

Fundamentals of SDRAM Memory



Table of Contents

Introduction	4-5	Memory System Design	16-18
DRAM Trends	4	Design Simulation	16
DRAM	5-7	Design Verification	16
SDRAM	7-10	Verification Strategy	17
DDR SDRAM	8	SDRAM Verification	17
DDR2 SDRAM	9	Glossary	19-23
DDR3 SDRAM	10		
DIMMs	10-15		
DIMM Physical Size	10		
DIMM Data Width	10		
DIMM Rank	10		
DIMM Memory Size & Speed	11		
DIMM Architecture	11		
Serial Presence Detect	14		

Introduction

DRAM (Dynamic Random Access Memory) is attractive to designers because it provides a broad range of performance and is used in a wide variety of memory system designs for computers and embedded systems. This DRAM memory primer provides an overview of DRAM concepts, presents potential future DRAM developments and offers an overview for memory design improvement through verification.

DRAM Trends

There is a continual demand for computer memories to be larger, faster, lower powered and physically smaller. These needs are the driving force in the advancement of DRAM technology. Mainstream DRAMs have evolved over the years through several technology enhancements, such as SDRAM (Synchronous DRAM), DDR (Double Data Rate) SDRAM, DDR2 (Double Data Rate 2) SDRAM, and DDR3 (Double Data Rate 3) SDRAM. This evolution has also been driven by how computer memories are used on DIMMs (Dual Inline Memory Modules). DIMM implementations have expanded from unregistered DIMMs to include registered DIMMs and FB-DIMMs (Fully Buffered DIMMs).

Computer memories are not the only systems that continue to demand larger, faster, lower powered and physically smaller memories. Embedded systems applications have similar requirements and can also use DRAMs.

However, memory systems are implemented differently in computers versus embedded systems. Typically, computer memories are mounted on pluggable DIMMs that are easily installed in the computer during assembly. The computer user may upgrade the computer memory by adding or replacing the DIMMs after the computer has been purchased. As a result, memories used in computers require a high level of compatibility with current and future computers, as well as current and future memories used in

conjunction with a DIMM. There are two major areas of compatibility. First, memory needs to be compatible with a wide variety of memory controller hubs used by the computer manufacturers. Secondly, memory needs to work when a mixture of different manufacturer's memories is used in the same memory system of the computer. Open memory standards are useful in helping to ensure memory compatibility.

On the other hand, embedded systems typically use a fixed memory configuration, meaning the user does not modify the memory system after purchasing the product. The embedded systems manufacturer then has total control over which memories from specific manufacturers are used in the embedded systems product. It is common to optimize an embedded system's performance and cost by using one specific memory from one memory manufacturer. As a result, it is less important in embedded systems, as compared to computer systems, to have a high level of multivendor memory interoperability.

The JEDEC (Joint Electron Device Engineering Council) has helped the memory industry by creating memory specifications in the form of JEDEC standards. JEDEC is a non-profit organization with members from memory manufacturers, computer manufacturers, test equipment manufacturers, etc. The open JEDEC standards define the required specifications that are needed for manufacturers to implement memory products that are to be interoperable with other manufacturers' memories and computer memory controller hubs. These standards cover physical characteristics, DIMM circuit board layouts, electrical signals, register definitions, functional operation, memory protocols, etc. Verifying and testing a memory conformance to the JEDEC specifications is a critical step to ensuring reliable and interoperable memory operation with other manufacturer's products.

New DRAM designs are meeting computer and embedded systems memory requirements to be larger, faster, lower powered and physically smaller. As a result, the following DRAMs changes are occurring: memory size is increasing, the numbers of banks are increasing, the burst length is increasing, the supply voltage is decreasing, the logic voltage swings are decreasing, the clock rates are increasing, the data rates are increasing, memory channels implementations are going from a large number of parallel signals to a reduced number of high speed serial signals, the number of memory channels are increasing, the circuit board density is increasing, etc. These trends are causing designers to use new techniques and tools to design, verify and debug their memory systems.

