AN4011

Choosing the Right Cypress Synchronous SRAM

Author: Sohelkhan Pathan/Pritesh Mandaliya
Associated Part Family: All Sync SRAMs
Related Application Notes: None

AN4011 provides an overview of Standard Synchronous, NoBL™, QDR®-II/II+, QDR-II+ Xtreme, DDR-II/II+, DDR-II+ Xtreme, and QDR-IV SRAM's. Cypress currently manufactures several major Synchronous SRAM architectures. The purpose of this application note is to provide a means to determine which architecture is right for a particular application. A brief description of each architecture and a comparison by address/data relationships and performance characteristics is also included.

Contents

1 Introduction

Cypress currently manufactures several major Synchronous SRAM architectures. These SRAMs all work on synchronous timing. Data, Address, and Control inputs are captured at a clock edge. Distinguishing features of several SRAM architectures are shown in Table 1.

Table 1. SRAM Architectures [1]

Architecture	Versions	Capability
Standard Synchronous	Flow-through	4-word burst user selectable; dual bus master; SAR; SDR; same-cycle Read result.
	Pipelined SCD	4-word burst user selectable; dual bus master; SAR; SDR; Read result after next CLK; same-cycle Read termination.
	Pipelined DCD	4-word burst user selectable; dual bus master; SAR; SDR; Read result after next CLK; Read termination after next CLK.
NoBL™	Flow-through	4-word burst user selectable; SAR; SDR; zero idle cycles for bus turnarounds; same-cycle Read result; data always one CLK cycle later than address.

Architecture	Versions	Capability
	Pipelined	4-word burst user selectable; SAR; SDR; zero idle cycles for bus turnarounds; Read result after next CLK; data always two CLK cycles later than address.
QDR-II/QDR-II+ /QDR-II+ Xtreme	2-word burst	2-word burst fixed; separate I/O; initiate Read and Write every clock cycle; DAR; DDR; Read result after next C or K; Read termination after data delivered; Write data begins with Write command; sustains 4 operands per clock cycle; features impedance matching ZQ circuitry; QDR-II features larger output data window and source synchronous clocks.
	4-word burst	4-word burst fixed; separate I/O; initiate Read or Write every clock cycle; SAR; DDR; Read result after next C or K; Read termination after data delivered; Write data begins next clock after Write command; sustains 4 operands per clock cycle; features impedance matching ZQ circuitry; QDR-II features larger output data window and source synchronous clocks.
DDR-II/DDR-II+ /DDR-II+ Xtreme[2] (common I/O)	2-word burst	2-word burst fixed; common I/O; SAR; DDR; Read result after next C or K; Read termination (common I/O) after data delivered; Write data begins next clock after Write command; sustains 2 operands per clock cycle; features impedance matching ZQ circuitry and source synchronous clocks; DDR-II features larger output data window.
	4-word burst	4-word burst fixed; common I/O; SAR; DDR; Read result after next C or K; Read termination after data delivered; Write data begins next clock after Write command; sustains 2 operands per clock cycle; features impedance matching ZQ circuitry and source synchronous clocks; DDR-II features larger output data window.
DDR-II/DDR-II+ (separate I/O)	2-word burst	2-word burst fixed; separate I/O; initiate Read or Write every clock cycle; SAR; DDR; (separate I/O) Read result after next C or K; Read termination after data delivered; Write data begins next clock after Write command; sustains 2 operands per clock cycle; features impedance matching ZQ circuitry and source synchronous clocks; DDR-II SIO features larger output data window.
QDR-IV	2-word burst	2-word burst fixed; two independent bidirectional data ports that support simultaneous read/write transactions; DAR; DDR; QDR-IV XP has write latency of 5 clock cycle and read latency of eight clock cycle while QDR-IV HP has three and five clock cycles as write latency and read latency respectively. Supports bus inversion, address bus parity, configurable on-die terminations, on chip error correction code (ECC) and de-skew training.

Note 1: SAR = Single Address Rate, SCD = Single Cycle Deselect, SDR = Single Data Rate, DAR = Double Address Rate, DCD = Double Cycle Deselect, DDR = Double Data Rate.

Note 2: DDR-II+ Xtreme devices have 2-word burst option only.

2 Standard Synchronous SRAM

The Standard Synchronous SRAM family was created especially for cache applications. These devices incorporate an internal 2-bit burst counter that supports a cache line size of four (that is, four bus transactions to fetch a cache line). The Standard Synchronous SRAM was designed to permit two bus masters; hence, it has two different master control inputs: $\overline{\text{ADSP}}$ that is generally controlled by the microprocessor, and $\overline{\text{ADSC}}$ that is generally controlled by the cache controller. $\overline{\text{ADSP}}$ operations are interpreted initially as READ cycles but can be turned into WRITE cycles at the next CLK rising edge. A single address can be allowed to operate on 1, 2, 3, or 4 words in response, controlled on the fly by advance ($\overline{\text{ADV}}$). Any $\overline{\text{ADV}}$ beyond four simply wraps around and replays from the beginning of the sequence.

Standard Synchronous SRAMs are the preferred choice for processors specifically designed around their control signals. In general, while this SRAM architecture is well suited for single data rate (SDR) cache applications, it is not well suited for applications in which frequent bus turnaround cycles occur.

Three different versions of the Cypress Standard Synchronous SRAMs are available (Flow-through, Single-cycle deselect [SCD] Pipelined, and Double-cycle deselect [DCD] pipelined). Choosing between the different versions is based upon a number of factors.

2.1 Choosing Pipelined versus Flow-through

Cypress offers two different pipelined versions and one Flow-through version of the Standard Synchronous SRAM as shown in Table 2. The pipelined versions have input registers and output registers. The Flow-through versions only have input registers. Because of the output register, pipelined parts have an extra cycle of delay before data is valid (during a READ) versus Flow-through devices.

For ASIC or DSP applications, the Flow-through device is frequently used. This is because many of these applications cannot tolerate the extra cycle of delay associated with pipelined parts when retrieving data. Pipelined parts do have a place in applications where maximizing bandwidth is critical. Users need to be aware that pipelined parts require a one-cycle delay when switching from Reads to Writes in order to prevent bus contention.

Table 2. Pipelined and Flow-Through Description

Version	Description
Flow-through	A Synchronous SRAM with registers on the input signals.
Pipelined with Single-Cycle Deselect (SCD)	A Synchronous SRAM with registers on both the input and output signals. Single-cycle deselect.
Pipelined with Double-Cycle Deselect (DCD)	A Synchronous SRAM with registers on both the input and output signals. Double-cycle deselect.

