Freescale Semiconductor
Application Note

Document Number: AN2922
Rev. 1, 12/2005

SDRAM and PowerQUICC™ II: An Introduction

by *Paul Chang and Randy Ryder*
NCSD Applications
Freescale Semiconductor, Inc.
Austin, TX

Introduction

Memory technology encompasses a wide range of options. The focus of this application note is memory called synchronous dynamic random access memory (SDRAM). Dynamic random access memory (DRAM) has evolved over the years into the most common form of memory called synchronous DRAM, or SDRAM.

These are further subdivided between single and double data rate memories (SDR/DDR). Earlier versions of DRAM were just known as conventional DRAMs. The first generation of memories were called page-mode DRAMs, the row address signal was kept active for multiple column address strobes (early form of bursting). The next form of DRAM was called extended data-out (EDO) DRAM, where an extra signal called output enable (OE) was added to control data outputs. The problem with early memory technologies was that they were asynchronous to the system clock, which meant the processor always had to wait on the read or write request based on the internal timing of the memory itself. SDRAM technology adds a separate clock signal as an input to the memory that synchronizes the data with the processor. SDR-SDRAM is the focus of this discussion.

Contents

freescale™ semiconductor

This application note is divided into the following sections:

- "Architecture," covers the basic internal memory organization.
- "Operation," describes the key concepts involved in programming the memory controller to communicate with the memory.
- "Pinout," explains the memory's external pins.
- "Bank Interleave Feature," describes the theory behind bank and page-based interleaving methods.
- "Physical Connections," describes the logical and physical relationship of the address and data lines.

Architecture

SDRAM memory is organized in a logical manner as a grid array. The basic components of storage within these arrays are cells which contain a capacitor and an FET as the control gate. Each cell stores 1 bit, a defining characteristic of DRAM architecture. Figure 1 shows 4-bit cells. The WL notation indicates a wordline, or row, which is connected to the control gate of the FET. The MOSFET opens and closes based on the voltage on the wordline, charging or draining the capacitor attached to it.

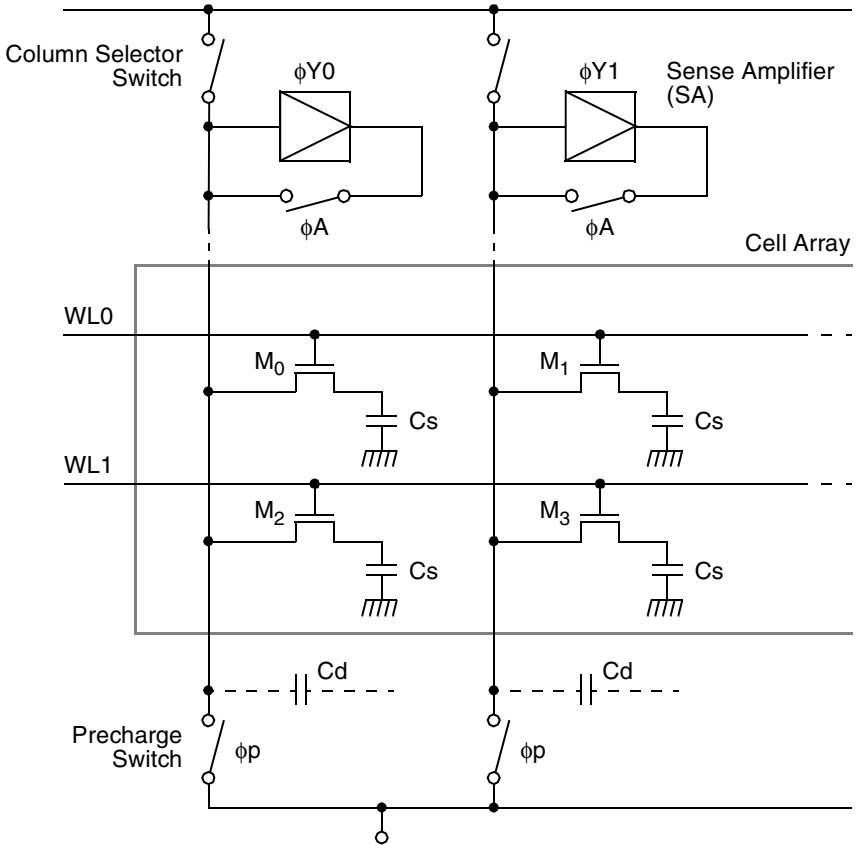

Figure 1. SRAM Bit Cells

After the appropriate wordline is charged, the column selector is switched to access the appropriate capacitor for the upcoming operation. Due to the natural eventual discharge of the capacitor, the cells have to be refreshed; this is the reason for the unique refresh command which is characteristic of SDRAMs. In a refresh cycle, the data stored in the bit cells is read out and then rewritten back. Additional circuitry, called sense amplifiers, is added to detect and amplify the logic levels being read out of memory.

This storage method is different from other architectures, such as static RAMs (SRAM), where multiple gates are used to hold the charge without the need for having their state refreshed. Figure 2 shows a RAM array with 8 row and column lines which form a 64-bit array. Remember, the row count is always greater than or equal to the column count due to a variety of design reasons.

