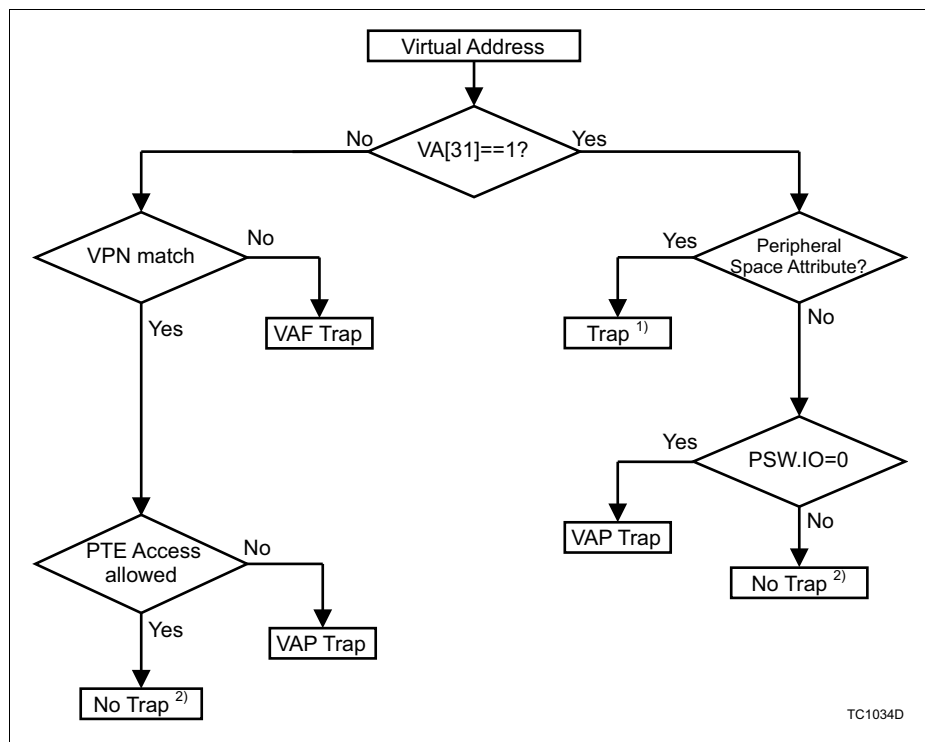
MMU traps belong to Trap Class Number (TCN) 0. The MMU can generate the following traps:

- VAF (Virtual Address Fill).
- VAP (Virtual Address Protection).

The VAF trap is generated if PTE-based translation is required for a virtual address and there is no PTE translation in the MMU. The VAP trap is generated if there is a matching PTE and the access is disallowed. The VAP trap can also occur when User-0 accesses upper segments in virtual mode.

The VAF trap is assigned a TIN (Trap Identification Number) of 0 while the VAP trap is assigned a TIN of 1. Both the VAF and VAP traps are synchronous traps.

The events that happen on an MMU trap are the same as those that happen on any other trap. In addition to context saving and control transfer, the virtual address is right shifted by  $10 + 2 * \min(\text{MMU\_CON.SZA}, \text{MMU\_CON.SZB})$ , and loaded into the Translation Fault Page Address register (MMU\_TFA). **Figure 37** shows Virtual mode traps for read, write and fetch accesses.



**Figure 37 Virtual Mode Traps for Read, Write and Fetch Accesses<sup>1) 2)</sup>**

Any User-0 mode access to a virtual address in the upper segments that does not have the peripheral space attribute results in a VAP trap. See [Trap Types, page 6-1](#) for more on traps.

<sup>1)</sup> In User-0 mode the MPP trap is for read/write accesses, and the PSE trap for fetch accesses. In User-1 and Supervisor modes, the PSE trap is for fetch accesses. There is no trap for read/write accesses in these modes.

<sup>2)</sup> Subject to constraints imposed by the Physical Memory Attributes. See [Physical Memory Attributes \(PMA\), page 8-1](#).

## **10.6 Virtual Mode Protection**

Memory protection is enforced using separate mechanisms for the two translation paths. These are described in this section.

### **10.6.1 Direct Translation**

Virtual memory protection for addresses that undergo direct translation is provided using standard TriCore range-based protection. The range-based protection mechanism provides support for protecting memory ranges from unauthorized read, write or instruction fetch accesses. Refer to the chapter [Memory Protection System, page 9-1](#).

### **10.6.2 Page Table Entry (PTE) Based Translation**

Virtual memory protection for addresses that undergo PTE-based translation is provided by properties of the PTE used for the address translation. The PTE provides support for protecting a virtual address from unauthorized read, write or instruction fetches by other processes. The following PTE bits are provided for this purpose:

- Execute Enable (XE) - allows instruction fetch from the virtual page.
- Write Enable (WE) - allows data writes to the virtual page.
- Read Enable (RE) - allows data reads from the virtual page.

For each of these bits; 1 = allows and 0 = disallows.

See [Translation Physical Address Register \(MMU\\_TPA\), page 10-17](#).

## **10.7 Cacheability**

A memory access is cacheable if both the virtual address is allowed to be cached and the physical memory attributes allow the access to be cached. The physical memory attributes (PMA) are described in [Physical Memory Attributes \(PMA\), page 8-1](#).

### **10.7.1 Direct Translation Virtual Address Cacheability**

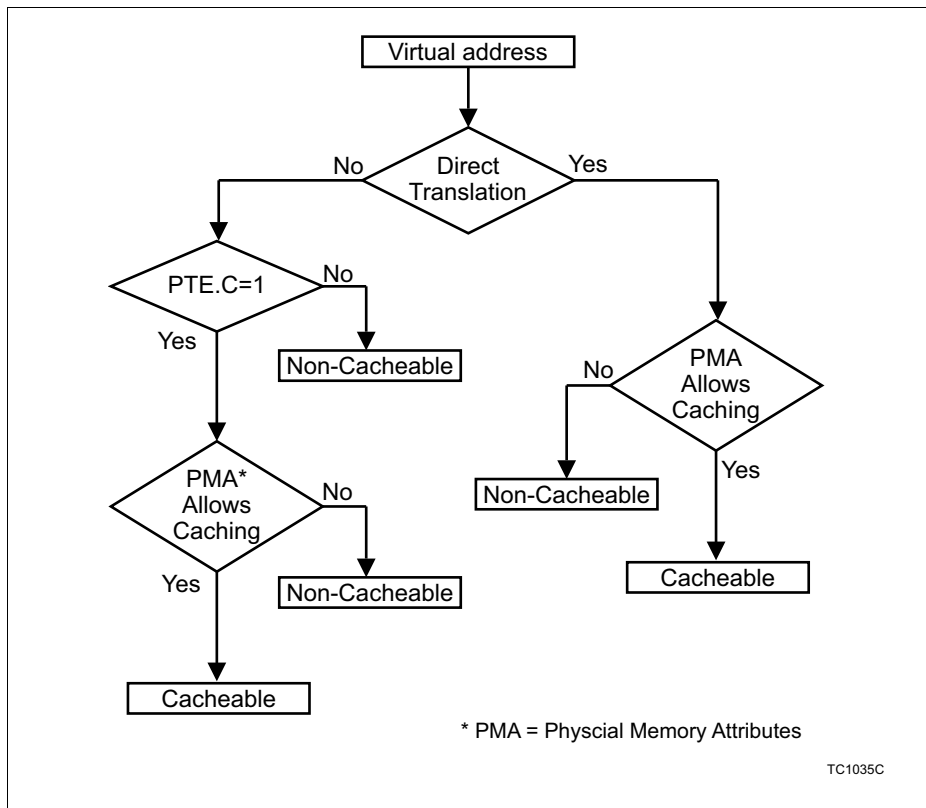
For all virtual addresses undergoing Direct Translation the virtual address is cacheable.

### **10.7.2 PTE Translation Cacheability**

For all virtual addresses undergoing PTE (Page Table Entry) Translation, the virtual address is cacheable if the PTE entry C bit is set and not cacheable if the PTE entry C bit is reset.

### 10.7.3 Cacheability of a Virtual Address Flow

The following figure describes the determination for cacheability of a memory access.



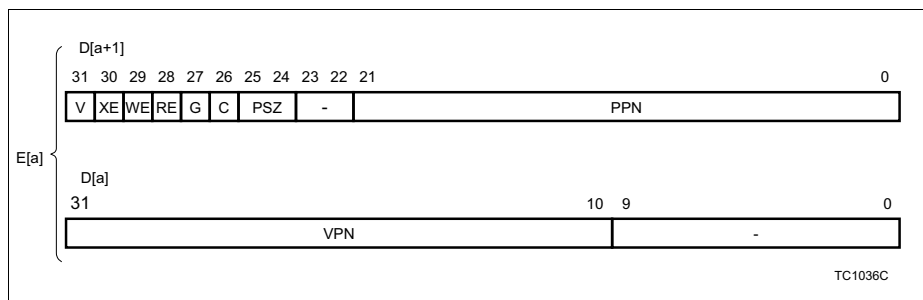
**Figure 38 Cacheability of a Virtual Address**

## 10.8 MMU Instructions

All MMU instructions are privileged instructions that require Supervisor mode for execution. If the MMU is physically present (`MMU_CON.NOMMU == 0`) the instructions execute normally, whether or not the MMU is enabled (`MMU_CON.V == 0` or `1`). If the MMU is not present (`MMU_CON.NOMMU == 1`), then all MMU instructions cause an unimplemented opcode (UOPC) trap.

### 10.8.1 TLBMAP (TLB Map)

The TLBMAP instruction is used to install a mapping in the MMU. The TLBMAP instruction takes an extended data register E[a] as a parameter. The data register D[a] contains the virtual address for the translation and the data register D[a+1] contains the page attributes and PPN (Physical Page Number). The ASI (Address Space Identifier) for the translation is obtained from the MMU\_ASI register. The page attributes are contained in the most significant byte of the odd register with the format as shown:



**Figure 39 TLBMAP E[a] Format**

Entering a mapping with any of the following properties results in undefined behaviour:

- A mapping with a virtual page that wholly or partly overlaps an existing virtual page mapping. A virtual page mapping will overlap with another virtual page mapping if the mappings have an equal or enclosing VPN and the same ASI or at least one of the mappings has its global bit (G) set.
- A mapping for a page size which is not one of the two valid page sizes.
- A mapping using a physical address in a memory region that has the peripheral space attribute.
- A mapping where the VPN is in the upper half of the address space.
- A mapping where the V bit is set to 0.

Undefined behaviour in the context of a TLBMAP instruction means that the MMU TLBs contain undefined PTEs (Page Table Entries). To restore defined behaviour both TLBs have to be flushed (see [TLBFLUSH \(TLB Flush\), page 10-10](#)), and the valid PTEs re-entered.


Entering a mapping when the two TLBs have identical page size settings results in the mapping being entered in one of the two TLBs, with the choice being implementation dependent.

Bits D[a][9:0] and D[a+1][23:22] are reserved and are 0. Bits D[a][15:10] and D[a+1][5:0] are reserved when unused, and therefore are 0 when unused by the page size. For example, if the page size (PSZ) is set to 4 KBytes (01<sub>B</sub>) then bits D[a][11:10] of the VPN are unused and must be set to 0. Similarly, bits D[a+1][1:0] of the PPN are also unused and must be set to 0.

Note that a TLBMAP instruction must be followed by an ISYNC instruction before attempting to use a new installed mapping.

### 10.8.2 TLBDEMAP (TLB Demap)

The TLBDEMAP instruction is used to uninstall a single mapping in the MMU. TLBDEMAP takes as a parameter a data register that contains the virtual address whose mapping is to be removed. The address space identifier for the demap operation is obtained from the Address Space Identifier (ASI) register (MMU\_ASI). Demapping a translation that is not present in the MMU is legal and results in a NOP.



**Figure 40 TLBDEMAP D[a] Format**

*Note: A TLBDEMAP instruction should be followed by an ISYNC before any access to an address in the demapped page is made.*

### 10.8.3 TLBFLUSH (TLB Flush)

The TLBFLUSH instructions are used to flush mappings from the MMU. There are two variants of the TLBFLUSH instruction:

- TLBFLUSH.A flushes all the mappings from TLB-A.
- TLBFLUSH.B flushes all the mappings from TLB-B.

*Note: A TLBFLUSH instruction should be followed by an ISYNC before any access to an address requiring PTE translation is made.*




### 10.8.4 TLBPROBE (TLB Probe)

The TLBPROBE instructions are TLBPROBE.A and TLBPROBE.I.

#### TLBPROBE.A (TLB Probe Address)


This instruction takes a data register D[a] as a parameter and is used to probe the MMU for a virtual address. The D[a] register contains the virtual address for the probe. The address space identifier for the probe is obtained from the ASI (Address Space Identifier) register.



**Figure 41 TLBPROBE.A**

#### TLBPROBE.I (TLB Probe Index)

This instruction takes a data register D[a] as a parameter and is used to probe the TLB at a given TLB index. The D[a] register contains the index for the probe. The index for the TLBs is implementation specific and there is no architecturally defined way to predict what TLB index value will be associated with a given address mapping. Bits D[a][31:8] are reserved and must be set to 0's.



**Figure 42 TLBPROBE.I**

The TLBPROBE instructions return the following:

- The ASI and VPN (Virtual Page Number) of the translation in the Translation Virtual Address register (TVA).
- The PPN (Physical Page Number) and attributes in the Translation Physical Address register (TPA).
- The TLB index of the translation in the Translation Page Index register (TPX).

The MMU\_TPA.V bit is set to zero if the TTE contained an invalid translation or an invalid index was used for the probe. All of the other bitfields are undefined.

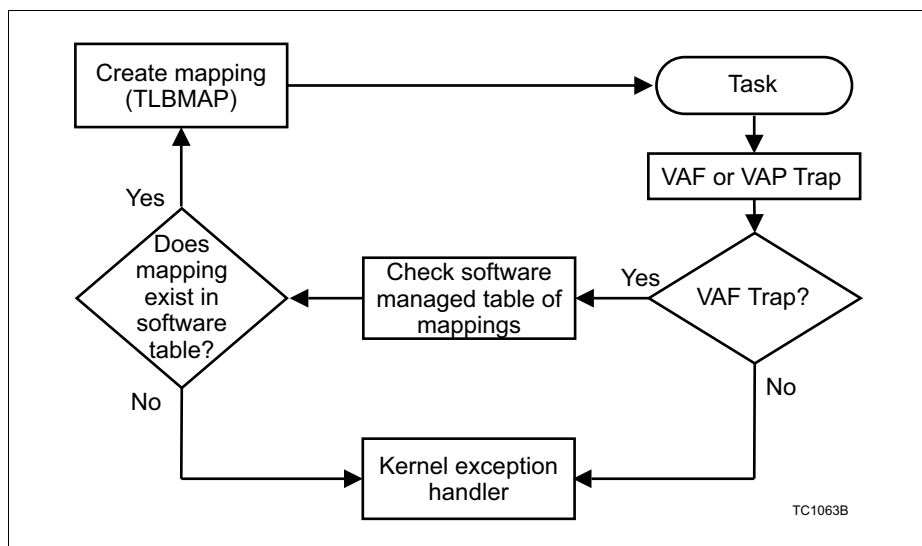
## 10.9 TLB Usage

The TriCore architecture does not specify any PTE replacement algorithm. Entry of a valid new mapping using the TLBMAP instruction does not guarantee the continued existence of a previously entered mapping even if the TLB has not been filled; i.e. the replacement algorithm may over-write a previous mapping. An implementation will always provide a means to ensure forward progress of instructions requiring multiple mappings to execute. Use of the MMU will therefore involve a software architecture with the same fundamental mechanisms as shown in [Figure 43](#).

When executing TLBDEMAP, TLBFLUSH.A or the TLBFLUSH.B instruction, an implementation may require additional operations to be performed in order to maintain coherency of the processors view of memory. For example, removing a mapping that has the cacheability bit set using the TLBDEMAP instruction may require the data and program caches to be flushed before a new mapping can be entered that uses an overlapping physical page with the cacheability bit clear.

An implementation may also impose additional restrictions on the PTEs that can be mapped in order to maintain coherency of the processors view of memory. For example, an implementation may require that the least significant *n* bits of cacheable PTEs, VPN and PPN, be identical to avoid aliasing in a virtually indexed cache.

Context saves and restores are always directly translated irrespective of the segment the context save area resides in.



**Figure 43 Using TLB (Translation Lookaside Buffer)**

## 10.10 MMU Core Special Function Registers

All MMU Core Special Function Registers can be read using the MFCR instruction. The MMU\_CON and MMU\_ASI registers are the only software-writeable registers. They are both written using the MTCR instruction. The MTCR instruction must be followed by an ISYNC instruction before any PTEs are used.

*Note: If MMU\_CON.NOMMU == 1 (MMU not present) then all other registers in the section do not exist and are undefined. If they are accessed no error occurs, but the read and write results are undefined.*

*Note: If no MMU is present (MMU\_CON.NOMMU == 1), then MMU instructions will cause a UOPC (Unimplemented Opcode) trap.*

All MMU registers other than the MMU\_CON register have undefined values at reset.

### 10.10.1 MMU Configuration Register (MMU\_CON)

A MTCR instruction that changes the SZA (Page Size A) bit field must be preceded by a TLBFLUSH.A instruction to ensure that TLB A remains coherent with the programmers view. Similarly a MTCR instruction that changes the SZB (Page Size B) bit field must be preceded by a TLBFLUSH.B instruction. MTCR instructions that change the MMU\_CON.V bit must only be executed from memory addresses undergoing direct translation or, in the case of disabling the MMU, be executed from a virtual address that translates to the exact same physical address, otherwise the MMU behaviour is undefined.

#### MMU\_CON

**Configuration Register**

**(8000<sub>H</sub>)**

**Reset Value: Implementation Specific**

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
<b>NO MMU</b>	-			<b>TSZ</b>						<b>SZB</b>		<b>SZA</b>		<b>V</b>	
r				r						rw		rw		rw	

Field	Bits	Type	Description
-	[31:16]	-	<b>Reserved Field</b>
NOMMU	15	r	<b>No MMU Available</b> 0 : MMU is present. 1 : MMU is not present (all other bits in MMU_CON undefined). Note that to enable the MMU when present (MMU_CON.NOMMU == 0), the MMU_CON.V bit must be set.
-	[14:12]	-	<b>Reserved Field</b>
TSZ	[11:5]	r	<b>TLB Size</b> Determines the size of each TLB. The entries of TLB-A are indexed 0 through TSZ while the entries of TLB-B are indexed 128 through 128+TSZ. Each TLB has a maximum of TSZ+1 entries.
SZB	[4:3]	rw	<b>Page Size B</b> Page size of the mappings in TLB-B. 00 <sub>B</sub> : 1 KByte. 01 <sub>B</sub> : 4 KByte. 10 <sub>B</sub> : 16 KByte. 11 <sub>B</sub> : 64 KByte.
SZA	[2:1]	rw	<b>Page Size A</b> Page size of the mappings in TLB-A. 00 <sub>B</sub> : 1 KByte. 01 <sub>B</sub> : 4 KByte. 10 <sub>B</sub> : 16 KByte. 11 <sub>B</sub> : 64 KByte.
V	0	rw	<b>Virtual mode</b> 0 : Physical mode. 1 : Virtual mode. This bit enables the MMU when the MMU is present (MMU_CON.NOMMU == 0).


### 10.10.2 Address Space Identifier Register (MMU\_ASI)

The Memory Management Unit (MMU), Address Space Identifier (ASI) register description.

#### MMU\_ASI

#### Address Space Identifier Register (8004<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
-	[31:5]	-	<b>Reserved Field</b>
ASI	[4:0]	rw	<b>Address Space Identifier</b> The ASI register contains the Address Space Identifier of the current process.

### 10.10.3 Translation Virtual Address Register (MMU\_TVA)


The MMU\_TVA register is used to return the ASI and VPN (Virtual Page Number) result of a TLBPROBE instruction.

#### MMU\_TVA

#### Translation Virtual Address Register

(800C<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
-	[31:29]	-	<b>Reserved Field</b>
ASI	[28:24]	r	<b>Address Space Identifier</b> The ASI field contains the Address Space Identifier of the PTE.
-	[23:22]	-	<b>Reserved Field</b>
VPN	[21:0]	r	<b>Virtual Page Number</b> The VPN of the PTE accessed by the last TLBPROBE instruction.


#### 10.10.4 Translation Physical Address Register (MMU\_TPA)

The MMU\_TPA register is used to return the PPN (Physical Page Number) and attributes of a translation by a TLBPROBE instruction.

##### MMU\_TPA

##### Translation Physical Address Register (8010<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
V	31	r	<b>Valid bit</b> Indicates that the TTE contains a valid mapping. 0 : Invalid. 1 : Valid.
XE	30	r	<b>Execute Enable</b> Enables instruction fetches to the page. 0 : Disabled. 1 : Enabled.
WE	29	r	<b>Write Enable</b> Enables data writes to the page. 0 : Disabled. 1 : Enabled.
RE	28	r	<b>Read Enable</b> Enables data reads from the page. 0 : Disabled. 1 : Enabled.
G	27	r	<b>Global</b> Indicates that the page is globally mapped therefore making it visible in all address spaces. 0 : Not globally mapped. 1 : Globally mapped.

**Memory Management Unit (MMU)**

Field	Bits	Type	Description
C	26	r	<b>Cacheability</b> Indicates that the page is cacheable. 0 : Not Cacheable. 1 : Cacheable.
PSZ	[25:24]	r	<b>Page Size</b> 00 <sub>B</sub> : 1 KByte. 01 <sub>B</sub> : 4 KByte. 10 <sub>B</sub> : 16 KByte. 11 <sub>B</sub> : 64 KByte.
0	[23:22]	-	<b>Reserved Field</b>
PPN	[21:0]	r	<b>Physical Page Number</b> Holds the PPN from the PTE.




### 10.10.5 Translation Page Index Register (MMU\_TPX)

The MMU\_TPX register is used to return the TLB (Translation Lookaside Buffer) index result of a TLBPROBE instruction.

#### MMU\_TPX

Translation Page Table Index Register (8014<sub>H</sub>)

Reset Value: Implementation Specific



Field	Bits	Type	Description
-	[31:8]	-	Reserved Field
INDEX	[7:0]	r	Translation Index


### 10.10.6 Translation Fault Page Address Register (MMU\_TFA)

The MMU\_TFA register contains the faulting virtual page number (where faulting refers to a VAP or VAF trap). It is the faulting virtual address, right shifted by  $10 + 2 * \min(\text{SZA}, \text{SZB})$  bits.

#### MMU\_TFA

##### Translation Fault Page Address Register(8018<sub>H</sub>)

**Reset Value: Implementation Specific**



Field	Bits	Type	Description
-	[31:22]	-	Reserved Field
FPN	[21:0]	r	<b>Faulting Page Number</b> VPN from the faulting Virtual Address.

## **11 Floating Point Unit (FPU)**

This chapter describes the TriCore® Floating Point Unit (FPU) architecture. The FPU is an optional component in TriCore configurations. It need not be present in every system that uses the core, and even when present it can be disabled.

The optional FPU is an IEEE-754 compatible floating-point unit to accompany the TriCore instruction set.

### **11.1 Functional Overview**

The FPU executes single precision IEEE-754 compatible floating-point arithmetic instructions and supports the following feature set:

- Floating-point add, subtract, multiply, MAC, and divide instructions.
- Conversion to or from IEEE-754 single precision format from or to TriCore signed and unsigned integers and 32-bit signed fractions (Q31 format).
- QSEED.F instruction used to obtain an approximate value intended for use in Newton-Raphson iterations to perform a square-root operation.
- Comparison of two floating-point numbers.
- All four IEEE-754 rounding modes are implemented.
- Asynchronous traps can be generated on selected IEEE-754 exceptions (TriCore 1.3.1).


### **Restrictions**

The FPU has the following restrictions and usage limitations:

- Only IEEE-754 single precision format is supported.
- IEEE-754 denormalized numbers are not supported for arithmetic operations.
- IEEE-754 compliant remainder function cannot be implemented using FPU instructions because of the effects of multiple rounding when using a sequence of individually rounded instructions.
- Fused multiply-and-accumulate operations (MACs) are not part of the IEEE-754 standard. Using FPU MAC operations can give different results from using separate multiply and accumulate operations because the result is only rounded once at the end of a MAC.
- Full compliance with the IEEE-754 standard is not achieved because denormal numbers are not supported.
- If no FPU is present, then FPU instructions will cause a UOPC (unimplemented opcode) trap.