2.2 Choosing the Right Pipelined Version

Because Cypress has two different types of pipelined SRAMs, it is important for designers to choose the version that best meets their needs. Figure 1 shows the difference between the DCD and SCD versions of the pipelined SRAMs. This diagram shows a Read cycle followed by three advance cycles and then a deselect cycle. The DCD device reads out all four words, but the SCD cuts off the last word. To read all four words out of the SCD device, the deselect command must be delayed one cycle.

Cypress designed the SCD version to be compatible with the Intel BSRAM specification. This specification was developed to ensure that there is no possibility of bus contention in a system with multiple banks of SRAM. This could occur when switching between two or more banks of SRAMs if parts from different vendors are used, or if the banks are separated by electrically long paths. This bus contention could occur during the Read transition if the turn-on and turn-off times for the DQ lines vary significantly among manufacturers.

For cache users it is important to select an SRAM and a chipset that are compatible and provide the required performance. For greatest availability of parts, a chipset might be chosen that supports both DCD and SCD. For best performance on a single bank only system, a DCD might be the best solution. It is important for cache designers to check with their chipset vendor to verify the type(s) of Standard Sync SRAM they support. Most chipsets only support the BSRAM standard and can only use SCD parts. Other chipsets may be able to use both DCD and SCD parts, but leave a dead cycle between bank switching, so even if you use the DCD versions, there is no performance advantage. Others may support both DCD and SCD, but give a performance advantage if the DCD parts are chosen.

2.3 Designing for Both DCD and SCD Versions

To have the widest supply base or when fastest performance is not an issue, users can create a design that can use both DCD and SCD. One way of doing this is to detect what type of SRAM (DCD or SCD) is in the system and then configure the Deselect cycle accordingly. During boot-up/power-up, the system can test to see if a DCD or SCD SRAM is present. This can be done by writing four sequential words into the SRAM and then reading them out in a burst, with a Deselect cycle occurring after the last ADV cycle (see Figure 1). If a DCD is in the system, all four words will be read correctly; if SCD, only three words will be read correctly.

Figure 1. Single-Cycle Deselect (SCD) versus Double-Cycle

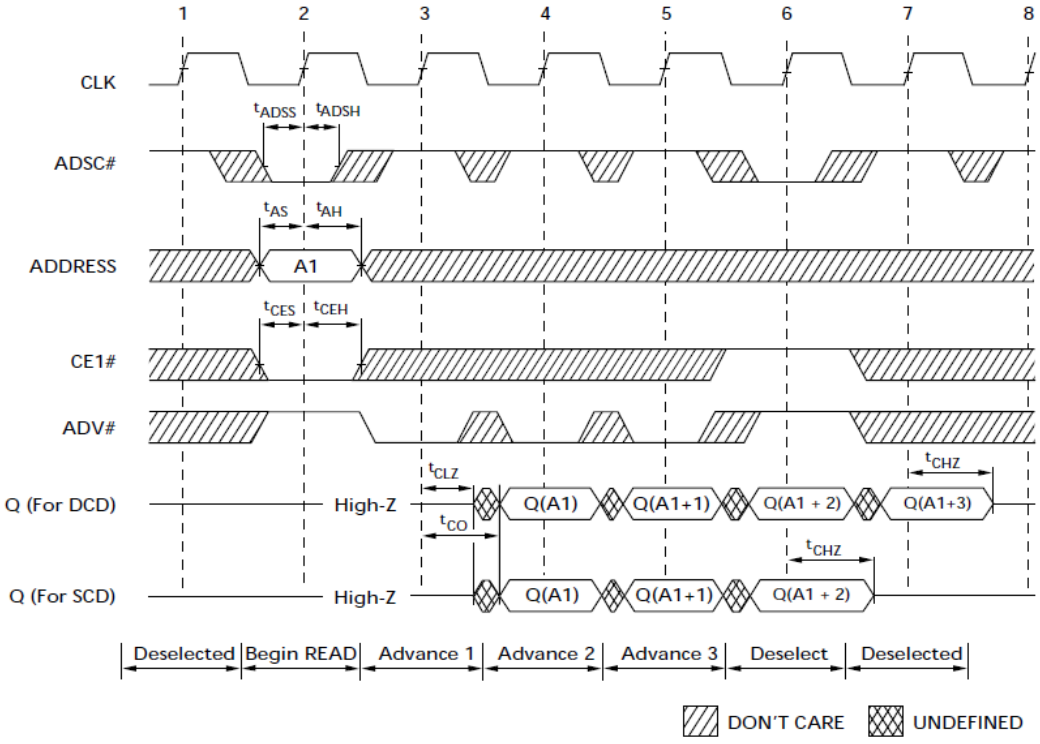

Deselect (DCD)