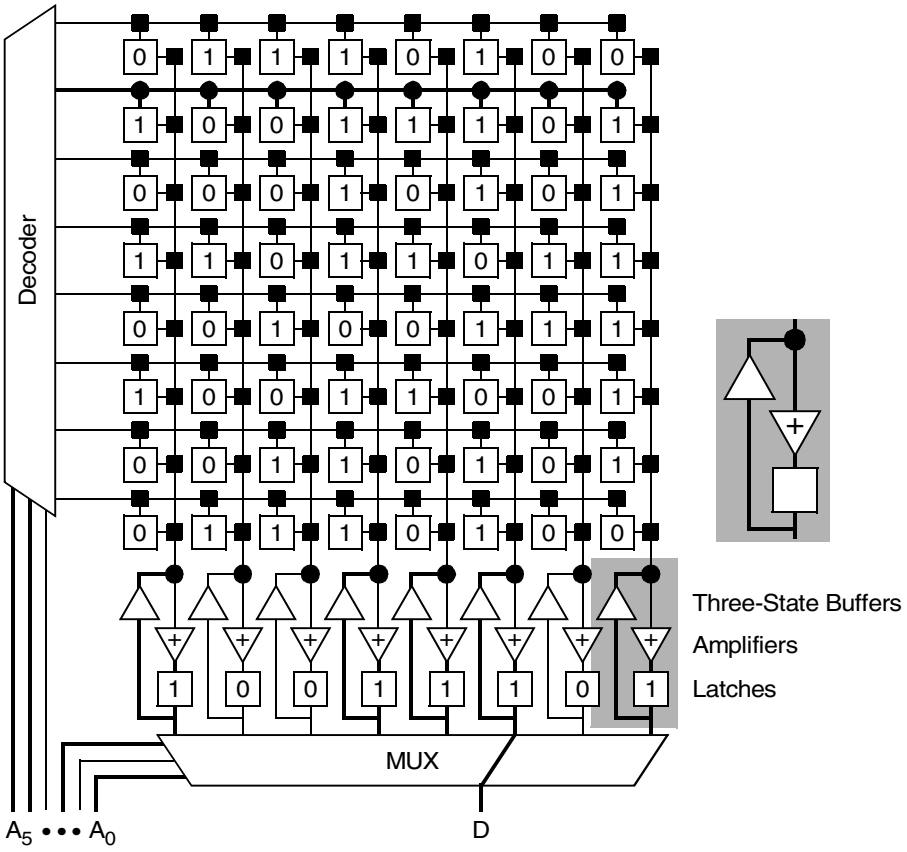

Figure 2. DRAM Array

This array holds 64 bits but can only read or write 1 bit of information at a time. For this reason, SDRAM has further segmentations in its architecture. To output a greater number of bits at once, an SDRAM has multiple arrays. An 8-bit port size SDRAM has 8 arrays, each the same size. When an access is requested, the memory only requires 1 address to access 1 bit from each of the 8 arrays to provide 8 bits of data.

The final significant concept in SDRAM technology is the concept of a bank. Internal banks are added to the internal structure of SDRAMs to meet the requirement for increased memory sizes. This type of segmentation provides several advantages, such as decreased power consumption and less cost due to fewer physical address lines being used. Early SDRAMs had a maximum block of 16 Mb in the form of eight 2 Mb arrays. As density grew to 32 Mb, 2 banks of 16 Mb total storage were designed. In other words, the 2 banks were comprised of eight 2 Mb arrays. Current densities of 64 Mb or greater are composed of 4 banks. The PowerQUICC™ II family supports SDRAMs with a maximum of 4 internal banks (not to be confused with DIMM banks on a motherboard which are sometimes referred to as ranks).

Figure 3 shows a 64 Mb SDRAM device with 12 address lines coming into the device, which are multiplexed for row and column. Also, a separate auto precharge line is used, as well as the bank activate lines. Two bank activate lines are used because 4 banks are available. The control signals, such as $\overline{RAS}$, $\overline{CAS}$, $\overline{WE}$, and DQM are all inputs to the SDRAM control logic. The control signals are discussed later. Within each bank are eight 4096×512 arrays which determine a port size of 8 for this device.

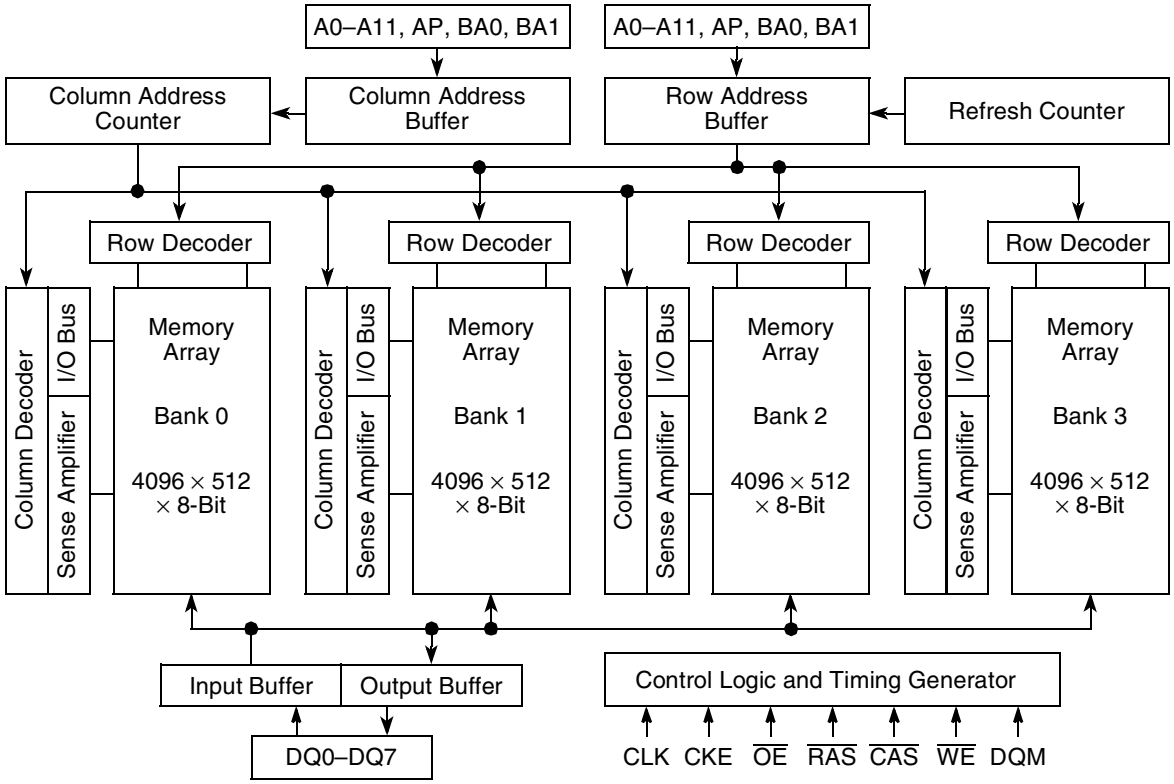

Figure 3. 64 Mb SDRAM Device Block Diagram

To understand the nomenclature used to describe the size of an SDRAM memory, for example, Micron's MT48LC16M8A2 device is described below (see "References," on page 17). This is a 128 Mb device with 4M × 8 × 4 banks. The part number, 16M8, indicates a 16 Mb device with an 8-bit internal port size. The arrangement is 4K rows by 1K column for each array and has 8 arrays per bank, with 4 internal banks per device. The 8 arrays per bank provide the 8-bit port width of this device. If each bank had 16 such arrays, its port width would be 16 bits, as is the case for the 8M16 device. Figure 4 shows an example of the internal bank architecture, focusing on how the 8 arrays comprise the 8 bits of data output on each bank.