## 11.2 IEEE-754 Compliance

### 11.2.1 IEEE-754 Single Precision Data Format



**Figure 44 Single Precision IEEE-754 Floating-Point Format**

The single precision IEEE-754 floating-point format has three sections: a sign bit, an 8-bit biased exponent, and a 23-bit fractional mantissa with an implied binary point before bit 22. For normal numbers the mantissa has an implied 1 immediately to the left of the binary point. **Table 14** shows the different types of number representation in IEEE-754 single precision format. In this table:

s = bit [31]: sign bit.

e = bits [30:23]: biased exponent.

f = bits [22:0]: fractional part of mantissa.

**Table 14 IEEE-754 Single Precision Representation Types**

Condition	Represented Value	Description
$0 < e < 255$	$(-1)^s \cdot 2^{(e-127)} \cdot 1.f$	Normal number.
$e == 0$ AND $f != 0$	$(-1)^s \cdot 2^{(-126)} \cdot 0.f$	Denormal number.
$e == 0$ AND $f == 0$	$(-1)^s \cdot 0$	Signed zero.
$s == 0$ AND $e == 255$ AND $f == 0$	$+\infty$	+ infinity.
$s == 1$ AND $e == 255$ AND $f == 0$	$-\infty$	- infinity.
$e == 255$ AND $f != 0$ AND $f[22] == 0$		Signalling NaN <sup>1)</sup> .
$e == 255$ AND $f != 0$ AND $f[22] == 1$		Quiet NaN <sup>1)</sup> .

<sup>1)</sup> IEEE-754 does not define how to distinguish between signalling NaNs and quiet NaNs, but bit[22] has become the standard way of doing this.

*Note: Both signed values of zero are always treated identically and never produce different results except different signed zeros.*

### 11.2.2 Denormal Numbers

Denormal numbers are not supported for arithmetic operations. With the exception of the CMP.F instruction, all instructions replace denormal operands with the appropriately signed zero before computation. Following computation, if a denormal number would otherwise be the result, it is replaced with the appropriately signed zero.

Conceptually, the conventional order for making IEEE-754 computations is:

1. Compute result to infinite precision.
2. Round to IEEE-754 format.

This is replaced with:

1. Substitute signed zero for all denormal operands.
2. Compute result to infinite precision.
3. Round to IEEE-754 format.
4. Substitute signed zero for all denormal results.

This procedure has a subtle effect on underflow; see [Round to Nearest: Denormals and Zero Substitution, page 11-7](#).

Denormal numbers are supported only by the CMP.F instruction which makes comparisons of denormal numbers in addition to identifying denormal operands.

### 11.2.3 NaNs (Not a Number)

NaNs (Not a Number) are bit combinations within the IEEE-754 standard that do not correspond to numbers. There are two types of NaNs: signalling and quiet. The FPU defines signalling NaNs to have bit 22 = '0', and quiet NaNs to have bit 22 = '1'.

When invalid operations are performed (including operations with a signalling NaN operand), FI is asserted and a quiet NaN is produced as the floating-point result. The quiet NaN contains information about the origin of the invalid operation; see [Invalid Operations and their Quiet NaN Results, page 11-10](#).

IEEE-754 suggests that quiet NaNs should be propagated so that the result of an instruction receiving a quiet NaN as an operand (with no signalling NaN operands) should be that quiet NaN. The FPU does not propagate quiet NaNs in this way. The result of an operation that has one (or more) quiet NaN operands and no signalling NaN operands is always the quiet NaN 7FC00000<sub>H</sub>.

### 11.2.4 Underflow

Underflow occurs when the result of a floating-point operation is too small to store in floating-point representation.

IEEE-754 requires two conditions to occur before flagging underflow:

- The result must be 'tiny'.
  - A result is 'tiny' if it is non-zero and its magnitude is  $< 2^{-126}$  (for single precision). IEEE-754 allows this to be detected either before or after rounding.
- There must be a loss of accuracy in the stored result.

Loss of accuracy can be detected in two ways: either as a denormalization loss, or an inexact result.

Denormalization loss occurs when the result is calculated assuming an unbounded exponent, but is rounded to a normalized number using 23 fractional bits. If this rounded result must be denormalized to fit into IEEE-754 format and the resultant denormalized number differs from the normalized result with unbounded exponent range, then a denormalization loss occurs.

An inexact result is one where the infinitely precise result differs from the value stored.

The FPU determines tininess before rounding and inexact results to determine loss of accuracy.

In the case of the FPU, even if a denormal result would produce no loss of accuracy, because it is replaced with a zero, accuracy is lost and underflow must be flagged.

Any tiny number that is detected must therefore result in a loss of accuracy since it will either be a denormal that is replaced with zero or rounded up. Therefore underflow detection can be simplified to tiny number detection alone; i.e. any non-zero unrounded number whose magnitude is  $< 2^{-126}$ .

### 11.2.5 Fused MACs

Fused multiply-and-accumulate operations (MACs) are not supported by the IEEE-754 standard. Using FPU MAC operations (MADD.F and MSUB.F) can give different results from using separate multiply (MUL.F) and accumulate (ADD.F or SUB.F) operations because the result is only rounded once at the end of a MAC.

### 11.2.6 Traps (TriCore 1.3.1)

For TriCore 1.3.1, IEEE-754 allows optional provision for synchronous traps to occur when exception conditions occur. Under these circumstances the results returned by arithmetic operations may differ from IEEE-754 requirements to allow intermediate results to be passed to the trap handling routines. These traps are provided to assist in debugging routines and operations.

FPU traps are asynchronous and therefore are not IEEE-754 compliant traps. Since IEEE-754 traps are optional this does not cause any IEEE-754 non compliance.

### 11.2.7 Software Routines

Operations required for IEEE-754 compliance, but not implemented in the FPU instruction set, are detailed in [Table 15](#).

**Table 15 IEEE-754 Operations Requiring Software Implementation**

IEEE-754 Operation	Suggested Implementation
Square root	Newton-Raphson using QSEED.F instruction.
Remainder	<p>FPU instructions cannot be used to implement the remainder function because of the errors that can occur from multiple rounding. For reference, the IEEE method for calculating remainder is given below. Note that rounding must only occur on the conversion to integer, and for the final result.</p> $\text{rem} = x - (d * (\text{FTOI}(x/d)^{1}))$ <p>rem: remainder  x: dividend  d: divisor</p>
Round to integer in Floating-point format	ITOF(FTOI(x)).
Convert between binary and decimal	-

<sup>1)</sup> Round to nearest.

## 11.3 Rounding

All four rounding modes specified in IEEE-754 are supported. The rounding mode is selected using the RM field of the PSW (PSW[25:24]).

**Table 16 Rounding Mode Definition (PSW.RM)**

Rounding Mode Value	Mode
00 <sup>1)</sup>	Round to nearest.
01	Round toward $+\infty$ .
10	Round toward $-\infty$ .
11	Round toward zero.

<sup>1)</sup> Round to nearest is the default rounding mode.

IEEE-754 defines the rounding modes in terms of representable results, in relation to the 'infinitely precise' result. The infinitely precise result is the mathematically exact result that would be computed by the operation, if the number of mantissa and exponent bits were unlimited.

- **Round to nearest** is defined as returning the representable value that is nearest to the infinitely precise result. This is the default rounding mode that should be selected when RTOS software initializes a task. See [Round to Nearest: Even, page 11-7](#), for further information.
- **Round toward  $+\infty$**  is defined as returning the representable value that is closest to and no less than the infinitely precise result.
- **Round toward  $-\infty$**  is defined as returning the representable value that is closest to and no greater than the infinitely precise result.
- **Round toward zero** is defined as returning the representable value that is closest to and no greater in magnitude than the infinitely precise result. It is equivalent to truncation.

The rounding mode can be changed by the UPDFL (Update Flags) instruction.

Rounding is performed at the end of each relevant FPU instruction, followed by the replacement of all denormal numbers with the appropriately signed 0.

IEEE-754 does not specify the MAC instructions (MADD.F and MSUB.F) that combine multiplication and addition in a single operation. The result from the multiply part of a MAC instruction is not rounded before it is used in the addition in the FPU. Instead the whole MAC is calculated with infinite precision and rounded at the end of the add. It is therefore possible that the result from a MADD.F instruction will differ from the result that would be obtained using the same operands in a MUL.F followed by an ADD.F.



## **Rounding Mode Restored (TriCore 1.3.1)**

In TriCore 1.3.1 the rounding mode is restored on a RET (Return from Call) instruction by default. This default behaviour may be inhibited by clearing the relevant bit in the COMPAT register. The rounding mode is also restored on an RFE (Return From Exception) instruction or on an RFM (Return From Monitor) instruction.

### **11.3.1 Round to Nearest: Even**

'Round to nearest' is defined as returning the representable value that is nearest to the infinitely precise result. If two representable values are equally close (i.e. the infinitely precise result is exactly half way between two representable values), then the one whose LSB (Least Significant Bit) is zero is returned. This is sometimes known as rounding to nearest even.

This is usually straight forward, but if the infinitely precise result is half way between two representable numbers with different exponents, the result with the larger exponent is always selected (the LSB of its mantissa is zero).

For example, if the infinitely precise result is:

1.111 1111 1111 1111 1111 1111 1000 0000 0000<sub>B</sub> \* 2<sup>0</sup>

This is half way between:

1.0000 0000 0000 0000 0000 0000<sub>B</sub> \* 2<sup>1</sup>

and:

1.111 1111 1111 1111 1111 1111<sub>B</sub> \* 2<sup>0</sup>

The result with the larger exponent is returned.

### **11.3.2 Round to Nearest: Denormals and Zero Substitution**

Following computation, results are first rounded to IEEE-754 representable numbers and then the appropriately signed zero is substituted for any denormal results that may have occurred. This produces some results that can seem counter intuitive.

Consider an infinitely precise result that has been computed and falls between the smallest representable positive IEEE-754 normal number (1.000 ... 000 \* 2<sup>-126</sup>) and the largest representable positive IEEE-754 denormal number (0.111 ... 111 \* 2<sup>-126</sup>).

- If the infinitely precise result is nearer to the normal number, or halfway between the two, then the result must be rounded to the normal number.
- If the infinitely precise result is nearer to the denormal number, then the result is rounded to the denormal value. Zero is then substituted for the denormal result.

The FPU architecture cannot produce denormal results, however the concept of denormal numbers is important to the FPU. It would be wrong to assume that the infinitely precise result should be rounded to the nearest FPU representable number, in this case (+1.000 ... 000 \* 2<sup>-126</sup>) or (0). Such an implementation would mean that all

unrounded results between  $(+1.000 \dots 000 * 2^{-126})$  and  $(+0.100 \dots 000 * 2^{-126})$  would be rounded to the smallest representable positive IEEE-754 normal number.

### 11.3.3 Round Towards $\pm \infty$ : Denormals and Zero Substitution

Following computation results are first rounded to IEEE-754 representable numbers, then the appropriately signed zero is substituted for any denormal results that may have occurred. See [Denormal Numbers, page 11-3](#).

According to the IEEE-754 definition of the rounding modes, when rounding towards  $+\infty$  ( $-\infty$ ) the rounded result should not be less than (greater than) the infinitely precise result. However if a positive (negative) result would otherwise be rounded to a denormal number, it is then substituted for a zero. Therefore the returned result of zero is less than (greater than) the infinitely precise result. The returned result appears to contradict the definition of these rounding modes in this case.

## 11.4 Exceptions

The FPU implements all five IEEE-754 exceptions (invalid operation, overflow, divide by zero, underflow, and inexact). When one of these exceptions occur the corresponding exception flag in the PSW is asserted.

### Asynchronous Traps (TriCore 1.3.1)

In TriCore 1.3.1 an asynchronous trap may optionally be taken when an exception occurs, however IEEE-754 compliant traps are not implemented, see [Section 11.5 Asynchronous Traps \(TriCore 1.3.1\) \( page 11-13\)](#).

### IEEE-754 Exception Flags

The IEEE-754 exception flags are stored as part of the PSW register as shown in the following table. In accordance with IEEE-754, each bit is sticky so that the FPU instructions in general assert these flags when an exception occurs and do not negate them when the exception does not occur. The UPDFL instruction can be used to clear the exception flags.

**Table 17 FPU Exception Flags**

ALU Flag	FPU Flag	FPU Exception	PSW Bit Position
C	FS	Some Exception.	31
V	FI	Invalid Operation.	30
SV	FV	Overflow.	29
AV	FZ	Divide by Zero.	28

**Table 17 FPU Exception Flags**

ALU Flag	FPU Flag	FPU Exception	PSW Bit Position
SAV	FU	Underflow.	27
-	FX	Inexact.	26

Since the IEEE-754 exception flags are sticky, it can be impossible to tell if an exception occurred on the last instruction if it was asserted before the last instruction executed. An additional, non sticky, exception flag (FS) is therefore implemented to identify if the last FPU instruction caused an IEEE-754 exception or not.

Note that the PSW bits used to store the exception flags are also used to store ALU flags as shown in the table above. When an ALU instruction updates these flags, the corresponding FPU exception flag is overwritten and lost.

The following conditions are true for all FPU operations asserting exception flags, with the exception of UPDFL.

- Any FPU operation can assert only one of the FI, FV, FZ or FU exception flags.
- FX can be asserted by any operation so long as FI and FZ are negated.
- When either FV or FU are asserted, FX is also asserted.

### **FS - Some Exception**

This bit is not sticky and is asserted or negated for all instructions that can cause IEEE-754 exceptions to occur. If any of the IEEE-754 exceptions (FI, FV, FZ, FU, FX) have occurred during that instruction, FS is also asserted.

*Note: UPDFL can assert IEEE-754 exceptions without asserting FS.*

### **FI - Invalid Operation**

FI is asserted in three circumstances:

- When a signalling NaN (see **NaNs (Not a Number), page 11-3**) is an operand for a FPU instruction.
- For invalid operations such as QSEED.F ( $\approx 1/\sqrt{x}$ ) of a negative number.
- Conversions from floating-point to other formats where the rounded result is outside the range of the target.

When an instruction that produces a floating-point result asserts FI as a result of a signalling NaN or invalid operation, the result is a quiet NaN.

**Table 18 Invalid Operations and their Quiet NaN Results**

<b>Invalid Operation</b>	<b>Quiet NaN</b>
Signalling NaN operand for arithmetic instructions. <sup>1)</sup>	7FC00000 <sub>H</sub> <sup>2)</sup>
Signalling NaN operand for CMP.F instruction.	n.a. <sup>5)</sup>
ADD.F with $+\infty$ and $-\infty$ as operands.	7FC00001 <sub>H</sub>
SUB.F with $(+\infty$ and $+\infty)$ or $(-\infty$ and $-\infty)$ as operands.	7FC00001 <sub>H</sub>
MADD.F if the result of the multiplication is $\pm\infty$ and the addend is the oppositely signed $\infty$ .	7FC00001 <sub>H</sub>
MSUB.F if the result of the multiplication is $\pm\infty$ and the minuend is the same signed $\infty$ .	7FC00001 <sub>H</sub>
MUL.F with 0 and $\pm\infty$ as multiplicands.	7FC00002 <sub>H</sub>
MADD.F with 0 and $\pm\infty$ as multiplicands.	7FC00002 <sub>H</sub>
MSUB.F with 0 and $\pm\infty$ as multiplicands.	7FC00002 <sub>H</sub>
QSEED.F with a negative operand <sup>3)</sup> .	7FC00004 <sub>H</sub>
DIV.F with 0 as both operands <sup>4)</sup> .	7FC00008 <sub>H</sub>
DIV.F with both operands being an $\infty$ of either sign.	7FC00008 <sub>H</sub>
FTOI, FTOU or FTOQ31 with rounded result outside the range of the target format.	n.a. <sup>5)</sup>
FTOIZ, FTOUZ or FTOQ31Z with rounded result outside the range of the target format. (TriCore 1.3.1).	n.a. <sup>5)</sup>
FTOI, FTOU or FTOQ31 with the input operand a quiet NaN, a signalling NaN or $\pm\infty$ .	n.a. <sup>5)</sup>
FTOIZ, FTOUZ or FTOQ31Z with the input operand a quiet NaN, a signalling NaN or $\pm\infty$ . (TriCore 1.3.1).	n.a. <sup>5)</sup>

<sup>1)</sup> Also see the FPU operation syntax description in the Instruction Set.

<sup>2)</sup> The quiet NaN (7FC00000<sub>H</sub>) is produced as the result of arithmetic operations that have any NaN as an operand. FI is only asserted when one of these NaNs is signalling. See **NaNs (Not a Number)**, page 11-3.

<sup>3)</sup> -0 is not negative, therefore QSEED.F of -0 is  $-\infty$ .

<sup>4)</sup> 0/0 is defined as being an invalid operation (FI) rather than a divide by zero (FZ).

<sup>5)</sup> The result is not in floating-point format and therefore cannot be a quiet NaN. Refer to the instruction description for what the result should be.

**FV - Overflow**

For operations that return a floating-point result, the FV flag is set as stated in IEEE-754; 'whenever the destination format's largest finite number is exceeded in magnitude by what would have been the rounded floating-point result, were the exponent range unbounded'.

The result returned is determined by the rounding mode and the sign of the unrounded result:

- Round to nearest carries all overflows to infinity, with the sign of the unrounded result.
- Round toward zero carries all overflows to the format's largest finite number with the sign of the unrounded result.
- Round toward minus infinity carries positive overflows to the format's largest finite number, and carries negative overflows to minus infinity.
- Round toward plus infinity carries negative overflows to the format's most negative finite number, and carries positive overflows to plus infinity.

When overflow is flagged (FV asserted), the returned result can not be exactly equal to the unrounded result. Therefore whenever FV is asserted FX is also asserted.

### **FZ - Divide by Zero**

The FZ flag is set by DIV.F if the divisor operand is zero and the dividend operand is a finite non zero number. The result is an infinity with sign determined by the usual rules.

Note that:

- 0/0 is defined as an invalid operation, so FI is asserted rather than FZ.
- All arithmetic with  $\pm \infty$  as an operand is defined as being exact, except for invalid operations where FI is asserted. Therefore for  $\pm \infty / \pm 0$  FZ is not asserted, the appropriately signed  $\infty$  is returned as the result with no other exceptions occurring.

### **FU - Underflow**

As discussed in [Underflow, page 11-4](#), underflow is detected and so FU is asserted, when the unrounded result is smaller in magnitude than the smallest representable normal number ( $2^{-126}$ ).

The Q31TOF instruction can cause an underflow as well as the arithmetic instructions ADD.F, SUB.F, MUL.F, MADD.F, MSUB.F, and DIV.F.

The return result for instructions flagging an underflow are complicated by the way that FPU treats denormal numbers. This is described in detail in [Round to Nearest: Denormals and Zero Substitution, page 11-7](#).

### **FX - Inexact**

If the rounded result of an operation is not exactly equal to the unrounded result, then the FX flag is set.

The result delivered is the rounded result, unless either overflow (FV) or underflow (FU) has also occurred during this instruction, when the overflow or denormalization return result rules are followed.

## **11.5 Asynchronous Traps (TriCore 1.3.1)**

In TriCore 1.3.1 the FPU can be configured such that a trap is signalled to the TriCore core when an FPU instruction causes an IEEE-754 exception. The trap generated is a Co-Processor Asynchronous Error (CAE), Trap Class 4 - TIN 4. FPU CAE traps should not be confused with the synchronous exception traps optional to IEEE-754 which allow software routines to correct arithmetic overflow or underflow.

The FPU CAE trap is intended for debug purposes only and has no effect on either the exceptional instruction or any other instruction which may be executing within the FPU. The result returned by an exceptional instruction causing a CAE trap is identical to that which would be returned if no trap were taken. The CAE trap is signalled after instruction completion.

The specific exception conditions which cause FPU CAE traps to be generated are under software control. To enable the trap generation for a specific exception type the appropriate enable bit in the FPU\_TRAP\_CON register must be asserted (FIE, FVE, FZE, FUE or FXE). Any number of these enable bits may be set to allow traps to be taken if any of a range of exceptions occur. FX is a regularly occurring condition, care should be taken in enabling this trap.

When an instruction causes one of the enabled exceptions, information about the exceptional instruction including the instruction PC, opcode and source operands are captured in the FPU special function registers. At the same time the Trap Status flag (TST) is set within the FPU\_TRAP\_CON register, denoting that the contents of the FPU trap capture registers are valid. In addition, so long as FPU\_TRAP\_CON.TST remains set, further FPU CAE trap generation is inhibited. This avoids multiple traps being generated from the same root problem and the original information being lost. Once the trap handler has interrogated the FPU to determine the cause of the trap, the FPU\_TRAP\_CON.TST bit may be cleared to enable further traps.

## 11.6 FPU CSFR Registers (TriCore 1.3.1)

The FPU CSFR registers are used to store the details of instructions causing traps.

The result of the exceptional instruction causing a trap is not stored in an FPU register. The result will be available in the instruction's destination register as long as it has not been overwritten before the asynchronous trap is taken.

### 11.6.1 FPU Trap Control Register

*Note: TriCore 1.3.1 architecture only.*

#### FPU\_TRAP\_CON

##### Trap Control Register

(A000<sub>H</sub>)

Reset value: 0000 0000<sub>H</sub>

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
-	FI	FV	FZ	FU	FX		-		FIE	FVE	FZE	FUE	FXE		-
	rh	rh	rh	rh	rh				rw	rw	rw	rw	rw		
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
			-				RM				-			TCL	TST
							rh	rh						w	rh

Field	Bits	Type	Description
-	31	-	<b>Reserved Field</b>
FI	30	rh	<b>Captured FI</b> Asserted if the captured instruction asserted FI. Only valid when TST is asserted.
FV	29	rh	<b>Captured FV</b> Asserted if the captured instruction asserted FV. Only valid when TST is asserted.
FZ	28	rh	<b>Captured FZ</b> Asserted if the captured instruction asserted FZ. Only valid when TST is asserted.
FU	27	rh	<b>Captured FU</b> Asserted if the captured instruction asserted FU. Only valid when TST is asserted.



**Floating Point Unit (FPU)**

Field	Bits	Type	Description
FX	26	rh	<b>Captured FX</b> Asserted if the captured instruction asserted FX. Only valid when TST is asserted.
-	[25:23]	-	<b>Reserved Field</b>
FIE	22	rw	<b>FI Trap Enable</b> When set, an instruction generating an FI exception will trigger a trap.
FVE	21	rw	<b>FV Trap Enable</b> When set, an instruction generating an FV exception will trigger a trap.
FZE	20	rw	<b>FZ Trap Enable</b> When set, an instruction generating an FZ exception will trigger a trap.
FUE	19	rw	<b>FU Trap Enable</b> When set, an instruction generating an FU exception will trigger a trap.
FXE	18	rw	<b>FX Trap Enable</b> When set, an instruction generating an FX exception will trigger a trap.
-	[17:10]	-	<b>Reserved Field</b>
RM	[9:8]	rh	<b>Captured Rounding Mode</b> The rounding mode of the captured instruction. Only valid when TST is asserted. Note that this is the rounding mode supplied to the FPU for the exceptional instruction. UPDFL instructions may cause a trap and change the rounding mode. In this case the RM bits capture the input rounding mode.
-	[7:2]	-	<b>Reserved Field</b>
TCL	1	w	<b>Trap Clear</b> 1 : Clears the trapped instruction (TST will be negated). 0 : Does nothing. Read: always reads as 0.

Field	Bits	Type	Description
TST	0	rh	<b>Trap Status</b> 0 : No instruction captured: The next enabled exception will cause the exceptional instruction to be captured. 1 : Instruction captured: No further enabled exceptions will be captured until TST is cleared.


## 11.6.2 FPU Trapping Instruction Program Counter Register

*Note: TriCore 1.3.1 architecture only.*

### FPU\_TRAP\_PC

#### Trapping Instruction Program Counter (A004<sub>H</sub>)

Reset value: Implementation Specific



Field	Bits	Type	Description
PC	[31:0]	rh	<b>Captured Program Counter</b> The program counter (virtual address) of the captured instruction. Only valid when FPU_TRAP_CON.TST is asserted.