Figure 2. Designing for Both DCD and SCD Versions

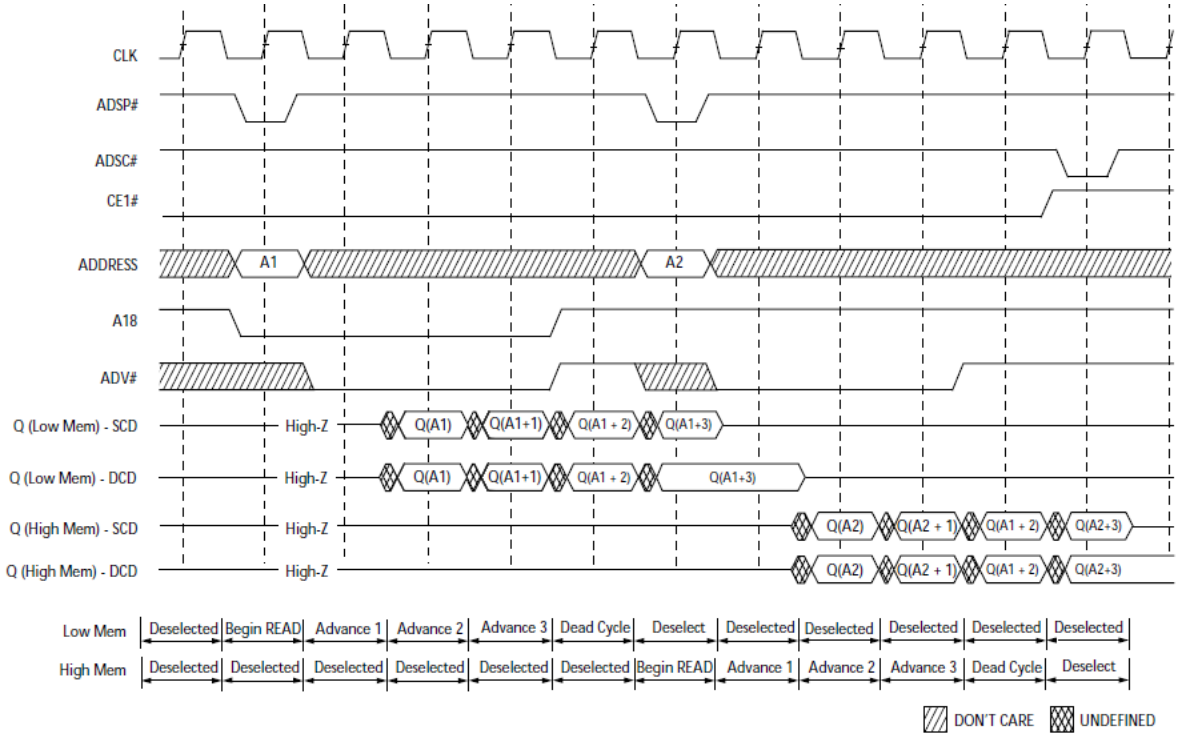

Another method is to design a memory interface that can accommodate either the DCD or SCD versions without any modifications during power-up or boot-up. This can be done by always inserting a dead cycle, an example of which is shown in Figure 2. This method permits the simplest and lowest pin count interface, but does take a performance hit when doing back-to-back READs.

3 No Bus Latency™ (NoBL™)

The NoBL SRAM was created in response to applications that required frequent bus turnarounds but could not afford the idle cycles needed by Standard Synchronous SRAMs. Like Standard Synchronous devices, NoBL incorporates an internal 2-bit burst counter to reduce the required address bus bandwidth. However, most applications do not make use of this feature since the bus turns around so often. This device operates at SDR and SAR.

One innovation of NoBL architecture is the realignment of write data such that the address-to-data relationship is identical whether reading or writing. For this reason, no idle cycles are needed when the bus is turned around, that is, when it transitions from Read to Write, or Write to Read. It should be noted that all versions of NoBL SRAMs use internal data registers and address snooping so that the correct data is always returned in response to a Read cycle, regardless of whether or not there has been time to write the data into the memory array. This is called data coherency.

4 NoBL Flow-through

Two versions of NoBL SRAM exist: Flow-through and Pipelined. NoBL Flow-through SRAMs always have a one clock cycle delay from address to data in the system. Reads provide data during the clock cycle for the one it was requested, and produce the result in time to capture by the next CLK rising edge. Write data is also expected at that next CLK rising edge. This device style minimizes the system latency—for a price. The SRAM must perform a data fetch and data delivery in time for the next CLK rising edge. This means that the frequency of operation is limited. This device is well suited for SDR applications that experience frequent bus turnarounds. It is used effectively when there is a need to operate on small data chunks, especially one-word chunks that require minimum data latency. In general, Flow-through SRAMs operate at a lower frequency than pipelined SRAMs because data fetch, data delivery, and subsequent data capture by the requestor must be accomplished in one clock cycle.

5 NoBL Pipelined

Pipelined NoBL SRAM is named after the additional pipelined register in the outputs of this device. This permits the internal Read to take one entire clock cycle. During the next clock cycle, data is delivered, that makes a two-stage pipeline. While data is being delivered, the memory array is free for another operation. The address-to-data relationship therefore lags the Flow-through version by one additional clock cycle. Write cycles, therefore, expect data two CLK rising edges after the address. Pipelined NoBL devices operate at higher frequencies than Flow-through and sustain greater system throughput. The one clock cycle penalty for priming the pipeline is incurred only once when the first operation is initiated, and then one new transaction can be sustained every cycle.

This device is well suited for SDR applications that experience frequent bus turnarounds, need to operate on small data chunks (especially one-word chunks), and need to operate at higher frequencies than permitted by the Flow-through version.

6 QDR-II/II+/II+ Xtreme™ and DDR-II/II+/II+ Xtreme

The following section discusses QDR-II/II+/II+ Xtreme.

QDR stands for Quad Data Rate. QDR-II SRAMs were developed to address network applications that require the low latency and full cycle utilization of NoBL SRAMs but also require a significantly higher operating frequency. An important factor was reducing ASIC pin count.

QDR-II SRAMs have separate Read and Write buses. This SRAM solution eliminates turnaround cycles and scales to any frequency allowed by the inherent speed of memory manufacturing technology. Control signals are few: $\overline{\text{RPS}}$ controls the Read port, and $\overline{\text{WPS}}$ controls the Write port. Width expansion is simply done with all $\overline{\text{RPS}}$ and $\overline{\text{WPS}}$ controls in parallel. Depth expansion is done by adding another $\overline{\text{RPS}}$ and $\overline{\text{WPS}}$ for each bank.

An important feature of QDR-II architecture is Double Data Rate (DDR) on each data pin. This permits smaller bus sizes and reduced ASIC pin counts. All cycles operate at the device burst length. Cycles cannot be terminated or interrupted. Parts of the Write cycle can optionally be masked using Byte Write control pins.

All versions require an input master clock pair, K and $\overline{K}$. Only the rising edges are utilized by the SRAM. $\overline{K}$ rising edge should ideally occur exactly one half clock cycle after K rising edge. This balances the output data so that each data word has the same valid time.

QDR-II and DDR-II SRAMs have a data output clock pair, C and $\overline{C}$, that can optionally be used to control when output data emerges from the device. This is very useful in systems in which multiple SRAMs are located at differing physical distances from the bus master. All output data can be aligned using C and $\overline{C}$, that the whole result can be captured and easily synchronized at the bus master simultaneously. If not used, C and $\overline{C}$ can be strapped HIGH, placing the device in single clock mode (K and $\overline{K}$ are used for both input and output registers).

QDR-II/DDR-II devices have optional-use source-synchronous output echo clocks, CQ and $\overline{CQ}$. These outputs are timed exactly same as the output data Q signals and can be used to trigger input registers. It is recommended to use CQ and $\overline{CQ}$ at clock speeds above 200 MHz.

QDRII+/DDR-II+/QDR-II+ Xtreme/DDR-II+ Xtreme devices do not use C and $\overline{C}$. The outputs are synchronized to CQ and $\overline{CQ}$. In addition there is a QVLD signal that goes high half cycle before data is output from the SRAM and goes low half cycle before the last data is output from the device. QDR-II+/DDR-II+/QDR-II+ Xtreme/DDR-II+ Xtreme have devices with and without on-die termination (ODT). The combination of mandatory and optional-use clocks makes this the most flexible SRAM architecture available today.