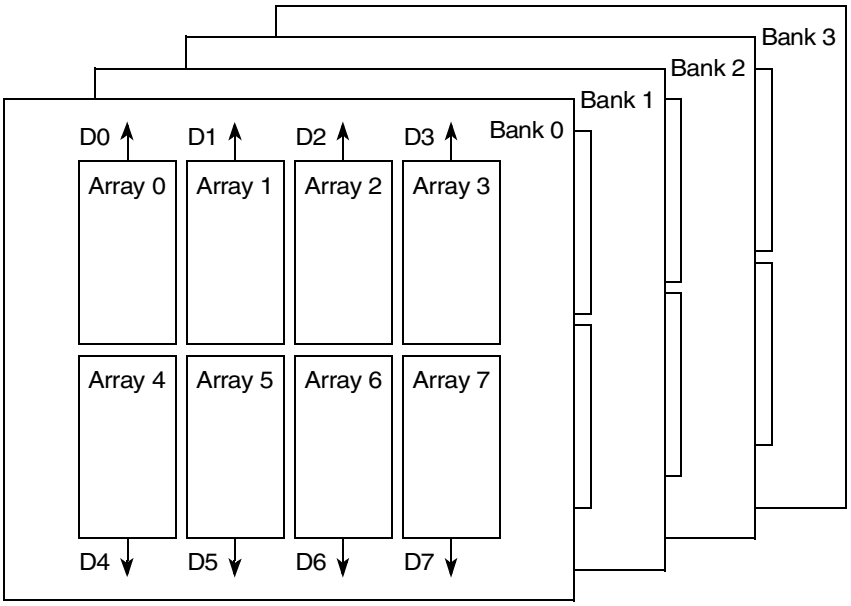

Figure 4. Internal Layout

Operation

SDRAM operation consists of combinations of signals that translate into specific types of SDRAM commands. In addition to the utilization of address lines and data lines, there are six memory-specific signals, three of which determine the command that is issued. The three command signals are write enable ($\overline{WE}$), row address strobe ($\overline{RAS}$), and column address strobe ($\overline{CAS}$). Using these three signals in parallel, all main SDRAM commands can be issued. The basic commands are read, write, precharge, and the refresh operation. Other commands can also be issued through a similar combination of $\overline{WE}$, $\overline{CAS}$, and $\overline{RAS}$. Table 1 shows a truth table with common SDRAM commands.

Three other important signals are chip select ($\overline{CS}$), input/output mask (DQM), and bank address (BA), also known as bank select (BS). The $\overline{CS}$ signal is active low and enables the command decoder inside the control logic of the memory. DQM gives the memory controller the ability to select only 1 byte at a time from memory in cases where multiple memories are configured for port sizes greater than 8 bits (as in a 64-bit data bus). DQM is active low and masks both read and write accesses when requested. The bank address signal selects which one of the internal banks the SDRAM command is being applied to.

Table 1. Command Truth Table

Command	$\overline{CS}$	$\overline{RAS}$	$\overline{CAS}$	$\overline{WE}$	ADDR
Command inhibit (NOP)	H	X	X	X	X
No operation (NOP)	L	H	H	H	X
Active (select bank and activate row)	L	L	H	H	Bank/row
READ (select bank and column, and start READ burst)	L	H	L	H	Bank/column
WRITE (select bank and column, and start WRITE burst)	L	H	L	L	Bank/column
PRECHARGE (deactivate row in bank or banks)	L	L	H	L	X
AUTO REFRESH or SELF REFRESH (enter self refresh mode)	L	L	L	H	X

Note that in order to do a read or write operation, the active (ACT) command has to be issued first. ACT will cause a specific row and bank to be activated before the read or write command is issued. Although most controllers control the refresh rate of the memory, a refresh command could also be issued by using the PSDMR register in the memory controller. The refresh command, whether issued by the user or done automatically by the controller, recharges the cells in each row by reading that row out and subsequently writing it back in. After the completion of one row, the next row is refreshed. The precharge command closes a particular row in one bank or multiple banks. When a row access is finished, a new row cannot be opened until the previous row has been closed.

There are more complex commands that can be issued if the controller supports such commands. Two of these commands are burst reads and burst writes. The basic premise behind a burst access is that the row address is kept latched in while the column address is automatically incremented internal to the memory, without the need to issue the $\overline{CAS}$ command. Most bursts can be programmed to occur in 2, 4, or 8 column locations, allowing for a maximum of 8 consecutive accesses on one row. It is during burst modes that concepts of page hits and misses occur. A page is another name for a row. A page hit means that the next access that the controller is requesting occurs on the same row that's already open. If the controller requests an access on a page or row that is currently not open, the access is referred to as a page miss and has to be retried.

SDRAM has a special sequence of operations that must occur in order for the memory to be initialized. All SDRAMs have a mode register which must be programmed before operation begins. The mode register contains configurable parameters. Figure 5 illustrates a typical mode register.

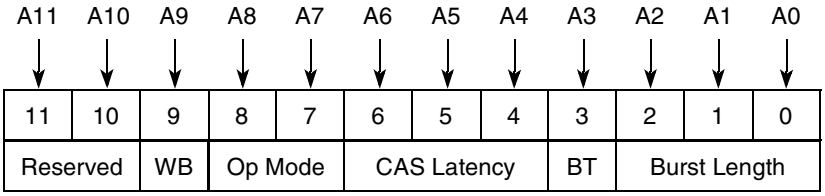

Figure 5. Typical Mode Register

WB (write burst mode)—Sets whether the burst length will be applicable to reads and writes or just reads.

Op mode (operating mode)—Usually programmed to a default value, in this case, 00.

CAS latency—The ratio of column access time to clock access time; for an example, check Micron's data sheet (see "References," on page 17) for available values, usually 2 or 3.

BT (burst type)—Determines whether the bursts are interleaved or sequential; usually defaulted to sequential.