As memory clock rates increase and logic voltage swings decrease, signal integrity has become more of an issue for reliable memory operation. As result, there are trends for new DRAM features to focus on improving signal integrity of the memory system. These features include dynamically controlled ODT (on-die termination), OCD (off-chip driver) calibration and Fully Buffered DIMMs with AMBs (Advanced Memory Buffers).

DRAM

An advantage of DRAM over other types of memory is its ability to be implemented with fewer circuits per memory cell on the IC (integrated circuit). The DRAM's memory cell is based on storing charge on a capacitor. A typical DRAM cell is built with one capacitor and one or three FET(s) (field-effect transistor). A typical SRAM (Static Random Access Memory) memory cell takes six FET devices, resulting in fewer memory cells per same size IC. SRAMs are simpler to use, easier to interface to and have faster data access times than DRAMs.

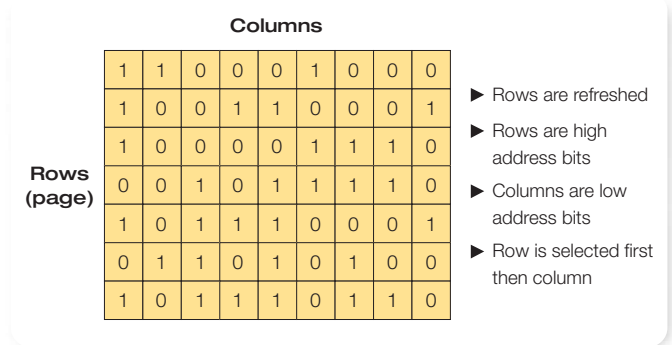


Figure 1. DRAMs memory cells organized into a two-dimensional array of rows and columns.

DRAMs core architecture consists of memory cells organized into a two-dimensional array of rows and columns (See Figure 1). To access a memory cell requires two steps. First, you address a specific row and then you address a specific column in the selected row. In other words, first an entire row is read internally in the DRAM IC and then the column address selects which column of the row is to be read or to be written to the DRAM IC I/O (Input/Output) pins.

DRAM reads are destructive, meaning the data in the row of memory cells are destroyed in the read operation. Therefore, the row data need to be written back into the same row after the completion of a read or write operation on that row. This operation is called precharge and is the last operation on a row. It must be done before accessing a new row and is referred to as closing an open row.

Fundamentals of SDRAM Memory

► Primer

Analysis of computer memory accesses show that reads of sequential memory addresses are the most common types of memory accesses. This is reasonable since reading computer instructions are typically more common than data read or writes. Also, most instruction reads are sequential in memory until an instruction branch or a jump to subroutine occurs.

A DRAM row is called a memory page and once the row is opened you can access multiple sequential or different column addresses in that row. This increases memory access speed and reduces memory latency by not having to resend the row address to the DRAM when accessing memory cells in the same memory page. As a result, the row address is the computer's higher order address bits and the column address is the lower order address bits. Since the row and column addresses are sent at different times, the row address and the column address are multiplexed on the same DRAM pins in order to reduce package pin count, cost and size. Typically the size of the row address is larger than the column address because the power usage is related to the number of columns.

Early DRAMs had control signals such as RAS# (Row Address Select active low) and CAS# (Column Address Select active low) to select the row and column addressing operation being performed. Additional DRAM control signals include WE# (Write Enable active low) for selecting write or read operation, CS# (Chip Select active low) for selecting the DRAM and OE# (output enable active low). The early

DRAMs had control signals that were asynchronous and had various timing specifications covering their sequence and time relationships to determine the DRAM operating mode.

The early DRAMs read cycle had four steps. First, RAS# goes low with a row address on the address bus. Secondly, CAS# goes low with a column address on the address bus. Third, OE# goes low and read data appears on DQ data pins. The time from the first step to the third step when data is available on DQ pins is called latency. The last step is RAS#, CAS# and OE# going high (inactive) and waiting for the internal precharge operation to complete restoration of the row data after the destructive read. The time from the first step to completion of the last step is the memory cycle time. Signal timing of the above signals is related to the sequence of edges and is asynchronous. There are no synchronous clock operations with these early DRAMs.

