CoreSDR v3.0

Handbook

Actel®

Actel Corporation, Mountain View, CA 94043

Printed in the United States of America

Part Number: 50200104-0

Release: June 2007

Trademarks

Table of Contents

Introduction

General Description

CoreSDR provides a high-performance interface to single-data-rate (SDR) synchronous dynamic random access memory (SDRAM) devices. CoreSDR accepts read and write commands using the simple local bus interface and translates these requests to the command sequences required by SDRAM devices. CoreSDR also performs all initialization and refresh functions.

CoreSDR uses bank management techniques to monitor the status of each SDRAM bank. Banks are only opened or closed when necessary, minimizing access delays. Up to four banks can be managed at one time. Access cascading is also supported, allowing read or write requests to be chained together. This results in no delay between requests, enabling up to 100% memory throughput for sequential accesses.

CoreSDR is provided with configurable memory settings (Row Bits, Column Bits, Bank Bits) and timing parameters (CAS latency, t_{RAS}, t_{RC}, t_{RFC},t_{RCD}, t_{RP}, t_{MRD}, t_{RRD}, t_{REFC}, t_{WR}). The memory settings are configured through the CoreConsole GUI. The timing parameters are configured by setting the CoreSDR inputs. This ensures compatibility with virtually any SDRAM configuration.

CoreSDR consists of the following primary blocks, as shown in Figure 1:

1. Control and Timing Block – Main controller logic
2. Initialization Control – Performs initialization sequence after reset_n is deactivated or sd_init is pulsed.
3. Address Generation – Puts out address, bank address, and chip select signals on SDRAM interface.
4. Bank Management – Keeps track of last opened row and bank to minimize command overhead.
5. Refresh Control – Performs automatic refresh commands to maintain data integrity.

For CoreSDR, the datapath is external to the core. The tristate buffer shown in Figure 1 resides in the I/O and its output enable is controlled by the core.

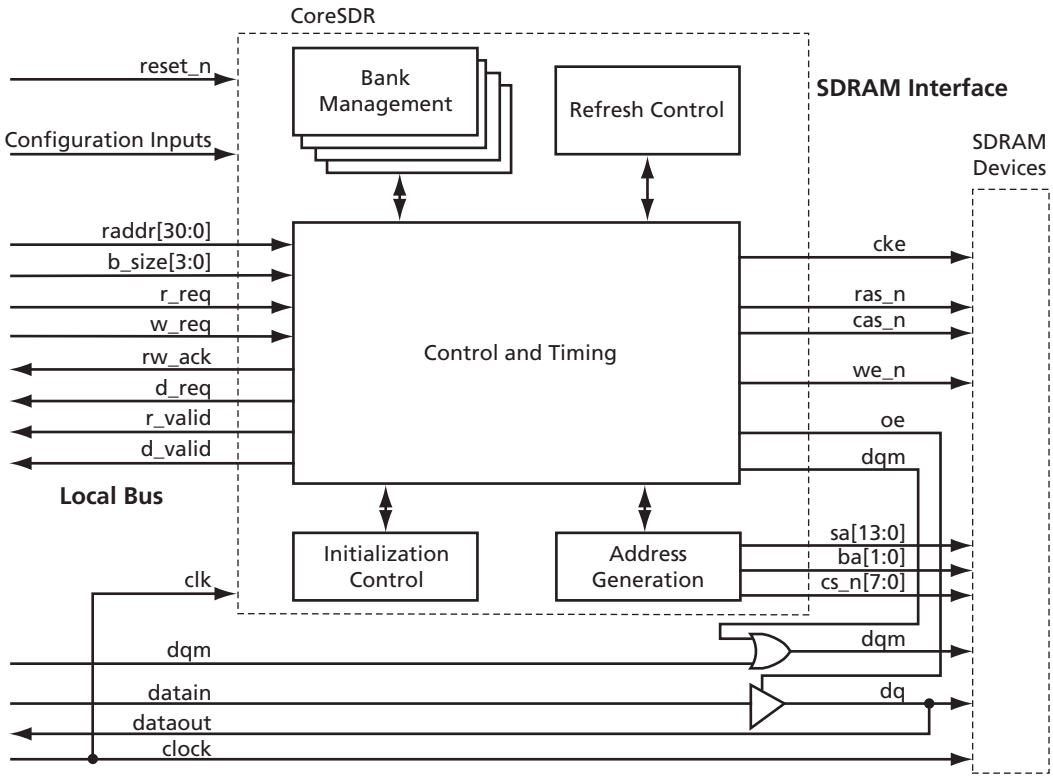

Figure 1 · CoreSDR Block Diagram

Core Versions

This handbook applies to CoreSDR v3.0. The release notes provided with the core list known discrepancies between this handbook and the core release associated with the release notes.

Device Utilization and Performance

CoreSDR has been implemented in several Actel device families. A summary of the implementation data is listed in Table 1.

Table 1 · CoreSDR Device Utilization and Performance

Family	Cells or Tiles			Utilization		Performance
	Sequential	Combinatorial	Total	Device	Total	
Fusion	492	858	1,350	AFS250-2	11%	133 MHz
IGLOO™/e	492	858	1,350	AGL250-2	11%	100 MHz
ProASIC®3/E	492	858	1,350	A3P250-2	11%	133 MHz
ProASIC PLUS®	752	1,046	1,798	APA075-STD	58%	91 MHz
Axcelerator®	505	642	1,147	AX125-2	57%	153 MHz
RTAX-S	505	642	1,147	RTAX250S-1	27%	95 MHz

Note: All data was obtained using a typical system configuration with the local bus tied to internal user logic and the controller configuration inputs hard-coded. CoreSDR requires 62 FPGA I/O pins when interfaced to a 32-bit SDRAM memory. RTAX-S performance data was obtained under military (MIL) conditions. All other performance data was obtained under commercial (COM) conditions.

1

Functional Block Description

Address Mapping

The mapping of the raddr bus at the local bus interface to the chip select, row, column, and bank addresses is shown in Figure 1-1. The exact bit positions of the mapping will vary depending on the rowbits and colbits configuration port settings. The chip select and row bits are mapped from the most significant bits of raddr, and the column bits are mapped from the least significant bits. The bank bits are mapped from the bits between the row and column bits. By mapping the bank bits from this location, long accesses to contiguous address space are more likely to take place without the need for a precharge.