Burst length—Available burst lengths; usually 2, 4, or 8 column locations are available for bursting.

Once the proper configuration has been chosen, the command is given via the SDRAM command register and a subsequent dummy write to the SDRAM is done. The dummy write contains the mode bits on the low order address bits. Following is an example of a typical programming sequence based on a Micron mode register, with the processor in 60x compatible mode:

Port size = 64 bits (PowerQUICC II must be programmed for 60x compatible mode)

BL = 4

BT = 0 (sequential)

CAS latency = 2

Op Mode = 00 (normal)

1. Issue precharge command to all banks
2. Issue 8 refresh cycles
3. Write to PSDMR register to issue mode register write command
4. Write 0x00 to 0x0000_0122

Note that in single mode, the mode register configuration does not have to be written onto the low order address bits. In single mode, setting PSDMR[OP] = 011 and doing a dummy write to memory suffices.

Pinout

Figure 6 shows a Micron, 128 Mb, 54-pin SDRAM in three different data port sizes:

- MT48LC32M4A2 (x4)—8 Meg × 4 × 4 banks
- MT48LC16M8A2 (x8)—4 Meg × 8 × 4 banks
- MT48LC8M16A2 (x16)—2 Meg × 16 × 4 banks

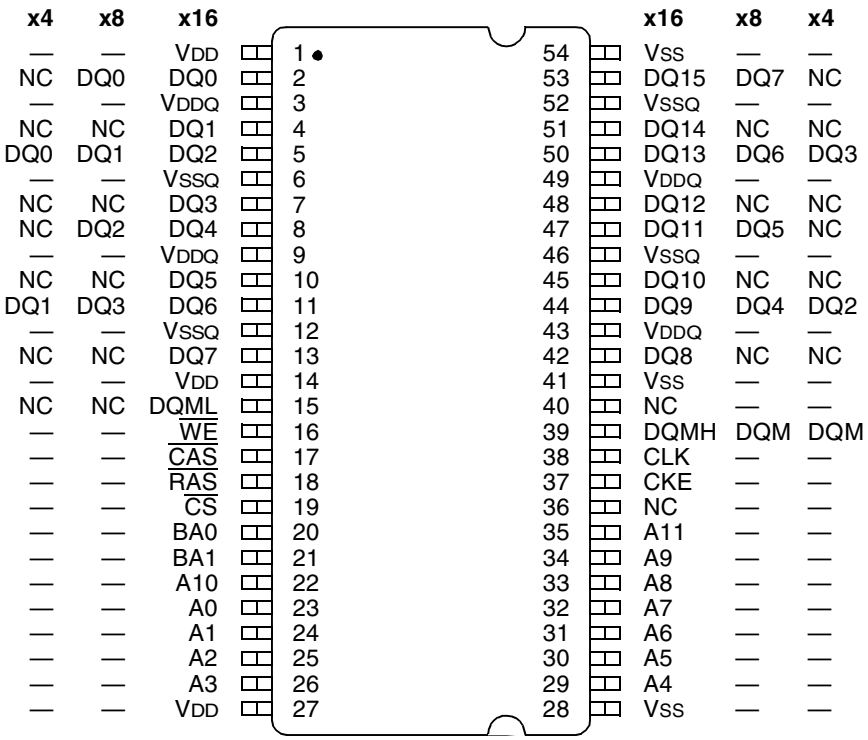

NOTE: A dash (—) indicates the x8 and x4 pin functions are the same as the x16 pin function.

	32 Meg × 4	**16 Meg × 8**	**8 Meg × 16**
Configuration	8 Meg × 4 × 4 banks	4 Meg × 8 × 4 banks	2 Meg × 16 × 4 banks
Refresh count	4K	4K	4K
Row addressing	4K (A0–A11)	4K (A0–A11)	4K (A0–A11)
Bank addressing	4 (BA0, BA1)	4 (BA0, BA1)	4 (BA0, BA1)
Column addressing	2K (A0–A9, A11)	1K (A0–A9)	512 (A0–A8)

Figure 6. Micron SDRAM

The above SDRAM device has 2 bank select lines to support a total of 4 banks, 12 row address lines, and various column address lines depending on the port size desired. The SDRAM device in the subsequent sections has 4 banks, 12 rows, and 9 columns.

SDRAM Physical Pin Description

Table 2 summarizes all the physical pins on the SDRAM device.

Table 2. SDRAM Physical Pin Functions

Pin	Function
V_{DD}	Power supply: +3.3 V ± 0.3 V
V_{DDQ}	Isolated DQ power on die for better noise immunity
V_{SSQ}	Isolated DQ GND on die for better noise immunity

Table 2. SDRAM Physical Pin Functions (continued)

Pin	Function
V_{SS}	GND
CLK	Input clock pin
CKE	Clock enable
NC	No connect
A[0:9, 11,12]	Address pins
A10/AP	Auto precharge pin (A10/SDA10 is internally mux'd with AP)
DQ[0:15]	Data pins
BA[0:1]	Bank selects
DQM_H/L	Byte selects Note: DQM_L is only available in x16 configuration
$\overline{WE}$	Write enable
$\overline{RAS}$	Row address strobe (latch for row)
$\overline{CAS}$	Column address strobe (latch for column)
$\overline{CS}$	Chip select

Bank Interleave Feature

In a typical SDRAM access, if a missed page (row) is in a different SDRAM bank than the current opened page, the memory controller must deactivate the previously accessed bank first and then activate the new bank. A READ or WRITE access follows activation of the new bank.

To improve performance, the PowerQUICC II SDRAM controller supports a feature called bank interleaving. This means that if a missed page is in a different SDRAM bank than the currently opened page, the PowerQUICC II SDRAM controller can issue an ACTIVATE command to the new bank prior to a DEACTIVATE command to the old bank. By using multi-tasking, this feature eliminates the wait time between deactivating, activating, and accessing sequences.

This feature can be performed if both pages reside on different SDRAM devices or on different internal SDRAM banks as shown in Figure 7.