The DRAM memory cell needs to refresh to avoid losing its data contents. This requires refresh of the capacitor before it loses its charge. Refreshing memory is the responsibility of the memory controller and the refresh time specification varies with different DRAM memories. The memory controller performs a refresh by doing a RAS# only cycle with the row address. At the end of the RAS# only cycle is the precharge operation of restoring the row data that was address in the RAS# only cycle. Typically, the memory controller would have a row counter that would sequentially generate all row addresses that were needed by the RAS# only refresh cycles.




Figure 2. DRAM refresh implementations include distributed refresh and burst refresh. (Image courtesy of Micron Technology, Inc.)

There are two refresh strategies (See Figure 2). The first strategy is for the memory controller to refresh all rows sequentially in a burst of refresh cycles and then return control of memory back to the processor for normal operation. The next burst of refresh operations occurs before reaching the maximum refresh time. The second refresh strategy is for the memory controller to interleave the refresh cycles with normal processor memory operations. This refresh method spreads out the refresh cycles over the maximum refresh time.

Early DRAMs evolved and implemented the refresh counter on the DRAM IC to take care of sequentially generated row addresses. Internally to the DRAM IC, the refresh counter is an input to a multiplexer that controls the memory array row address. The other multiplexer input is from the row address from the external address input pins. This internal refresh counter eliminated the need for an external refresh counter circuit in the memory controller. Some of these DRAMs supported a CAS# before RAS# cycle to initiate a refresh cycle using the internally generated row address.

SDRAM

The DRAM's asynchronous operation caused many design challenges when interfacing it to a synchronous processor. SDRAM (Synchronous DRAM) was designed to synchronize the DRAM operation to the rest of the computer system and to eliminate defining all the different modes of memory operations based on the sequence of CE# (Chip Enable active low), RAS#, CAS# and WE# edge transitions. SDRAM added a clock signal and the concept of memory commands. The type of memory command is determined by the state of CE#, RAS#, CAS# and WE# signals at the rising edge of the SDRAM clock. Data sheets describe the memory commands in table form based on the state of CE#, RAS#, CAS# and WE# signals.

For example, an Activate command sends a row address to the SDRAM to open a row (page) of memory. Next is a sequence of Deselect commands to satisfy timing requirements before sending the Read or Write command with the column address. Once the row (page) of memory is opened with an Activate command, several Read and Write commands can operate on the data in that row (page) of memory. A Precharge command is required to close the row before another row can open.

Fundamentals of SDRAM Memory

► Primer

DDR SDRAM

DDR (Double Data Rate) SDRAMs increased the memory data rate performance by increasing clock rates, bursting of data and transferring two data bits per clock cycle (See Table 1). DDR SDRAMs burst multiple memory locations in a single read or single write command. A read memory operation entails sending an Activate command followed by a Read command. The memory responds after its latency with a burst of two, four, or eight memory locations at a data rate of two memory locations per clock cycle. Therefore, four memory locations are read from or written to in two consecutive clock cycles.

DDR SDRAMs have multiple banks to provide multiple interleaved memory access, which increases memory bandwidth. A bank is one array of memory, two banks are two arrays of memory, four banks are four arrays of memory, etc (See Figure 3). Four banks require two bits for bank address (BA0 & BA1).

For example, a DDR SDRAM with four banks operates in the following manner. First, an Activate command opens a row in the first bank. A second Activate command opens a row in the second bank. Now any combinations of Read or Write commands can be sent to either the first bank or the second bank with their open rows. When Read and Write operations on the bank are completed, a Precharge command closes the row and the bank is ready for an Activate command to open a new row.