### 11.6.3 FPU Trapping Instruction Opcode Register


*Note: TriCore 1.3.1 architecture only.*

#### FPU\_TRAP\_OPC

**Trapping Instruction Opcode**

**(A008<sub>H</sub>)**

**Reset value: Implementation Specific**



Field	Bits	Type	Description
-	[31:20]	-	<b>Reserved Field</b>
DREG	[19:16]	rh	<b>Captured Destination Register</b> The destination register of the captured instruction. 0 <sub>H</sub> : Data general purpose register 0. ... F <sub>H</sub> : Data general purpose register 15. Only valid when FPU_TRAP_CON.TST is asserted.
-	[15:9]	-	<b>Reserved Field</b>
FMT	8	rh	<b>Captured Instruction Format</b> The format of the captured instruction's opcode. 0 : RRR. 1 : RR. Only valid when FPU_TRAP_CON.TST is asserted.
OPC	[7:0]	rh	<b>Captured Opcode</b> The secondary opcode of the captured instruction. When FPU_TRAP_OPC.FMT=0 only bits [3:0] are defined. OPC is valid only when FPU_TRAP_CON.TST is asserted.

## 11.6.4 FPU Trapping Instruction Operand SRC1 Register


Note: TriCore 1.3.1 architecture only.

### FPU\_TRAP\_SRC1

Trapping Instruction Operand

(A010<sub>H</sub>)

Reset value: Implementation Specific



Field	Bits	Type	Description
SRC1	[31:0]	rh	<b>Captured SRC1 Operand</b> The SRC1 operand of the captured instruction. Only valid when FPU_TRAP_CON.TST is asserted.

## 11.6.5 FPU Trapping Instruction Operand SRC2 Register


*Note: TriCore 1.3.1 architecture only.*

### FPU\_TRAP\_SRC2

Trapping Instruction Operand

(A014<sub>H</sub>)

Reset value: Implementation Specific



Field	Bits	Type	Description
SRC2	[31:0]	rh	<b>Captured SRC2 Operand</b> The SRC2 operand of the captured instruction. Only valid when FPU_TRAP_CON.TST is asserted.

## 11.6.6 FPU Trapping Instruction Operand SRC3 Register


Note: TriCore 1.3.1 architecture only.

### FPU\_TRAP\_SRC3

Trapping Instruction Operand

(A018<sub>H</sub>)

Reset value: Implementation Specific



Field	Bits	Type	Description
SRC3	[31:0]	rh	<b>Captured SRC3 Operand</b> The SRC3 operand of the captured instruction. Only valid when FPU_TRAP_CON.TST is asserted.

### 11.6.7 FPU Identification Register

*Note: TriCore 1.3.1 architecture only.*


The FPU Identification Register identifies the FPU type and revision.

#### FPU\_ID

**FPU Module Identification**

**(A020<sub>H</sub>)**

**Reset Value: Implementation Specific**



Field	Bits	Type	Description
MOD	[31:16]	r	<b>Module Identification Number</b> Used for module identification.
MOD_32B	[15:8]	r	<b>32-Bit Module Enable</b> A value of C0 <sub>H</sub> in this field indicates a 32-bit module with a 32-bit module ID register.
MOD_REV	[7:0]	r	<b>Module Revision Number</b> Used for revision numbering. The value of the revision starts at 01 <sub>H</sub> (first revision) up to FF <sub>H</sub> .



## **12 Core Debug Controller (CDC)**

The TriCore® debug functionality is an interface of architecture, implementation and software tools, so users are advised that mechanisms may differ in subsequent architecture generations.

The Core Debug Controller (CDC) is designed to support real-time systems that require non-intrusive debugging. Most of the architectural state in the CPU Core and Core on-chip memories can be accessed through the system Address Map.

Access to the CDC is typically provided via the On-Chip Debug Support (OCDS) of the system containing the CPU.

### **CDC Features**

CDC features are aimed predominantly at the software development environment. It offers real-time run control and internal visibility of resources such as data and memories. Features include:

- Real-time run control (Halt and Restart the CPU).
- Access and update internal registers and core local memory.
- Setting breakpoints and watchpoints with complex trigger conditions.

### **Enabling the CDC**

To enable the CDC, the system containing the core must set the Debug Enable bit (DE) in the Debug Status Register (DBGSR). The CDC is disabled when DBGSR.DE == 0, and enabled when DBGSR.DE == 1. How the DBGSR.DE bit is controlled and how the CDC is enabled or disabled, is system dependent.

## **12.1 Run Control Features**

Real-time run control functions are accessed and controlled by address mapped reads and writes, typically by the OCDS or by any other bus master that has the appropriate authorization. The CDC provides hardware hooks into the core allowing the detection of Debug Events which result in Debug Actions.

Four signals are provided by the CDC for communication with the OCDS:

- Core Break-In.
  - An indication from the OCDS to the Core of a condition of interest.
- Core Break-Out.
  - An indication from the Core to the OCDS of a condition of interest.
- Core Suspend-In.
  - An indication from the OCDS to the Core to enter Halt mode.
- Core Suspend-Out.

- An indication from the Core to the OCDS of the state of the Debug Status register (DBGSR) SUSP field (DBGSR.SUSP). This signal can be controlled by writes to the Debug Status register, whereas the Core Break-Out signal can not.

## Features

- Single-Step support in hardware.
- Debug Events that can cause a Debug Action:
  - Assertion of the external Core Break-In signal to the core.
  - Execution of the DEBUG instruction.
  - Execution of the MTCR (Move To Control Register) or the MFCR (Move From Control Register) instruction.
  - Events raised by the Trigger Event Unit (see [Trigger Event Unit, page 12-4](#)).
- Debug Actions can be one or more of the following:
  - Update Debug Status register.
  - Indicate event on Core Break-Out signal and/or Core Suspend-Out signal.
  - Halt CPU execution.
  - Take Breakpoint Trap.
  - Raise Breakpoint Interrupt.
  - Control performance counters.
- Real-time features:
  - Read and write of core memory and register while the core is running, with minimum intrusion (may steal cycles).
  - The service of high priority interrupt routines by use of the Breakpoint Interrupt Debug Action.

*Note: The reading and writing of other system memory while the CPU is running can be intrusive, depending on the number of cycles that are required to perform the operation. When this happens, cycle stealing occurs.*

The programming of Debug Events and Debug Actions can occur while the CPU is running with little or no intrusion. The detection of Debug Events has no effect on real-time execution.

## **12.2 Debug Events**

When the CDC is enabled, a Debug Event can be generated by:

- Core Break-In signal.
  - See [External Debug Event, page 12-3](#).
- Execution of a DEBUG instruction.
  - See [Debug Instruction, page 12-3](#).
- Execution of the MTCR or MFCR instruction.
  - See [MTCR and MFCR Instructions, page 12-3](#).
- A hardware Event generation unit.
  - See [Trigger Event Unit, page 12-4](#).

### **12.2.1 External Debug Event**

An External Debug Event is not correlated in any way to the instruction flow, but it provides the ability to stop and gain control of the CPU without having to reset. It may take several clocks for the Debug Event to be recognized by the CPU if it is currently executing a multi-cycle, non-cancellable instruction (such as a context save and restore for example).

The Debug Action taken on the assertion of the Core Break-In signal is specified in the EXEVT (External Event) register (see [EXEVT, page 12-17](#)).

### **12.2.2 Debug Instruction**

TriCore supports a User mode DEBUG instruction which can generate a Debug Event when the CDC is enabled. When the CDC is disabled it is treated as a NOP (No Operation). Both 16-bit and 32-bit forms of the DEBUG instruction are provided. This feature facilitates software debug, which allows a jump to a monitor program and provides a relatively inexpensive software instrumentation and interrogation mechanism. The Debug Action taken on the Debug Event is specified in the SWEVT (Software Debug Event) register (See [SWEVT, page 12-19](#)).

### **12.2.3 MTCR and MFCR Instructions**

A Debug Event is raised when a MTCR (Move To Control Register) or MFCR (Move From Control Register) instruction is used to read or modify a user Core Special Function Register (CSFR). This gives the debug software the ability to monitor, detect and modify changes to CSFRs. A Debug Event is not raised when code reads or modifies one of the dedicated Debug SFRs (Special Function Registers):

- Debug Status Register ([DBGSR, page 12-15](#)).
- Core Register Access Event Register ([CREVT, page 12-18](#)).
- Software Debug Event Register ([SWEVT, page 12-19](#)).
- External Event Register ([EXEVT, page 12-17](#)).

- Trigger Event Register (TRnEVT) ([TR0EVT, page 12-33](#) and [TR1EVT, page 12-33](#)).
- Debug Monitor Start Register ([DMS, page 12-37](#)).
- Debug Context Pointer Register ([DCX, page 12-24](#)).

#### Additional TriCore 1.3.1 Registers

- Counter Control Register - [Counter Control Register, page 12-45](#).
- CPU Clock Count Register - [CPU Clock Cycle Count Register, page 12-47](#).
- Instruction Count Register - [Instruction Count Register, page 12-48](#).
- Multi-Count Register 1 - [Multi-Count Register 1, page 12-49](#).
- Multi-Count Register 2 - [Multi-Count Register 2, page 12-50](#).
- Multi-Count Register 3 - [Multi-Count Register 3, page 12-51](#).

The Debug Action taken when the Debug Event is raised is specified in the CREVT register (See [CREVT, page 12-18](#)).

#### 12.2.4 Trigger Event Unit

The Trigger Event Unit is responsible for generating Debug Events when a programmable set of Debug Triggers are active. Debug Triggers come from the protection system and are either:

- Code Addresses.
- Data Accesses.

*Note: Compared addresses are virtual addresses.*

These Debug Triggers provide the inputs to a programmable block of logic which produces Debug Events as its output (see [Debug Triggers, page 12-5](#) for more details on programmable combinations).

The Debug Action taken when the Debug Event is raised, is specified in the Trigger Event register (TRnEVT). See [Trigger Event Registers, page 12-33](#) for the register definition.

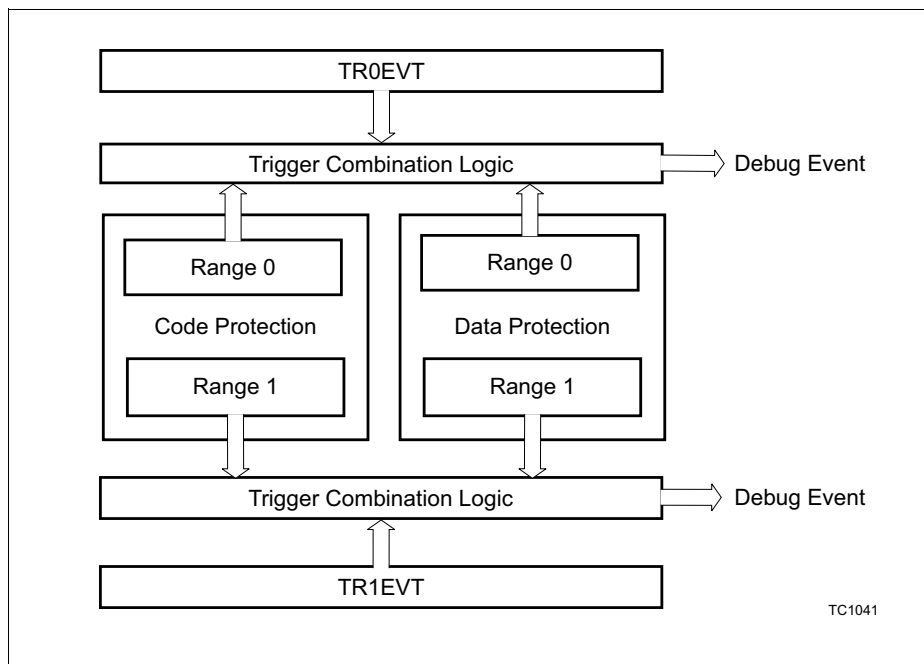
## 12.3 Debug Triggers

The CDC can generate the following types of Debug Triggers:

- Execution of an instruction at a specific address.
- Execution of an instruction within a range of addresses.
- Loading a value from a specific address.
- Loading a value from within a range of addresses.
- Storing a value to a specific address.
- Storing a value to within a range of addresses.

The Debug Trigger Event Unit takes inputs from the Protection mechanism Range Table Entries (RTEs) and combines them as specified by the TR $n$ EVT registers to produce a Debug Event. Range Table Entries that do not have a corresponding TR $n$ EVT register can not be used to generate a Debug Event.

The RTEs which provide Debug Trigger information are those selected by the PRS (Protection Register Set) field in the PSW (Program Status Word) register (PSW.PRS). If it is not known which PRS is active, the RTEs in all potentially used PRSs should be programmed in the same way.



**Figure 45 An Implementation Combination Example, Using two TR $n$ EVT Registers**

### 12.3.1 Combining Debug Triggers

The CDC allows Code and Data triggers to be combined to create a Debug Event. The combination is specified by the Trigger Event Register (TRnEVT).

The Trigger Event Unit can generate a number of Trigger Debug Events by combining four Debug Triggers for each Trigger Debug Event. The Debug Triggers are generated by the memory protection system. The four Triggers used to generate the Triggern Debug Event come from the Coden and Datan Range Table Entries (RTE).

**Table 19 Debug Triggers Generated by the Memory Protection System**

Trigger	Description
D <sub>U</sub>	Data read or write access to the upper bound of the Data RTE <sub>n</sub> , as enabled in the Data Protection Mode (DPM) register.
D <sub>LR</sub>	Data read or write access to the lower bound or range of the Data RTE <sub>n</sub> , as enabled in the Data Protection Mode (DPM) register.
C <sub>U</sub>	Code execution from the upper bound address of the Code RTE <sub>n</sub> , as enabled in the Code Protection Mode (CPM) register.
C <sub>LR</sub>	Code execution from the lower bound address or address range of the Code RTE <sub>n</sub> , as enabled in the Code Protection Mode (CPM) register.

The combinations of Debug Triggers that generate a Debug Event are controlled by bits in the TRnEVT register. The possible combinations are given in [Table 20](#).

**Table 20 Debug Trigger Combinations that Generate a Debug Event**

TRnEVT [11:8]	Data (D) and Code (C), Upper(U) and Lower(L) Bound, Combinations
0000	D <sub>U</sub> or D <sub>LR</sub> or C <sub>U</sub> or C <sub>LR</sub>
0001	(D <sub>LR</sub> and C <sub>LR</sub> ) or D <sub>U</sub> or C <sub>U</sub>
0010	(D <sub>LR</sub> and C <sub>U</sub> ) or D <sub>U</sub> or C <sub>LR</sub>
0011	(D <sub>LR</sub> and C <sub>LR</sub> ) or (D <sub>LR</sub> and C <sub>U</sub> ) or D <sub>U</sub>
0100	(D <sub>U</sub> and C <sub>LR</sub> ) or D <sub>LR</sub> or C <sub>U</sub>
0101	(D <sub>LR</sub> and C <sub>LR</sub> ) or (D <sub>U</sub> and C <sub>LR</sub> ) or C <sub>U</sub>
0110	(D <sub>LR</sub> and C <sub>U</sub> ) or (D <sub>U</sub> and C <sub>LR</sub> )
0111	(D <sub>LR</sub> and C <sub>LR</sub> ) or (D <sub>LR</sub> and C <sub>U</sub> ) or (D <sub>U</sub> and C <sub>LR</sub> )
1000	(D <sub>U</sub> and C <sub>U</sub> ) or D <sub>LR</sub> or C <sub>LR</sub>
1001	(D <sub>LR</sub> and C <sub>LR</sub> ) or (D <sub>U</sub> and C <sub>U</sub> )
1010	(D <sub>LR</sub> and C <sub>U</sub> ) or (D <sub>U</sub> and C <sub>U</sub> ) or C <sub>LR</sub>

**Table 20**      **Debug Trigger Combinations that Generate a Debug Event**

TRnEVT [11:8]	Data (D) and Code (C), Upper(U) and Lower(L) Bound, Combinations
1011	(D <sub>LR</sub> and C <sub>LR</sub> ) or (D <sub>LR</sub> and C <sub>U</sub> ) or (D <sub>U</sub> and C <sub>U</sub> )
1100	(D <sub>U</sub> and C <sub>LR</sub> ) or (D <sub>U</sub> and C <sub>U</sub> ) or D <sub>LR</sub>
1101	(D <sub>LR</sub> and C <sub>LR</sub> ) or (D <sub>U</sub> and C <sub>LR</sub> ) or (D <sub>U</sub> and C <sub>U</sub> )
1110	(D <sub>LR</sub> and C <sub>U</sub> ) or (D <sub>U</sub> and C <sub>LR</sub> ) or (D <sub>U</sub> and C <sub>U</sub> )
1111	(D <sub>LR</sub> and C <sub>LR</sub> ) or (D <sub>LR</sub> and C <sub>U</sub> ) or (D <sub>U</sub> and C <sub>LR</sub> ) or (D <sub>U</sub> and C <sub>U</sub> )

*Note: DBGSR.EVTSRC, DBGSR.PREVSUSP and DBGSR.SUSP are updated for all Debug Actions except 000 (None; disabled) and 101/110/111 reserved (same behaviour as 000).*

## 12.4      **Debug Actions**

When a Debug Event occurs, one or more of the following Debug Actions are taken depending upon the programming of the relevant Event Register:

- [Update Debug Status Register \(DBGSR\), page 12-7.](#)
- [Indicate on Core Break-Out Signal, page 12-8.](#)
- [Indicate on Core Suspend-Out Signal, page 12-8.](#)
- [Halt, page 12-8.](#)
- [Breakpoint Trap, page 12-8.](#)
- [Breakpoint Interrupt, page 12-10.](#)
- [Suspend Out, page 12-11.](#)
- [Performance Counter Start/Stop \(TriCore 1.3.1\), page 12-11.](#)
- [None \(TriCore 1.3.1\), page 12-11.](#)
- [Disabled, page 12-12.](#)
- [Suspend In Halt \(TriCore 1.3.1\), page 12-12.](#)

### 12.4.1      **Update Debug Status Register (DBGSR)**

When a Debug Event occurs the EVTSRC (Event Source), PEVT (Posted Event), PREVSUSP (Previous State of Suspend Signal) and SUSP (Current State of Suspend Signal) fields of the Debug Status Register (DBGSR) are always updated.

The PREVSUSP field is updated from the contents of the SUSP field.

SUSP is updated from the EVTA field of the register that prompted the Debug Event (EXEVT, CREVT, SWEVT or TRnEVT).

### **12.4.2 Indicate on Core Break-Out Signal**

A Debug Event can indicate to the OCDS that the Event has occurred. Note that it is implementation dependent whether or not this signal is connected to an external pin.

### **12.4.3 Indicate on Core Suspend-Out Signal**

On a Core Suspend-Out action, the value of the SUSP field in the Debug Status Register (DBGSR) is copied to the PREVSUSP field (DBGSR.PREVSUSP).

The DBGSR.SUSP field is updated with the contents of the SUSP field from the register that prompted the Debug Event (EXEVT, CREVT, SWEVT or TRnEVT).

Modification of the DBGSR.SUSP bit will be reflected in the Core Suspend-Out Signal. When writing to the DBGSR.SUSP bit, PREVSUSP is not updated.

When a debug event causes a breakpoint interrupt to be posted, DBGSR.SUSP, DBGSR.PREVSUSP and the Core Suspend-Out signal remain unchanged.

### **12.4.4 Halt**

The Debug Action Halt, causes the Halt mode to be entered. Halt mode performs a cancel of:

- All instructions after and including the instruction that caused the breakpoint if Break Before Make (BBM) is set.
- All instructions after the instruction that caused the breakpoint if BBM is clear.

Once these instructions have been cancelled the CPU enters Halt mode, where no more instructions are fetched or executed. Halt mode is entered when the DBGSR.HALT bit field is set to 01<sub>B</sub>. On entering Halt mode the DBGSR.EVTSRC bit field is updated.

Once in Halt mode the external Debug system is used to interrogate the target through the mapping of the architectural state into the FPI address space.

While halted, the CPU does not respond to any interrupts and only resumes execution once the Debug Status register HALT bit is clear (DBGSR.HALT). The bit is cleared by writing 10<sub>B</sub> to the HALT field.

### **12.4.5 Breakpoint Trap**

The Breakpoint Trap enters a Debug Monitor without using any user resource. It relies upon the following emulator resources:

- A Debug Monitor which is executed commencing at the address defined in the DMS (Debug Monitor Start Address) register.
- A 4-word area of RAM is available at the address defined in the DCX (Debug Context Save Area Pointer) register. This is used to store the critical state during the Debug Monitor entry sequence.



When a Breakpoint Trap is taken, the following actions are performed:

- Write PSW to DCX + 4<sub>H</sub>
- Write PCXI to DCX + 0<sub>H</sub>
- Write A[10] to DCX + 8<sub>H</sub>
- Write A[11] to DCX + C<sub>H</sub>
- A[11] = PC
- PCXI.PIE = ICR.IE
- PCXI.PCPN = ICR.CCPN
- PC = DMS
- PSW.PRS = 0<sub>H</sub>
- PSW.IO = 2<sub>H</sub>
- PSW.GW = 0<sub>H</sub>
- PSW.IS = 1<sub>H</sub>
- PSW.CDE = 0<sub>H</sub>
- PSW.CDC = 0000000<sub>B</sub>
- ICR.IE = 0<sub>H</sub>
- DBGTCR.DTA = 1<sub>H</sub> (TriCore 1.3.1)

The corresponding return sequence is provided through the privileged instruction RFM (Return From Monitor).

This provides an automated route into the Debug Monitor which does not take any User resource. The RFM (Return From Monitor) instruction is then used to return control to the original task. The RFM instruction is a NOP (No Operation) when the CDC is disabled (i.e. DBGSR.DE == 0).

### **Multiple Breakpoint Traps (TriCore 1.3.1)**

On taking a breakpoint trap TriCore saves a debug context (PCX, PSW, A10, A11) at the location indicated by the DCX register. At the end of the debug trap handler an RFM instruction is used to restore this state.

The DCX location is only able to store a single debug context. Problems therefore arise if multiple breakpoint traps are triggered. Only the state saved by the final breakpoint trap is retained, all state from the previous breakpoint traps is lost.

To prevent this situation occurring the breakpoint trap entry sequence sets the Debug Trap Active (DTA) bit in the Debug Trap Control Register (DBGTCR). This bit is used to inhibit further breakpoint traps.

The DTA bit is cleared on an RFM instruction and set on a breakpoint trap (It may also be set and cleared by MTCR).

A breakpoint trap may only be taken in the condition DTA==0. Taking a breakpoint trap sets the DTA bit to one. Further breakpoint traps are therefore disabled until such time as the breakpoint trap handler clears the DTA bit or until the breakpoint trap handler terminates with a RFM or on a debug reset.

### 12.4.6 Breakpoint Interrupt

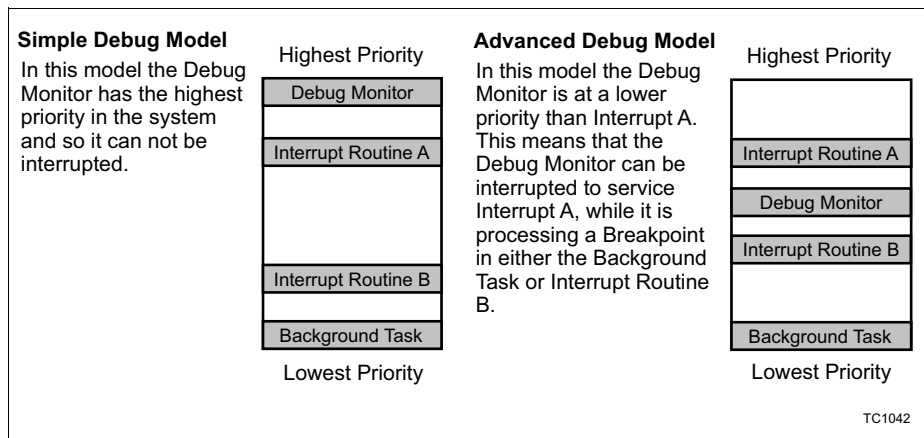
One of the possible Debug Actions to be taken on a Debug Event, is to raise a Breakpoint Interrupt. The interrupt priority is programmable and is defined in the control register associated with the breakpoint interrupt.

The architecture allows a Debug Event to raise one of four Breakpoint Interrupts, each of which can have its own interrupt priority. The number of Breakpoint Interrupts is implementation dependant.