7 QDR-II Burst of 2

The 2-word burst QDR SRAM version can accept two addresses during each clock cycle. Hence the burst of 2 QDR is Double Address Rate. This enables the application to utilize the full bandwidth on the bus. $\overline{RPS}$ and $\overline{WPS}$ are latched at the same K rising edge. Externally, it appears as though both Read and Write are initiated at the same time.

The Write cycle is shown in Figure 3. The first burst of data is captured beginning at the same K rising edge when $\overline{WPS}$ is asserted low. The second burst is captured at the next $\overline{K}$ rising edge. The address is captured at the $\overline{K}$ rising edge. The data is written into the array during the time between the $\overline{K}$ rising edge and K rising edge.

Figure 3. QDR SRAM Write Cycle Data Placement

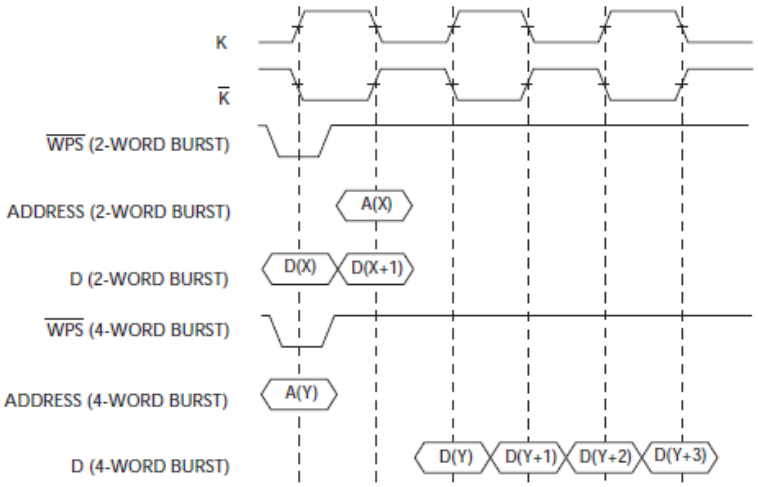

The Read cycle is shown in Figure 4. $\overline{RPS}$ is asserted low at the rising edge of K to initiate a read. The address is supplied at the same time. Internally, data is read from K rising edge to $\overline{K}$ rising edge. It is registered, and then data appears at the output during the next clock cycle for two consecutive half clock cycles.

Figure 4. QDR SRAM Read Cycle

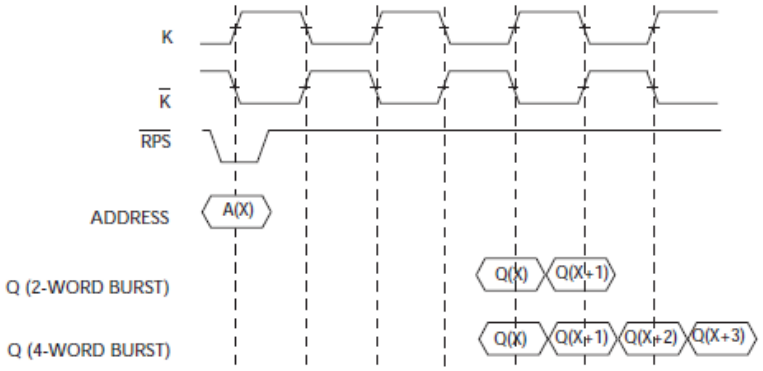

Output data is coherent. If a Read and a Write cycle are initiated during the same clock cycle (at K rising edge), $\overline{RPS}$ and $\overline{WPS}$ are both LOW. A Read is requested from address 1; address 1 is provided. At $\overline{K}$ rising edge, address 2 is provided for the Write. In this example, addresses 1 and 2 are the same. The Read data provided is the same data presented for the Write to address 2.

QDR burst of 2 is ideal for applications that require small data chunks and in that the near-term ratio between Read and Write operations is close to one. (This ratio is discussed in more depth later.) The DAR and double data rate operation permit minimum pin count on any devices that interface with the SRAMs. Frequency is limited by two factors: system loading and SRAM internal memory array speed. The controller must be able to provide the address at the same signaling frequency as the clock. Since the address fans out to all devices on the bus, and often data is just point-to-point, the address path is the most difficult to satisfy for AC timing. The SRAM internal memory array speed limitation arises because the memory array is accessed twice per clock cycle. SRAM speed is limited by the ability of the memory array to keep up with incoming requests.

8 QDR-II Burst of 4

The long burst QDR-II SRAM requires a single address be supplied during each clock cycle to maintain full bus utilization [single address rate (SAR)]. Only one command is received at a K rising edge: $\overline{RPS}$ or$\overline{WPS}$. Externally, it appears as though Read and Write operations toggle. That is, a Read slot occurs during one clock cycle, a Write slot occurs during the next, and so on. Indeed, this is the case. Quad data rate operation occurs because the four data words encompass two clock cycles. At full device utilization, read and write bus transactions overlap and two operations are in simultaneous flight.

The Write cycle in Figure 3 shows that data is captured beginning at the next K rising edge that initiated the cycle and continues for three more K/$\overline{K}$ rising edges. The address is accepted at K rising edge. The actual internal Write to the memory data array does not occur until the next write request. Any Read from the pending location is forwarded correctly from the holding registers, and full data coherency is maintained.

The Read cycle in Figure 4 shows that the address is supplied at the same K rising edge that initiated the cycle. Internally, data is read from K rising edge to the next K rising edge. It is registered, and then data delivery begins during the next clock cycle and continues for four consecutive half clock cycles.

This device is ideal for applications that require 4-word data chunks and in that the near-term ratio between Read and Write cycles is near one. The SAR and double data rate operation permit minimum pin count on any devices that interface with the SRAMs. The address bus only needs to be updated once each clock period to maintain 100 percent bus utilization on both buses.

9 QDR-II versus QDR-II+/II+ Xtreme

Details on QDR-II/II+/II+ Xtreme SRAM operation are explained in the Application Note AN4065. The main difference between QDR-II and QDR-II+/II+ Xtreme is the read latency. Also QDR-II+/II+ Xtreme does not use C/$\overline{C}$ and has a QVLD signal to indicate read data availability. The control signals remain the same. A brief outline of differences and how those differences may influence SRAM selection follows.