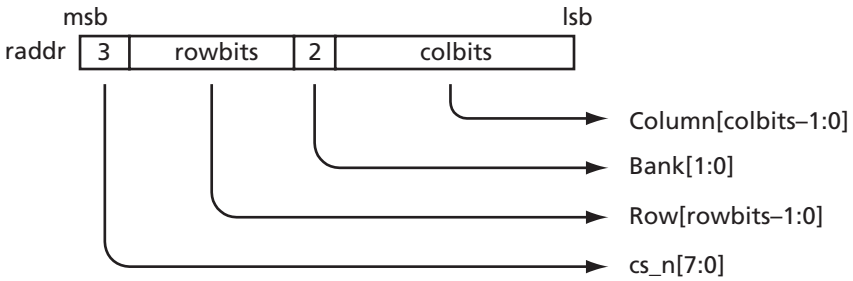

Figure 1-1 · Mapping raddr

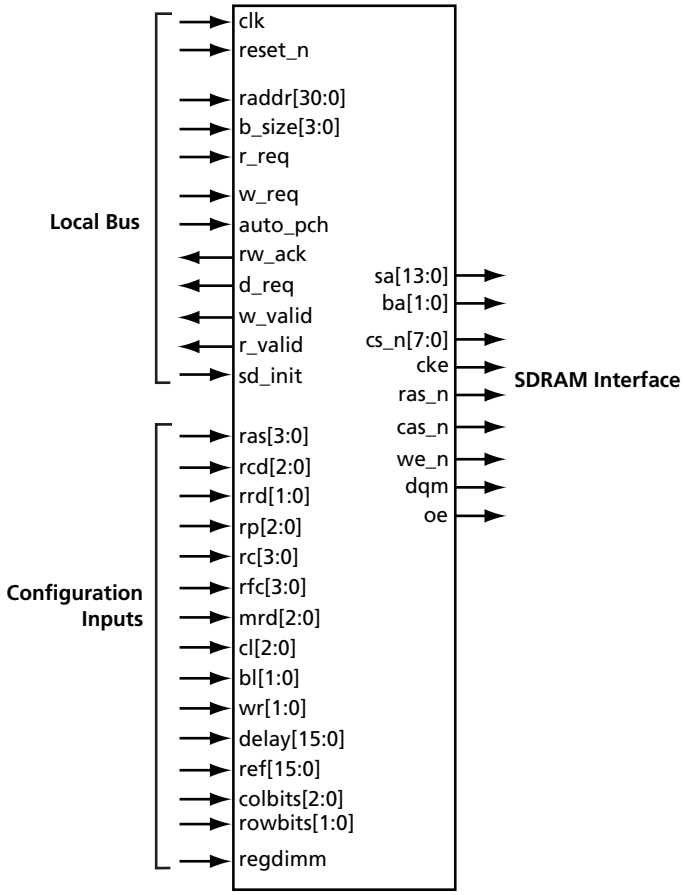

Figure 1-2 · I/O Signal Diagram

CoreSDR Operation

The synchronous interface and fully pipelined internal architecture of SDRAM allow extremely fast data rates if used efficiently. SDRAM is organized in banks of memory addressed by row and column. The number of row and column address bits depends on the size and configuration of the memory.

SDRAM is controlled by bus commands that are formed using combinations of the ras_n, cas_n, and we_n signals. For instance, on a clock cycle where all three signals are HIGH, the associated command is a *no operation* (NOP). A NOP is also indicated when the chip select is not asserted. The standard SDRAM bus commands are shown in Table 1-1.

Table 1-1 · SDRAM Bus Commands

Command	ras_n	cas_n	we_n
NOP	H	H	H
Active	L	H	H
Read	H	L	H
Write	H	L	L
Burst Terminate	H	H	L
Precharge	L	H	L
Auto-Refresh	L	L	H
Load Mode Register	L	L	L

SDRAM devices are typically divided into four banks. These banks must be opened before a range of addresses can be written to or read from. The row and bank to be opened are registered coincident with the *active* command. When a new row on a bank is accessed for a read or a write, it may be necessary to first close the bank and then reopen it to the new row. The *precharge* command closes a bank. Opening and closing banks costs memory bandwidth, so CoreSDR has been designed to monitor and manage the status of the four banks simultaneously. This enables the controller to intelligently open and close banks only when necessary.

When the *read* or *write* command is issued, the initial column address is presented to the SDRAM devices. In the case of SDR SDRAM, the initial data is presented concurrent with the *write* command. For the *read* command, the initial data appears on the data bus 1–4 clock cycles later. This is known as CAS latency and is due to the time required to physically read the internal DRAM and register the data on the bus. The CAS latency depends on the speed grade of the SDRAM and the frequency of the memory clock. In general, the faster the clock, the more cycles of CAS latency required. After the initial *read* or *write* command, sequential reads and writes will continue until the burst length is reached or a *burst terminate* command is issued. SDRAM devices support a burst length of up to eight data cycles. CoreSDR is capable of cascading bursts to maximize SDRAM bandwidth.

SDRAM devices require periodic refresh operations to maintain the integrity of the stored data. CoreSDR automatically issues the *auto-refresh* command periodically. No user intervention is required.

The *load mode register* command is used to configure the SDRAM operation. This register stores the CAS latency, burst length, burst type, and write burst mode. CoreSDR supports a *sequential burst type* and *programmed-length write burst mode*. The SDR controller only writes to the base mode register. Consult the SDRAM device specification for additional details on these registers.

To reduce pin count, SDRAM row and column addresses are multiplexed on the same pins. Table 1-2 lists the number of rows, columns, banks, and chip selects required for various standard discrete SDR SDRAM devices. CoreSDR will support any of these devices.

Table 1-2 · Standard SDR SDRAM Device Configurations

Chip Size	Configuration	Rows	Columns	Banks
64 Mb	16M×4	12	10	4
64 Mb	8M×8	12	9	4
64 Mb	4M×16	12	8	4
64 Mb	2M×32	11	8	4
128 Mb	32M×4	12	11	4
128 Mb	16M×8	12	10	4
128 Mb	8M×16	12	9	4
128 Mb	4M×32	12	8	4
256 Mb	64M×4	13	11	4
256 Mb	32M×8	13	10	4
256 Mb	16M×16	13	9	4
512 Mb	128M×4	13	12	4
512 Mb	64M×8	13	11	4
512 Mb	32M×16	13	10	4
1,024 Mb	256M×4	14	12	4
1,024 Mb	128M×8	14	11	4
1,024 Mb	64M×16	14	10	4

SDRAM is typically available in dual in-line memory modules (DIMMs), small outline DIMMs (SO-DIMMs), and discrete chips. The number of row and column bits for a DIMM or SO-DIMM configuration can be found by determining the configuration of the discrete chips used on the module. This information is available in the module datasheet.

Instruction Timing

Initialization

After reset_n is deasserted or sd_init is pulsed, CoreSDR performs the following sequence:

1. NOP command is issued for 200 µs (period controlled by the delay port parameter).
2. *Precharge-all* command
3. Eight *auto-refresh* commands
4. *Load mode register* command – This causes the SDRAM mode register to be loaded with the proper burst length (bl) and CAS latency (cl) values.

CoreSDR initialization timing is shown in Figure 1-3.