The Breakpoint Interrupt allows a flexible Debug environment to be defined which is capable of satisfying many of the requirements for efficient debugging of a real-time system. For example, the execution of safety critical code can be preserved while the debugger is active.

Breakpoint Interrupts can be used to provide the conventional Debug Model available in traditional microcontrollers, where a Breakpoint stops the processor, by simply assigning the highest interrupt priority level to the Debug Monitor or by ensuring interrupts are disabled in the Debug Monitor. It also provides the flexibility for critical interrupts to be programmed with a higher priority than the Debug Monitor. The advantages of this are that:

- The Debug Monitor can be interrupted in an identical manner to any other interrupt by a higher level interrupt. This allows the CPU to service critical interrupts while the Debug Monitor is running.
- Any Debug Events posted in a critical routine are postponed until the CPU priority drops below that of the Debug Monitor.



**Figure 46 Debug Monitor - Simple and Advanced Models**

## **Posted Breakpoint Interrupts**

The situation needs to be considered where a Breakpoint Interrupt targeted at the CPU is at an interrupt priority level below the current CPU priority. In the Advanced Model in [Figure 46](#) for example, if a Breakpoint Interrupt is set in Interrupt Routine 'A' it is a problem, because the Debug Monitor is programmed to be at a lower priority than the current Task.

This scenario is indicated by posting a software interrupt at the interrupt level associated with the Breakpoint. Therefore, when the CPU interrupt priority level falls below that of the Debug Monitor, the Debug Monitor routine is entered. In order to indicate to the Monitor routine that the Breakpoint was postponed, the Posted Event bit (PEVT) in the Debug Status register is set when the software interrupt is posted. It is the responsibility of the Breakpoint Interrupt handler to check this bit in the Debug Status register and to subsequently clear that bit if necessary.

*Note: DBGSR.SUSP and DBGSR.PREVSUSP are not updated when a breakpoint interrupt is posted.*

*Note: DBGSR.EVTSRC is always updated regardless of whether or not a breakpoint interrupt is posted.*

## **Interrupts to Other Targets**

As well as being targeted at the CPU, a breakpoint interrupt can be targeted at other cores in the system.

### **12.4.7 Suspend Out**

The suspend out signal will either be asserted or negated when a debug event occurs. The previous state of the suspend out signal is recorded in DBGSR.PREVSUSP.

### **12.4.8 Performance Counter Start/Stop (TriCore 1.3.1)**

When the performance counter is operating in task mode, the counters are started and stopped by debug actions. All event registers allow the counters to either be started or stopped.

The trigger event registers also allow the mode to be toggled to active (start) or inactive (stop). This allows a single RTE to be used to control the performance counter, in certain applications.

### **12.4.9 None (TriCore 1.3.1)**

No action is implemented through the EVTA field of the event's register, however the suspend out signal, performance count and DBGSR register updates still occur as normal for an event.

#### **12.4.10 Disabled**

The event is disabled and no actions occur: the suspend out signal, performance counter control and DBGSR register ignore the event.

#### **12.4.11 Suspend In Halt (TriCore 1.3.1)**

When the Suspend In signal is asserted, halt mode is always entered so long as debug is enabled. The CPU remains in halt mode so long as Suspend In is asserted. When Suspend In is negated, the CPU is released from halt.

This facility is implemented so that in a multi core system, several cores can be halted and released from halt simultaneously.

### **12.5 Priority of Debug Events**

When two or more Debug Events occur on the same instruction, priorities are used to determine which Debug Event occurs. This section describes how those priorities are determined.

All Debug Events can be linked to specific instruction except for the external condition (EXEVT) Debug Event which is not linked to any instruction being executed. The latency of the EXEVT Debug Event is not defined.

When linking Debug Events to a specific instruction, they can be generated either logically before or logically after the execution of the instruction that they are linked with. This is known as Break Before Make (BBM) and Break After Make (BAM) respectively, and is controlled by the BAM<sup>1)</sup> bit in the various Debug Event registers.

*Note: Data access and data/code combination access triggers can only create BAM Debug Events. When triggers occur, TRnEVT.BBM is ignored.*

There are two types of Debug Action: those that change the program flow immediately by either halting the processor or redirecting the program flow (pipeline Debug Actions), and those that do not change the program flow immediately (non-pipeline Debug Actions). Halt, breakpoint trap and taken breakpoint interrupts are the pipeline Debug Actions. Indicate on core breakout and posted breakpoint interrupts are the non-pipeline Debug Actions.

There are four classes of priority to determine the relative priority of two Debug Events.

Priority (high to low):

1. Pipeline / Non-pipeline.
2. Instruction order.
3. BBM/BAM.
4. Debug Event Priorities.

<sup>1)</sup> EXEVT.BAM has no effect on the latency of external condition Debug Events.

Pipeline debug events change the program flow and must always be taken, so if a non-pipeline debug event occurs in the same cycle as a pipeline debug event, the pipeline debug event takes priority. This may occur because a BBM non-pipeline debug event and BAM pipeline debug event occur on the same instruction. Because of TriCore's superscalar architecture, it may also occur if a pipeline debug event occurs on an instruction immediately following an instruction with a non-pipeline debug event.

When the CDC is setup to create more than one BBM Debug Event on the same instruction, only one Debug Event is generated. Similarly, when the CDC is setup to create more than one BAM Debug Event on the same instruction, only one Debug Event is generated. In both cases the Debug Event priorities are used to determine which Debug Event is generated.

**Table 21      Debug Event Priorities**

Priority (high to low)	Type of Debug Event
1	Debug Instruction (SWEVT) / Core Register Access (CREVT).
2	Trigger 0 (TR0EVT).
3	Trigger 1 (TR1EVT).
4	Trigger <i>n</i> (TR <i>n</i> EVT). Note that the number of TR <i>n</i> EVT registers is implementation dependent.

*Note: The external condition may not generate a Debug Event even when programmed to do so, if a BAM Debug Event is generated in the same cycle. To avoid potential loss of EXEVT Debug Events, BBM Debug Events should be used.*

## 12.6      Call Tracing

The tracing of subroutine calls in a TriCore system is performed using the PSW based call depth counter and the CDO trap handler.

The sequence followed for call tracing is as follows:

1. The PSW based Call Depth Counter is set so as to generate a CDO trap on every subroutine call. (PSW.CDC = 1111110<sub>B</sub>).
2. The Call Depth counting system is enabled. (PSW.CDE = 1).
3. When the next CALL is attempted, a CDO trap will be taken instead of the subroutine call.
4. The CDO trap handler then performs the required trace function.
5. The CDO trap handler clears the PSW.CDE bit of the trapping context in memory.
6. The CDO trap handler executes a Return from Exception (RFE). This restores the trapping context from memory, this time with the call depth tracing disabled. (PSW.CDE=0).

7. The original CALL is executed. As the call depth tracing system is now disabled (PSW.CDE=0) the subroutine call will be successful.
- Whenever the PSW is saved by a CALL instruction the CDE bit is forced to "1".
  - The state of the PSW.CDE bit at the start of a subroutine is "1".

Refer to the CALL instruction in Volume 2: Instruction Set.

Therefore in a Call Tracing sequence the PSW.CDE bit has a "one-shot" operation, being disabled for a single subroutine call after being cleared by the CDO trap.

## **12.7 The CDC Control Registers**

The Debug Status Register (DBGSR) contains information about the current status of the Core Debug Controller (CDC) hardware in the CPU core:

- A bit to indicate whether the CDC is enabled.
- The source of the last Debug Event..

Each source of a Debug Event has an associated register which defines the Debug Actions to be taken when the Debug Event is raised. These registers may contain extra information about the criteria that must be met for the Debug Event to be raised, such as the combination of Debug Triggers for example.

TriCore 1.3 and TriCore 1.3.1 have different register fields. Both versions are described in this chapter. TriCore 1.3 registers are described from [page 12-15](#). TriCore 1.3.1 registers are described from [page 12-25](#).

## 12.8 CDC Control Registers (TriCore 1.3)

### 12.8.1 DBGSR Debug Status Register

Note: TriCore 1.3 Architecture only.

#### DBGSR

#### Debug Status Register

(FD00<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub> (Boot Execute)

0000 0002<sub>H</sub> (Boot Halt)

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
-			EVTSRC					P EVT	PSU SP	-	SU SP	-	HALT	DE	
			rh					rwh	rh		rwh		rwh	rh	

Field	Bits	Type	Description
-	[31:13]	-	<b>Reserved Field</b>
EVTSRC	[12:8]	rh	<b>Event Source</b> 0 : EXEVT. 1 : CREVT. 2 : SWEVT. 16 + n TR <sub>n</sub> EVT (n = 0, 1). other = Reserved.
PEVT	7	rwh	<b>Posted Event</b> 0 : No posted event. 1 : Posted event.
PREVSUSP	6	rh	<b>Previous State of Core Suspend-Out Signal</b> 0 : Previous core suspend-out inactive. 1 : Previous core suspend-out active. Updated when a Debug Event causes a hardware update of DBGSR.SUSP. This field is not updated for writes to DBGSR.SUSP.
-	5	-	<b>Reserved Field</b>

Field	Bits	Type	Description
SUSP	4	rwh	<b>Current State of the Core Suspend-Out Signal</b> 0 : Core suspend-out inactive. 1 : Core suspend-out active.
-	3	-	<b>Reserved Field</b>
HALT	[2:1]	rwh	<b>CPU Halt Request / Status Field</b> HALT can be set or cleared by software. HALT[0] is the actual Halt bit. HALT[1] is a mask bit to specify whether or not HALT[0] is to be updated on a software write. HALT[1] is always read as 0. HALT[1] must be set to 1 in order to update HALT[0] by software (R: read; W: write). 00 <sub>B</sub> R: CPU running. W: HALT[0] unchanged. 01 <sub>B</sub> R: CPU halted. W: HALT[0] unchanged. 10 <sub>B</sub> R: Not Applicable. W: reset HALT[0]. 11 <sub>B</sub> R: Not Applicable. W: If DBGSR.DE == 1 (The CDC is enabled), set HALT[0]. If DBGSR.DE == 0 (The CDC is not enabled), HALT[0] is left unchanged.
DE	0	rh	<b>Debug Enable</b> Indicates whether the CDC is enabled. 0 : The CDC disabled. 1 : The CDC enabled.



## 12.8.2 External Event Register

*Note: TriCore 1.3 Architecture only.*

### EXEVT

**External Event Register (FD08<sub>H</sub>)**      **Reset Value: 0000 0000<sub>H</sub>**

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
-										SU SP	-	BBM	EVTA		
										rw		rw	rw		

Field	Bits	Type	Description
-	[31:6]	-	<b>Reserved Field</b>
SUSP	5	rw	<b>CDC Suspend-Out Signal State</b> Value to be assigned to the CDC suspend-out signal when the Debug Event is raised.
-	4	-	<b>Reserved Field</b>
BBM	3	rw	<b>Break Before Make (BBM) or Break After Make (BAM) Selection</b> 0 : Break after make (BAM). 1 : Break before make (BBM).
EVTA	[2:0]	rw	<b>Event Associated</b> Debug Action associated with the Debug Event: 000 <sub>B</sub> : None; disabled. 001 <sub>B</sub> : Pulse BRKOUT signal. 010 <sub>B</sub> : Halt and pulse BRKOUT signal. 011 <sub>B</sub> : Breakpoint trap and pulse BRKOUT signal. 100 <sub>B</sub> : Breakpoint interrupt 0 and pulse BRKOUT signal. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 and pulse BRKOUT signal <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 and pulse BRKOUT signal <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 and pulse BRKOUT signal <sup>1)</sup> .

<sup>1)</sup> If not implemented, None.

### 12.8.3 Core Register Access Event Register

*Note: TriCore 1.3 Architecture only.*

#### CREVT

**Core Register Access Event (FD0C<sub>H</sub>)**      **Reset Value: 0000 0000<sub>H</sub>**

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
-										SU SP	-	BBM	EVTA		
										rw		rw	rw		

Field	Bits	Type	Description
-	[31:6]	-	<b>Reserved Field</b>
SUSP	5	rw	<b>CDC Suspend-Out Signal State</b> Value to be assigned to the CDC suspend-out signal when the Debug Event is raised.
-	4	-	<b>Reserved Field</b>
BBM	3	rw	<b>Break Before Make (BBM) or Break After Make (BAM) Selection</b> 0 : Break after make (BAM). 1 : Break before make (BBM).
EVTA	[2:0]	rw	<b>Event Associated</b> Debug Action associated with the Debug Event: 000 <sub>B</sub> : None; disabled. 001 <sub>B</sub> : Pulse BRKOUT signal. 010 <sub>B</sub> : Halt and pulse BRKOUT signal. 011 <sub>B</sub> : Breakpoint trap and pulse BRKOUT signal. 100 <sub>B</sub> : Breakpoint interrupt 0 and pulse BRKOUT signal. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 and pulse BRKOUT signal <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 and pulse BRKOUT signal <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 and pulse BRKOUT signal <sup>1)</sup> .

<sup>1)</sup> If not implemented, None.

## 12.8.4 Software Debug Event Register

*Note: TriCore 1.3 Architecture only.*

### SWEVT

**Software Debug Event (FD10<sub>H</sub>)** **Reset Value: 0000 0000<sub>H</sub>**

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
-										SU SP	-	BBM	EVTA		
										rw		rw	rw		

Field	Bits	Type	Description
-	[31:6]	-	<b>Reserved Field</b>
SUSP	5	rw	<b>CDC Suspend-Out Signal State</b> Value to be assigned to the CDC suspend-out signal when the event is raised.
-	4	-	<b>Reserved Field</b>
BBM	3	rw	<b>Break Before Make (BBM) or Break After Make (BAM) Selection</b> 0 : Break after make (BAM). 1 : Break before make (BBM).
EVTA	[2:0]	rw	<b>Event Associated</b> Debug Action associated with the Debug Event: 000 <sub>B</sub> : None; disabled. 001 <sub>B</sub> : Pulse BRKOUT signal. 010 <sub>B</sub> : Halt and pulse BRKOUT signal. 011 <sub>B</sub> : Breakpoint trap and pulse BRKOUT signal. 100 <sub>B</sub> : Breakpoint interrupt 0 and pulse BRKOUT signal. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 and pulse BRKOUT signal <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 and pulse BRKOUT signal <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 and pulse BRKOUT signal <sup>1)</sup> .

<sup>1)</sup> If not implemented, None.

## 12.8.5 Trigger Event Registers

*Note: TriCore 1.3 Architecture only.*

### TR0EVT

Trigger Event 0

(FD20<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>

### TR1EVT

Trigger Event 1

(FD24<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
-											ASI				
rw															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
ASI_EN	-			DU_U	DU_LR	DLR_U	DLR_LR	-			SU_SP	-	BBM	EVTA	
rw				rw	rw	rw	rw				rw		rw	rw	

Field	Bits	Type	Description
-	[31:21]	-	<b>Reserved Field</b>
ASI	[20:16]	rw	<b>Address Space Identifier</b> The ASI of the Debug Trigger process. (Not implemented in TriCore 1.2)
ASI_EN	15	rw	<b>Enable ASI Comparison</b> 0 : No ASI comparison performed. Debug Trigger is valid for all processes. 1 : Enable ASI comparison. Debug Events are only triggered when the current process ASI matches TRnEVT.ASI. Field should be set to 0 for implementations without an MMU. (Not implemented in TriCore 1.2)
-	[14:12]	-	<b>Reserved Field</b>

Field	Bits	Type	Description
DU_U	11	rw	<b>Controls combinations of D<sub>U</sub> and C<sub>U</sub></b> Note: Refer to <a href="#">Table 20, page 12-6</a> for clarification of trigger conditions that generate a Debug Event. 0 : D <sub>U</sub> triggers event unless DU_LR == 1, where C <sub>LR</sub> is also required. C <sub>U</sub> triggers event unless DLR_U == 1, where D <sub>LR</sub> is also required. 1 : D <sub>U</sub> and C <sub>U</sub> only trigger an event when they are both present.
DU_LR	10	rw	<b>Controls combination of D<sub>U</sub> and C<sub>LR</sub></b> Note: Refer to <a href="#">Table 20, page 12-6</a> for clarification of trigger conditions that generate a Debug Event. 0 : D <sub>U</sub> triggers event unless DU_U == 1. Where C <sub>U</sub> is also required. C <sub>LR</sub> triggers event unless DU_U == 1. Where D <sub>U</sub> is also required. 1 : D <sub>U</sub> and C <sub>LR</sub> only trigger an event when they are both present.
DLR_U	9	rw	<b>Controls combination of D<sub>LR</sub> and C<sub>U</sub></b> Note: Refer to <a href="#">Table 20, page 12-6</a> for clarification of trigger conditions that generate a Debug Event. 0 : D <sub>LR</sub> triggers event unless DLR_LU == 1. Where C <sub>LR</sub> is also required. C <sub>U</sub> triggers event unless DU_U == 1. Where D <sub>U</sub> is also required. 1 : D <sub>LR</sub> and C <sub>U</sub> only trigger an event when they are both present.
DLR_LR	8	rw	<b>Controls combination of D<sub>LR</sub> and C<sub>LR</sub></b> Note: Refer to <a href="#">Table 20, page 12-6</a> for clarification of trigger conditions that generate a Debug Event. 0 : D <sub>LR</sub> triggers event unless DLR_LU == 1. Where C <sub>U</sub> is also required. C <sub>LR</sub> triggers event unless DU_U == 1. Where D <sub>U</sub> is also required. 1 : D <sub>LR</sub> and C <sub>LR</sub> only trigger an event when they are both present.

**Core Debug Controller (CDC)**

Field	Bits	Type	Description
SUSP	5	rw	<b>CDC Suspend-Out Signal State</b> Value to be assigned to the CDC suspend-out signal when the Debug Event is raised.
BBM	3	rw	<b>Break Before Make (BBM) or Break After Make (BAM) Selection</b> Code triggers BBM or BAM selection. 0 : Code only triggers Break After Make (BAM). 1 : Code only triggers Break Before Make (BBM). Note that data access and data/code combination access triggers can only create BAM Debug Events. When these triggers occur, TRnEVT.BBM is ignored.
EVTA	[2:0]	rw	<b>Event Associated</b> Specifies the Debug Action associated with the Debug Event: 000 <sub>B</sub> : None; disabled. 001 <sub>B</sub> : Pulse BRKOUT signal. 010 <sub>B</sub> : Halt and pulse BRKOUT signal. 011 <sub>B</sub> : Breakpoint trap and pulse BRKOUT signal. 100 <sub>B</sub> : Breakpoint interrupt 0 and pulse BRKOUT signal. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 and pulse BRKOUT signal <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 and pulse BRKOUT signal <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 and pulse BRKOUT signal <sup>1)</sup> .

<sup>1)</sup> If not implemented, None.

## 12.8.6 Debug Monitor Start Address Register


*Note: TriCore 1.3 Architecture only.*

The DMS reset value is DE00 00n0<sub>H</sub>, where 'n' is Core ID.

### DMS

**Debug Monitor Start Address (FD40<sub>H</sub>)**

**Reset Value: DE00 00n0<sub>H</sub>**



Field	Bits	Type	Description
DMS Value	[31:1]	rw	<b>Debug Monitor Start Address</b> The address at which monitor code execution begins when a breakpoint trap is taken.
-	0	-	<b>Reserved Field</b>


## 12.8.7 Debug Context Save Area Pointer Register

*Note: TriCore 1.3 Architecture only.*

### DCX

**Debug Context Save Area Pointer (FD44<sub>H</sub>)**

**Reset Value: DE80 0000<sub>H</sub>**



Field	Bits	Type	Description
DCX Value	[31:4]	rw	<b>Debug Context Save Area Pointer</b> Address where the debug context is stored following a breakpoint trap.
-	[3:0]	-	<b>Reserved Field</b>



## 12.9 CDC Control Registers (TriCore 1.3.1)

### 12.9.1 Debug Status Register

*Note: TriCore 1.3.1 Architecture only.*

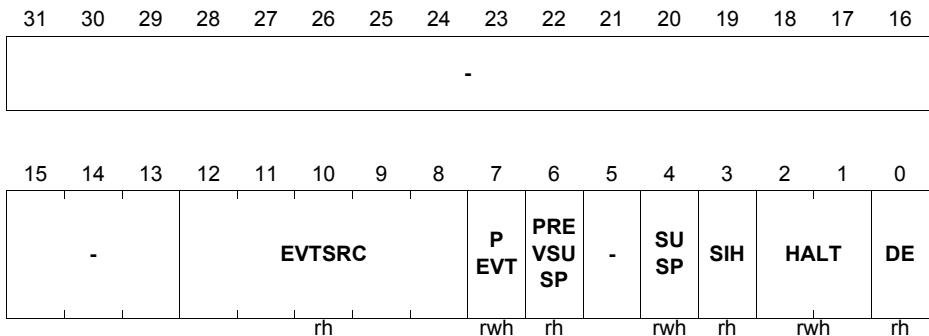
#### DBGSR

#### Debug Status Register

(FD00<sub>H</sub>)

**Reset Value: 0000 0000<sub>H</sub> (Boot Execute)**

**0000 0002<sub>H</sub> (Boot Halt)**



Field	Bits	Type	Description
-	[31:13]	-	<b>Reserved Field</b>
EVTSRC	[12:8]	rh	<b>Event Source</b> 0 : EXEVT. 1 : CREVT. 2 : SWEVT. 16 + n TRnEVT (n = 0, 1). other = Reserved.
PEVT	7	rwh	<b>Posted Event</b> 0 : No posted event. 1 : Posted event.
PREVSUSP	6	rh	<b>Previous State of Core Suspend-Out Signal</b> 0 : Previous core suspend-out inactive. 1 : Previous core suspend-out active. Updated when a Debug Event causes a hardware update of DBGSR.SUSP. This field is not updated for writes to DBGSR.SUSP.
-	5	-	<b>Reserved Field</b>

Field	Bits	Type	Description
SUSP	4	rwh	<b>Current State of the Core Suspend-Out Signal</b> 0 : Core suspend-out inactive. 1 : Core suspend-out active.
SIH	3	rh	<b>Suspend-in Halt</b> State of the Suspend-In signal. 1 : The Suspend-In signal is asserted. The CPU is in Halt Mode. 0 : The Suspend-In signal is negated. The CPU is not in Halt Mode, (except when the Halt mechanism is set following a Debug Event or a write to DBGSR.HALT).
HALT	[2:1]	rwh	<b>CPU Halt Request / Status Field</b> HALT can be set or cleared by software. HALT[0] is the actual Halt bit. HALT[1] is a mask bit to specify whether or not HALT[0] is to be updated on a software write. HALT[1] is always read as 0. HALT[1] must be set to 1 in order to update HALT[0] by software (R: read; W: write). 00 <sub>B</sub> R: CPU running. W: HALT[0] unchanged. 01 <sub>B</sub> R: CPU halted. W: HALT[0] unchanged. 10 <sub>B</sub> R: Not Applicable. W: reset HALT[0]. 11 <sub>B</sub> R: Not Applicable. W: If DBGSR.DE == 1 (The CDC is enabled), set HALT[0]. If DBGSR.DE == 0 (The CDC is not enabled), HALT[0] is left unchanged.
DE	0	rh	<b>Debug Enable</b> Indicates whether CDC is enabled. 0 : The CDC disabled. 1 : The CDC enabled.

## 12.9.2 External Event Register


*Note: TriCore 1.3.1 Architecture only.*

### EXEVT

#### External Event Register

(FD08H)

Reset Value: 0000 0000H



Field	Bits	Type	Description
-	[31:8]	-	<b>Reserved Field</b>
CSP	7	rw	<b>Counter Stop</b> When this event occurs, in addition to the event's action, stop the performance counters when they are in task mode.
CST	6	rw	<b>Counter Start</b> When this event occurs, in addition to the event's action, start the performance counters when they are in task mode.
SUSP	5	rw	<b>CDC Suspend-Out Signal State</b> Value to be assigned to the CDC suspend-out signal when the Debug Event is raised.
BOD	4	rw	<b>Breakout Disable</b> 0 : BRKOUT signal asserted according to the Debug Action specified in the EVTA field. 1 : BRKOUT signal not asserted. This takes priority over any assertion generated by the EVTA field.