Table 3. Summary - QDR-II versus QDR-II+/II+ Xtreme

	QDR-II	QDR-II+	QDR-II+ Xtreme
Max frequency	Burst of 2: 333 MHz Burst of 4: 333 MHz	Burst of 2: 333 MHz Burst of 4: 550 MHz	Burst of 2: 450 MHz Burst of 4: 663 MHz
Initial latency	1.5 clock cycles	2.0/2.5 cycles	2.5 cycles
Echo clocks	Yes	Yes	Yes
Density	18 Mb/36 Mb/72 Mb	18 Mb/36 Mb/72 Mb/144 Mb	36 Mb/72 Mb/
Power supply	1.8 V	1.8 V	1.8 V
ODT enabled device option	No	Yes	Yes
Data valid signal - QVLD	No	Yes	Yes

10 DDR-II Common I/O

Some SRAM applications require a Read followed by a Write favoring QDR-II SRAM operation. Other applications may need data streaming (for example, 16 Reads, then 16 Writes), in that case the near-term balance between Read and Write operations is 100 percent Read or 100 percent Write. In this case, one of the QDR-II SRAM buses is not utilized half of the time. These two latter cases resulted in the development of DDR-II common I/O SRAMs in that the input and output data share the same bus. This reduces the number of signals to be routed on the board compared to QDR-II. Bus turnaround cycles reduce available bandwidth; however, for some systems this results in a better average bus utilization than QDR architecture could provide.

The DDR-II SRAMs have common Read and Write ports. Bus turnaround cycles are required during the transition from SRAM Read to Write. The number of clock idle cycles during this transition varies with frequency: one clock cycle is needed at approximately 166 MHz and below; two clock cycles are needed at approximately 200 MHz and above. Simulations need to be done to determine exactly how many cycles are required to avoid excessive bus contention. This solution scales to higher frequencies and is limited only by the inherent speed of memory manufacturing technology. Control signals are slightly different than the QDR-II device control signals. R/$\overline{W}$ controls the Read/Write operation, and $\overline{LD}$ acts same as a chip enable and causes the SRAM to initiate a new cycle. Width expansion is simply accomplished with all R/$\overline{W}$ and $\overline{LD}$ controls in parallel. Depth expansion is accomplished by adding another $\overline{LD}$ for each bank.

All other features are the same as the QDR-II device.

11 DDR-II Burst of 2

The DDR-II SRAM can perform one operation during one clock cycle by asserting $\overline{LD}$ LOW and setting the R/$\overline{W}$ at the rising edge of K.

The Write cycle is shown in Figure 5. A write operation is initiated by setting $\overline{LD}$ and R/$\overline{W}$ low at the rising edge of K. The address is also captured at the same K rising edge. The two bursts of data are captured one clock cycle later - the first burst at the rising edge of K and second at the rising edge of $\overline{K}$. The actual Write to the memory array occurs during the next Write cycle. If a Read to the pending address occurs before data is written into the memory array, data is forwarded from the holding register so that supplied data is always coherent.

The Read cycle is shown in Figure 6. A read operation is initiated by setting $\overline{LD}$ low and R/$\overline{W}$ high at the rising edge of K. The address is captured at the same K rising edge. Data appears at the outputs after 1.5 cycles. The first burst appears at the rising edge of $\overline{K}$ and the second burst appears at the next rising edge of K.

Figure 5. DDR-II Write Cycle

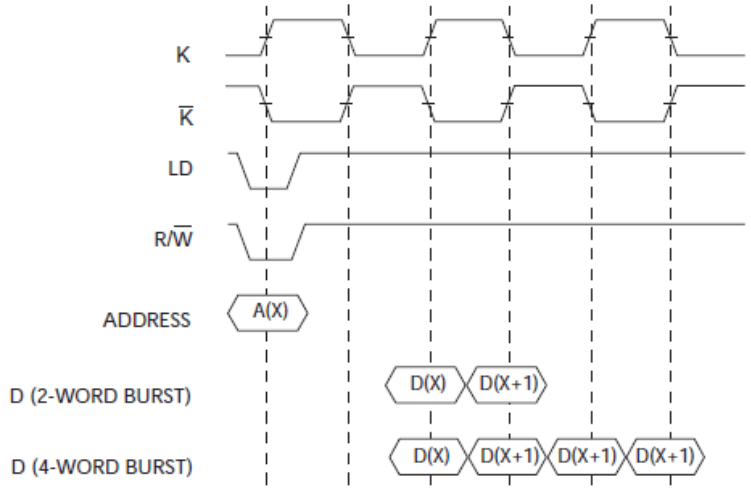

Figure 6. DDR-II Read Cycle

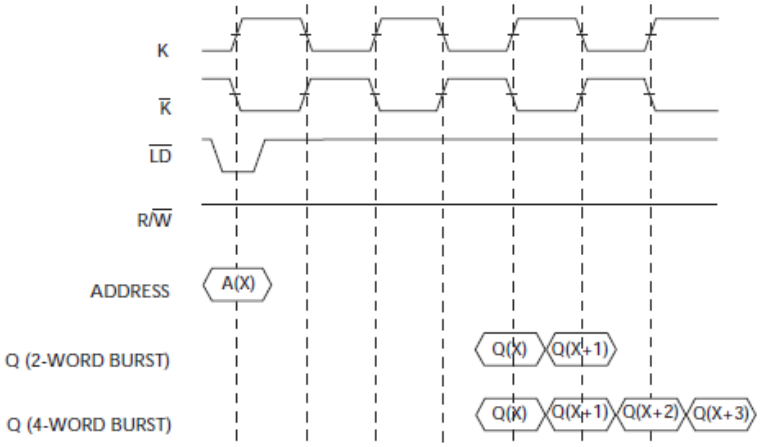

The low address rate and double data rate operation permit minimum pin count on any device that interfaces with the SRAMs, and they simplify high-frequency operation. This device is best utilized in applications requiring data streaming or in that periods of unidirectional bus operation is sustained.

12 DDR-II Burst of 4

The DDR-II SRAM requires that a single address be supplied every two clock cycles to maintain full bus utilization, making this the easiest version to use for address generation. Only one operation can be requested during one clock cycle by asserting $\overline{LD}$ LOW and setting the R/$\overline{W}$ direction at K rising edge.