Note that the power required by the DDR SDRAM is related to the number of banks with open rows. More open rows require more power and larger row sizes require more power. Therefore, for low power applications one should open only one row at a time in each bank and not have multiple banks each with open rows.

DDR SDRAM	Data Rate	Memory Clock
DDR-266	266 Mb/s/pin	133 MHz
DDR-333	333 Mb/s/pin	166 MHz
DDR-400	400 Mb/s/pin	200 MHz

Table 1. DDR SDRAM data rates and clock speeds.

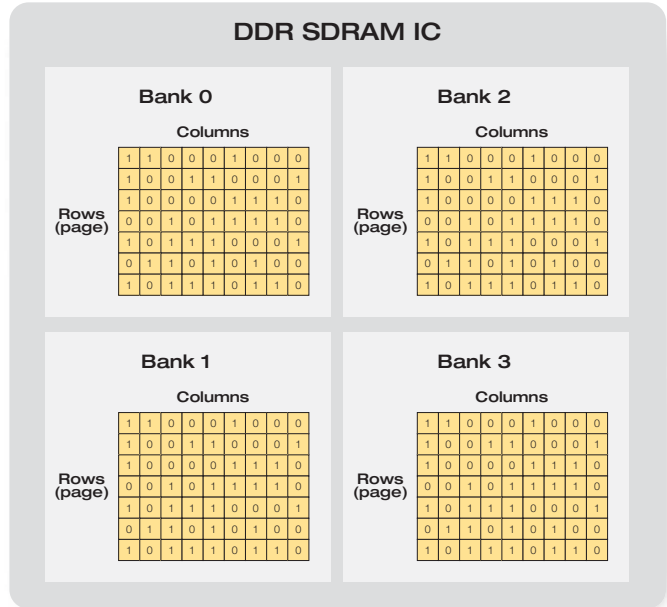


Figure 3. Multiple memory banks in a DDR SDRAM provide increased access flexibility and improved performance.

Interleaving consecutive memory words in consecutive memory banks is supported when the bank address bits are connected to the lower order address bits in the memory system. Consecutive memory words are in the same memory bank when the bank address bits are connected to the higher order address bits in the memory system.

DDR2 SDRAM

DDR2 SDRAM has several improvements over DDR SDRAM. DDR2 SDRAM clock rates are higher, thus increasing the memory data rates (See Table 2). Signal integrity becomes more important for reliable memory operation as the clock rates increase. As clock rates increase, signal traces on the circuit boards become transmission lines and proper layout and termination at the end of the signal traces becomes more important.

Termination of the address, clock and command signals are somewhat straightforward because these signals are unidirectional and are terminated on the circuit boards. The data signals and data strobes are bidirectional. The memory controller hub drives them during a write operation and the DDR2 SDRAM drives them during a read operation. To add to the complexity, multiple DDR2 SDRAMs are connected to the same data signals and data strobes. These multiple DDR2 SDRAMs can be on the same DIMM and on different DIMMs in the memory system. As a result, the data and data strobe drivers and receivers are constantly changing depending upon the read/write operation and which DDR2 SDRAM is being accessed.

DDR2 SDRAM improves the signal integrity of data signals and data strobes by providing ODT (On-Die Termination), an ODT signal to enable the on-die termination and the ability to program the on-die termination values (75 ohms, 150 ohms, etc.) with the DDR2 SDRAM extended mode register. The on-die termination value and operation is controlled by the memory controller hub and are a function of a DDR2 SDRAM DIMM's location and type of memory operation (reads or writes). ODT operation results in better signal integrity by creating a larger eye diagram for the data valid window with increased voltage margins, increased slew rates, reduced overshoot and reduced ISI (Inter-Symbol Interference).

DDR2 SDRAM reduces memory system power by operating at 1.8 volts, which is 72% of DDR SDRAM's 2.5 volts. In some implementations, the number of columns in a row has been reduced, resulting in lower power when a row is activated for read or writes.