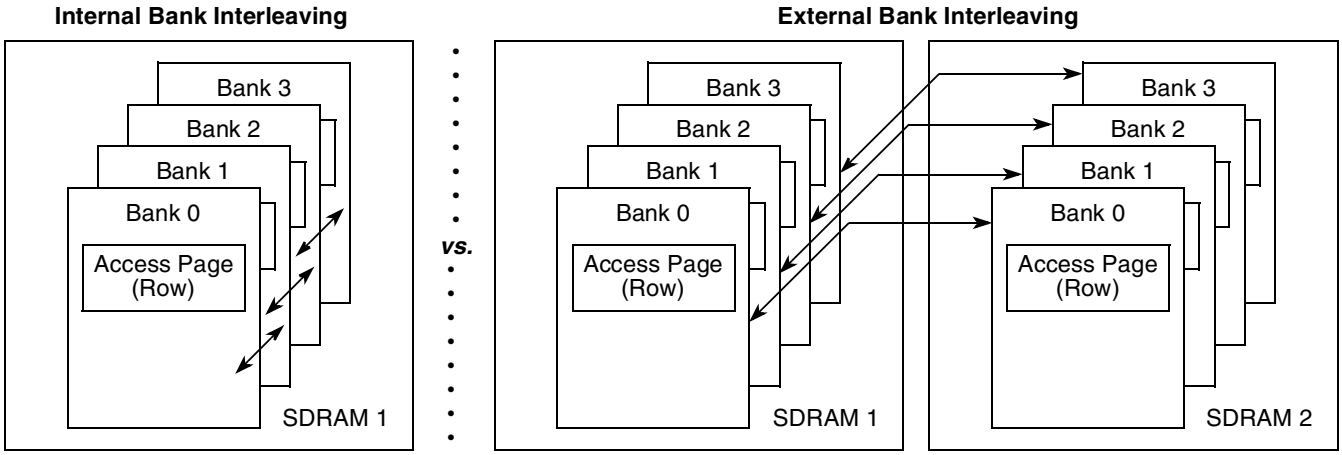

Figure 7. Internal versus External Bank Interleaving

Within the internal bank interleaving there are two types of interleaving modes that take advantage of the bank-interleaving feature on PowerQUICC II: page-based interleaving (PBI) and bank-based interleaving (BBI).

PBI and BBI are two different ways of reading and storing bits in a SDRAM device. Before explaining how BBI and PBI work, the concept of address partitioning must be understood.

NOTE

Page-based and bank-based interleaving modes are not available in external bank interleaving.

Address Multiplexing

An innovation of SDRAM devices is that they do not physically implement 32 address pins. A typical SDRAM device has only 9–13 physical address pins, whereas the 60x bus provides 32 address pins. What makes it possible to interface between SDRAM and PowerQUICC II when there is an uneven number of pins is through an internal address multiplexing via the PowerQUICC II SDRAM controller in single mode. Address multiplexing can also be done externally via an external address MUX chip when running in 60x mode.

From a 32-bit logical address, the SDRAM controller partitions the logical address into either BBI or PBI format. The SDRAM controller then extracts row and column information from the partitioned address where this information is used to map out certain sections on the SDRAM device for memory accesses. Row and column are latched into the SDRAM device in two separate phases with the assertion of $\overline{\text{RAS}}$ followed by $\overline{\text{CAS}}$.

NOTE

In single mode, the user would not see the actual 32-bit logical address displayed on the 60x bus. Instead, a 32-bit address containing either row + bank select or column + bank select information is displayed.

To illustrate with an example, assume that PowerQUICC II is configured for single mode and PBI. In addition, the data port size is 64 bits and the logical address to be accessed is 0x1000. Given the above scenario, address 0x1000 is partitioned into the binary format shown in Figure 8.

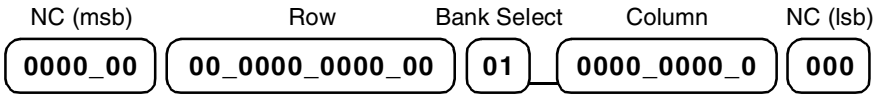

Figure 8. PBI Partitioning Example

A typical memory access requires the execution of an ACTIVATE command and a READ/WRITE command.

During an ACTIVATE command when row information is latched into the SDRAM, the address on the address bus looks like the diagram shown in Figure 9.

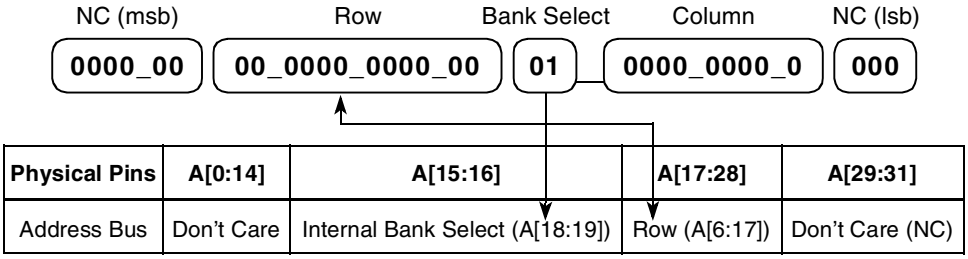

Figure 9. ACTIVATE Command

The 12 row addresses shown in the PBI bus partition (see Figure 8) correspond to the 12 physical address pins on the SDRAM. These 12 bits of row along with 2 bits of bank select information are latched into SDRAM when $\overline{\text{RAS}}$ is asserted.

Similarly, when column information is latched into the SDRAM during a READ/WRITE command, the address on the address bus looks like the diagram shown in Figure 10.

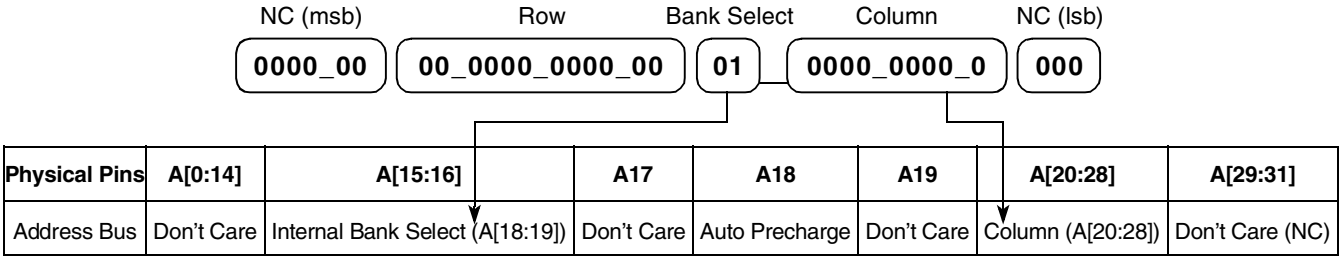

Figure 10. READ/WRITE Command

The 9 bits of column information along with bank select information are latched into SDRAM with the assertion of $\overline{\text{CAS}}$. Note that row and column are sharing the12 SDRAM physical address pins.