The Write cycle is shown in Figure 5. To initiate a write $\overline{LD}$ and R/$\overline{W}$ are set to low at the rising edge of K. The write address is captured at the same time. The first burst of Data is captured at the next rising edge of K. The second burst is captured at the $\overline{K}$ rising edge, the third burst at K rising edge, and the fourth burst at the $\overline{K}$ rising edge. Thus it takes two clocks to register all pieces of data required for the write to complete. The actual Write proceeds during the next Write cycle. If a Read to the pending address occurs before data is written into the memory array, data is forwarded from the holding register; therefore, supplied data is always coherent.

The Read cycle is shown in Figure 6. To initiate a read, $\overline{LD}$ is set to low and R/$\overline{W}$ is set to high at the rising edge of K. The address is captured at the same K rising edge. Internally, data is read from the memory array during the next clock cycle. It is registered, and then data is output 1.5 cycles after the read was initiated. Each burst of data is output at consecutive K and $\overline{K}$ rising edges.

This device is ideal for applications that require 4-word data chunks and in that the near-term ratio between Read and Write cycles is not near one, as discussed previously. The low address rate and double data rate operation permit minimum pin count on any device that interfaces with the SRAMs, and simplify high-frequency operation. This device is best utilized in applications requiring data streaming or in that periods of unidirectional bus operation is sustained.

13 DDR-II versus DDR-II+/II+ Xtreme

Details of DDR-II/II+/II+ Xtreme operation are explained in Application Note AN4065.

Table 4. Summary - DDR-II versus DDR-II+/II+ Xtreme

	DDR-II	DDR-II+	DDR-II+ Xtreme
Max frequency	Burst of 2: 333 MHz Burst of 4: 333 MHz	Burst of 2: 550 MHz Burst of 4: 550 MHz	Burst of 2: 663 MHz
Initial latency	1.5 clock cycles	2.0/2.5 cycles	2.5 cycles
Echo clocks	Yes	Yes	Yes
Density	18 Mb/36 Mb/ 72 Mb/144 Mb	18 Mb/36 Mb/ 72 Mb/144 Mb	36 Mb/ 72 Mb
Power supply	1.8 V	1.8 V	1.8 V
ODT enabled device option	No	Yes	Yes
Data valid signal - QVLD	No	Yes	Yes

14 DDR-II Separate I/O

DDR-II Separate I/O is an architecture that is a hybrid between QDR-II and DDR-II devices. DDR-II SIO has separate input and output bus so this part is very similar to the QDR-II option. The only difference is that the DDR-II SIO can only perform one operation per clock cycle. Unlike the QDR-II burst of 2, DDR-II SIO cannot perform internal data forwarding where it can write and immediately read from the same address. The controller has to wait for the write data to be committed to the memory array before reading from the same address.

This architecture was created to overcome the two factors limiting the frequency of QDR-II 2-word burst device-signaling frequency of the address bus and the memory array response time. DDR-II separate I/O relieves those frequency limitations but incurs a new penalty: one half of the average bus utilization of QDR-II 2-word burst. In some applications, a constant address request rate must be maintained to preserve system throughput, and no bus turn-around penalties can be tolerated. DDR-II separate I/O was created for exactly this purpose. The average collective data pin utilization is always 50 percent, regardless of Read to Write ratio. This performance level is superior to the older architectures (NoBL and Standard Synchronous) under all bus conditions, and equal to those architectures for unidirectional data transactions.

15 QDR-IV

The QDR-IV SRAM is a high-performance memory device optimized to maximize the number of random transactions per second by the use of two independent bidirectional data ports. These ports are equipped with DDR interfaces and designated as port A and port B respectively. Accesses to these two data ports are concurrent and independent of each other. Access to each port is through a common address bus running at DDR. The control signals are running at SDR and determine if a read or write should be performed.

QDR-IV has three types of differential clocks CK/CK# for address and command clocking, DKA/DKA#/DKB/DKB# for data input clocking and QKA/QKA#/QKB/QKB# for data output clocking.

Addresses for port A are latched on the rising edge of the input clock (CK), and addresses for port B are latched on the falling edge of the input clock (CK).

The QDR-IV SRAM read and write commands are driven by the control inputs (LDA#, LDB#, RWA#, and RWB#) and the Address Bus.

The port A control inputs (LDA# and RWA#) are sampled at the **rising** edge of the input clock. The port B control inputs (LDB# and RWB#) are sampled at the **falling** edge of the input clock.

*For port A, w*hen LDA# = 0 and RWA# = 1, a read operation is initiated. When LDA# = 0 and RWA# = 0, a write operation is initiated.

*For port B, w*hen LDB# = 0 and RWB# = 1, a read operation is initiated. When LDB# = 0 and RWB# = 0, a write operation is initiated.

QDR-IV also supports features such as bus inversion, address bus parity, ECC, de-skew, configurable ODT and impedance,

The QDR-IV family includes the following:

- QDR-IV High Performance (HP) SRAM: A two-word burst architecture device with two accesses for each cycle at a maximum frequency of 667 MHz and with a read latency of five clock cycles.
- QDR-IV Xtreme Performance (XP) SRAM: A banked two-word burst architecture device with two accesses for each cycle at a maximum frequency of 1066 MHz and with a read latency of eight clock cycles.

Table 5. Operational Modes

	QDR-IV HP SRAM		QDR-IV XP SRAM	
Clock Frequency	600 MHz	667 MHz	933 MHz	1066 MHz
Read Latency	5 cycles	5 cycles	8 cycles	8 cycles
	8.33 ns	7.5 ns	8.57 ns	7.5 ns
Write Latency	3 cycles	3 cycles	5 cycles	5 cycles
Bus Width	x18, x36			
I/O Type	1.1 V and 1.2 V POD 1.2 V and 1.25 V HSTL/SSTL			
Package	361 FCBGA			
Port configuration	Bidirectional R/W ports			
Density	144 Mb, 72 Mb			

As Figure 7 and Figure 8 show, the Port A read data comes out of the data pins exactly five Read Latency (RL) clock cycles later in the case of the QDR-IV HP SRAM or eight RL clock cycles later in the case of the QDR-IV XP SRAM. The data is available after the number of RL clock cycles from the **rising** edge of the CK signal when the READ command was issued.

The Port A write data is supplied to the data pins exactly three Write Latency (WL) clock cycles later in the case of the QDR-IV HP SRAM or five WL clock cycles later in the case of QDR-IV XP SRAM. The data comes after the number of WL clock cycles from the **rising** edge of the CK signal when the WRITE command was issued.

Port B works same way as port A but data is available after the number of RL clock cycles from the **falling** edge of the CK signal when the READ command was issued and the data comes after the number of WL clock cycles from the **falling** edge of the CK signal when the WRITE command was issued.