DDR2 SDRAM	Data Rate	Memory Clock
DDR2-400	400 Mb/s/pin	200 MHz
DDR2-533	533 Mb/s/pin	266 MHz
DDR2-667	667 Mb/s/pin	333 MHz
DDR2-800	800 Mb/s/pin	400 MHz

Table 2. DDR2 SDRAM data rates and clock speeds.

Another benefit of the lower operating voltages is the lower logic voltage swings. For the same slew rate, the reduced voltage swings increase logic transition speeds to support faster clock rates. In addition, the data strobe can be programmed to be a differential signal. Using differential data strobe signals reduces noise, crosstalk, dynamic power consumption and EMI (Electromagnet Interference) and increases noise margin. Differential or single-end data strobe operation is configured with the DDR2 SDRAM extended mode register.

A new feature introduced with DDR2 SDRAM is additive latency, which provides the memory controller hub the flexibility to send the Read and Write commands sooner after the Activate command. This optimizes memory throughput and is configured by programming the additional latency using the DDR2 SDRAM extended mode register.

DDR2 SDRAM improves data bandwidth of 1Gb and 2Gb DDR2 SDRAMs by using eight banks. The eight banks increase the flexibility of accessing large memory DDR2 SDRAMs by interleaving different memory bank operations. Also, for large memories, DDR2 SDRAM supports a burst length up to eight.

DDR2 SDRAM data sheets are over 100 pages and the above DDR2 SDRAM features are highlights of its key features. Refer to DDR2 SDRAM data sheets for their complete features and details of operation.

Fundamentals of SDRAM Memory

► Primer

DDR3 SDRAM

DDR3 SDRAM is a performance evolution beyond DDR2 SDRAM. DDR3 SDRAMs will support the next level of faster data rates and clock speeds (See Table 3). Other expected changes include reducing the DDR3 SDRAM operating voltage to 1.5 volts, which is 83% of DDR2 SDRAM's 1.8 volts. DDR3 SDRAM is the memory that will be used by FB-DIMM2 (Fully Buffered DIMM2) implementations. DDR3 SDRAM specifications are under development and are subject to change until they are approved by JEDEC.

You can join the JEDEC DDR3 SDRAM subcommittee at www.jedec.org to help in the development of the DDR3 SDRAM specifications and you can monitor the JEDEC web site for posting of approved DDR3 SDRAM specifications.

DIMMs

Dual inline memory modules (DIMMs) are plug-in memory modules for computers. DIMMs vary in physical size, memory data width, ranks, memory sizes, memory speeds and memory architectures. JEDEC has defined DIMMs standards and continues to work on defining new DIMMs based on new memory types and memory architectures.

DIMM Physical Size

The standard DIMM size is used in desktops, workstations and servers. SO-DIMMs (Small Outline DIMMs) are small size DIMMs used in laptops and other space constant implementations. The Butterfly configuration refers to two SO-DIMMs parallel to the computer motherboard that have their edge connectors next to each other. Think of the two edge connectors as the butterfly body and the SO-DIMMs as the open butterfly wings. Mini-DIMMs (Miniature DIMMs) are smaller than SO-DIMMs and are used in single board computers. VLP-DIMMs (Very Low Profile DIMMs) are shorter in height and are used in blade servers.

DDR3 SDRAM	Data Rate	Memory Clock
DDR3-800	800 Mb/s/pin	400 MHz
DDR3-1066	1066Mb/s/pin	533 MHz
DDR3-1333	1333Mb/s/pin	667 MHz
DDR3-1600	1600 Mb/s/pin	800 MHz

Table 3. Expected DDR3 SDRAM data rates and clock speeds.

DIMM	4 I/Os per IC	8 I/Os per IC	16 I/Os per IC
1 Rank	16 ICs	8 ICs	4 ICs
2 Ranks	32 ICs	16 ICs	8 ICs
4 Ranks	64 ICs	32 ICs	16 ICs

Table 4. Number of memory ICs per DIMM without ECC based on ranks on the DIMM and number of data I/Os per memory IC.