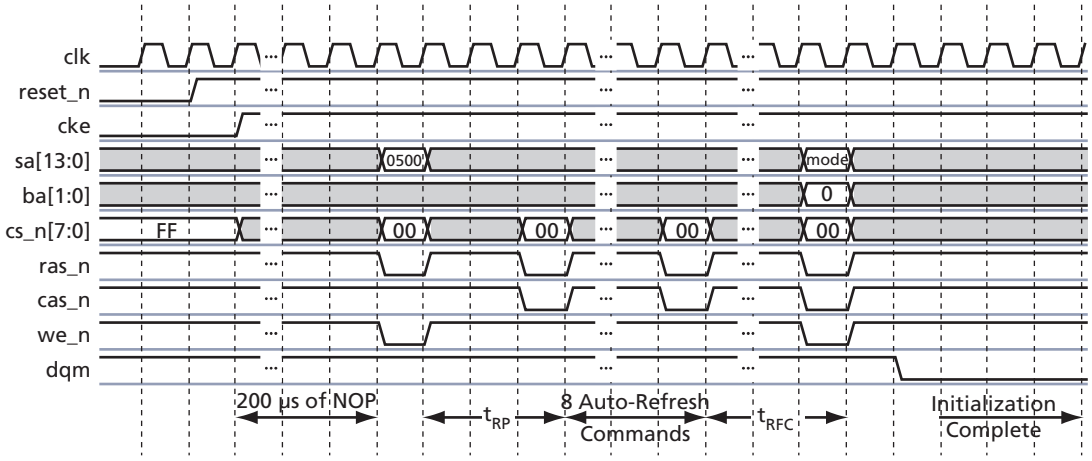

Figure 1-3 · SDR Initialization Sequence

Auto-Refresh

SDRAM devices require periodic *auto-refresh* commands to maintain data integrity. CoreSDR will automatically issue periodic *auto-refresh* commands to the SDRAM device(s) without user intervention. The refresh period configuration port (ref) specifies the period between refreshes, in clock cycles. Figure 1-4 shows an example of two refresh commands. The first refresh sequence occurs when one or more banks have been left open as a result of a *read without precharge* or *write without precharge* operation. All open banks are closed using the *precharge-all* command (ras_n, we_n asserted with sa[10] and sa[8]) prior to the refresh command. In Figure 1-4, a refresh occurs again after the refresh period has elapsed, as determined using the ref configuration port. The refresh will never interrupt a read or write in the middle of a data burst. However, if the controller determines that the refresh period has elapsed at a point concurrent with or prior to a read or write request, the request may be held off (rw_ack will not get asserted) until after the refresh has been performed.

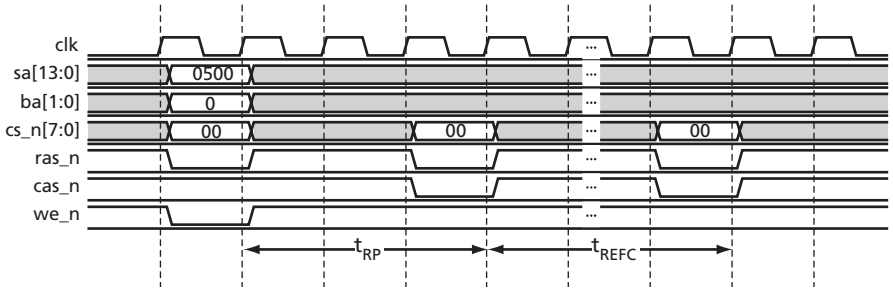

Figure 1-4 · Refresh Timing

Bank Management

CoreSDR incorporates bank management techniques to minimize command overhead. For each bank, the controller records the last opened row and whether the bank has been closed or not. When a local bus interface read or write request occurs, CoreSDR checks to determine if the requested bank is already opened and whether the request is for the same row as the one the bank is already opened with. If the bank is already opened with the requested row, CoreSDR performs the function immediately. If the bank is opened to a different row, the controller closes the bank (using the *precharge* command) and reopens the bank (using the *active* command) to the requested row. If the bank is already closed, the controller opens the bank to the requested row (using the *active* command).

Requests to the controller can be issued as *read with auto-precharge*, *write with auto-precharge*, *read without auto-precharge*, and *write without auto-precharge*. Commands are issued with auto-precharge if the auto_pch signal is set concurrent with the read request (r_req) or write request (w_req) signals. After a read with auto-precharge or write with auto-precharge, the accessed bank is automatically closed internally by the SDRAM device(s). After a read without auto-precharge or write without auto-precharge, the accessed bank is left open until closing is required. Closing will occur whenever a request is issued to a row different from the row a bank is already open to, or during the next refresh sequence. The refresh sequence will close all the banks (using the *precharge-all* command) if all banks are not already closed.

The default configuration of the controller tracks the status of four banks at a time. This means that an access to row *a* on bank *a* on chip select *a* is treated differently from an access to row *a* on bank *a* on chip select *b*. Therefore, a close and open sequence is performed when switching between these rows.

2 Tool Flows

Licenses

CoreSDR is licensed in three ways. Depending on your license, tool flow functionality may be limited.

Evaluation

Pre-compiled simulation libraries are provided, allowing the core to be instantiated in CoreConsole and simulated within Actel Libero® Integrated Design Environment (IDE), as described in the "CoreConsole" section. Using the Evaluation version of the core, it is possible to create and simulate the complete design in which the core is being included. The design may not be synthesized, as source code is not provided.

Obfuscated

Complete RTL code is provided for the core, enabling the core to be instantiated with CoreConsole. Simulation, synthesis, and layout can be performed with Actel Libero IDE. The RTL code for the core is obfuscated, and some of the testbench source files are not provided. They are pre-compiled into the compiled simulation library instead.

RTL

Complete RTL source code is provided for the core and testbenches.

CoreConsole

CoreSDR is preinstalled in the CoreConsole IP Deployment Platform. To use the core, click and drag it from the IP core list into the main window. The CoreConsole project can be exported to Actel Libero IDE at this point, providing access to the core only. Alternatively, IP blocks can be interconnected, allowing the complete system to be exported from CoreConsole to Libero IDE. The core can be configured using the configuration GUI within CoreConsole, as shown in Figure 2-1.

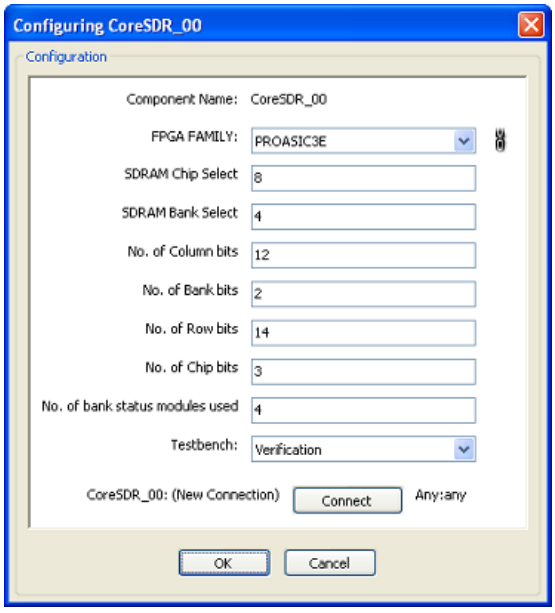

Figure 2-1 · CoreSDR Configuration Within CoreConsole

After configuring the core, Actel recommends you use the top-level Auto Stitch function to connect all the core interface signals to the top level of the CoreConsole project. Once the core is configured, invoke the Generate function in CoreConsole. This will export all the required files to the project directory in the *LiberoExport* directory. This is in the CoreConsole installation directory by default.

Importing into Libero IDE

After generating and exporting the core from CoreConsole, the core can be imported into Actel Libero IDE. Create a new project in Libero IDE and import the CoreConsole project from the *LiberoExport* directory. Libero IDE will then install the core and the selected testbenches, along with constraints and documentation, into its project.

Note: If two or more DirectCores are required, they can both be included in the same CoreConsole project and imported into Actel Libero IDE at the same time.