Field	Bits	Type	Description
BBM	3	rw	<b>Break Before Make (BBM) or Break After Make (BAM) Selection</b> 0 : Break after make (BAM). 1 : Break before make (BBM).
EVTA	[2:0]	rw	<b>Event Associated</b> Debug Action associated with the Debug Event: <b>When field BOD = 0</b> 000 <sub>B</sub> : Disabled. 001 <sub>B</sub> : Pulse BRKOUT Signal. 010 <sub>B</sub> : Halt and pulse BRKOUT Signal. 011 <sub>B</sub> : Breakpoint trap and pulse BRKOUT Signal. 100 <sub>B</sub> : Breakpoint interrupt 0 and pulse BRKOUT Signal. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 and pulse BRKOUT Signal <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 and pulse BRKOUT Signal <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 and pulse BRKOUT Signal <sup>1)</sup> . <b>When field BOD = 1</b> 000 <sub>B</sub> : Disabled. 001 <sub>B</sub> : None. 010 <sub>B</sub> : Halt. 011 <sub>B</sub> : Breakpoint trap. 100 <sub>B</sub> : Breakpoint interrupt 0. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 <sup>1)</sup> .


<sup>1)</sup> If not implemented, None

### 12.9.3 Core Register Access Event Register

*Note: TriCore 1.3.1 Architecture only.*

#### CREVT

**Core Register Access Event (FD0C<sub>H</sub>)** **Reset Value: 0000 0000<sub>H</sub>**



Field	Bits	Type	Description
-	[31:8]	-	<b>Reserved Field</b>
CSP	7	rw	<b>Counter Stop</b> When this event occurs, in addition to the event's action, stop the performance counters when they are in task mode.
CST	6	rw	<b>Counter Start</b> When this event occurs, in addition to the event's action, start the performance counters when they are in task mode.
SUSP	5	rw	<b>CDC Suspend-Out Signal State</b> Value to be assigned to the CDC suspend-out signal when the Debug Event is raised.
BOD	4	rw	<b>Breakout Disable</b> 0 : BRKOUT signal asserted according to the action specified in the EVTA field. 1 : BRKOUT signal not asserted. This takes priority over any assertion generated by the EVTA field.

Field	Bits	Type	Description
BBM	3	rw	<b>Break Before Make (BBM) or Break After Make (BAM) Selection</b> 0 : Break after make (BAM). 1 : Break before make (BBM).
EVTA	[2:0]	rw	<b>Event Associated</b> Debug Action associated with the Debug Event: <b>When field BOD = 0</b> 000 <sub>B</sub> : Disabled. 001 <sub>B</sub> : Pulse BRKOUT Signal. 010 <sub>B</sub> : Halt and pulse BRKOUT Signal. 011 <sub>B</sub> : Breakpoint trap and pulse BRKOUT Signal. 100 <sub>B</sub> : Breakpoint interrupt 0 and pulse BRKOUT Signal. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 and pulse BRKOUT Signal <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 and pulse BRKOUT Signal <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 and pulse BRKOUT Signal <sup>1)</sup> . <b>When field BOD = 1</b> 000 <sub>B</sub> : Disabled. 001 <sub>B</sub> : None. 010 <sub>B</sub> : Halt. 011 <sub>B</sub> : Breakpoint trap. 100 <sub>B</sub> : Breakpoint interrupt 0. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 <sup>1)</sup> .

<sup>1)</sup> If not implemented, None

## 12.9.4 Software Debug Event Register

*Note: TriCore 1.3.1 Architecture only.*

### SWEVT

#### Software Debug Event

(FD10<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>

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
-								CSP	CST	SU SP	BOD	BBM	EVTA		
								rw	rw	rw	rw	rw	rw		

Field	Bits	Type	Description
-	[31:8]	-	<b>Reserved Field</b>
CSP	7	rw	<b>Counter Stop</b> When this event occurs, in addition to the event's action, stop the performance counters when they are in task mode.
CST	6	rw	<b>Counter Start</b> When this event occurs, in addition to the event's action, start the performance counters when they are in task mode.
SUSP	5	rw	<b>CDC Suspend-Out Signal State</b> Value to be assigned to the CDC suspend-out signal when the event is raised.
BOD	4	rw	<b>Breakout Disable</b> 0 : BRKOUT signal asserted according to the action specified in the EVTA field. 1 : BRKOUT signal not asserted. This takes priority over any assertion generated by the EVTA field.

Field	Bits	Type	Description
BBM	3	rw	<b>Break Before Make (BBM) or Break After Make (BAM) Selection</b> 0 : Break after make (BAM). 1 : Break before make (BBM).
EVTA	[2:0]	rw	<b>Event Associated</b> Debug Action associated with the Debug Event: <b>When field BOD = 0</b> 000 <sub>B</sub> : Disabled. 001 <sub>B</sub> : Pulse BRKOUT Signal. 010 <sub>B</sub> : Halt and pulse BRKOUT Signal. 011 <sub>B</sub> : Breakpoint trap and pulse BRKOUT Signal. 100 <sub>B</sub> : Breakpoint interrupt 0 and pulse BRKOUT Signal. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 and pulse BRKOUT Signal <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 and pulse BRKOUT Signal <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 and pulse BRKOUT Signal <sup>1)</sup> . <b>When field BOD = 1</b> 000 <sub>B</sub> : Disabled. 001 <sub>B</sub> : None. 010 <sub>B</sub> : Halt. 011 <sub>B</sub> : Breakpoint trap. 100 <sub>B</sub> : Breakpoint interrupt 0. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 <sup>1)</sup> .

<sup>1)</sup> If not implemented, None



## 12.9.5 Trigger Event Registers

*Note: TriCore 1.3.1 Architecture only.*

### TR0EVT

**Trigger Event 0**

**(FD20<sub>H</sub>)**

**Reset Value: 0000 0000<sub>H</sub>**

### TR1EVT

**Trigger Event 1**

**(FD24<sub>H</sub>)**

**Reset Value: 0000 0000<sub>H</sub>**

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
											ASI				
rw															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
ASI_EN	-			DU_U	DU_LR	DLR_U	DLR_LR	CNT		SU SP	BOD	BBM	EVTA		
rw				rw		rw	rw	rw		rw	rw	rw	rw		

Field	Bits	Type	Description
-	[31:21]	-	<b>Reserved Field</b>
ASI	[20:16]	rw	<b>Address Space Identifier</b> The ASI of the Debug Trigger process.
ASI_EN	15	rw	<b>Enable ASI Comparison</b> 0 : No ASI comparison performed. Debug Trigger is valid for all processes. 1 : Enable ASI comparison. Debug Events are only triggered when the current process ASI matches TRnEVT.ASI. Field should be set to 0 for implementations without an MMU.
-	[14:12]	-	<b>Reserved Field</b>

Field	Bits	Type	Description
DU_U	11	rw	<b>Controls combinations of D<sub>U</sub> and C<sub>U</sub></b> Note: Refer to <a href="#">Table 20, page 12-6</a> for clarification of trigger conditions that generate a Debug Event. 0 : D <sub>U</sub> triggers event unless DU_LR == 1, where C <sub>LR</sub> is also required. C <sub>U</sub> triggers event unless DLR_U == 1, where D <sub>LR</sub> is also required. 1 : D <sub>U</sub> and C <sub>U</sub> only trigger an event when they are both present.
DU_LR	10	rw	<b>Controls combination of D<sub>U</sub> and C<sub>LR</sub></b> Note: Refer to <a href="#">Table 20, page 12-6</a> for clarification of trigger conditions that generate a Debug Event. 0 : D <sub>U</sub> triggers event unless DU_U == 1. Where C <sub>U</sub> is also required. C <sub>LR</sub> triggers event unless DU_U == 1. Where D <sub>U</sub> is also required. 1 : D <sub>U</sub> and C <sub>LR</sub> only trigger an event when they are both present.
DLR_U	9	rw	<b>Controls combination of D<sub>LR</sub> and C<sub>U</sub></b> Note: Refer to <a href="#">Table 20, page 12-6</a> for clarification of trigger conditions that generate a Debug Event. 0 : D <sub>LR</sub> triggers event unless DLR_LU == 1. Where C <sub>LR</sub> is also required. C <sub>U</sub> triggers event unless DU_U == 1. Where D <sub>U</sub> is also required. 1 : D <sub>LR</sub> and C <sub>U</sub> only trigger an event when they are both present.
DLR_LR	8	rw	<b>Controls combination of D<sub>LR</sub> and C<sub>LR</sub></b> Note: Refer to <a href="#">Table 20, page 12-6</a> for clarification of trigger conditions that generate a Debug Event. 0 : D <sub>LR</sub> triggers event unless DLR_LU == 1. Where C <sub>U</sub> is also required. C <sub>LR</sub> triggers event unless DU_U == 1. Where D <sub>U</sub> is also required. 1 : D <sub>LR</sub> and C <sub>LR</sub> only trigger an event when they are both present.

Field	Bits	Type	Description
CNT	[7:6]	rw	<b>Counter</b> When this event occurs adjust the control of the performance counters in task mode as follows: 00 : No change. 01 : Start the performance counters. 10 : Stop the performance counters. 11 : Toggle the performance counter control (i.e. start it if it is currently stopped, stop it if it is currently running).
SUSP	5	rw	<b>CDC Suspend-Out Signal State</b> Value to be assigned to the CDC suspend-out signal when the Debug Event is raised.
BOD	4	rw	<b>Breakout Disable</b> 0 : BRKOUT signal asserted according to the action specified in the EVTA field. 1 : BRKOUT signal not asserted. This takes priority over any assertion generated by the EVTA field.

Field	Bits	Type	Description
BBM	3	rw	<b>Break Before Make (BBM) or Break After Make (BAM) Selection</b> Code triggers BBM or BAM selection. 0 : Code only triggers Break After Make (BAM). 1 : Code only triggers Break Before Make (BBM). Note that data access and data/code combination access triggers can only create BAM Debug Events. When these triggers occur, TRnEVT.BBM is ignored.
EVTA	[2:0]	rw	<b>Event Associated</b> Specifies the Debug Action associated with the Debug Event: <b>When field BOD = 0</b> 000 <sub>B</sub> : Disabled. 001 <sub>B</sub> : Pulse BRKOUT Signal. 010 <sub>B</sub> : Halt and pulse BRKOUT Signal. 011 <sub>B</sub> : Breakpoint trap and pulse BRKOUT Signal. 100 <sub>B</sub> : Breakpoint interrupt 0 and pulse BRKOUT Signal. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 and pulse BRKOUT Signal <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 and pulse BRKOUT Signal <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 and pulse BRKOUT Signal <sup>1)</sup> . <b>When field BOD = 1</b> 000 <sub>B</sub> : Disabled. 001 <sub>B</sub> : None. 010 <sub>B</sub> : Halt. 011 <sub>B</sub> : Breakpoint trap. 100 <sub>B</sub> : Breakpoint interrupt 0. 101 <sub>B</sub> : If implemented, breakpoint interrupt 1 <sup>1)</sup> . 110 <sub>B</sub> : If implemented, breakpoint interrupt 2 <sup>1)</sup> . 111 <sub>B</sub> : If implemented, breakpoint interrupt 3 <sup>1)</sup> .

<sup>1)</sup> If not implemented, None

## 12.9.6 Debug Monitor Start Address Register


*Note: TriCore 1.3.1 Architecture only.*

The DMS reset value is  $DE00\ 0nn0_H$ , where 'nn' is the 4-bit Core ID in bits [9:6].

### DMS

**Debug Monitor Start Address (FD40<sub>H</sub>)**

**Reset Value: DE00 0nn0<sub>H</sub>**



Field	Bits	Type	Description
DMS Value	[31:1]	rw	<b>Debug Monitor Start Address</b> The address at which monitor code execution begins when a breakpoint trap is taken.
-	0	-	<b>Reserved Field</b>

## 12.9.7 Debug Context Save Area Pointer Register


*Note: TriCore 1.3.1 Architecture only.*

The reset value of the DCX register is DE80 0nn0<sub>H</sub>, where 'nn' is the 4-bit Core ID in bits [9:6].

### DCX

**Debug Context Save Area Pointer (FD44<sub>H</sub>)**

**Reset Value: DE80 0nn0<sub>H</sub>**



Field	Bits	Type	Description
DCX Value	[31:4]	rw	<b>Debug Context Save Area Pointer</b> Address where the debug context is stored following a breakpoint trap.
-	[3:0]	-	<b>Reserved Field</b>

## 12.9.8 Debug Trap Control Register

*Note: TriCore 1.3.1 Architecture only.*


The Debug Trap Control Register contains the DTA (Debug Trap Active) bit. The reset value of DTA is zero. The DTA bit is defined as being cleared on an RFM instruction and set on a breakpoint trap. It may also be set and cleared by MTCR.

### DBGTCR

**Debug Trap Control Register**

**(FD48<sub>H</sub>)**

**Reset Value: 0000 0000<sub>H</sub>**



Field	Bits	Type	Description
-	[31:1]		<b>Reserved Field</b>
DTA	0	rwh	<b>Debug Trap Active Bit</b> 1: A breakpoint Trap is active. 0: No breakpoint trap is active.  A breakpoint trap may only be taken in the condition DTA==0. Taking a breakpoint trap sets the DTA bit to one. Further breakpoint traps are therefore disabled until such time as the breakpoint trap handler clears the DTA bit or until the breakpoint trap handler terminates with a RFM.

### 12.9.9 Software Breakpoint Service Request Control Register

The Software Breakpoint Service Request Control Register (SBSRC $n$ ) defines the interrupt request parameters for a breakpoint interrupt, where  $n = 0, 1, 2$  or  $3$ . SBSRC1, 2 and 3 are optional and may not be implemented

Software Breakpoint Service Request Control Registers are located in the address range of the CPU slave interface (CPS).

#### SBSRC $n$ ( $n = 0$ to $3$ )

##### Software Breakpoint Service Request Control Register

(FFBC $_H - n * 4_H$ )

Reset Value: 0000 0000 $_H$

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
SET R	CLR R	SRR	SRE	TOS		-		SRPN							
w	w	rh	rw	rw				rw							

Field	Bits	Type	Description
-	[31:16]	-	<b>Reserved Field</b>
SETR	15	w	<b>Service Request Set</b> SETR is required to set SRR. 0 : No action. 1 : Set SRR. Written value is not stored. Read always returns 0. No action if CLRR is also set.
CLRR	14	w	<b>Service Request Clear</b> CLRR is required to clear SRR. 0 : No action. 1 : Clear SRR. Written value is not stored. Read always returns 0. No action if SETR is also set.
SRR	13	rh	<b>Service Request Flag</b> 0 : No Breakpoint Interrupt Service Request is pending. 1 : A Breakpoint Interrupt Service Request is pending.
SRE	12	rw	<b>Service Request Enable</b> 0 : Breakpoint Interrupt Service Request is disabled. 1 : Breakpoint Interrupt Service Request is enabled.



Field	Bits	Type	Description
TOS	[11:10]	rw	<b>Type Of Service Control</b> 00 <sub>B</sub> : Service Provider 0 - Typically CPU service is initiated. 01 <sub>B</sub> : Service Provider 1 - Implementation Specific. 10 <sub>B</sub> : Service Provider 2 - Implementation Specific. 11 <sub>B</sub> : Service Provider 3 - Implementation Specific.
-	[9:8]	-	<b>Reserved Field</b>
SRPN	[7:0]	rw	<b>Service Request Priority Number</b> 00 <sub>H</sub> : Breakpoint Interrupt Service Request is never serviced. 01 <sub>H</sub> : Breakpoint Interrupt Service Request, lowest priority. ... FF <sub>H</sub> : Breakpoint Interrupt Service Request, highest priority.

## **12.10 Core Performance Measurement and Analysis (TriCore 1.3.1)**

Real-time measurement of core performance provides useful insights to system developers, architects, compiler developers, application developers, OS developers, and so on.

TriCore includes the ability to measure different performance aspects of the processor without any real-time effect on its execution. The performance measurement hardware is configured so that only a subset of performance measurements can be taken simultaneously.

The performance measurement block can be used to measure basic parameters:

- CPU Clocks.
- Instruction Count.
- Instruction Cache Hit / Miss.
- Data Cache Hit / Miss (clean or dirty).

The performance counters can be used in a free running manner, enabled to acquire aggregate information. Alternatively they can be used in conjunction with the debug event logic to control 'windows' of operation for an individual task, for example starting and stopping the counters dynamically to filter the measured information on some desired event.

### **Performance Counter Overview**

The Performance counters are controlled in the Counter Control Register (CCTRL).

The performance counters can be enabled or disabled by writing the appropriate value to the counter enable CCTRL.CE bit.

Typically two parameters are always counted for base line measurement;

- The clock count.
- The number of instructions issued.

One of:

- Instruction Cache Hits.
- Data Cache Hits.

One of:

- Instruction Cache Misses.
- Data Cache Clean Misses.

Additionally:

- Data Cache Dirty Misses (cache write-back / eviction was required).

*Note: Counters can only be written when they are disabled (i.e. not in 'counting mode'). Any attempt to write during counting-mode will have no effect.*

*Note: The counters are free running incrementors once enabled, and will roll over to zero after the maximum value is reached.*

The grouping of counter functions allows typical measurements to be clustered; i.e. Data Cache performance and Instruction Cache performance.

These can all be measured against the background statistics of clock cycles and instructions issued.

The start of counters is not precisely synchronized to any pipeline stage. For example, once the instruction counter is enabled to count, it starts counting all retiring instructions from that clock cycle onward. Similarly, once the instruction cache miss counter is started, it will count all the instruction cache misses from that clock cycle onward.

There are two ways to enable counters: Normal mode and Task mode (CCTRL.CM).

Normal (default mode) or Task mode are configured by CCTRL.CM:

- Normal mode - The counters start counting as soon as they are enabled, and will keep counting until they are disabled.
- Task mode - The counters will only count if the processor detected a debug event with the action to start the performance counters.

### **Writing of the Counters**

Counters can be read any time, but they can only be written when they are not actively counting (i.e. when they are disabled). If the counters are disabled, then they are not considered to be in counting mode and so they can be written.

A counter is said to be in the counting mode if:

- The Normal or Task mode is selected.
- The mode is active (Normal mode is always active).
- The counter enable CE bit (in the Counter Control register - CCTRL) is enabled.

### **Counter Modes**

The Counter Mode (CM) bit in the Counter Control CSFR (i.e. CCTRL.CM) determines the operating mode of all the counters.

In the Normal mode of operation the counter increments on their respective triggers if the Count enable bit in the CCTRL is set (CCTRL.CE). In Task mode there is additional gating control from the debug unit which allows the data gathered in the performance counters to be filtered by some specific criteria, such as a single task for example.

### **Wrapping of the counters / Sticky bit**

The performance counters give the user some indication that the counters had wrapped (by use of a sticky bit.) This helps to tell whether the counter has wrapped between two measured values.

- All performance counters are 31bit counters with free wrapping operation.
- Bit 31 of each counter is sticky. It gets set when bits 30:0 wrap. It stays set until written by software.

For example:

```
if (counter_event && counters_en)
begin
    counter[30:0] <= counter[30:0] + 1;
    if (counter[30:0] == 31'hFFFFFFF)
        counter[31] <= 1;
end
else if (count_we)
    counter[31:0] <= write_data;
```

## 12.11 Performance Counter Registers (TriCore 1.3.1)

The performance counter registers used are:

**Table 22 OCDS Control Registers**

Register	Description	Offset Address	Reference
CCTRL	Counter Control Register.	FC00 <sub>H</sub>	<a href="#">page 12-45</a>
CCNT	CPU Clock Count Register.	FC04 <sub>H</sub>	<a href="#">page 12-47</a>
ICNT	Instruction Count Register.	FC08 <sub>H</sub>	<a href="#">page 12-48</a>
M1CNT	Multi Count Register 1.	FC0C <sub>H</sub>	<a href="#">page 12-49</a>
M2CNT	Multi Count Register 2.	FC10 <sub>H</sub>	<a href="#">page 12-50</a>
M3CNT	Multi Count Register 3.	FC14 <sub>H</sub>	<a href="#">page 12-51</a>

### 12.11.1 Counter Control Register

*Note: TriCore 1.3.1 Architecture only.*

#### CCTRL

##### Counter Control

(FC00<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>

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
			-				<b>M3</b>			<b>M2</b>			<b>M1</b>	<b>CE</b>	<b>CM</b>
							rw			rw			rw	rw	rw

Field	Bits	Type	Description
-	[31:11]	-	<b>Reserved Field</b>
M3	[10:8]	rw	<b>M3CNT configuration</b> 000 : Reserved. 001 : Reserved. 010 : Reserved. 011 : Data Cache Dirty Miss Count. 100 : Reserved. 101 : Reserved. 110 : Reserved. 111 : Reserved.
M2	[7:5]	rw	<b>M2CNT configuration</b> 000 : Reserved. 001 : Instruction Cache Miss Count. 010 : Reserved. 011 : Data Cache Clean Miss Count. 100 : Reserved. 101 : Reserved. 110 : Reserved. 111 : Reserved.


Field	Bits	Type	Description
M1	[4:2]	rw	<b>M1CNT configuration</b> 000 : Reserved. 001 : Reserved. 010 : Instruction Cache Hit Count. 011 : Data Cache Hit Count. 100 : Reserved. 101 : Reserved. 110 : Reserved. 111 : Reserved.
CE	1	rw	<b>Count Enable</b> 0 : Disable the counters: CCNT, ICNT, M1CNT, M2CNT, M3CNT. 1 : Enable the counters: CCNT, ICNT, M1CNT, M2CNT, M3CNT.
CM	0	rw	<b>Counter Mode</b> 0 : Normal Mode. 1 : Task Mode.

## 12.11.2 CPU Clock Cycle Count Register

*Note: TriCore 1.3.1 Architecture only.*

### CCNT

**CPU Clock Cycle Count (FC04<sub>H</sub>)** **Reset Value: 0000 0000<sub>H</sub>**




Field	Bits	Type	Description
SOvf	31	rw	<b>Sticky Overflow bit</b> This bit is set by hardware when count value [30:0] = 31'h7FFF_FFFF. It can only be cleared by software.
Count Value	[30:0]	rw	<b>Count Value</b> Current Count of the CPU Clock Cycles.

### 12.11.3 Instruction Count Register

*Note: TriCore 1.3.1 Architecture only.*

#### ICNT

**Instruction Count (FC08<sub>H</sub>)** **Reset Value: 0000 0000<sub>H</sub>**



Field	Bits	Type	Description
SOvf	31	rw	<b>Sticky Overflow bit</b> This bit is set by hardware when count value [30:0] = 31'h7FFF_FFFF. It can only be cleared by software.
Count Value	[30:0]	rw	<b>Count Value</b> Count of the Instructions Executed.



### 12.11.4 Multi-Count Register 1


*Note: TriCore 1.3.1 Architecture only.*

#### M1CNT

#### Multi-Count Register 1

(FC0C<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>



Field	Bits	Type	Description
SOvf	31	rw	<b>Sticky Overflow bit</b> This bit is set by hardware when count value [30:0] = 31'h7FFF_FFFF. It can only be cleared by software.
Count Value	[30:0]	rw	<b>Count Value</b> Count of the Selected Event.

### 12.11.5 Multi-Count Register 2


*Note: TriCore 1.3.1 Architecture only.*

#### M2CNT

#### Multi-Count Register 2

(FC10<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>



Field	Bits	Type	Description
SOvf	31	rw	<b>Sticky Overflow bit</b> This bit is set by hardware when count value [30:0] = 31'h7FFF_FFFF. It can only be cleared by software.
Count Value	[30:0]	rw	<b>Count Value</b> Count of the Selected Event.

### 12.11.6 Multi-Count Register 3


Note: TriCore 1.3.1 Architecture only.

#### M3CNT

#### Multi-Count Register 3

(FC14<sub>H</sub>)

Reset Value: 0000 0000<sub>H</sub>



Field	Bits	Type	Description
SOvf	31	rw	<b>Sticky Overflow bit</b> This bit is set by hardware when count value [30:0] = 31'h7FFF_FFFF. It can only be cleared by software.
Count Value	[30:0]	rw	<b>Count Value</b> Count of the Selected Event.



## **13 TriCore 1.3.1 Architectural Extensions**

The TriCore® 1.3.1 CPU contains a small number of extensions to the existing TriCore 1.3 Architecture to support the required feature set.