DIMM Data Width

DIMM data width depends upon support for the ECC (Error Correction Code). ECC is eight check bits used for error detection and correction. The standard DIMM data width is 64 bits without ECC and 72 bits with the eight ECC bits.

DIMM Rank

A rank is a complete group of memory devices on a DIMM to support 64 data bits or 72 bits with ECC. A rank of two is two groups of memory devices on a DIMM. Rank of four is four groups of memory devices on a DIMM. Table 4 shows how many memory ICs are on a DIMM that supports a data width of 64 bits without ECC. At some point there is not enough room on both sides of the DIMM for all the memory ICs. To solve this problem, memory ICs are stacked on top of each other.

DIMM Memory Size & Speed

DIMM memory size depends on size of memory ICs used and DIMM configuration. A 512Mb (Meg bit) memory IC can be designed as different configurations (See Table 5). DIMM speed depends on the clock speed supported by the DDR, DDR2 and DDR3 SDRAMs used on the DIMM.

DIMM Architecture

There are three major DIMM architectures: UDIMMs, RDIMMs and FB-DIMMs. Each DIMM architecture has advantages and limitations.

UDIMM is an unregistered DIMM. UDIMM has no buffering of the DDR, DDR2 and DDR3 SDRAMs signals on the DIMM (See Figure 4). UDIMMs were the first implementation of DIMMs. For a single or dual DIMM memory system the UDIMMs are the fastest and lowest cost. The memory controller hub directly controls all the DRAM signals. No buffers or registers delay the signals between the memory controller hub and the SDRAM on the UDIMM. The number of UDIMMs that can be on a memory controller hub memory channel is limited by signal integrity. Signal integrity is decreased by the following factors: increased memory clock speed, increased trace lengths, increased number of UDIMMs on a memory channel, and increased number of ranks on a UDIMM. The memory controller hub sees every connector, every trace, every trace branch and every SDRAM pin. The impedance problems of the tree stub architecture limit the clock frequency and the number of UDIMMs that a memory channel can reliably operate.

Page Size	I/Os Pins	Banks
32 Meg bit	4	4
16 Meg bit	8	4
8 Meg bit	16	4

Table 5. Examples of different 512Mb (Meg bit) memory IC configurations.

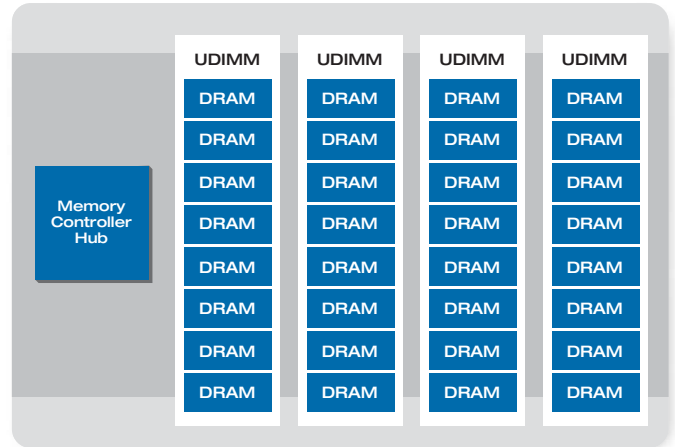


Figure 4. UDIMM has no buffering of the DRAM signals on the DIMM.

Memory controller hubs that have separate memory channels are one way to increase the number of UDIMMs in a memory system. Two separate memory channels can support two high speed UDIMMs with one UDIMM per memory channel.