Simulation Flows

To run simulations, select the required testbench flow within CoreConsole and run **Save & Generate** from the Generate pane. Select the required testbench through the core configuration GUI in CoreConsole. The following simulation environments are supported:

- Full CoreSDR verification environment (Verilog only)
- Simple testbench (this VHDL testbench is a subset of the Verilog testbench)

When CoreConsole generates the Libero IDE project, it will install the appropriate testbench files. To run the testbenches, **simply set the design root to the CoreSDR instantiation in the Actel Libero IDE file manager** and click the **Simulation** icon in Libero IDE. This will invoke Model*Sim*® and automatically run the simulation.

Synthesis in Libero IDE

To run Synthesis on the core with parameters set in CoreConsole, set the design root to the top of the project imported from CoreConsole. This is a wrapper around the core that sets all the generics appropriately. Make sure the required timing constraints files are associated with the synthesis tool. Click the **Synthesis** icon in Libero IDE. The synthesis window appears, displaying the Synplicity® project. To run Synthesis, click the **Run** icon.

Place-and-Route in Libero IDE

Having set the design route appropriately and run Synthesis, click the **Layout** icon in Libero IDE to invoke Designer. CoreSDR requires no special place-and-route settings.

3 CoreSDR

Generics

Customers can define the generics listed in Table 3-1 as required in the source code.

Table 3-1 · CoreSDR Generics

Generic	Default Setting	Description
sdram_rasize*	31	Local address bus size
sdram_chips	8	Number of chip selects
sdram_banks	4	Number of SDRAM banks
sdram_colbits	12	Maximum number of SDRAM column bits
sdram_bankbits	2	Maximum number of SDRAM bank bits
sdram_rowbits	14	Maximum number of SDRAM row bits
sdram_chipbits	3	Number of encoded chip select bits
sdram_bankstatmodules	4	Number of bank status modules used (in multiples of four)
Family	16	Must be set to match the supported FPGA family. 11 – Axcelerator 12 – RTAX-S 14 – ProASIC^{PLUS} 15 – ProASIC3 16 – ProASIC3E 17 – Fusion 20 – IGLOO 21 – IGLOOe

*Note: *This parameter value is fixed at 31 through CoreConsole.*

Controller Configuration Ports

CoreSDR is configured using runtime-programmable configuration ports. These ports can be tied off by the user to fixed values or programmed at runtime. These values should not change after reset_n is deasserted. Table 3-2 lists the configuration ports.

Table 3-2 · SDRAM Controller Parameters

Parameter	Port Bits	Valid Values	Description
ras	4	1–10	SDRAM active to precharge (t_{RAS}), specified in clock cycles
rcd	3	2–5	SDRAM active to read or write delay (t_{RCD}), specified in clock cycles
rrd	2	2–3	SDRAM active bank *a* to active bank *b* (t_{RRD}), specified in clock cycles
rp	3	1–4	SDRAM *precharge* command period (t_{RP}), specified in clock cycles
rc	4	3–12	SDRAM *active* to *active/auto-refresh* command period (t_{RC}), specified in clock cycles
rfc	4	2–14	*Auto-refresh* to *active/auto-refresh* command period (t_{RFC}), specified in clock cycles
mrd	3	1–7	SDRAM *load mode register* command to *active* or *refresh* command (t_{MRD}), specified in clock cycles
cl	3	1–4	SDRAM CAS latency, specified in clock cycles
bl	2	0–3	SDRAM maximum burst length (encoded). Values are decoded as follows: 0: 1 transfer/burst 1: 2 transfers/burst 2: 4 transfers/burst 3: 8 transfers/burst (valid for SDR only)
wr	2	1–3	SDRAM write recovery time (t_{WR})
delay	16	10–65,535	Delay after a reset event that the controller waits before initializing the SDRAM, specified in clock cycles. Per JEDEC standards, SDR devices require this delay to be a minimum of 200 µs.
ref	16	10–65,535	Period between *auto-refresh* commands issued by the controller, specified in clock cycles REF = auto refresh interval / t_{CK}
colbits	7	3–7	Number of bits in the column address (encoded). Values are decoded as follows: 3: 8 column bits 4: 9 column bits 5: 10 column bits 6: 11 column bits 7: 12 column bits
rowbits	3	0–3	Number of bits in the row address (encoded). Values are decoded as follows: 0: 11 row bits 1: 12 row bits 2: 13 row bits 3: 14 row bits
regdimm	1	0–1	Set when using registered/buffered DIMMs. Causes adjustment in local bus interface timing to synchronize with SDRAM command timing delayed by register/buffer on DIMM.

Example settings for the timing-related parameters are shown in Table 3-3. These settings are based on the speed grade of the SDRAM devices and the desired operating frequency. Consult the datasheet for the SDRAM device you are using for the specific timing values of that device.

Table 3-3 · Example Controller Parameter Values for CoreSDR

Parameter	100 MHz (10 ns period)[1]		133 MHz (7.5 ns period)[2]	
	Specification	Value	Specification	Value
ras	44.0 ns	5	37.0 ns	6
rcd	20.0 ns	2	15.0 ns	3
rrd	15.0 ns	2	14.0 ns	2
rp	20.0 ns	2	15.0 ns	3
rc	66.0 ns	7	60.0 ns	8
rfc	66.0 ns	7	66.0 ns	9
mrd	2 clks	2	2 clks	2
cl	–	2	–	2
wr	15.0 ns	2	14.0 ns	2
delay	200 µs	20,000	200 µs	26,667
ref	7.8125 µs	781	7.8125 µs	1,041

Notes:

1. *Values based on Micron MT48LC32M8A2-75*
2. *Values based on Micron MT48LC32M8A2-7E*

4

Core Interfaces

The port signals for CoreSDR are defined in Table 4-1 below, Table 4-2 on page 20, and Table 1 on page 6. The port signals are also illustrated in Figure 1 on page 5. All signals are designated either Input (input-only) or Output (output-only).

Local Bus Signals

The user interface to CoreSDR is referred to as the local bus interface. The local bus signals are shown in Table 4-1.

Table 4-1 · Local Bus Signals

Signal	Name	I/O	Description
clk	Clock	Input	System clock. All local bus signals are synchronous to this clock.
reset_n	Reset	Input	System reset
raddr[30:0]	Memory Address	Input	Local bus address
b_size[3:0]	Burst Size	Input	Local bus burst length. Valid values are 1 through BL, where BL is the programmed burst length. (Refer to Table 3-2 on page 16 for discussion of the BL parameter.)
r_req	Read Request	Input	Local bus read request
w_req	Write Request	Input	Local bus write request
auto_pch	Auto-Precharge Request	Input	When asserted in conjunction with r_req or w_req, causes command to be issued as *read with auto-precharge* or *write with auto-precharge*, respectively.
rw_ack	Read/Write Acknowledge	Output	Acknowledgement of read or write request
d_req	Data Request	Output	Requests data on the local bus write data bus (datain) during a write transaction. Asserts one clock cycle prior to when data is required.
w_valid	Write Data Valid	Output	Frames the active data being written to SDRAM. Mimics d_req, except that it is delayed by one clock cycle. This signal is typically not used and is retained for legacy compatibility.
r_valid	Read Data Valid	Output	Indicates that the data on the local bus read data bus (dataout) is valid during a read cycle.
sd_init	Initialization Strobe	Input	Causes CoreSDR to reissue the initialization sequence to SDRAM devices. CoreSDR will always issue the initialization sequence (including the startup delay) after reset, regardless of the SD_INIT state. This signal can be tied LOW if runtime reinitialization is not required.