### **13.1 TriCore 1.3.1 Architectural Extensions - Trap System**

The following trap types are introduced:

#### **CAE Coprocessor Trap Asynchronous Error (TIN 4)**

This asynchronous trap is generated by a coprocessor to report an error. Examples of typical errors that can cause a CAE trap are unimplemented coprocessor instructions and arithmetic errors (as found in the Floating Point Unit for example).

CAE is shared amongst all coprocessors in a given system. A trap handler must therefore inspect all coprocessors to determine the cause of a trap.

#### **PIE Program Memory Integrity Error (TIN 5)**

The PIE trap is raised whenever an uncorrectable memory integrity error is detected in an instruction fetch.. The trap is synchronous to the erroneous instruction. The trap is of Class 4 and TIN 5.

A PIE trap is raised if any element within the fetch group contains an unrecoverable error. Hardware is not required to localise the error to a particular instruction.

An implementation may provide additional registers that can be interrogated to determine the source of the error more precisely. Refer to the User manual for a specific Tricore implementation for more details.

#### **DIE Data Memory Integrity Error (TIN 6)**

The DIE trap is raised whenever an uncorrectable memory integrity error is detected in a data access. The trap is of Class 4 and TIN 6.

Implementations may choose to implement the DIE trap as either an asynchronous or synchronous trap.

A DIE trap is raised if any element accessed by a load or store contains an uncorrectable error. Hardware is not required to localise the error to the access width of the operation.

An implementation may provide additional registers that can be interrogated to determine the source of the error more precisely. Refer to the User manual for a specific Tricore implementation for more details.

### **Trap Priority**

The DIE trap has a priority one lower than the DAE trap when the exception is taken as an asynchronous trap.

The DIE trap has a priority one lower than the DSE trap when the exception is taken as a synchronous trap

The PIE trap has the lowest priority of the program side instruction fetch traps. (Between PSE and IOPC in the priority table).

## **13.2 TriCore 1.3.1 Architectural Extensions - Core Registers**

A number of Core Special Function Registers (CSFRs) have been introduced to the TriCore 1.3.1 architecture in order to fully support functional enhancements. These are:

**Table 13-1 New CSFR Registers**

<b>Register Name</b>	<b>Description</b>	<b>Page</b>
SMACON	SIST Mode Access Control Register	<a href="#">page 3-20</a>
BMACON	BIST Mode Access Control Register	<a href="#">page 3-19</a>
MIECON	Memory Integrity Error Control	<a href="#">page 7-9</a>
CCPIER	Count of Corrected Program Integrity Errors	<a href="#">page 7-3</a>
CCDIER	Count of Corrected Data Integrity Errors	<a href="#">page 7-4</a>
PIEAR	Program Integrity Error Address Register	<a href="#">page 7-6</a>
PIETR	Program Integrity Error Trap Register	<a href="#">page 7-5</a>
DIEAR	Data Integrity Error Address Register	<a href="#">page 7-8</a>
DIETR	Data Integrity Error Trap Register	<a href="#">page 7-7</a>
FPU_TRAP_CON	FPU Trap Control Register	<a href="#">page 11-14</a>
FPU_TRAP_PC	FPU Trapping Instruction Program Count	<a href="#">page 11-17</a>
FPU_TRAP_OPC	FPU Trapping Instruction Opcode Register	<a href="#">page 11-18</a>
FPU_TRAP_SRC1	FPU Trapping Instruction Operand Register	<a href="#">page 11-19</a>
FPU_TRAP_SRC2	FPU Trapping Instruction Operand Register	<a href="#">page 11-20</a>
FPU_TRAP_SRC3	FPU Trapping Instruction Operand Register	<a href="#">page 11-21</a>
FPU_ID	FPU Identification Register	<a href="#">page 11-22</a>
COMPAT	Compatibility Control Register	<a href="#">page 3-18</a>
DBGTCR	Debug Trap Control Register	<a href="#">page 12-39</a>
CCTRL	Counter Control Register	<a href="#">page 12-45</a>
CCNT	CPU Clock Count Register	<a href="#">page 12-47</a>
ICNT	Instruction Count Register	<a href="#">page 12-48</a>
M1CNT	Multi Count Registers 1	<a href="#">page 12-49</a>
M2CNT	Multi Count Registers 2	<a href="#">page 12-50</a>
M3CNT	Multi Count Registers 3	<a href="#">page 12-51</a>

### **13.3 TriCore 1.3.1 Architectural Extensions - Instruction Set**

The following instructions are introduced:

- CACHEI.W and CACHEI.WI
- FPU Conversion Instructions. FTOIZ, FTOQ31Z and FTOUZ.

#### **Cachei.w and Cachei.wi**

These cache index instructions are used for efficient flushing without knowing the cache contents and preferably knowing very little about the cache itself (the total cache size and the cache line size). This helps in debugging and coherence in flushing data structures to optimize performance.

#### **FPU Conversion Instructions**

These instructions convert from floating point to other formats and always use round towards zero rounding irrespective of the current rounding mode.



## **13.4 TriCore 1.3.1 - Documentation References**

This table references all sections of the manuals that contain TriCore 1.3.1 specific information.

**Table 23 TriCore 1.3.1 Documentation References - Volume 1**

---

### **General Purpose and System Registers**

---

[ENDINIT Protection, page 3-1](#)

---

[Compatibility Mode Register \(COMPAT\), page 3-18](#)

---

[Floating Point Registers \(TriCore 1.3.1\), page 3-21](#)

---

### **Trap System**

---

[CAE - Coprocessor Trap Asynchronous Error \(TIN 4\) \(TriCore 1.3.1\), page 6-13](#)

---

[PIE - Program Memory Integrity Error \(TIN 5\) \(TriCore 1.3.1\), page 6-13](#)

---

[DIE - Data Memory Integrity Error \(TIN 6\) \(TriCore 1.3.1\), page 6-13](#)

---

[Synchronous Trap Priorities, page 6-15](#)

---

[Asynchronous Trap Priorities, page 6-16](#)

---

### **Memory Integrity Error Mitigation**

---

[Memory Integrity Error Mitigation \(TriCore 1.3.1\), page 7-1](#)

---

### **Physical Memory Attributes**

---

[Scratchpad RAM \(TriCore 1.3.1\), page 8-4](#)

---

[BIST Mode Access Control Register \(BMACON\), page 3-19](#)

---

[SIST Mode Access Control Register \(SMACON\), page 3-20](#)

---

### **Core Debug Controller**

---

[Breakpoint Trap, page 12-8](#)

---

[Multiple Breakpoint Traps \(TriCore 1.3.1\), page 12-9](#)

---

[Performance Counter Start/Stop \(TriCore 1.3.1\), page 12-11](#)

---

[None \(TriCore 1.3.1\), page 12-11](#)

---

[Suspend In Halt \(TriCore 1.3.1\), page 12-12](#)

---

[CDC Control Registers \(TriCore 1.3.1\), page 12-25](#)

---

[Core Performance Measurement and Analysis \(TriCore 1.3.1\), page 12-42](#)

---

[Performance Counter Registers \(TriCore 1.3.1\), page 12-44](#)

---

Table 23 TriCore 1.3.1 Documentation References - Volume 1

---

**Floating Point Unit**

---

[Functional Overview, page 11-1](#)

---

[Traps \(TriCore 1.3.1\), page 11-4](#)

---

[Rounding Mode Restored \(TriCore 1.3.1\), page 11-7](#)

---

[Invalid Operations and their Quiet NaN Results, page 11-10](#)

---

[Asynchronous Traps \(TriCore 1.3.1\), page 11-13](#)

---

[FPU CSFR Registers \(TriCore 1.3.1\), page 11-14](#)

---

Table 24 TriCore 1.3.1 Documentation References - Volume 2

---

[CACHEI.W](#)

---

[CACHEI.WI](#)

---

[DVINIT DVINIT.U DVINIT.B DVINIT.BU DVINIT.H DVINIT.HU](#)

---

[RET](#)

---

[RFM](#)

---

[FTOIZ](#)

---

[FTOQ31Z](#)

---

[FTOUZ](#)

---

## 14 Core Register Table

The following tables list all the TriCore® CSFRs and GPRs. The memory protection system is modular and the actual number of registers is implementation-specific.

**Table 25 General Purpose Registers (GPR)**

Register Name	Description	Address Offset
D[0]	Data Register 0.	FF00 <sub>H</sub> <sup>1)</sup>
D[1]	Data Register 1.	FF04 <sub>H</sub>
D[2]	Data Register 2.	FF08 <sub>H</sub>
D[3]	Data Register 3.	FF0C <sub>H</sub>
D[4]	Data Register 4.	FF10 <sub>H</sub>
D[5]	Data Register 5.	FF14 <sub>H</sub>
D[6]	Data Register 6.	FF18 <sub>H</sub>
D[7]	Data Register 7.	FF1C <sub>H</sub>
D[8]	Data Register 8.	FF20 <sub>H</sub>
D[9]	Data Register 9.	FF24 <sub>H</sub>
D[10]	Data Register 10.	FF28 <sub>H</sub>
D[11]	Data Register 11.	FF2C <sub>H</sub>
D[12]	Data Register 12.	FF30 <sub>H</sub>
D[13]	Data Register 13.	FF34 <sub>H</sub>
D[14]	Data Register 14.	FF38 <sub>H</sub>
D[15]	Data Register 15 - Implicit Data Register.	FF3C <sub>H</sub>
A[0]	Address Register 0 - Global Address Register.	FF80 <sub>H</sub> <sup>1)</sup>
A[1]	Address Register 1 - Global Address Register.	FF84 <sub>H</sub>
A[2]	Address Register 2.	FF88 <sub>H</sub>
A[3]	Address Register 3.	FF8C <sub>H</sub>
A[4]	Address Register 4.	FF90 <sub>H</sub>
A[5]	Address Register 5.	FF94 <sub>H</sub>
A[6]	Address Register 6.	FF98 <sub>H</sub>
A[7]	Address Register 7.	FF9C <sub>H</sub>
A[8]	Address Register 8 - Global Address Register.	FFA0 <sub>H</sub>
A[9]	Address Register 9 - Global Address Register.	FFA4 <sub>H</sub>
A[10] (SP)	Address Register 10 - Stack Pointer Register.	FFA8 <sub>H</sub>
A[11] (RA)	Address Register 11 - Return Address Register.	FFAC <sub>H</sub>
A[12]	Address Register 12.	FFB0 <sub>H</sub>
A[13]	Address Register 13.	FFB4 <sub>H</sub>
A[14]	Address Register 14.	FFB8 <sub>H</sub>
A[15]	Address Register 15 - Implicit Address Register.	FFBC <sub>H</sub>

<sup>1)</sup> These address offsets are not used by the MTCR instruction.

**Table 26 Core Special Function Registers (CSFR)**

Register Name	Description	Address Offset
PCXI	Previous Context Information Register.	FE00 <sub>H</sub>
PCX	Previous Context Pointer Register.	
PSW	Program Status Word Register.	FE04 <sub>H</sub>
PC	Program Counter Register.	FE08 <sub>H</sub>
SYSCON	System Configuration Register.	FE14 <sub>H</sub>
CPU_ID	CPU Identification Register (Read Only).	FE18 <sub>H</sub>
BIV <sup>1)</sup>	Base Address of Interrupt Vector Table Register.	FE20 <sub>H</sub>
BTIV <sup>1)</sup>	Base Address of Trap Vector Table Register.	FE24 <sub>H</sub>
ISP <sup>1)</sup>	Interrupt Stack Pointer Register.	FE28 <sub>H</sub>
ICR	ICU Interrupt Control Register.	FE2C <sub>H</sub>
FCX	Free Context List Head Pointer Register.	FE38 <sub>H</sub>
LCX	Free Context List Limit Pointer Register.	FE3C <sub>H</sub>
COMPAT <sup>1)</sup>	Compatibility Mode Register (TriCore 1.3.1)	9400 <sub>H</sub>

**Memory Protection Registers**

DPR0_0L	Data Segment Protection Register 0, Set 0, Lower.	C000 <sub>H</sub>
DPR0_0U	Data Segment Protection Register 0, Set 0, Upper.	C004 <sub>H</sub>
DPR0_1L	Data Segment Protection Register 1, Set 0, Lower.	C008 <sub>H</sub>
DPR0_1U	Data Segment Protection Register 1, Set 0, Upper.	C00C <sub>H</sub>
DPR0_2L	Data Segment Protection Register 2, Set 0, Lower.	C010 <sub>H</sub>
DPR0_2U	Data Segment Protection Register 2, Set 0, Upper.	C014 <sub>H</sub>
DPR0_3L	Data Segment Protection Register 3, Set 0, Lower.	C018 <sub>H</sub>
DPR0_3U	Data Segment Protection Register 3, Set 0, Upper.	C01C <sub>H</sub>
DPR1_0L	Data Segment Protection Register 0, Set 1, Lower.	C400 <sub>H</sub>
DPR1_0U	Data Segment Protection Register 0, Set 1, Upper.	C404 <sub>H</sub>
DPR1_1L	Data Segment Protection Register 1, Set 1, Lower.	C408 <sub>H</sub>
DPR1_1U	Data Segment Protection Register 1, Set 1, Upper.	C40C <sub>H</sub>
DPR1_2L	Data Segment Protection Register 2, Set 1, Lower.	C410 <sub>H</sub>
DPR1_2U	Data Segment Protection Register 2, Set 1, Upper.	C414 <sub>H</sub>
DPR1_3L	Data Segment Protection Register 3, Set 1, Lower.	C418 <sub>H</sub>
DPR1_3U	Data Segment Protection Register 3, Set 1, Upper.	C41C <sub>H</sub>

**Table 26 Core Special Function Registers (CSFR)**

Register Name	Description	Address Offset
DPR2_0L	Data Segment Protection Register 0, Set 2, Lower.	C800 <sub>H</sub>
DPR2_0U	Data Segment Protection Register 0, Set 2, Upper.	C804 <sub>H</sub>
DPR2_1L	Data Segment Protection Register 1, Set 2, Lower.	C808 <sub>H</sub>
DPR2_1U	Data Segment Protection Register 1, Set 2, Upper.	C80C <sub>H</sub>
DPR2_2L	Data Segment Protection Register 2, Set 2, Lower.	C810 <sub>H</sub>
DPR2_2U	Data Segment Protection Register 2, Set 2, Upper.	C814 <sub>H</sub>
DPR2_3L	Data Segment Protection Register 3, Set 2, Lower.	C818 <sub>H</sub>
DPR2_3U	Data Segment Protection Register 3, Set 2, Upper.	C81C <sub>H</sub>
DPR3_0L	Data Segment Protection Register 0, Set 3, Lower.	CC00 <sub>H</sub>
DPR3_0U	Data Segment Protection Register 0, Set 3, Upper.	CC04 <sub>H</sub>
DPR3_1L	Data Segment Protection Register 1, Set 3, Lower.	CC08 <sub>H</sub>
DPR3_1U	Data Segment Protection Register 1, Set 3, Upper.	CC0C <sub>H</sub>
DPR3_2L	Data Segment Protection Register 2, Set 3, Lower.	CC10 <sub>H</sub>
DPR3_2U	Data Segment Protection Register 2, Set 3, Upper.	CC14 <sub>H</sub>
DPR3_3L	Data Segment Protection Register 3, Set 3, Lower.	CC18 <sub>H</sub>
DPR3_3U	Data Segment Protection Register 3, Set 3, Upper.	CC1C <sub>H</sub>
CPR0_0L	Code Segment Protection Register 0, Set 0, Lower.	D000 <sub>H</sub>
CPR0_0U	Code Segment Protection Register 0, Set 0, Upper.	D004 <sub>H</sub>
CPR0_1L	Code Segment Protection Register 1, Set 0, Lower.	D008 <sub>H</sub>
CPR0_1U	Code Segment Protection Register 1, Set 0, Upper.	D00C <sub>H</sub>
CPR0_2L	Code Segment Protection Register 2, Set 0, Lower.	D010 <sub>H</sub>
CPR0_2U	Code Segment Protection Register 2, Set 0, Upper.	D014 <sub>H</sub>
CPR0_3L	Code Segment Protection Register 3, Set 0, Lower.	D018 <sub>H</sub>
CPR0_3U	Code Segment Protection Register 3, Set 0, Upper.	D01C <sub>H</sub>
CPR1_0L	Code Segment Protection Register 0, Set 1, Lower.	D400 <sub>H</sub>
CPR1_0U	Code Segment Protection Register 0, Set 1, Upper.	D404 <sub>H</sub>
CPR1_1L	Code Segment Protection Register 1, Set 1, Lower.	D408 <sub>H</sub>
CPR1_1U	Code Segment Protection Register 1, Set 1, Upper.	D40C <sub>H</sub>
CPR1_2L	Code Segment Protection Register 2, Set 1, Lower.	D410 <sub>H</sub>
CPR1_2U	Code Segment Protection Register 2, Set 1, Upper.	D414 <sub>H</sub>
CPR1_3L	Code Segment Protection Register 3, Set 1, Lower.	D418 <sub>H</sub>
CPR1_3U	Code Segment Protection Register 3, Set 1, Upper.	D41C <sub>H</sub>

**Table 26 Core Special Function Registers (CSFR)**

<b>Register Name</b>	<b>Description</b>	<b>Address Offset</b>
CPR2_0L	Code Segment Protection Register 0, Set 2, Lower.	D800 <sub>H</sub>
CPR2_0U	Code Segment Protection Register 0, Set 2, Upper.	D804 <sub>H</sub>
CPR2_1L	Code Segment Protection Register 1, Set 2, Lower.	D808 <sub>H</sub>
CPR2_1U	Code Segment Protection Register 1, Set 2, Upper.	D80C <sub>H</sub>
CPR2_2L	Code Segment Protection Register 2, Set 2, Lower.	D810 <sub>H</sub>
CPR2_2U	Code Segment Protection Register 2, Set 2, Upper.	D814 <sub>H</sub>
CPR2_3L	Code Segment Protection Register 3, Set 2, Lower.	D818 <sub>H</sub>
CPR2_3U	Code Segment Protection Register 3, Set 2, Upper.	D81C <sub>H</sub>
CPR3_0L	Code Segment Protection Register 0, Set 3, Lower.	DC00 <sub>H</sub>
CPR3_0U	Code Segment Protection Register 0, Set 3, Upper.	DC04 <sub>H</sub>
CPR3_1L	Code Segment Protection Register 1, Set 3, Lower.	DC08 <sub>H</sub>
CPR3_1U	Code Segment Protection Register 1, Set 3, Upper.	DC0C <sub>H</sub>
CPR3_2L	Code Segment Protection Register 2, Set 3, Lower.	DC10 <sub>H</sub>
CPR3_2U	Code Segment Protection Register 2, Set 3, Upper.	DC14 <sub>H</sub>
CPR3_3L	Code Segment Protection Register 3, Set 3, Lower.	DC18 <sub>H</sub>
CPR3_3U	Code Segment Protection Register 3, Set 3, Upper.	DC1C <sub>H</sub>
DPM0	Data Protection Mode Register 0.	E000 <sub>H</sub>
DPM1	Data Protection Mode Register 1.	E080 <sub>H</sub>
DPM2	Data Protection Mode Register 2.	E100 <sub>H</sub>
DPM3	Data Protection Mode Register 3.	E180 <sub>H</sub>
CPM0	Code Protection Mode Register 0.	E200 <sub>H</sub>
CPM1	Code Protection Mode Register 1.	E280 <sub>H</sub>
CPM2	Code Protection Mode Register 2.	E300 <sub>H</sub>
CPM3	Code Protection Mode Register 3.	E380 <sub>H</sub>
<b>Memory Management Registers</b>		
MMU_CON	Memory Management Unit Configuration Register.	8000 <sub>H</sub>
MMU_ASI	MMU Address Space Identifier Register.	8004 <sub>H</sub>
MMU_TVA	MMU Translation Virtual Address Register.	800C <sub>H</sub>
MMU_TPA	MMU Translation Physical Address Register.	8010 <sub>H</sub>
MMU_TPX	MMU Translation Physical Index Register.	8014 <sub>H</sub>
MMU_TFA	MMU Translation Fault Address Register.	8018 <sub>H</sub>
BMACON <sup>(1)</sup>	BIST Mode Control Register. (TriCore 1.3.1)	9004 <sub>H</sub>
SMACON <sup>(1)</sup>	SIST mode Control Register. (TriCore 1.3.1)	900C <sub>H</sub>

**Table 26 Core Special Function Registers (CSFR)**

<b>Register Name</b>	<b>Description</b>	<b>Address Offset</b>
DIEAR	Data Integrity Error Address Register. (TriCore 1.3.1)	9020 <sub>H</sub>
DIETR	Data Integrity Error Trap Register. (TriCore 1.3.1)	9024 <sub>H</sub>
CCDIER	Count Corrected Data Integrity Errors. (TriCore 1.3.1)	9028 <sub>H</sub>
MIECON <sup>1)</sup>	Memory Integrity Error Control Register (TriCore 1.3.1)	9044 <sub>H</sub>
PIEAR	Program Integrity Error Address Register. (TriCore 1.3.1)	9210 <sub>H</sub>
PIETR	Program Integrity Error Trap Register. (TriCore 1.3.1)	9214 <sub>H</sub>
CCPIER	Count Corrected Program Integrity Errors. (TriCore 1.3.1)	9218 <sub>H</sub>

**Debug Registers**

DBGSR	Debug Status Register.	FD00 <sub>H</sub>
EXEVT	External Event Register.	FD08 <sub>H</sub>
CREVT	Core Register Event Register.	FD0C <sub>H</sub>
SWEVT	Software Event Register.	FD10 <sub>H</sub>
TR0EVT	Trigger Event 0 Register.	FD20 <sub>H</sub>
TR1EVT	Trigger Event 1 Register.	FD24 <sub>H</sub>
DMS	Debug Monitor Start Address Register.	FD40 <sub>H</sub>
DCX	Debug Context Save Address Register.	FD44 <sub>H</sub>
DBGTCR	Debug Trap Control Register. (TriCore 1.3.1)	FD48 <sub>H</sub>
CCTRL	Counter Control Register (TriCore 1.3.1)	FC00
CCNT	CPU Clock Count Register (TriCore 1.3.1)	FC04
ICNT	Instruction Count Register (TriCore 1.3.1)	FC08
M1CNT	Multi Count Register 1 (TriCore 1.3.1)	FC0C
M2CNT	Multi Count Register 2 (TriCore 1.3.1)	FC10
M3CNT	Multi Count Register 3 (TriCore 1.3.1)	FC14

**Floating Point Registers**

FPU_TRAP_CON	Trap Control Register. (TriCore 1.3.1)	A000 <sub>H</sub>
FPU_TRAP_PC	Trapping Instruction Program Control Register. (TriCore 1.3.1)	A004 <sub>H</sub>
FPU_TRAP_OPC	Trapping Instruction Opcode Register. (TriCore 1.3.1)	A008 <sub>H</sub>

**Core Register Table**

**Table 26 Core Special Function Registers (CSFR)**

<b>Register Name</b>	<b>Description</b>	<b>Address Offset</b>
FPU_TRAP_SRC1	Trapping Instruction SRC1 Operand Register. (TriCore 1.3.1)	A010 <sub>H</sub>
FPU_TRAP_SRC2	Trapping Instruction SRC2 Operand Register. (TriCore 1.3.1)	A014 <sub>H</sub>
FPU_TRAP_SRC3	Trapping Instruction SRC3 Operand Register. (TriCore 1.3.1)	A018 <sub>H</sub>
FPU_ID	FPU Identification Register. (TriCore 1.3.1)	A020 <sub>H</sub>

<sup>1)</sup> These registers are ENDINIT protected.

**Table 27 Special Function Registers Associated with the Core<sup>1)</sup>**

CPU_SRC0	CPU Service Request Control Register 0.	FFFC <sub>H</sub>
CPU_SRC1	CPU Service Request Control Register 1.	FFF8 <sub>H</sub>
CPU_SRC2	CPU Service Request Control Register 2.	FFF4 <sub>H</sub>
CPU_SRC3	CPU Service Request Control Register 3.	FFF0 <sub>H</sub>
CPU_SBSRC0	CPU Software Break Service Request Control Register 0.	FFBC <sub>H</sub>
CPU_SBSRC1 <sup>2)</sup>	CPU Software Break Service Request Control Register 1.	FFB8 <sub>H</sub>
CPU_SBSRC2 <sup>2)</sup>	CPU Software Break Service Request Control Register 2.	FFB4 <sub>H</sub>
CPU_SBSRC3 <sup>2)</sup>	CPU Software Break Service Request Control Register 3.	FFB0 <sub>H</sub>

<sup>1)</sup> These address offsets are calculated from a different base address to core registers. These registers cannot be accessed using the MTCR and MFCR instructions.