Refer to the *SDRAM Configuration Examples* chapter in various PowerQUICC II reference manuals for the bank-based interleaving (BBI) example.

Bank-Based Interleaving (BBI) versus Page-Based Interleaving (PBI)

In comparison, both internal interleaving modes have column addresses on the lowest order bits as shown in partition (Table 3 and Table 4). The difference in the addressing scheme is that the PBI bank select lines come before the row addresses, and the BBI row addresses come before the bank select lines.

Table 3. BBI Address Partition (64-Bit Port Size)

A[0:5]	A[6:7]	A[8:19]	A[20:28]	A[29:31]
msb of start address	Internal bank select	Row	Column	lsb

Table 4. PBI Address Partition (64-Bit Port Size)

A[0:5]	A[6:17]	A[18:19]	A[20:28]	A[29:31]
msb of start address	Row	Bank select	Column	lsb

Page-based interleaving has better overall performance than bank-based interleaving. By moving bank select ahead of column and below row addresses, it will guarantee that bank interleaving will happen more often. The sequence of deactivating a row, activating a new row, and then accessing it consumes more time than interleaving between different banks. More frequent use of the internal bank-interleaving feature means less overall SDRAM access time and, thus, yields better overall performance.

Physical Connections

Address Pins Connection

Figure 11 and Figure 12 illustrate the address connections between PowerQUICC II and a SDRAM device based on different data port size requirements. Important concepts about these connections are discussed in the following sections.

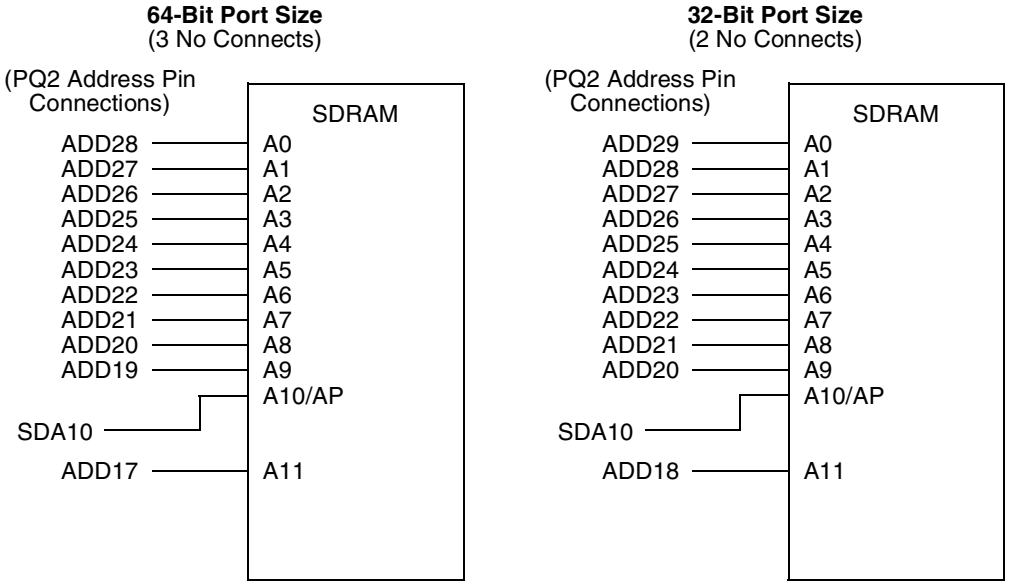

Figure 11. 32-/64-Bit Port Size Addressing

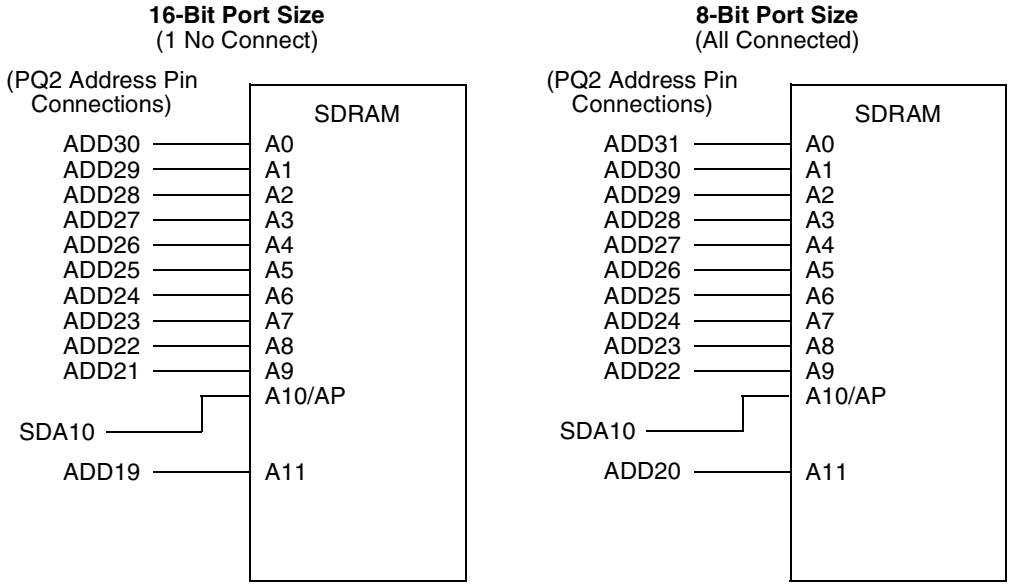

Figure 12. 8-/16-Bit Port Size Addressing

Big and Small Endianess

The PPC architecture implements big endian as opposed to little endian. The difference in endianess affects how address lines are connected between PowerQUICC II and the SDRAM device. In the PowerPC™ architecture, the most significant bit (msb) is logical address bit 0 and the least significant bit (lsb) is logical address bit 31. The addressing scheme in the SDRAM device is the opposite, where address bit 31 is msb and address bit 0 is lsb. Thus, the address connection between PowerQUICC II and SDRAM is reversed, such that address bit 0 of SDRAM is connected to the low-order address bit of PowerQUICC II and address bit 11 of SDRAM is connected to the high-order bit as shown in Figure 11 and Figure 12.