Notes:

1. *All control signals are active high except reset_n.*
2. *All local bus signals are synchronous to clk.*

SDR SDRAM Interface Signals

The external interface to SDRAM devices is referred to as the SDRAM interface. The SDRAM interface signals are shown in Table 4-2.

Table 4-2 · SDR SDRAM Interface Signals

Signal	Name	I/O	Description
sa[13:0]	Address Bus	Output	Sampled during the *active*, *precharge*, *read*, and *write* commands. This bus also provides the mode register value during the *load mode register* command.
ba[1:0]	Bank Address	Output	Sampled during *active*, *precharge*, *read*, and *write* commands to determine which bank command is to be applied to.
cs_n[7:0]	Chip Selects	Output	SDRAM chip selects
cke	Clock Enable	Output	SDRAM clock enable. Held LOW during reset to ensure SDRAM dq and dqs outputs are in the high-impedance state.
ras_n	Row Address Strobe	Output	Command input
cas_n	Column Address Strobe	Output	Command input
we_n	Write Enable	Output	Command input
dqm	Data Mask	Output	SDRAM data mask asserted by controller during SDRAM initialization and during burst terminate. User may sum with user data mask bits.
oe	Output Enable	Output	Tristate control for DQ data

5 Timing Diagrams

SDRAM Writes

The user requests writes at the local bus interface by asserting the w_req signal and driving the starting address and burst size on raddr and b_size, respectively. The auto_pch may also be asserted with w_req to cause the write to be issued as a *write with auto-precharge*.

The rules for write requests at the local bus interface are as follows:

1. Once w_req is asserted, it must remain asserted until rw_ack is asserted by CoreSDR. After rw_ack is asserted by CoreSDR, w_req may remain asserted to request a follow-on write transaction. The w_req signal may remain asserted over any number of rw_ack pulses to generate any number of cascaded write bursts. The only time w_req may be deasserted is during the clock cycle immediately following the rw_ack pulse from CoreSDR.
2. The signals raddr, b_size, and auto_pch must maintain static values from the point when w_req becomes asserted until CoreSDR asserts rw_ack. The raddr, b_size, and auto_pch signals may only change values in the clock cycle immediately following the rw_ack pulse from CoreSDR or when w_req is deasserted.
3. The w_req signal may not be asserted while the read request signal (r_req) is asserted.
4. The data request signal (d_req) will assert one clock prior to when the user must present data at the datain bus.
5. The timing relationship between an initial w_req assertion and an rw_ack assertion, or between rw_ack pulses as a result of multiple cascaded writes, will vary depending on the status of the banks being accessed, configuration port settings, refresh status, and initialization status. The user logic should not rely on any fixed timing relationship between w_req and rw_ack.

Example CoreSDR Write Sequence

Figure 5-1 on page 22 shows an example CoreSDR write sequence. In this sequence, two writes are requested, both to the same bank and row. The first write request is for a burst size of eight; the second is for a burst size of five. The write request signal (w_req) is first asserted with the starting address (raddr) and the burst size (b_size). As a result of this request, CoreSDR asserts the row address (sa), bank address (ba), and chip select (cs_n) with the *active* command to open the bank to the requested row. Next, CoreSDR requests data from the user at the local bus interface using the d_req signal and acknowledges the write request by asserting rw_ack. At the next clock, CoreSDR begins writing the data to the SDRAM devices. Eight data cycles are transferred.

The local bus interface write request (w_req) remains asserted after rw_ack is asserted by CoreSDR to request an additional write burst. After the rw_ack pulse, the local bus interface changes the address and changes the burst size to five. As soon as the eight write cycles from the previous request are processed, the next *write* command is issued by CoreSDR, and the five data cycles for the new burst begin. In the case of this example, the second write was to the same bank and row as the first write. If the second write had been to a different bank or row, CoreSDR would have issued *precharge* and/or *activate* commands prior to the *write* command.

Since the burst size (b_size) of the second write request is less than the programmed SDR SDRAM burst length, the burst must be terminated using the *burst terminate* command. CoreSDR does this automatically, asserting we_n five clock cycles after the *write* command was issued. The dqm signal is also asserted during the *burst terminate* command to comply with JEDEC specifications.

As demonstrated in Figure 5-1 on page 22, the output enable signal (oe) goes active one clock cycle before the data is actually required at the dq outputs. This is to accommodate long t_{z-x} delays that may exist in PLD or ASIC tristate drivers. The oe signal stays asserted until the last data element is written.

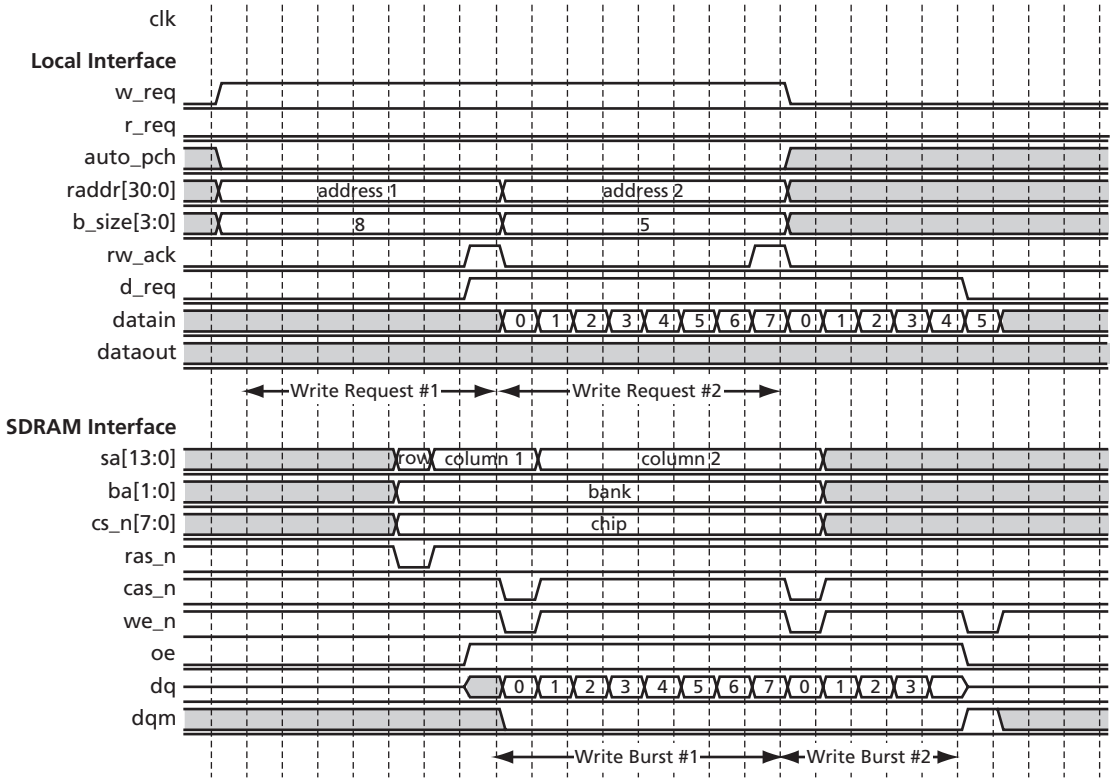

Figure 5-1 · CoreSDR Burst Write

Note: For the case shown, rcd = 3, bl = 2, and regdimm = 0.