Fundamentals of SDRAM Memory

► Primer

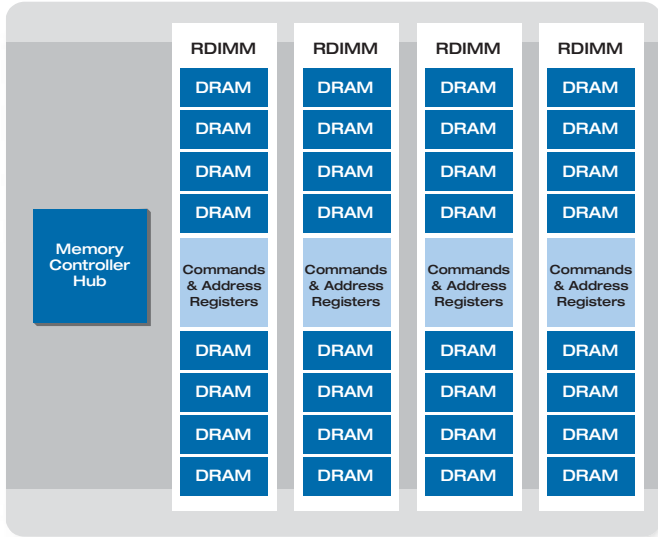


Figure 5. RDIMM buffers DRAM clock, command signals and address signals on the DIMM.

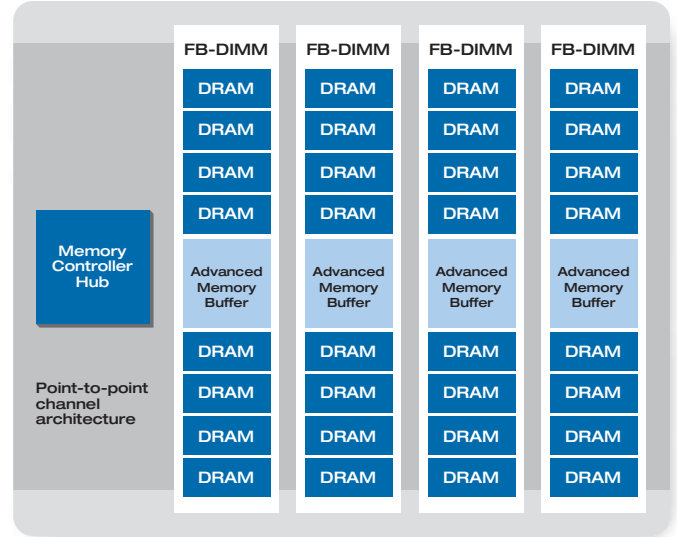


Figure 6. FB-DIMM buffers DDR2 SDRAM signals on the FB-DIMM.

RDIMM is a registered DIMM. RDIMM reduces some of the problems of the tree stub architecture by buffering the RDIMM SDRAMs clock, command signals and address signals on the RDIMM (See Figure 5). The clock signal is buffered with the Phase Lock Loop (PLL) and the command signals and addressing signals are buffered with register latches. A typical registered DIMM is implemented with a PLL IC and two ICs with registers. The memory controller hub clock, command signals and address signals see the impedances of the motherboard traces, DIMM connectors, RDIMM registers and RDIMM PLL. This reduced tree stub architecture allows for more RDIMMs to be used on a memory channel, making it faster. There is no buffering or reduced signal loading benefits for the bidirectional DQ data lines and DQS data strobe lines. Also, RDIMMs memory access times are one clock cycle slower than UDIMM because one clock cycle is required to latch the commands and address signals into the registers on a RDIMM.

FB-DIMM is a fully buffered DIMM. FB-DIMM uses DDR2 SDRAMs and FB-DIMM2 uses DDR3 SDRAMs. All DDR2