<sup>2)</sup> If implemented.



## 15 List of Registers

A[10](SP) .....	3-14	M1CNT .....	12-49
An (n = 0-15) .....	3-3	M2CNT .....	12-50
BIV .....	6-19	M3CNT .....	12-51
BMACON .....	3-19	MIECON .....	7-9
BTv .....	6-20	MMU_ASI .....	10-15
CCDIER .....	7-4	MMU_CON .....	10-13
CCNT .....	12-47	MMU_TFA .....	10-20
CCPIER .....	7-3	MMU_TPA .....	10-17
CCTRL .....	12-45	MMU_TPX .....	10-19
COMPAT .....	3-18	MMU_TVA .....	10-16
CPMx .....	9-12	module_SRCn .....	5-3
CPRx_nL .....	9-9	PC .....	3-5
CPRx_nU .....	9-8	PCX .....	4-15
CPU_ID .....	3-17	PCXI, PCX .....	3-12
CREVT .....	12-18	PIEAR .....	7-6
DBGSR .....	12-15	PIETR .....	7-5
DBGSR .....	12-25	PSW .....	3-6
DBGTCR .....	12-39	SBSRCn (n = 0 to 3) .....	12-40
DCX .....	12-24	SMACON .....	3-20
DIEAR .....	7-8	SWEVT .....	12-19
DIETR .....	7-7	SWEVT .....	12-31
DMS .....	12-23	SYSCON .....	3-16
DMS .....	12-37	TR0EVT .....	12-20
Dn (n = 0-15) .....	3-3	TR0EVT .....	12-33
DPMx .....	9-10	TR1EVT .....	12-20
DPRx_mL .....	9-7	TR1EVT .....	12-33
DPRx_mU .....	9-6		
EXEVT .....	12-17		
EXEVT .....	12-27		
FCX .....	4-14		
FPU_ID .....	11-22		
FPU_TRAP_CON .....	11-14		
FPU_TRAP_OPC .....	11-18		
FPU_TRAP_PC .....	11-17		
FPU_TRAP_SRC1 .....	11-19		
FPU_TRAP_SRC2 .....	11-20		
FPU_TRAP_SRC3 .....	11-21		
ICNT .....	12-48		
ICR .....	6-17		
ISP .....	3-15		
LCX .....	4-16		



## Index

### A

A0, A1, A8, A9

System Global Registers

GPRs 3-2

Overview 1-4

A0-A15

Address Registers 14-1

A10

Stack Pointer 3-13

Absolute

Addressing 2-8

Access Privilege 3-7

Address

Absolute 2-14

Array 2-12

Base Address of Vector Table 6-19

Code 2-14

Definition 2-2

Displacement 2-6

Effective 2-12, 4-13

General Purpose Registers 3-2

Half-word 6-20

Map 1-5

Physical Memory Attributes 8-3

Mapping 1-4

Multiple Address Spaces 10-5

Physical Memory 4-5

Range 9-16

Ranges 4-5

Register 2-14

Use with GPRs 3-2

Register A10 3-13

Return Address A11 3-2

Space 1-1, 1-4

Space Identifier (ASI) 10-1, 10-5

Spaces 10-2

Width 2-6

Address Offset

List of offsets 14-1

Address Register

Definition 3-14

Address Translation 10-3

Context Pointers 10-3

MMU\_CON 10-3

PPN 10-3

PTE 10-3

VPN 10-3

Addressing

Absolute 2-8

Address Register 2-9

Base + Offset 2-8

Bit Indexed 2-13

Bit Reverse 2-12

Circular 2-9

Indexed 2-13

Modes 1-6, 2-7

Programming Model 2-1, 2-7

Synthesized 2-13

PC-relative 2-14

Post-decrement 2-9

Post-increment 2-9

Pre-decrement 2-9

Pre-Increment 2-9

Synthesized 2-13

ADDSC.A Instruction

Indexed Addressing 2-13

ADDSC.AT 2-13

Alignment

Requirements 2-4

Rules 2-4

Trap 2-11

ALN Trap

Data Address Alignment 6-9

Arbitration

Scheme 5-8

Architectural Registers 1-3

Architecture

Addressing Data 2-14

Overview 1-1

Traps 6-1

Array

Index 2-12

ASI

Address Space Identifier 10-1, 10-5  
 Field in MMU\_ASI Register 10-15  
 Field in MMU\_TVA Register 10-16  
 Field in TRnEVT Register 12-20, 12-33  
 ASI\_EN  
 Field in TRnEVT Register 12-20, 12-33  
 Assertion Traps 6-14  
 Associativity (of TLB) 10-4  
 Asynchronous Traps 6-3, 11-1, 11-13  
 Atomic Operations 2-7  
 Automatic Switch  
 Stack Management 3-13

## **B**

BAM Trap  
 Break After Make 6-14  
 Priority of Debug Events 12-12  
 Base  
 + Offset Addressing 2-8  
 Address 2-12  
 Register 2-14  
 Base + Offset Addressing 2-8  
 BBM  
 Debug Halt Action 12-8  
 Field in CREVT Register 12-18, 12-30  
 Field in EXEVT Register 12-17, 12-28  
 Field in SWEVT Register 12-19, 12-32  
 Field in TRnEVT Register 12-22, 12-36  
 Priority of Debug Events 12-12  
 Trap  
 Break Before Make 6-14  
 BISR 4-4  
 BISR Instruction  
 Context Switching with Interrupts 4-7  
 Bit  
 Bit-Reverse Addressing 2-12  
 Enable and Disable 6-17  
 String 2-1  
 Type Abbreviations 1-2  
 Bit Type  
 Abbreviations in Tables  
 Definitions 1-2  
 Bit-Reverse Addressing 2-12

Bit-Reversed Order 2-12  
 BIV  
 Interrupt Vector Table Location 5-11  
 Register  
 Address Offset 14-2  
 Definition 6-19  
 Interrupt and Trap Handling 6-17

## **BL**

Field in CPMx Register 9-13  
 BMACON 3-19

Address Offset 14-4

## **Boolean**

Programming Model 2-1

## **Breakpoint**

CDC Features 12-1  
 Interrupt Debug Action 12-10  
 Trap 12-8

## **BTV**

Base Trap Vector Table Pointer 6-20  
 Register  
 Address Offset 14-2  
 Definition 6-20  
 Interrupt and Trap Handling 6-17

## **BU**

Field in CPMx Register 9-13

## **Buffer**

Aligned to a 64-bit Boundary 2-11  
 Size 2-13  
 Start 2-11

## **Byte**

Definition 1-2  
 Indices 2-13  
 Offset 2-10  
 Ordering 2-5

## **C**

Cacheability 10-7  
 Cacheability Bit (C)  
 TLB Table Entry Contents 10-5  
 Cacheable (C)  
 Physical Memory Address Properties  
 8-1  
 Cacheable Memory

Physical Memory Attribute 8-3  
 Call Depth Counter  
   CSAs and Context Lists 4-6  
 CALL Instruction  
   Context Switching & Calls 4-8  
 Calling and Called Functions 4-8  
 CCDIER  
   Address Offset 14-5  
 CCNT 12-47  
   Address Offset 14-5  
 CCPIER 7-3  
   Address Offset 14-5  
 CCPN  
   CPU Priority  
     Interrupt Priority Groups 5-12  
   Current CPU Priority Number 6-17  
   Field in ICR Register 6-18  
 CCTRL 12-42, 12-45  
   Address Offset 14-5  
 CCTRL.CM 12-43  
 CDC  
   Combining Debug Triggers 12-6  
   Control Registers 12-14  
   Core Debug Controller 12-1  
   Debug Triggers 12-5  
   Enabling 12-1  
   Features 12-1  
   Memory Protection System 9-1  
 CDE  
   Field in PSW Register 3-8  
 CDO Trap  
   Call Depth Overflow 6-11  
 CDU Trap  
   Call Depth Underflow 6-11  
 Circular  
   Addressing 2-9, 2-10  
   Buffer  
     Circular Addressing 2-9, 2-10  
     End Case 2-11  
     Restrictions 2-11  
 CLRR  
   Description 5-4  
   Field in SBSRC Register 12-40

Field in SRC Register 5-3  
 Code  
   Address 2-14  
   Fetch (F)  
     Physical Memory Address  
     Properties 8-2  
   Protection Mode (CPM) Register 12-6  
 Code Protection Mode (CPM) Register  
   Address Offset 14-4  
 COMPAT  
   Compatibility Register 14-2  
 Compatibility Mode Register 3-18  
 Context  
   Current 4-7  
   Information Register 3-12  
 List  
   Context Restore 4-11  
   Description 4-5  
   Previous 4-5  
 List Management  
   CTYP Trap 6-11  
 Lower 4-1  
 Lower Context  
   Context Restore 4-12  
   PCXI Register Field 3-12  
   Registers 3-4  
   Task Switching Operation 4-4  
 Management Registers 4-13  
 Management Traps 6-10  
 Of Task 1-7, 3-11  
 Pointers  
   Address Translation 10-3  
 Restore  
   CTYP Trap 6-11  
   Example 4-9  
   Operation 4-11  
 Save 4-9  
   Example 4-9  
   FCU Trap 6-11  
   Operation 4-6, 4-9  
 Switching 1-7  
   With Function Calls 4-8  
   With Interrupts 4-7

- Upper 4-1
- Upper Context
  - Registers 3-4
  - Task Switching Operation 4-4
  - UL Field in PCXI Register 3-12
- Context Save Area (CSA)
  - Context Lists 4-5
  - Context Management Registers 4-13
  - Description 4-3
  - Lower Context 1-7
  - Upper and Lower Contexts 4-1
- Coprocessor 1-11
- Core
  - Break-Out Signal 12-8
  - Debug Controller (CDC) 12-1
    - Registers 3-21
  - Special Function Registers (CSFRs)
    - Core Registers 1-4, 3-1
  - Suspend-Out Signal 12-8
- Core Register Table 14-1
- Core Special Function Registers 14-2
- Corrected Memory Integrity Errors 7-3
- Counters
  - Normal Mode 12-43
  - Task Mode 12-43
- CPM
  - Code Protection Mode Register 9-12
  - Combining Debug Triggers 12-6
- CPR
  - Code Segment Protection (CPR)
    - Register
      - Address Offset 14-3
- CPRx\_nL
  - Code Segment Protection Register
    - Lower Bound 9-9
- CPRx\_nU
  - Code Segment Protection Register
    - Upper Bound 9-8
- CPU
  - Current Priority Number 5-9
- CPU\_ID
  - CPU Identification Register
    - Address Offset 14-2
- CPU\_SBSRC
  - CPU Software Break Service Request
    - Control Register
      - Definition 12-40
- CPU\_SBSRC0
  - Address Offset 14-6
- CPU\_SBSRC1
  - Address Offset 14-6
- CPU\_SBSRC2
  - Address Offset 14-6
- CPU\_SBSRC3
  - Address Offset 14-6
- CPU\_SRC0
  - Address Offset 14-6
- CPU\_SRC1
  - Address Offset 14-6
- CPU\_SRC2
  - Address Offset 14-6
- CPU\_SRC3
  - Address Offset 14-6
- CREVT
  - Address Offset 14-5
  - Core Register Access Event Register
    - Definition 12-18, 12-29
- CSA
  - Context Lists 4-5
  - Context Save Area
    - Description 4-3
    - Lower Context 1-7
    - Upper and Lower Contexts 4-1
  - Effective Address 4-3
  - List Head Pointer 4-13
  - List Limit Pointer 4-13
  - List Underflow 4-16
- CSA memory location 4-17
- CSFR
  - Core Registers 1-4, 3-1
  - MMU 10-13
  - Register Table 14-1
- CSU Trap
  - Call Stack Underflow 6-11
- CTYP Trap
  - Context Type 6-11

## D

- D0-D15 Data Registers 14-1
- DAE Trap
  - Data Access Asynchronous Error 6-12
- Data
  - Data Registers (D0 to D15) 3-2
  - DPR Data Segment Protection Register
    - Address Offset 14-2
  - Formats 2-1, 2-2
  - General Purpose Registers 3-2
  - Memory 2-14
  - Protection Mode Register (DPM) 12-6
  - Size 2-11
  - Types 2-1
    - List of 1-4
  - Values
    - Circular Addressing 2-9
- Data Access
  - Cacheable and Speculative Properties 8-2
  - Physical Memory Address Properties 8-2
- Data Protection Mode Register 9-10
- Data Protection Mode Register (DPM)
  - Address Offset 14-4
- Data Segment Protection Registers 9-15
- DBGSR
  - Address Offset 14-5
  - Debug Status Register
    - CDC Control Registers 12-14
    - Definition 12-15, 12-25
  - Enabling CDC 12-1
- DBGTCR
  - Address Offset 14-5
- DCX
  - Address Offset 14-5
  - Debug Context Save Area Pointer Register
    - Definition 12-24, 12-38
  - Value
    - Field in DCX Register 12-24

## DE

- Field in DBGSR Register 12-16, 12-26
- Debug
  - Monitor Start Address Register (DMS) Breakpoint Trap 12-8
  - System 1-11
  - Traps 6-14
- Debug Action
  - Description 12-7
  - EXEVT 12-7
  - Halt 12-8
  - Run Control Features 12-1
  - TRnEVT 12-4
- Debug Event 12-1
  - Description 12-3
  - External 12-3
  - MTCR and MFCR 12-3
  - Priority 12-12
- DEBUG Instruction 12-2, 12-3
- Debug Monitor Start Address Register (DMS) 12-8
- Debug Registers 14-5
- Debug Triggers 12-5
  - Combining 12-6
- Debugging
  - Registers that support 3-21
- Denormal Numbers 11-3
- DIEAR
  - Address Offset 14-5
- DIETR
  - Address Offset 14-5
- Direct Memory Access (DMA) 1-8
- Direct Translation
  - Description 1-10
  - Memory Protection System 9-2
  - MMU 10-1
  - Permitted Versus Valid Accesses 8-6
  - Virtual Mode Protection 10-7
- DLR\_LR
  - Field in TRnEVT Register 12-21, 12-34
- DLR\_U
  - Field in TRnEVT Register 12-21, 12-34
- DMA

- Direct Memory Access 1-8
- DMS 12-23, 12-37
  - Address Offset 14-5
  - Debug Monitor Start Address Register
  - Breakpoint Trap 12-8
  - Value
    - Field in DMS Register 12-23
- Double-word
  - Accesses 2-4
  - Definition 1-2
- DPM
  - Data Protection Mode Register
    - Combining Debug Triggers 12-6
    - Definition 9-10
- DPR
  - Data Segment Protection Register 14-2
    - Definition 9-6, 9-7
  - DPRx Register 9-7
- DSE Trap
  - Data Access Synchronous Error 6-12
- DSPR
  - Data scratchpad RAM 8-4
  - Data Scratchpad Register 3-19
- DSYNC 4-17
- DU\_LR
  - Field in TRnEVT Register 12-21, 12-34
- DU\_U
  - Field in TRnEVT Register 12-21, 12-34
- E**
- EA
  - Effective Address 4-3
- Effective Address
  - Context Save Area (CSA) 4-3, 4-13
- Emulator Space
  - Physical Memory Attribute 8-3
- ENABLE Instruction 5-9
- Endianess 2-5
- ENDINIT protected 14-6
- EVT 12-39
- EVTa
  - Field in CREVT Register 12-18, 12-30
- Field in EXEVT Register 12-17, 12-28
- Field in SWEVT Register 12-19, 12-32
- Field in TRnEVT Register 12-22, 12-36
- EVTSRC
  - Field in DBGSR Register 12-15, 12-25
- Exceptions
  - Floating Point Exception Flags 11-8
- EXEVT
  - Address Offset 14-5
  - Register Definition 12-17, 12-27
- Extended-Size Registers 3-2
- EXTR.U 2-13
- F**
- FCD Trap 4-16
  - Free Context List Depletion 6-10
- FCU Trap
  - Free Context List Underflow 6-11
- FCX
  - Context Management Register 4-13
  - Context Restore 4-11
  - CSAs and Context Lists 4-5
  - Free Context List
    - Context Save Description 4-9
  - Free CSA List Head Pointer Register
    - Definition 4-14
  - Offset Address 4-14
  - Pointer 4-14
  - Register
    - Address Offset 14-2
    - Definition 4-14
    - FCU Trap 6-11
  - Segment Address Field 4-14
- FCXO
  - FCX Offset Address
    - Field in FCX Register 4-14
- Feature Summary
  - TriCore 1-2
- FFT
  - Algorithms 2-12
  - Bit-Reverse Addressing 2-13
- FI
  - FPU Exception Flag 11-9



- Filter Calculations 2-9
- Floating Point
  - Denormal Numbers 11-3
  - Exception Flags 11-8
  - Registers 3-2
  - Unit (FPU) 11-1
- Floating Point Registers 14-5
- Floating Point Unit (FPU) 11-1
- FPN
  - Field in MMU\_TFA Register 10-20
- FPU 11-1
  - Denormal Numbers 11-3
  - Exception Flags 11-8
  - Exceptions 11-8
  - FI Exception Flag 11-9
  - Floating Point Unit 11-1
  - FS Exception Flag 11-9
  - FU Exception Flag 11-12
  - FV Exception Flag 11-11
  - FX Exception Flag 11-12
  - FZ Exception Flag 11-12
  - Identification Register 11-22
  - IEEE-754 11-1
  - Invalid Operations 11-10
  - NaN 11-3
  - Rounding 11-6
  - Trap Control Register 11-14
- FPU\_ID
  - Address Offset 14-6
- FPU\_TRAP\_CON
  - Address Offset 14-5
- FPU\_TRAP\_OPC
  - Address Offset 14-5
- FPU\_TRAP\_PC
  - Address Offset 14-5
- FPU\_TRAP\_SCR1
  - Address Offset 14-6
- FPU\_TRAP\_SCR2
  - Address Offset 14-6
- FPU\_TRAP\_SCR3
  - Address Offset 14-6
- Free Context Depletion
  - CSA List Underflows 4-16
- Free Context List
  - Available CSA 4-5
  - Context Restore 4-11
  - Context Save 4-9
  - FCD Trap 6-10
  - Free CSA 4-6
- FS
  - FPU Exception Flag 11-9
- FU
  - FPU Exception Flag 11-12
- Function Call 4-8
  - Context Switching 4-8
- FV
  - FPU Exception Flag 11-11
- FX
  - FPU Exception Flag 11-12
- FZ
  - FPU Exception Flag 11-12
- G**
- GByte
  - Definition 1-2
- General Purpose Registers 3-1, 14-1, 14-2
- Global
  - Data 2-8
  - Register Write Permission 5-9
  - Registers 3-8
- Global bit
  - TLBMAP 10-9
- Global bit (G)
  - TLB Table Entry Contents 10-5
- GPR
  - 16-bit Instructions 3-2
  - Core Registers 1-3
  - General Purpose Registers
    - Data Formats 2-2
    - Overview 1-3
  - Register Table 3-4, 14-1
- GRWP Trap
  - Global Register Write Protection 6-8
- GW
  - Field in PSW Register 3-8

## H

- h
  - Definition 1-2
- Half-word
  - Boundary
    - Alignment Requirements 2-4
  - Definition 1-2
- HALT
  - Field in DBGSR Register 12-16, 12-26
- Halt
  - Debug Action 12-8
- Hardware Traps 6-3

## I

- ICNT 12-48
  - Address Offset 14-5
- ICR
  - Initial State upon a Trap 6-6
  - Interrupt Control Register
    - Address Offset 14-2
    - Definition 6-17
    - Description 5-7
- ICU
  - Interrupt Control Unit
    - Description 5-6
    - Interrupt Priority 1-8
    - Operation 5-7
- ID Registers 3-17
- IEEE-754 2-2
  - FPU 11-1
  - Single Precision Floating Point
    - Number 2-2
- Implicit
  - Address Register 1-3
  - Data Register 1-3
- INDEX
  - Field in MMU\_TPX Register 10-19
- Index
  - Algorithm 2-10
  - Array 2-12
  - Bit-Reverse 2-13
  - Modifier 2-12

- Indexed
  - Addressing 2-13
  - Arrays 2-13
- Indexed Addressing
  - Scaled Data Register 2-13
- Indexes
  - Table Indexes
    - GPRs 3-2
- Instruction
  - Load Double-word 2-11
  - Load Word 2-11
  - On-chip
    - PC-Relative Addressing 2-14
  - Word 2-10
- Instruction Fetch 9-3
- Instruction Formats 2-8
- Instruction Set Architecture (ISA)
  - Features 1-2
- Integers 2-2
- Internal Buffer
  - Context Restore 4-11
- Interrupt
  - Control Register 6-17
    - Definition 6-17
  - Enable 4-7
  - Enable/Disable Bit 5-7
  - Handler 4-4, 4-7
  - Interrupt Control Unit (ICU)
    - Interrupt Priority 1-8
  - Nested 1-8
  - Priority 1-8
  - Priority Groups 5-12
  - Register A11 3-2
  - Request
    - Priority Numbers 5-13
  - Requests 5-1
    - Priority 5-6
  - Service Routine (ISR) 1-7, 3-11, 3-13, 4-7
  - Signal 5-1
  - Software-Posted Interrupts 5-15
  - Stack Management 3-13
  - Stack Pointer 3-13

- Vector Table 6-17, 6-19
- Interrupt Control Register 5-7
  - Context Switching with Interrupts 4-7
- Interrupt Control Unit (ICU) 5-6
- Interrupt Service
  - Request 5-8
  - Request Node 5-1
- Interrupt Service Routine (ISR)
  - Dividing into Priorities 5-14
  - Entering an ISR 5-8
  - Exiting an ISR 5-9
  - Stack Management 3-13
- Interrupt Stack Control 3-7
- Interrupt System
  - Chapter 5-1
  - Description 1-8
  - Service Request Enable 5-5
  - Service Request Flag (SRR) 5-4
  - Service Request Priority Number (SRPN) 5-6
  - Type-of-Service Control (TOS) 5-5
  - Typical Block Diagram 5-2
  - Using the Interrupt System 5-12
- Interrupt-1 5-15
- IO
  - I/O Privilege
    - Field in PSW Register 3-7
- IOPC Trap
  - Illegal Opcode 6-8
- IS
  - Interrupt Stack Control
    - Field in PSW Register 3-7
- ISA
  - Feature Summary 1-2
- ISP
  - Initialize 3-13
  - Interrupt Stack Pointer Register
    - Address Offset 14-2
  - Interrupt Stack Pointer Register
    - Definition 3-15
- ISR
  - Entering an ISR 5-8
  - Exiting an ISR 5-9

- Splitting on to Different Priorities 5-14
- Stack Management 3-13
- Tasks and Contexts 1-7, 3-11
- ISYNC Instruction
  - Entering an ISR 5-9
  - TLBMAP 10-10
- ISYNC instruction 3-22

## J

- Jump and Link
  - Instruction
    - PC-Relative Addressing 2-14

## K

- KByte
  - Definition 1-2

## L

- LCX
  - Context Management Registers 4-13
  - FCD Trap 6-10
  - Free CSA List Limit Pointer Register
    - Address Offset 14-2
    - Definition 4-16
    - Offset 4-16
    - Segment Address 4-16
- LDMST 2-7
- LDMST Instruction
  - Alignment Requirements 2-4
- LEA
  - Load Effective Address
    - PC-Relative Addressing 2-14
- Link Word
  - Context Restore Example 4-11
  - Context Save Areas (CSAs) 4-5
  - Context Save Example 4-10
  - CSA 4-3
  - Lower Context and CSAs 1-7
- Little-Endian 2-5
- Load
  - Effective Address (LEA)
    - PC-Relative Addressing 2-14
  - Task Switching Operations 4-4