SDRAM Reads

The user requests reads at the local interface by asserting the r_req signal and driving the starting address and burst size on raddr and b_size, respectively. The auto_pch may also be asserted with r_req to cause the read to be issued as a *read with auto-precharge*.

The rules for read requests at the local interface are as follows:

1. Once r_req is asserted, it must remain asserted until rw_ack is asserted by CoreSDR. After rw_ack is asserted by CoreSDR, r_req may remain asserted to request a follow-on write transaction. r_req may remain asserted over any number of rw_ack pulses to generate any number of cascaded read bursts. The only time r_req may be deasserted is during the clock cycle immediately following the rw_ack pulse from CoreSDR.
2. The signals raddr, b_size, and auto_pch must maintain static values from the point when r_req is asserted until CoreSDR asserts rw_ack. raddr, b_size, and auto_pch may only change values in the clock cycle immediately following the rw_ack pulse from CoreSDR, or when r_req is deasserted.
3. The r_req signal may not be asserted while the write request signal (w_req) is asserted.
4. The read data valid signal (r_valid) will assert when valid data is available at the dataout bus.
5. The timing relationship between an initial r_req assertion and an rw_ack assertion, or between rw_ack pulses as a result of multiple cascaded reads, will vary depending on the status of the banks being accessed, configuration port settings, refresh status, and initialization status. The user logic should not rely on any fixed timing relationship between r_req and rw_ack.

Example CoreSDR Read Sequence

Figure 5-2 shows an example CoreSDR read sequence. In this sequence, two reads are requested, both to the same bank and row. The first read request is for a burst size of eight, the second is for a burst size of five. The read request signal (r_req) is first asserted with the starting address (raddr) and the burst size (b_size). As a result of this first request, CoreSDR asserts the row address (sa), bank address (ba), and chip select (cs_n) with the *active* command to open the bank to the requested row. Next, CoreSDR acknowledges the read request by asserting rw_ack. Since the CAS latency is set at two in this case, read data appears at dataout two clock cycles after CoreSDR issues the *read* command. CoreSDR asserts the r_valid signal to indicate valid data at the dataout bus. Eight data cycles are transferred.

The local bus interface read request (r_req) remains asserted after rw_ack is asserted by CoreSDR to request an additional read burst. After the rw_ack pulse, the local bus interface changes the address and changes the burst size to five. CoreSDR issues the next read command to the SDRAM devices as soon as is possible to avoid interruption in the data flow. In the case of this example, the second read was to the same bank and row as the first read. If the second read had been to a different bank or row, CoreSDR would have issued *precharge* and/or *activate* commands prior to the *read* command. Since the burst size (b_size) of the second read request is less than the programmed SDR SDRAM burst length, the burst must be terminated using the *burst terminate* command. CoreSDR does this automatically, asserting we_n five clock cycles after the *read* command was issued.

CoreSDR never asserts the dqm signal while read data is transacting. User logic must never assert dqm while read data is transacting, as this will cause the SDRAM devices to drive high-impedance instead of valid data.

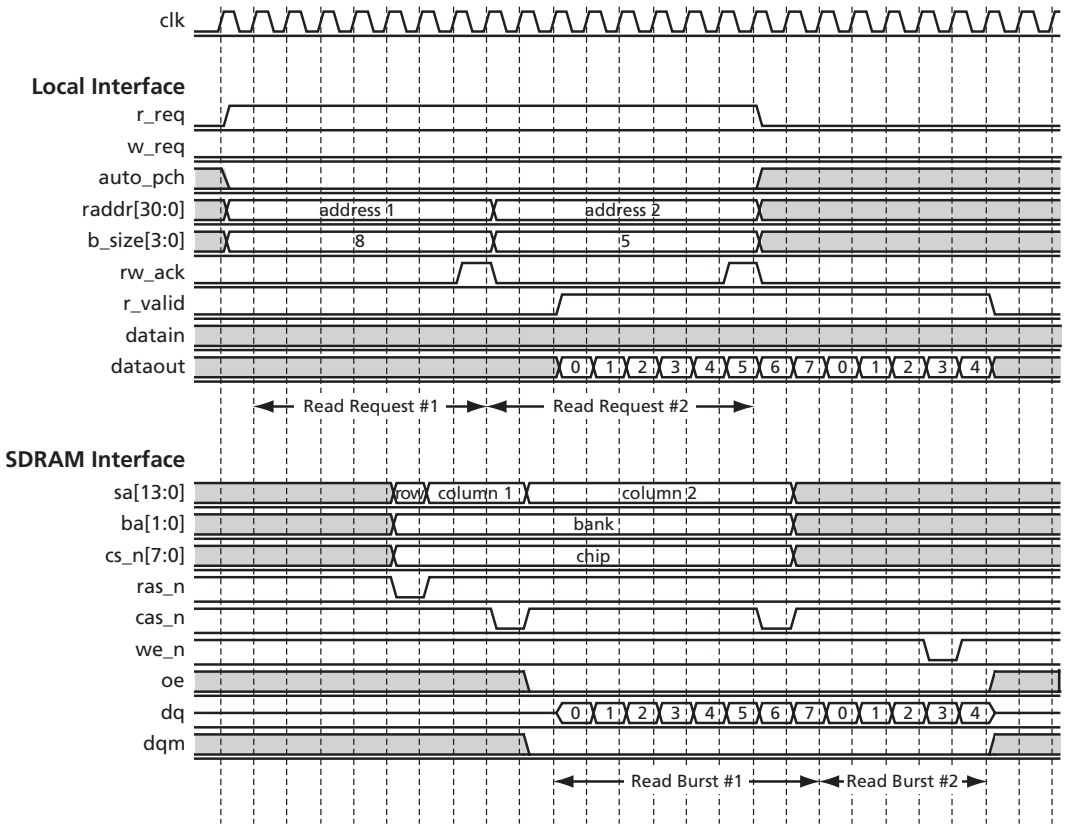

Figure 5-2 · CoreSDR Burst Read

Note: For the case shown, bl = 3, rcd = 3,wr = 2, rp = 3, cl = 2, and regdimm = 0.

Auto-Precharge

Read commands can be issued to the SDRAM devices as *read with auto-precharge* or *read without auto-precharge*. Likewise, *write* commands can be issued to the SDRAM devices as *write with auto-precharge* or *write without auto-precharge*. If the auto-precharge option is used, the SDRAM device will automatically close (precharge) banks being read from or written to at the end of the transaction. Any subsequent reads or writes to these banks will not require an explicit *precharge* command from CoreSDR. The user selects whether *read* or *write* commands are issued with auto-precharge using the auto_pch signal. If auto_pch is asserted along with w_req or r_req, the command will be issued to the SDRAM with auto-precharge. The auto-precharge option is useful in situations where the requested read or write addresses tend to be random. With random address sequences, banks are seldom left open with the exact row required by a subsequent request. If the auto-precharge was not used for the previous access to a bank, subsequent transactions to that bank first require the bank to be closed (precharged), causing a delay in the transaction. If auto-precharge was used for the previous access, the bank is already closed and ready to be opened to the desired row. Figure 5-3 shows example transactions using the auto-precharge feature for CoreSDR. In the example, two write requests are issued, the first using auto-precharge and the second not using auto-precharge. Both requests are to the same bank, but may be to different rows. CoreSDR issues the first command as a *write with auto-precharge* by driving bit sa[10] HIGH during the *write* command. CoreSDR issues the second command as a *write without auto-precharge* by driving bit sa[10] LOW during the *write* command. Since the first and second requests are to the same bank, CoreSDR must wait before reopening the bank to meet the SDRAM t_{WR} and t_{RP} requirements. If the first and second requests were to different banks, the second request would follow the first request such that there would be no interruption in data flow.