No Connect (NC) in Relation to Data Port Size

Some of the address pins on PowerQUICC II are left unconnected depending on the data port size. For different data port sizes, such as 64, 32, 16, and 8 bits, the number of unconnected pins are 3, 2, 1, and 0, respectively. For example, when working with the 64-bit data port size, the last 3 low-order address pins on PowerQUICC II are don't care and, therefore, are not connected. Figure 11 shows that the starting PowerQUICC II address pin ADD28 is connected to A0 of SDRAM, followed by ADD27 to A1, and so forth until all 12 SDRAM pins are connected.

Auto Precharge on A10

SDA10 is a multi-purpose pin on the memory controller which connects to A10 on the SDRAM. It is used for the auto precharge command and to provide addressing through the use of internal muxing. Refer to Figure 11 and Figure 12 for connection information.

Chip Select Connection

Multiple chip selects ($\overline{CS}$) can be programmed to support multiple SDRAM devices. There is no limitation on the number of chip selects that can be programmed for SDRAM ($\overline{CS}$[1:11] can be allocated for SDRAM). If multiple chip selects are configured to support SDRAM on a single bus, each SDRAM device must have the same port size and run at the same frequency.

Byte Select Connection

Byte select is used to access an individual byte within a SDRAM device. If the SDRAM has x16 arrays, then connections for both DQM_L and DQM_H are used to select between the upper and lower byte. However, if the device has x8 arrays or lower, then only DQM_H is required and DQM_L is NC.

Figure 13 shows a scenario using x8 arrays SDRAM devices where all 8 DQM lines are used to form a 64-bit data port size. In the diagram, all 8 SDRAM devices are tied to the same $\overline{CS}$ and each device is tied to a different DQM.

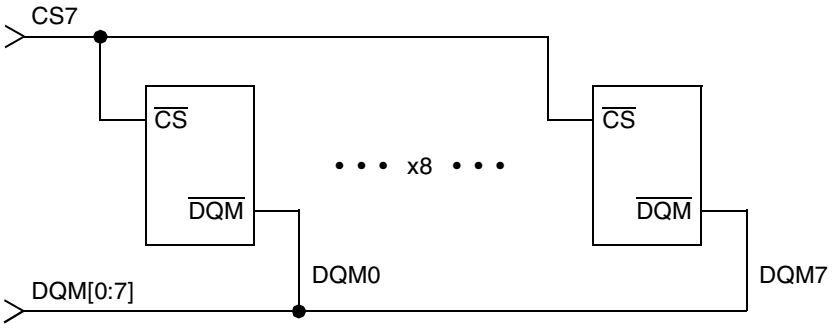

Figure 13. $\overline{CS}$/DQM Connects Example

Data Pins Connection

PowerQUICC II can utlitize up to 64 data pins in 60x compatible mode. In the following illustrations, each SDRAM device provides only 16 data pins. Two SDRAM devices can be tied together to form a 32-bit

data port size. Similarly, 4 SDRAM devices can be tied together to form a 64-bit data port size. Figure 14 and Figure 15 illustrate the connection of data pins based on different port size requirements:

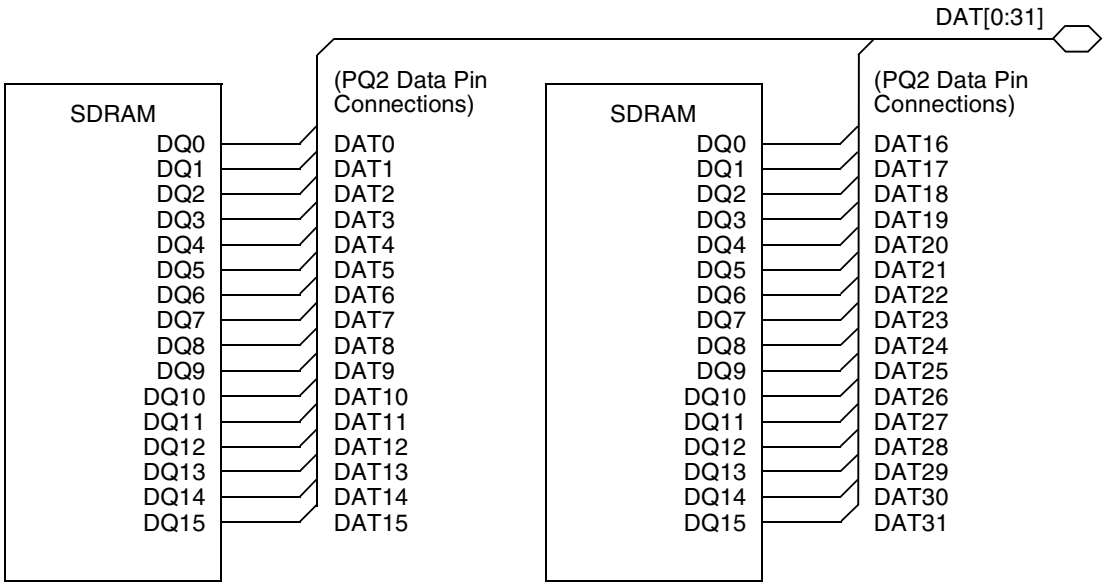

Figure 14. 32-Bit Data Port Connections

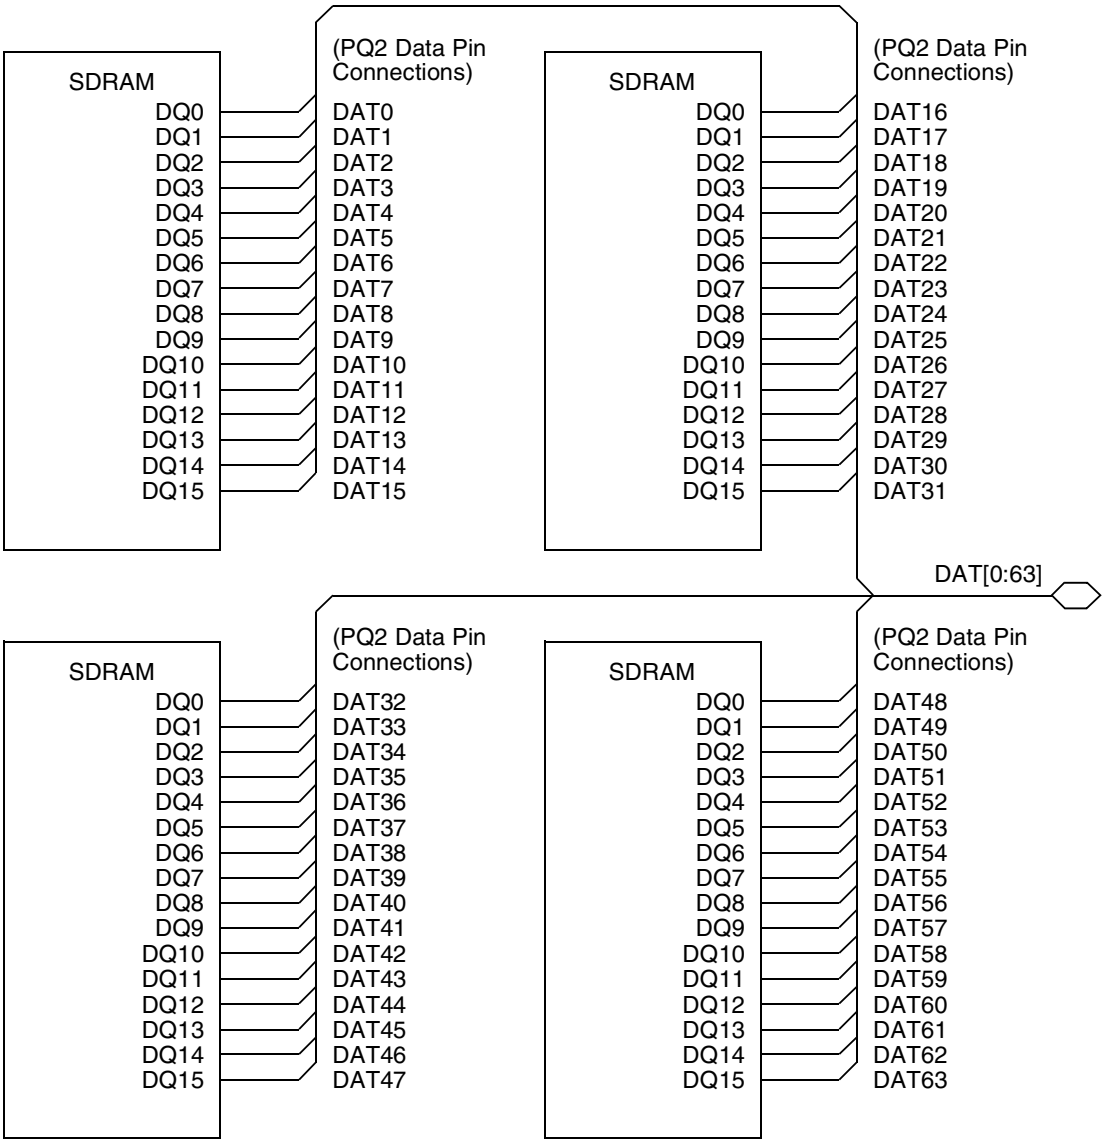

Figure 15. 64-Bit Data Port Connections

Definitions

Bank-based interleaving (BBI)—Process of accessing two banks of memory at once; the second bank of memory is read during the first bank's latter access cycle. Each SDRAM chip has two banks of memory cells, thus allowing simultaneous accesses as in interleaving.

Burst mode—The accessing of multiple words of memory at once, in a sequential manner, unlike reading or writing one word at a time. The block of data is located in a sequential address space.

Dual inline memory module (DIMM)—A 168- or 184-pin module with 64-bit data bus width. A newer module technology than the older SIMM technology.

Error correction code (ECC)—An algorithm that detects and corrects single bit errors as well as detecting, but not correcting double bit errors within a byte.

Fast page mode (FPM)—Older DRAM design that is able to access the same row multiple times without additional row access times; every access costs three clock cycles.

Page of data—Common row address, that is, the whole row of cells is a page in the array

Parity error checking—SDRAM chip checks the number of 1s in an incoming word; if the number of 1s is odd, the parity bit is set to 1, else 0 (even parity). For odd parity, when data is read back out of the chip, the parity bit is compared to a new count to see if it is equal. Not equal means an error and 8 bits are dropped; there is no error correction. Also, if two bits are swapped, no error will even be detected.

Registered memory modules versus unregistered—A module which includes a clock register to help buffer the address lines which allows device capacity to increase with the same amount of address lines. Registered memory modules induce an extra clock cycle of delay, which is negligible on system performance due to an increase in capacity.

Synchronous DRAM—Memory that is synchronized to the system clock vs. having a separate clock (asynchronous); done by having a register which typically keeps track of the number bytes being requested.

References

Synchronous DRAM, 128MB: x4, x8, x16 SDRAM, Rev. K, Micron Technology, Inc., May 2005.

For a current list of PowerQUICC II documentation, refer to the website listed on the back cover of this document.

Revision History

Table 5. Document Revision History

Rev. Number	Date	Editor/ Writer	Substantive Change(s)
1	12/2005	NB/PC/RR	Document—made editorial changes. Page 7—changed programming sequence BL = 1 to BL = 4. Changed step 4. Page 13—Updated Figures 11 and 12. Page 14—Replaced paragraph under “Auto Precharge on A/D.”
0	08/2005	NB/PC/RR	Initial release.

THIS PAGE INTENTIONALLY LEFT BLANK

THIS PAGE INTENTIONALLY LEFT BLANK

How to Reach Us:

Home Page:
www.freescale.com

email:
support@freescale.com

USA/Europe or Locations Not Listed:
Freescale Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
(800) 521-6274
480-768-2130
support@freescale.com

Europe, Middle East, and Africa:
Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
support@freescale.com

Japan:
Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064, Japan
0120 191014
+81 2666 8080
support.japan@freescale.com

Asia/Pacific:
Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate,
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:
Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
(800) 441-2447
303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor
@hibbertgroup.com

Document Number: AN2922
Rev. 1
12/2005

freescale™ semiconductor