Word 2-10

Local

- Variables 2-8

LOWBND

- Field in CPRx\_nL Register 9-9
- Field in DPRx\_nL Register 9-7

Lower Context 4-1

- PCXI Register Field 3-12
- Registers 3-4
- Task Switching Operation 4-4

**M**

M1CNT 12-49

- Address Offset 14-5

M2CNT 12-50

- Address Offset 14-5

M3CNT 12-51

- Address Offset 14-5

MByte

- Definition 1-2

MEM Trap

- Invalid Local Memory Address 6-9

Memory

- Access
  - Circular Addressing 2-10
  - Permitted versus Valid 8-5
- Management 3-21
  - TLB Description 10-4
- Management Unit (MMU)
  - Architecture Overview 1-10
- Management Unit Registers 3-21
- Memory Protection Enable (SYSCON.PROTEN) 3-16
- Model 1-5, 2-6
  - Description 1-4, 2-6
  - Programming Model Overview 2-1
- Protection
  - Model 9-7
- Protection Model 9-6
- Protection Register Sets 9-3
- Protection Registers 9-2
  - Active Set 3-6
  - Overview 3-21

PSW.PRS Field 3-6

Protection System 9-1

- Using 9-16

Memory Integrity Error

- Classification 7-1
- Data 7-2
- Mitigation 7-1
- Program 7-2

Memory Management Registers 14-4

Memory Management Unit

- MMU Chapter 10-1, 11-1

Memory Protection Registers 14-2

- Description 3-21

Memory Protection System 9-1

MFCR Instruction

- Debug Events 12-3
- Reading MMU CSFRs 10-13
- Run-Control Features 12-2

MHz

- Definition 1-2

MIECON

- Address Offset 14-5

MMU 10-1

- Architecture Overview 1-10
- Instructions 10-8
- Physically Present 10-8
- Protection System 1-10, 9-1
- Traps 10-5

MMU Configuration Register 10-13

MMU Traps 6-7

MMU\_ASI

- Address Offset 14-4
- Address Space Identifier Register
- Definition 10-15

MMU\_CON 10-13

- Address Offset 14-4
- Address Translation 10-3
- MMU Configuration Register 10-13

MMU\_TFA

- Address Offset 14-4
- Translation Fault Page Address
- Register Definition 10-20

MMU\_TPA

- Address Offset 14-4
- Translation Physical Address Register Definition 10-17
- MMU\_TPX
  - Address Offset 14-4
  - Translation Page Index Register Definition 10-19
- MMU\_TVA
  - Address Offset 14-4
  - Translation Virtual Address Register Definition 10-16
- Mode
  - Supervisor 1-7, 3-11
  - User-0 1-7, 3-11
  - User-1 1-7, 3-11
- Module Identification Number
  - CPU\_ID.MOD Field 3-17, 3-18, 11-22
- MPN Trap
  - Memory Protection Null Address 6-8
- MPP Trap
  - Memory Protection Access 6-8
- MPR Trap 9-17
  - Memory Protection Read 6-7
- MPW Trap 9-17
  - Memory Protection Write 6-8
- MPX Trap 9-17
  - Memory Protection Execute 6-8
- MTCR Instruction
  - Debug Events 12-3
  - ICR.CCPN Update 6-18
  - MMU CSFRs 10-13
  - Modifying ICR.IE and ICR.CCPN 5-9
  - Run Control Features 12-2
  - Writing to the BIV Register 5-11
- MTCR instruction 10-13
- MTCR update 3-22
- N**
  - Negative Logic
    - Text Conventions 1-2
  - NEST Trap
    - Nesting Error Description 6-12
  - Nesting
    - Ranges PRS Usage Example 9-14
  - NMI
    - Asynchronous Traps 6-3
    - Description 6-14
    - Non-Maskable Interrupt
      - Trap Class 6-3
    - Trap
      - Non-Maskable Interrupt 6-14
    - Trap System
      - Architecture Overview 1-9, 9-1
      - Trap System Overview 6-1
  - NOMMU
    - Field in MMU\_CON Register 10-14
  - Non-Cacheable Memory
    - Physical Memory Attribute 8-3
  - Normal Mode 12-43
- O**
  - OCDS
    - Control Registers 12-14
  - On-Chip Instruction
    - PC-Relative Addressing 2-14
  - OPD Trap
    - Invalid Operand 6-9
  - Overflow
    - Arithmetic Overflow
      - OVF Trap 6-2
  - OVF Trap
    - Arithmetic Overflow 6-14
- P**
  - Packed
    - Arithmetic in DSP 2-4
  - Page Mapping
    - TLB Map 10-9
  - Page Table Entry (PTE)
    - Memory Protection System 9-2
    - Virtual Address Translation 1-10, 10-1
  - PC
    - Program Counter Register
      - Address Offset 14-2

- Architectural Registers 1-3
- Architecture Overview 1-3
- Definition 3-5
- Register A11 3-2
- Relative Addressing 2-14
- PCX
  - Context Management Registers 4-13
  - Context Restore 4-11
  - Context Save 4-9
  - CSU Trap 6-11
  - Offset 4-15
  - Previous Context Pointer Register 14-2
    - Definition 4-15
  - Segment Address 4-15
- PCXI
  - Architectural Registers 1-3
  - Architecture Overview 1-3
  - Exiting an Interrupt Service Routine 5-9
  - Previous Context Information Register
    - Address Offset 14-2
    - Definition 3-12
  - Task Switching Operation 4-4
- PCXO
  - Previous Context Pointer Offset
    - Field in PCXI Register 3-12
- PCXS
  - PCX Segment Address
    - Field in PCXI Register 3-12
- Pending
  - Interrupt Priority Number (PIPn)
    - Context Switching - Interrupts 4-7
    - Entering an ISR 5-9
    - Interrupt Control Register 6-17
- Peripheral
  - Registers 2-8
- Peripheral Space
  - Physical Memory Attribute 8-3
- PEVT
  - Field in DBGSR Register 12-15, 12-25
- Physical Address Space
  - Memory Management Unit 10-1
  - Memory Model 2-6
  - Physical Memory Attributes 8-3
- Physical Memory Attributes (PMA) 8-1
- Physical Memory Attributes for all Segments 8-3
- Physical Memory Properties (PMP)
  - Cacheable (C) 8-1
  - Code Fetch (F) 8-1
  - Data Access (D) 8-1
  - Description 8-1
  - Privileged Peripheral (P) 8-1
  - Speculative (S) 8-1
- PIEAR
  - Address Offset 14-5
- PIETR
  - Address Offset 14-5
- PIPn
  - Field in ICR Register 6-17
  - ICU Operation 5-7
  - Used with BIV Register 6-19
- PMA
  - Description 8-1
  - Memory Protection System 9-16
  - Physical Memory Attributes 8-3
- PMP
  - Description 8-1
- Pointer
  - Interrupt Vector Table 6-17
  - Trap Vector Table 6-17
- Posted Software Events
  - Debug Actions 12-11
- Post-Increment Addressing 2-9
- PPN
  - Address Spaces 10-2
  - Field in MMU\_TPA Register 10-18
  - Page Table Entry Translation 1-10, 10-1
  - Physical Page Number
    - TLB Table Entry Contents 10-5
- Pre-Decrement Addressing 2-9
- Pre-Increment Addressing 2-9
- Previous Context
  - CSAs and Context Lists 4-6
- Previous Context Information (PCXI)
  - Register Definition 3-12

- Previous Context Pointer (PCX)
  - Context Management Registers 4-13
  - Context Restore 4-11
  - Context Save 4-9
  - Register Definition 4-15
- Previous CPU Priority Number (PCPN)
  - Field in PCXI Register 3-12
- Previous Interrupt Enable (PIE)
  - Field in PCXI Register 3-12
- PREVSUSP
  - Field in DBGSR Register 12-15, 12-25
- Priority
  - Debug Events 12-12
- Priority Number
  - CPU 4-7
  - of Interrupt Task 3-12
  - Pending Interrupt
    - Context Switching - Interrupts 4-7
  - Previous CPU 4-7
- PRIV Trap
  - Privilege Violation 6-7
- Privilege Level 3-7
- Privileged Peripheral (P)
  - Physical Memory Address 8-1
- Program
  - Counter
    - Architectural Registers 1-3
    - Register A11 3-2
  - Memory 2-14
  - State Information 3-5
- Programming Model 2-1
  - Address Data Type 2-2
  - Bit String 2-1
  - Boolean 2-1
  - IEEE-754
    - Single Precision Floating Point Number 2-2
  - Signed Fraction 2-2
  - Signed/Unsigned Integers 2-2
- Protection
  - Boundaries
    - Crossing Boundaries 9-17
  - I/O Privilege Level 1-9, 9-1
- Internal Protection Traps 6-7
- Memory Protection System 1-9
- Page-Based 1-10, 9-1
- Range-Based 1-10
- Register Set 9-6, 9-7, 9-16
  - Data 9-14
  - Using 9-5
- System 1-9, 9-1
- Trap System 1-9, 9-1
- Virtual Mode 10-7
- Protection Register Set 3-6
- Protection Register Set Example 9-4
- PRS
  - Field in PSW Register 3-6
  - Protection Register Set
    - Debugger Triggers 12-5
- PSE Trap
  - Fetch Synchronous Error 6-12
- PSPR
  - Program scratchpad RAM 8-4
- PSW
  - Architectural Registers 1-3
  - Architecture Overview 1-3
  - Example Register Set Usage 9-14
  - FPU Exceptions 11-8
  - Initial State upon a Trap 6-6
  - Interrupt Service Routine 5-8
  - Memory Protection 9-2
  - Processor Status Word 1-7
  - Program Status Word Register
    - Address Offset 14-2
    - CDC Field 4-6
    - Definition 3-6
  - Supervisor Mode 3-7
  - Task Switching Operation 4-4
  - User Status Bits 3-10
    - Definition 3-10
    - USB Field in PSW Register 3-6
  - User-0 Mode 3-7
  - User-1 Mode 3-7
- PSZ
  - Field in MMU\_TPA Register 10-18
- PTE

Execute Enable (XE) bit 10-7  
 Page Table Entry Based Translation  
     Description 10-7  
     Overview 1-10  
 Read Enable (RE) bit 10-7  
 Write Enable (WE) bit 10-7

## **Q**

Q31 format  
 Floating Point Overview 11-1

## **R**

r  
     Definition of 1-2  
 RA  
     Return Address 3-2  
     Task Switching Operation 4-4  
 Range Entry  
     Debugging 9-4  
 Range Table Entry  
     Mode Register 3-21  
     Modes of Use 9-4  
     Segment Protection 3-21  
 RBL  
     Field in DPMx Register 9-11  
 RBU  
     Field in DPMx Register 9-11  
 RE  
     Field in DPMx Register 9-10  
     Field in MMU\_TPA Register 10-17  
     Read Enable  
         TLB Table Entry Contents 10-5  
 Real Time Operating System (RTOS)  
     Tasks and Functions 4-1  
 Record Elements  
     Base + Offset Addressing 2-8  
 Register  
     A10(SP) 3-14  
     BIV 6-19  
     BMACON 3-19  
     BTv 6-20  
     CCDIER 7-4  
     CCNT 12-47

CCPIER 7-3  
 CCTRL 12-45  
 CDC 3-21  
 COMPAT 3-18  
 Context Management 4-13  
 CPMx\_n 9-12  
 CPRx\_nL  
     Code Segment Protection Register  
     Lower Bound 9-9  
 CPRx\_nU  
     Code Segment Protection Register  
     Upper Bound 9-8  
 CREVT 12-18, 12-29  
 CSFR 3-1  
 Data Registers (D0 to D15) 3-2  
 DBGSR 12-15, 12-25  
 DCX 12-24, 12-38  
 DIEAR 7-8  
 DIETR 7-7  
 DMS 12-23, 12-37  
 DPMx\_n 9-10  
 DPRx\_nL  
     Data Segment Protection Register  
     Lower Bound 9-7  
 DPRx\_nU  
     Data Segment Protection Register  
     Upper Bound 9-6  
 EXEVT 12-17, 12-27  
 Extended-Size 3-2  
 FCX 4-14  
 Floating Point 3-2  
 Global 3-8  
 GPR 3-1  
 ICNT 12-48  
 ICR 6-17  
 LCX 4-16  
 Lower Registers 1-4  
 M1CNT 12-49  
 M2CNT 12-50  
 M3CNT 12-51  
 Memory Protection Overview 3-21  
 MIECON 7-9  
 MMU\_ASI 10-15



- MMU\_CON 10-13
- MMU\_TFA 10-20
- MMU\_TPA 10-17
- MMU\_TPX 10-19
- MMU\_TVA 10-16
- Mode 3-21
- Overview of Registers 1-3
- PCX 4-15
- PCXI 3-12
- PIEAR 7-6
- PIETR 7-5
- SBSRCn 12-40
- Scaled Data 2-13
- SRC 5-3
- SWEVT 12-19, 12-31
- SYSCON 3-16
- System Global Registers (A0, A1, A8, A9) 1-4
- TR0EVT 12-20, 12-33
- TR1EVT 12-20, 12-33
- Reserved Field (-)
  - Definition 1-2
- Reset Values 3-1
- Restore
  - Task Switching Operation 4-4
- Return Address (RA) 3-2, 6-4
  - Register A11
  - GPR Overview 1-3
  - Trap System 6-4
- Return From Call (RET)
  - Context Switching - Function Calls 4-8
  - Task Switching 4-4
- Return From Exception (RFE)
  - Exiting an ISR 5-9
  - Interrupt Priority Groups 5-12
  - Task Switching 4-4
- Revision History of this Document 1-4
- RISC
  - Architecture Overview 1-1
- RM
  - Field in PSW 11-6
  - Floating Point Rounding 11-6
- Rounding
  - Floating-Point 11-6
- RS
  - Field in DMPx Register 9-10
- RTOS
  - Context Switching with Interrupts 4-7
  - Service Request Notes (SRNs) 5-1
  - Software-Posted Interrupts 5-15
- Run-control Features
  - Core Debug Controller (CDC) 12-2
- rw
  - Definition of 1-2
- rwh
  - Definition of 1-2
- S**
  - SBSRCn 12-40
  - Scale Factor
    - Indexed Addressing 2-13
  - Scaled
    - Data Register
      - Indexed Addressing 2-13
  - Scratchpad RAM
    - Physical Memory Attributes 8-4
  - Segments 2-6
    - 0 to 7
      - MMU 1-10
    - 8 to 15
      - MMU 1-10
    - Address Space 1-4
    - Memory Model
      - Address Space 2-6
      - Physical Memory Attributes 8-3
  - Semaphores 2-7
  - Service Providers
    - Interrupt System 5-1
  - Service Request Control Register (SRC)
    - Definition 5-3
    - Interrupt Registers 3-21
    - Interrupt System 5-1
  - Service Request Node (SRN)
    - Interrupt System 1-8
    - Overview 5-1
  - Service Request Priority Number (SRPN)

- Interrupt Priority 1-8
- Service Requests
  - Interrupt Priority 1-8
- SETR
  - Description 5-4
  - Field in SBSRC Register 12-40
  - Field in SRC Register 5-3
- Signed Fraction
  - Programming Model 2-2
- Signed/Unsigned Integers
  - Programming Model 2-2
- SMACON 3-20
  - Address Offset 14-4
- SMT
  - CSU Trap 6-11
  - Software Managed Task
    - Tasks and Functions 4-1
  - Software Managed Tasks
    - Tasks and Contexts 1-7
- Software Managed Tasks (SMT)
  - Overview 1-7, 3-11
- SOVF Trap
  - Sticky Arithmetic Overflow
    - Assertion Traps 6-14
- SP
  - A10 Register
    - Task Switching Operation 4-4
  - Stack Pointer 3-14
  - Stack Pointer A10 Register
    - General Purpose Registers 3-2
- Spanned Service Routine
  - Spanning ISRs 5-12
- Speculative (S)
  - Physical Memory Properties 8-1
- SRC
  - Service Request Control Register
    - Definition 5-3
- SRE
  - Description 5-5
  - Field in SBSRC Register 12-40
  - Field in SRC Register 5-4
- SRN
  - Interrupt System Introduction 5-1
- Service Request Node
  - Interrupt System 1-8
  - Overview 5-1
- Software-Posted Interrupts 5-15
- SRPN
  - Description 5-6
  - Different Priorities for the same
    - Interrupt Source 5-14
  - Field in SBSRC Register 12-41
  - Field in SRC Register 5-4
  - Fields 5-6
  - Service Request Priority Number 1-8
- SRR
  - Description 5-5
  - Field in SBSRC Register 12-40
  - Field in SRC Register 5-4
- Stack
  - Pointer A10
    - Architecture Register Overview 1-3
  - Pointer Register 10
    - General Purpose Registers 3-2
- Stack Management
  - Description 3-13
- State Information 3-19, 3-20
  - PCXI Register 3-12
  - Program Counter (PC) 3-5
- Static Data
  - Base + Offset Addressing 2-8
- Sticky Overflow
  - SOVF
    - Supported Traps 6-2
- STLCX 4-5
- STUCX 4-5
- Supervisor Mode 3-7
  - Overview 1-7, 3-11
- SUSP
  - Field in CREVT Register 12-18, 12-29
  - Field in DBGSR Register 12-16, 12-26
  - Field in EXEVT Register 12-17, 12-27
  - Field in SWEVT Register 12-19, 12-31
  - Field in TRnEVT Register 12-22, 12-35
- SVLCX 4-4
- SVLCX Instruction

Context Switching with Interrupts 4-7

SWAP Instruction

Alignment Requirements 2-4

SWAP.W 2-7

SWEVT

Address Offset 14-5

SWEVT Register

Debug Action 12-3

Software Debug Event Register

Definition 12-19, 12-31

Synchronous Trap

Overview 6-3

Synthesized

Addressing Modes 1-6, 2-7

SYS Trap

System Call Trap

Description 6-14

SYSCALL Instruction

SYS Trap Description 6-14

SYSCON

Free Context List Depletion Trap 6-10

Register 3-16

Address Offset 14-2

Memory Protection System 9-16

System

Global Registers (A0, A1, A8, A9) 3-2

System Call - SYS Trap

Supported Traps 6-2

System Call Traps 6-14

SZA

Field in MMU\_CON Register 10-14

SZB

Field in MMU\_CON Register 10-14

## **T**

Table Indexes

General Purpose Registers 3-2

Task

Context

Current 4-7

Switching 4-4

Task Mode 12-43

Tasks

Software Managed Tasks (SMTs)

Overview 4-1

Tasks and Functions

Overview 4-1

Text Conventions

Used in this Document 1-2

TIN

SYS Trap (System Call) 6-14

TIN-0 6-7

TIN-1 6-7

TIN-2 6-7

Trap Identification Number

Trap System 1-9

Trap Types 6-1

TLB 10-5

TTE Contents 10-5

TLB (Translation Lookaside Buffer)

Description 10-4

Hardware Traps 6-3

Usage 10-12

VAF Trap 6-7

TLBDEMAP Instruction

Follow by ISYNC 10-10

TLB Usage 10-12

Use in MMU 10-10

TLBFLUSH Instruction

Description in MMU 10-10

TLBMAP Instruction

Description 10-9

TLBPROBE Instruction

Description 10-11

TLBPROBE.I Instruction

Description 10-11

TOS

Description 5-5

Field in SBSRC Register 12-41

Field in SRC Register 5-4

TR0EVT

Register Definition 12-20, 12-33

TR1EVT

Register Definition 12-20, 12-33

Translation

Direct 10-1

- PTE
  - Description 10-7
  - MMU Overview 10-1
- Translation Virtual Address (TVA) Register
- TLBPROBE.I 10-11
- Trap
  - Accessing the Trap Vector Table 6-4
  - ALN - Data Address Alignment 6-9
  - Assertion 6-14
  - Asynchronous 6-3
  - BAM - Break After Make 6-14
  - Base Trap Vector Table Pointer (BTVP)
  - Register Definition 6-20
  - BBM - Break Before Make 6-14
  - Class 0 6-7
  - Class 1 6-7
  - Class 2 6-8
  - Class 3 6-10
  - Class 4 6-12
  - Class 5 6-14
  - Class 6 6-14
  - Class 7 6-14
  - Class Number 6-4
  - Classes 1-9, 6-20
  - Context Management 6-10
  - CSU - Call Stack Underflow 6-11
  - CTYP - Context Type 6-11
  - Debug 6-14
  - Descriptions 6-7
  - DIE 7-2
  - FCD - Context List Depletion 6-10
  - FCU - Context List Underflow 6-11
  - Handler Vector 6-4
  - Identification Number (TIN)
    - Trap System Overview 1-9
    - Trap Types 6-1
  - Initial State 6-6
  - Internal Protection 6-7
  - Memory Protection Traps 9-17
  - MPP - Memory Protection Peripheral Access 6-8
  - MPR - Memory Protection Read 6-7
  - MPW - Memory Protection Write 6-8
  - MPX - Memory Protection Execute 6-8
  - NEST - Nesting Error 6-12
  - NMI - Non-Maskable Interrupt 6-14
  - OPD - Invalid Operand 6-9
  - OVF - Arithmetic Overflow 6-14
  - PCXI Register
    - UL Field 3-12
  - PIE 7-2
  - Priorities 6-15
  - PRIV - Privilege Violation 6-7
  - Register A11 (RA) use with Traps 3-2
  - Return Address 6-4
  - SOVF - Arithmetic Overflow 6-14
  - Synchronous Overview 6-3
  - SYS - System Call 6-14
  - System 1-9
  - System Call (SYS) 6-14
  - Trap Handler 6-1
  - Trap System 6-1
  - Types 6-1
  - VAF - Virtual Address Fill 6-7
  - VAP - Virtual Address Protection 6-7
  - Vector Table Pointer 6-17
- Trap Registers 3-21
- Trap system
  - Trap vector table 6-5
- Traps
  - FPU 11-4
  - MMU 6-7
- TRAPSV Instruction
  - SOVF Trap 6-14
- TRAPV Instruction
  - OVF Trap 6-14
- Trigger Event Register (TRnEVT)
  - Definition 12-20, 12-33
- Trigger Event Unit
  - Description 12-4
- TRnEVT
  - Address Offset 14-5
  - Debug Action 12-4
  - Register Definition 12-20, 12-33
- TSZ
  - Field in MMU\_CON Register 10-14

TTE  
 TLB Table Entry Contents 10-5  
 Translation Lookaside Buffer (TLB)  
 Description 10-4  
 Type-of-Service Control (TOS)  
 Field in SRC Register 5-4, 5-5

## U

UOPC Trap  
 Unimplemented Opcode 6-8  
 UPDFL Instruction  
 Changing the Rounding Mode 11-6  
 UPBND  
 Field in CPRx\_nU Register 9-8  
 Field in DPRx\_nU Register 9-6  
 Upper Context 4-1  
 Registers 3-4  
 Task Switching Operation 4-4  
 UL Field in PCXI Register 3-12  
 User Status Bits 3-6, 3-10  
 User-0 Mode 3-7  
 Description 1-7, 3-11  
 User-1 Mode 3-7  
 Description 1-7, 3-11

## V

V  
 Field in MMU\_CON Register 10-14  
 Field in MMU\_TPA Register 10-17  
 VAF Trap  
 Hardware Traps 6-3  
 MMU Traps 10-5  
 Virtual Address Fill 6-7  
 Valid bit (V) 10-5  
 VAP Trap  
 Hardware Traps 6-3  
 MMU Traps 10-5  
 Virtual Address Protection 6-7  
 Vector Table  
 Base Address 6-19  
 Virtual  
 Address Space 10-2  
 Multiple Address Spaces 10-5

Addressing 1-1  
 MMU Address 1-10  
 Translation 8-6  
 Virtual Address 10-3  
 VPN  
 Address Spaces 10-2  
 Field in MMU\_TVA Register 10-16  
 MMU Page Table Entry Translation 1-10  
 TLB Table Entry (TTE) Contents 10-5

## W

w  
 Definition of 1-2  
 Watchpoints  
 CDC Features 12-1  
 WBL  
 Field in DPMx Register 9-11  
 WBU  
 Field in DPMx Register 9-11  
 WE  
 Field in DPMx Register 9-10  
 Field in MMU\_TPA Register 10-17  
 Write Enable 10-5  
 Word  
 Definition 1-2  
 Wrap Around Behaviour  
 Circular Addressing 2-10  
 WS  
 Field in DPMx Register 9-10

## X

XE  
 Execute Enable 10-5  
 Field in CPMx Register 9-12  
 Field in MMU\_TPA Register 10-17  
 XS  
 Field in CPMx Register 9-12





[www.infineon.com](http://www.infineon.com)

Ordering No. B158-H8581-G2-X-7600

Published by Infineon Technologies AG