Figure 7. Read Timing

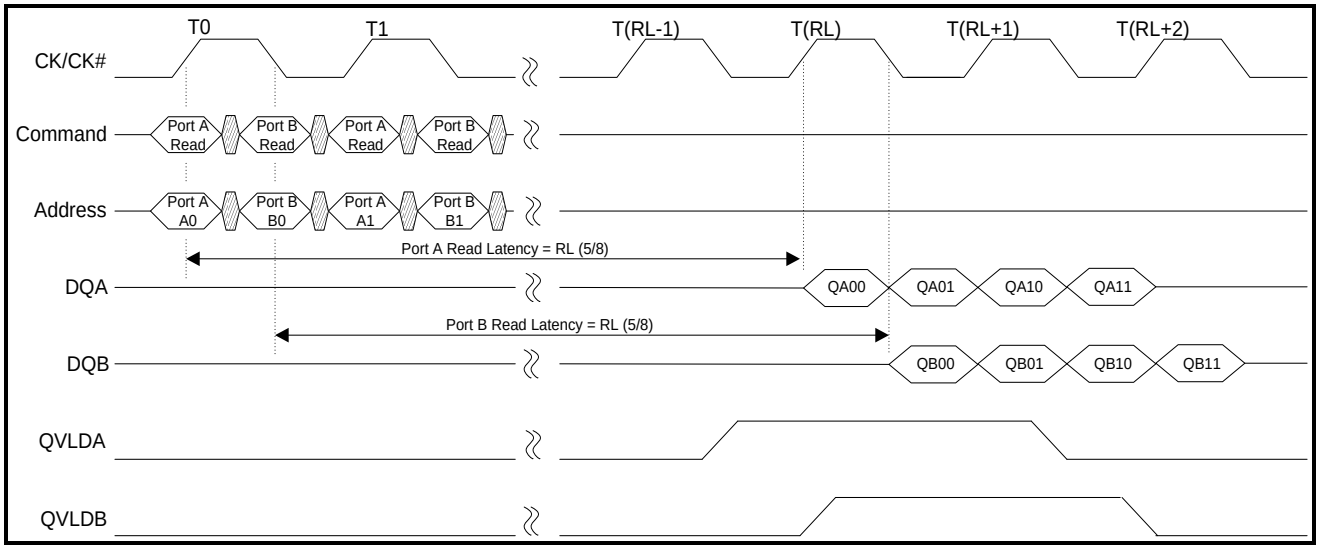

Figure 8. Write Timing

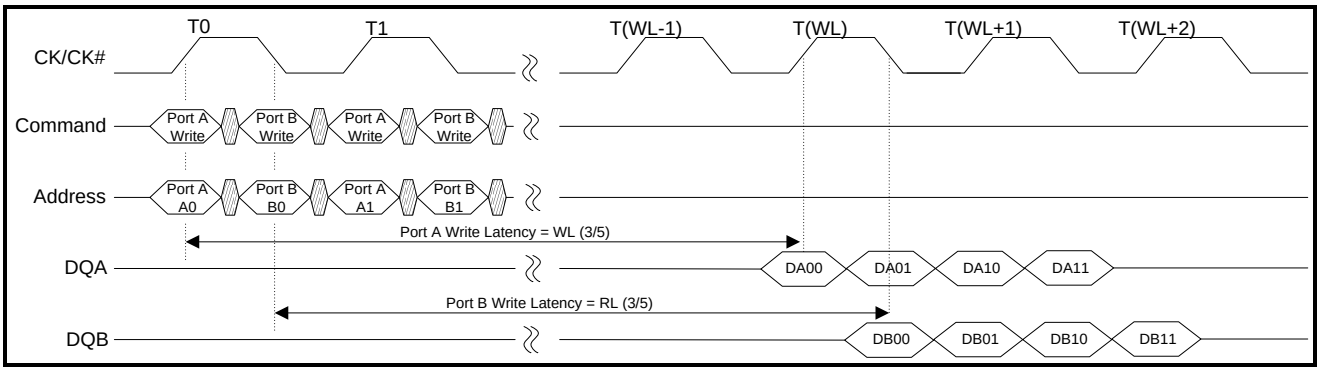

The QDR-IV XP SRAM is designed with eight internal banks. The lower three address pins (A2, A1, and A0) select the bank that will be accessed during a read or write operation.

The only banking restriction is that a particular bank can be accessed only once for every clock cycle. This banking restriction applies only to the QDR-IV XP SRAM. The QDR-IV HP SRAM does not have any banking restriction.

If a banking violation occurs, the read/write operation of Port B is denied. Because the Port A address is sampled on the rising edge of the input clock, there are no restrictions to Port A access.