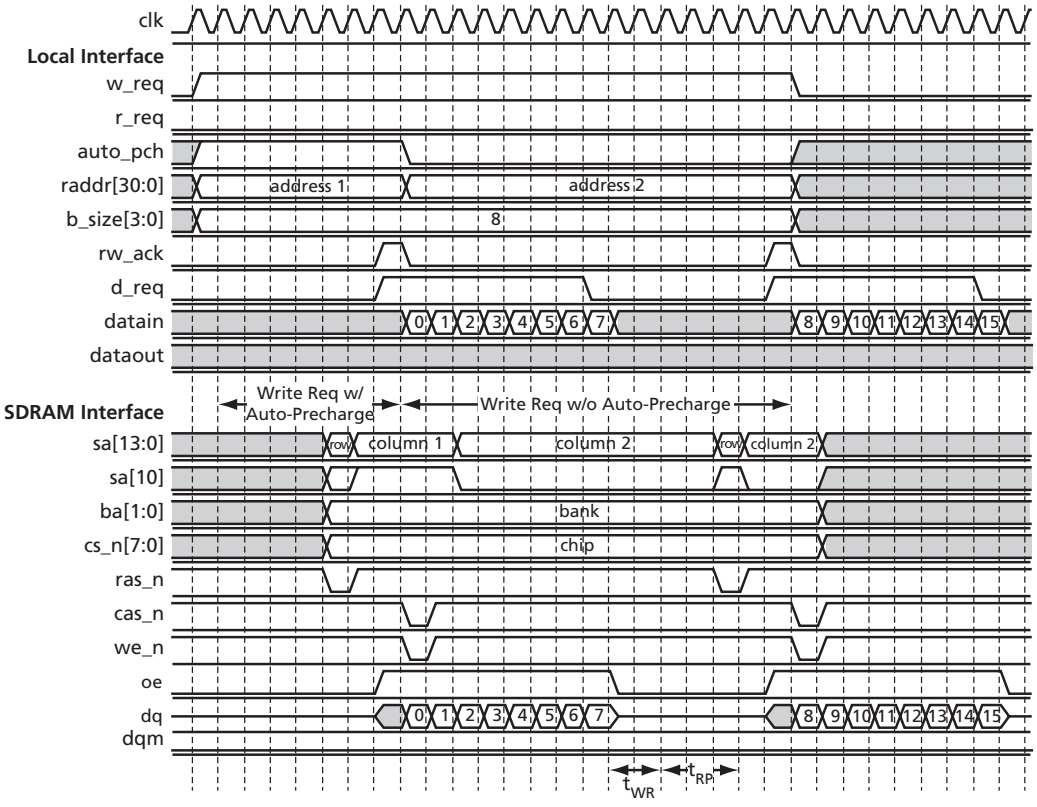

Figure 5-3 · CoreSDR Burst Writes with Auto-Precharge Option

Note: For the case shown, bl = 3, rcd = 3, wr = 2, rp = 3, cl = 2, and regdimm = 0.

6 Testbench Operation

Two testbenches are provided with CoreSDR:

- Verilog testbench: Testbench that verifies core operation. This testbench exercises all the features of the core.
- VHDL testbench: Testbench that verifies core operation. This testbench is a subset of Verilog testbench.

Testbench Description

Included with the Obfuscated and RTL releases of CoreSDR is a testbench that verifies operation of the CoreSDR macro. A simplified block diagram of the testbench is shown in Figure 6-1. By default, the Verilog version, *coresdr_tb.v*, instantiates a Micron 256 Mbit SDRAM model (*MT48LC16M16A2.v*, 4M×16 × 4 banks). The VHDL version, *coresdr_tb.vhd*, instantiates a Micron 64 Mbit SDRAM model (*MT48LC4M16A2.vhd*, 1M×16 × 4 banks). The testbench instantiates the DUT (design under test), which is the CoreSDR macro, the SDRAM model, as well as the test vector modules that provide stimuli sources for the DUT. A procedural testbench controls each module and applies the sequential stimuli to the DUT.

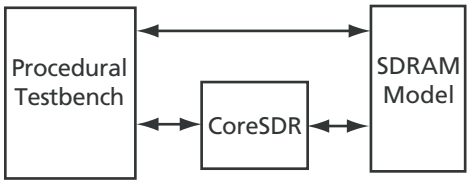

Figure 6-1 · CoreSDR Testbench

Verilog Testbench

Several tests comprise the test suite of the provided Verilog testbench. The various types of tests used in the Verilog testbench are listed in Table 6-1.

Table 6-1 · CoreSDR Verilog Tests

Test	Description
Address	Writes bursts to power-of-two address sequence (8, 16, 32, 64, etc.) to all of address space, then reads values back.
Partial Burst	Test done with and without auto-precharge: 1. Write FF, 1, 2, 3, FF, FF, 6, 7, FF, FF, FF, 11. 2. Write 0 to addr 0 with SIZE of 1. 3. Write 4, 5 to addr 4 with SIZE of 2. 4. Write 8, 9, 10 to addr 8, with SIZE of 3. 5. Read back 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11.
Partial Burst Read	Test of partial burst reads with and without auto-precharge: 1. Write counting pattern. 2. Read using all bsize lengths up to BLEN.
Back-to-Back Partial Burst Read	Partial burst reads using back-to-back reads with and without auto-precharge: 1. Write counting pattern. 2. Read using all bsize lengths up to BLEN (does not check read-back data).

Table 6-1 · CoreSDR Verilog Tests (continued)

Test	Description
Back-to-Back Partial Burst Write	Partial burst writes using back-to-back writes with and without auto-precharge: 1. Write burst of BSIZE 1, 2, 3, 4, 5, 6, 7, then 8 (or less if limited by BLEN), all back-to-back 2. Read back to verify.
Data Mask Walking Byte Test	1. Each write_data location is loaded with all ones, except for a count value, which is loaded in the byte that will be masked in the following write. This burst is written to memory. 2. Each write_data location is loaded with a counting pattern, except for the byte position that will be masked, which is loaded with FF. This burst is written to memory. 3. All locations are read back. Counting pattern is expected in each location.
Long Write and Long Read Test	Long write followed by read tests: 1. Long write to memory (W_REQ held HIGH continuously) 2. Long read from memory (R_REQ held HIGH continuously)
Alternating Write and Read Bursts	Fast switching between single write and read bursts to the same bank/row
Alternating Read and Write Bursts	Fast switching between single read and write bursts to the same bank/row
Back-to-Back Writes	Writes with W_REQ held HIGH, using various combinations of auto_pch and bank accesses: Auto-precharge sequence (first command -> second command): • No auto_pch -> no auto_pch • No auto_pch -> auto_pch • auto_pch -> no auto_pch • auto_pch -> auto_pch Bank access sequence (first command -> second command): • Precharge not required for second command (auto-precharged, same bank/row) • Precharge required for second command (same bank, new row)