Figure 9. QDR-IV XP SRAM – Write/Read Operation

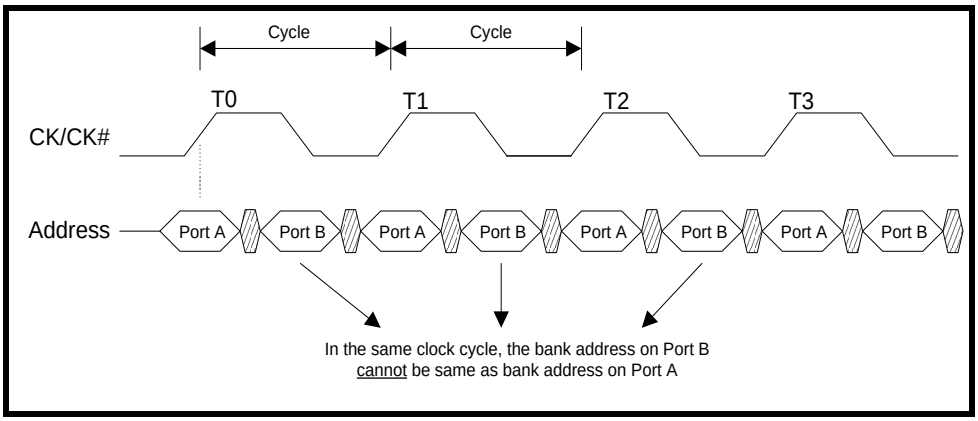

Figure 10. QDR-IV HP SRAM – Write/Read Operation

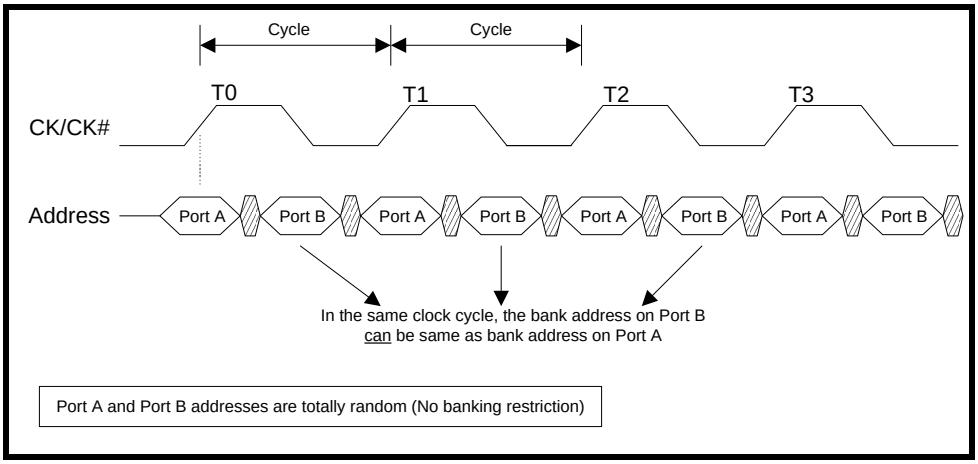

QDR-IV is the highest performance, standards-based memory solution available on the market. Its high performance, coupled with differentiated features such as dual bidirectional ports, ECC, bus inversion, ODT, and address parity, make it the optimal solution for networking systems. The advantages of QDR-IV also apply to other systems that require high-speed performance and signal integrity, such as high-performance computing and image processing.

16 Making the Optimal Choice

If the SRAM interfaces with an older processor, the choices are limited because older processor architectures were designed for Standard Synchronous SRAMs. Newer processors are designed for DDR-II/II+/II+ Xtreme SRAMs. Some network processors are designed for NoBL, QDR-II/II+, QDR-II+ Xtreme, DDR-II/II+, DDR-II+ Xtreme and QDR-IV SRAMs.

Network applications require maximum pin utilization. This strongly favors QDR-II/II+/II+ Xtreme, DDR-II/II+/II+ Xtreme and QDR-IV SRAMs due to double data rate operation on all data pins. Hence DDR-II/II+/II+ Xtreme, QDRII/II+/II+ Xtreme or QDR-IV SRAM is recommended for new designs with higher bandwidth requirement.

17 Device Cross Reference

Use the Competitor Memory Device Cross Reference Tool to determine Cypress memory parts compatible with devices from other vendors.

18 Summary

New designs should utilize the SRAM that best satisfies the requirements of a given application.

Cypress's high-bandwidth QDR-IV SRAMs are designed for high-speed performance and they satisfy demanding network functions, such as updating statistics, tracking flow states, scheduling packets, and performing table lookups.

QDR-II/II+ SRAMs are ideal for systems with back-to-back Read-Write operations. DDR-II/II+ SRAMs are preferred when there are consecutive reads or writes. The appropriate choice is determined by the actual read to write ratio and desired burst length, shown in the figures. QDR-II/II+ and DDR-II/II+ SRAMs are also recommended when pin count must be optimized and is facilitated by double data rate operation of the buses. Maximum clocking design flexibility is provided with up to three sets of clock pins.

The older SRAM architectures such as NoBL and Standard Synchronous are recommended when compatibility with an existing controller or processor is required or when the low latency of Flow-through operation is desired. NoBL architecture is preferred over Standard Synchronous where possible, since Standard Synchronous was designed for microprocessor cache applications in processors that are now being phased out of production.

Making the optimal SRAM choice is a matter of price versus performance. For cache users, the chipset choice and direct processor interface dictates the SRAM selection. Network applications are experiencing unprecedented new performance levels, thanks to QDR-II/II+/IV and DDR-II/II+ devices. Cypress offers a wide variety of synchronous SRAMs to meet the needs of all types of systems and is committed to continuing our leadership role in recognizing our customers' needs.

To know more about the Cypress Synchronous SRAM products and to get technical assistance, initiate a customer support case at www.cypress.com/support.

Document History

Document Title: AN4011 - Choosing the Right Cypress Synchronous SRAM

Document Number: 001-15488

Revision	ECN	Orig. of Change	Submission Date	Description of Change
**	1051060	SFV	05/16/2007	New Application note.
*A	1787965	VIDB	12/12/2007	Update copyright. Add source disclaimer, revision disclaimer. This note had no technical updates. There is an associated project but it was not updated.
*B	3112578	VIDB	12/16/2010	No technical updates. Updated copyright and revision. Added Document History heading.
*C	3356801	OSN	08/30/2011	Changed table 1 to feature list format from part listing format. Removed references to QDRI Added QDRII+/DDRII+ device descriptions and comparison table. Inserted figures 1 and 2 Added Table 2 and modified description under Standard Synchronous SRAMs Removed references to QDR-II/DDR-II Application Notes in Table 1.
*D	3735731	SKAP	09/06/2012	Updated template. Minor text edits.
*E	4227355	PRIT	12/20/2013	Updated in new template. Completing Sunset Review.
*F	4346172	PRIT	04/14/2014	Updated document with details of QDR-IV.
*G	4676098	DEVM	03/13/2015	Updated document with device cross reference and technical support link Updated template
*H	5149565	DEVM	02/24/2016	Updated Device Cross Reference link. Updated template
*I	5585611	DEVM	01/13/2017	Updated template
*J	5826167	AESATMP9	07/20/2017	Updated logo and copyright.

Worldwide Sales and Design Support

Cypress maintains a worldwide network of offices, solution centers, manufacturer's representatives, and distributors. To find the office closest to you, visit us at Cypress Locations.

Products

ARM® Cortex® Microcontrollers	cypress.com/arm
Automotive	cypress.com/automotive
Clocks & Buffers	cypress.com/clocks
Interface	cypress.com/interface
Internet of Things	cypress.com/iot
Memory	cypress.com/memory
Microcontrollers	cypress.com/mcu
PSoC	cypress.com/psoc
Power Management ICs	cypress.com/pmic
Touch Sensing	cypress.com/touch
USB Controllers	cypress.com/usb
Wireless Connectivity	cypress.com/wireless

PSoC® Solutions

PSoC 1 | PSoC 3 | PSoC 4 | PSoC 5LP | PSoC 6

Cypress Developer Community

Forums | WICED IOT Forums | Projects | Videos | Blogs | Training | Components

Technical Support

cypress.com/support

Cypress Semiconductor
198 Champion Court
San Jose, CA 95134-1709