Table 6-1 · CoreSDR Verilog Tests (continued)

Test	Description
Back-to-Back Write Followed by Read	Write immediately followed by read with various combinations of auto_pch and bank accesses. For each combination of auto_pch and bank access, the sequence is as follows: 1. Simple *write* command (using write task); this will be read/verified by the test read. 2. *Write* command 3. *Read* command (read data written in step 1) 4. Simple *read* command (using read task) to verify data written in step 2 Auto-precharge sequence (first command -> second command): • No auto_pch -> no auto_pch • No auto_pch -> auto_pch • auto_pch -> no auto_pch • auto_pch -> auto_pch Bank access sequence (first command -> second command): • Precharge not required for second command (auto-precharged, same bank/row) • Precharge required for second command (same bank, new row)
Back-to-Back Read Followed by Write	Back-to-back read followed by write. Validate read immediately followed by write with various combinations of auto_pch and bank accesses. For each combination of auto_pch and bank access, sequence is as follows: 1. Simple *write* command (using write task); this will be read/verified by the test read. 2. *Read* command (verify data from write in #500) 3. *Write* command 4. Simple *read* command (using read task) to verify data written in step 3 Auto-precharge sequence (first command -> second command): • No auto_pch -> no auto_pch • No auto_pch -> auto_pch • auto_pch -> no auto_pch • auto_pch -> auto_pch Bank access sequence (first command -> second command): • Precharge not required for second command (auto-precharged, same bank/row) • Precharge required for second command (same bank, new row)

Table 6-1 · CoreSDR Verilog Tests (continued)

Test	Description
Back-to-Back Read	Validate read immediately followed by another read with various combinations of auto_pch and bank accesses. For each combination of auto_pch and bank access, the sequence is as follows: 1. Simple *write* command (using write task); this will be read/verified by first test read. 2. Simple *write* command (using write task); this will be read/verified by second test read. 3. Test *read* command (verify data from write in #500) 4. Test *read* command (verify data from write in #2) Auto-precharge sequence (first command -> second command): • No auto_pch -> no auto_pch • No auto_pch -> auto_pch • auto_pch -> no auto_pch • auto_pch -> auto_pch Bank access sequence (first command -> second command): • Precharge not required for second command (auto-precharged, same bank/row) • Precharge required for second command (same bank, new row)
Row Sequence	Addresses SDRAM in row increments. Precharges typically required in this case to close bank and open new row. 1. Burst writes, row by row 2. Burst reads, row by row, data verified 3. Burst reads, R_REQ asserted continuously, data verified
Bank Sequence	Addresses all SDRAM banks in sequence. • Burst writes to each bank • Burst reads from each bank
Bus Turnaround	Tests situations where chip selects are changed during a burst read and bus turnaround is required to prevent contention between SDRAM device outputs. This tests case where chips are changed and one clock cycle of turnaround appears on bus. • Write 16 values, – 8 on last row, last bank of current chip – 8 on first row, first bank of next chip • Results need to be verified visually on waveform.

VHDL Testbench

The VHDL testbench implements a subset of the functionality tested in the Verilog testbench, described in Table 6-1 on page 25. By default, the VHDL testbench instantiates a Micron 64 Mbit SDRAM model (*MT48LC4M16A2.vhd*, 1M×16 × 4 banks).

The VHDL testbench is provided in pre-compiled Model*Sim* format and in VHDL source code for all releases (Evaluation, Obfuscated, and RTL) for you to examine and modify to suit your needs. The source code for the VHDL testbench is provided to ease the process of integrating the CoreSDR macro into your design and verifying its functionality.

A

Product Support

Actel backs its products with various support services including Customer Service, a Customer Technical Support Center, a web site, an FTP site, electronic mail, and worldwide sales offices. This appendix contains information about contacting Actel and using these support services.

Customer Service

Contact Customer Service for non-technical product support, such as product pricing, product upgrades, update information, order status, and authorization.

From Northeast and North Central U.S.A., call **650.318.4480**
From Southeast and Southwest U.S.A., call **650. 318.4480**
From South Central U.S.A., call **650.318.4434**
From Northwest U.S.A., call **650.318.4434**
From Canada, call **650.318.4480**
From Europe, call **650.318.4252** or **+44 (0) 1276 401 500**
From Japan, call **650.318.4743**
From the rest of the world, call **650.318.4743**
Fax, from anywhere in the world **650.318.8044**

Actel Customer Technical Support Center

Actel staffs its Customer Technical Support Center with highly skilled engineers who can help answer your hardware, software, and design questions. The Customer Technical Support Center spends a great deal of time creating application notes and answers to FAQs. So, before you contact us, please visit our online resources. It is very likely we have already answered your questions.

Actel Technical Support

Visit the Actel Customer Support website (www.actel.com/custsup/search.html) for more information and support. Many answers available on the searchable web resource include diagrams, illustrations, and links to other resources on the Actel web site.

Website

You can browse a variety of technical and non-technical information on Actel's home page, at www.actel.com.

Contacting the Customer Technical Support Center

Highly skilled engineers staff the Technical Support Center from 7:00 A.M. to 6:00 P.M., Pacific Time, Monday through Friday. Several ways of contacting the Center follow:

Email

You can communicate your technical questions to our email address and receive answers back by email, fax, or phone. Also, if you have design problems, you can email your design files to receive assistance. We constantly monitor the email account throughout the day. When sending your request to us, please be sure to include your full name, company name, and your contact information for efficient processing of your request.

The technical support email address is tech@actel.com.

Phone

Our Technical Support Center answers all calls. The center retrieves information, such as your name, company name, phone number and your question, and then issues a case number. The Center then forwards the information to a queue where the first available application engineer receives the data and returns your call. The phone hours are from 7:00 A.M. to 6:00 P.M., Pacific Time, Monday through Friday. The Technical Support numbers are:

650.318.4460
800.262.1060

Customers needing assistance outside the US time zones can either contact technical support via email (tech@actel.com) or contact a local sales office. Sales office listings can be found at www.actel.com/contact/offices/index.html.

Index

For more information about Actel's products, visit our website at http://www.actel.com

Actel Corporation *• 2061 Stierlin Court • Mountain View, CA 94043 USA*
Customer Service: 650.318.1010 • Customer Applications Center: 800.262.1060

Actel Europe Ltd*. • River Court, Meadows Business Park • Station Approach, Blackwater • Camberley, Surrey GU17 9AB• United Kingdom*
Phone +44 (0) 1276 609 300 • Fax +44 (0) 1276 607 540

Actel Japan *• EXOS Ebisu Bldg. 4F • 1-24-14 Ebisu Shibuya-ku • Tokyo 150 • Japan*
Phone +81.03.3445.7671 • Fax +81.03.3445.7668 • www.jp.actel.com

Actel Hong Kong *• Suite 2114, Two Pacific Place • 88 Queensway, Admiralty Hong Kong*
Phone +852 2185 6460 • Fax +852 2185 6488 • www.actel.com.cn

50200104-0 /6.07