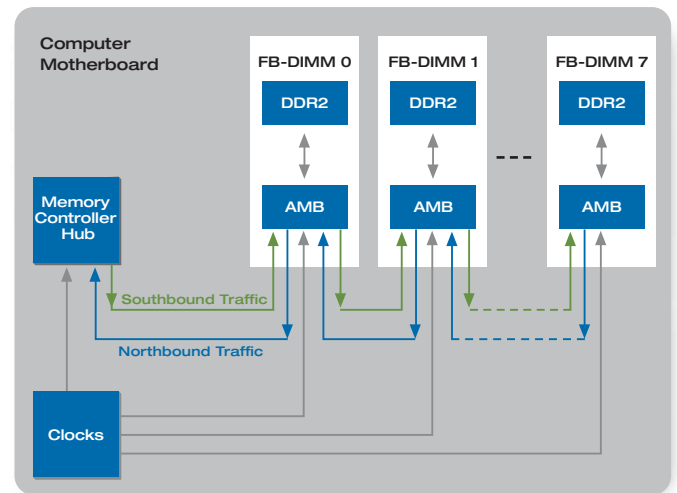


Figure 7. FB-DIMM point-to-point, high-speed serial architecture.

SDRAMs and DDR3 SDRAMs signals are buffered from the memory system with the AMB (Advanced Memory Buffer) IC on the FB-DIMM and FB-DIMM2 (See Figure 6). Different AMBs are used for FB-DIMM and FB-DIMM2. The first

FB-DIMM next to the memory controller hub uses up to 24 high-speed differential signals to communicate to the memory controller hub and uses up to 24 high-speed differential signals to communicate to the adjacent FB-DIMM (See Figure 7). These signals are unidirectional, point-to-point and use high-speed serial techniques to send commands and data. Ten signals carry the Southbound traffic from the memory controller hub toward the AMB, and 12 to 14 signals carry the Northbound traffic from the AMB back toward the memory controller hub. The FB-DIMM architecture is following modern high-speed digital design trends by replacing wide parallel buses with their tree stub architectures with a few point-to-point, high-speed serial lanes.

FB-DIMM architecture is designed to transfer data at full memory speed from the DDR2 memories on the FB-DIMM through the AMB and then on to the serial Northbound point-to-point signals to the memory controller hub. A serial frame is 12 bits long and its width varies from 10 Southbound signal lanes to 14 Northbound signal lanes. The 12 serial bits are transferred in a single DDR2 clock cycle. Therefore, you multiply the DDR2 clock frequency by 12 to obtain the data rate of the high-speed serial lanes. Alternatively, you multiply the DDR2 data rate by six to obtain the data rate of the high-speed serial lanes (See Table 6). The computer motherboard memory system clock is one half of the DDR2 clock. The AMB provides the PLL function to drive the DDR2 SDRAM clocks twice as fast. The number of Southbound lanes is fixed at 10 and the number of Northbound lanes varies from 12 to 14. You multiply the lanes by 12 serial bits to obtain the number of

DDR2 SDRAM	Data Rate	DDR2 Clock	Serial Data Rate
DDR2-533	533 Mb/s/pin	266 MHz	3.2 Gb/s
DDR2-667	667 Mb/s/pin	333 MHz	4.0 Gb/s
DDR2-800	800 Mb/s/pin	400 MHz	4.8 Gb/s

Table 6. The serial data rate of the FB-DIMM point-to-point serial lanes depends upon the speed of the DDR2 SDRAM on the FB-DIMM.

Lanes	Bits in a Serial Frame
10 Southbound Lanes	120
12 Northbound Lanes	144
13 Northbound Lanes	156
14 Northbound Lanes	168

Table 7. The size of the FB-DIMM serial data Northbound frame depends upon the number of Northbound lanes used.

bits in a frame (See Table 7). The Southbound lanes are fixed at 10 and therefore have 120 bits per frame. The Northbound lanes vary depending on whether or not ECC is used on the FB-DIMM and the number of serial CRC bits used.

For example, without ECC, two 64 bit words ($2 \times 64 = 128$ bits total bits) need to be transferred in a Northbound frame. 12 Northbound lanes with 144 total bits per frame are capable of handling this data rate. On the other hand, with ECC, two 72-bit words ($2 \times 72 = 144$ bits total bits) need to be transferred in a Northbound frame. 13 or 14 Northbound lanes are needed because 12 Northbound lanes with 144 total bits per frame do not have enough bits for 144 data bits plus the serial CRC bits.

Fundamentals of SDRAM Memory

► Primer

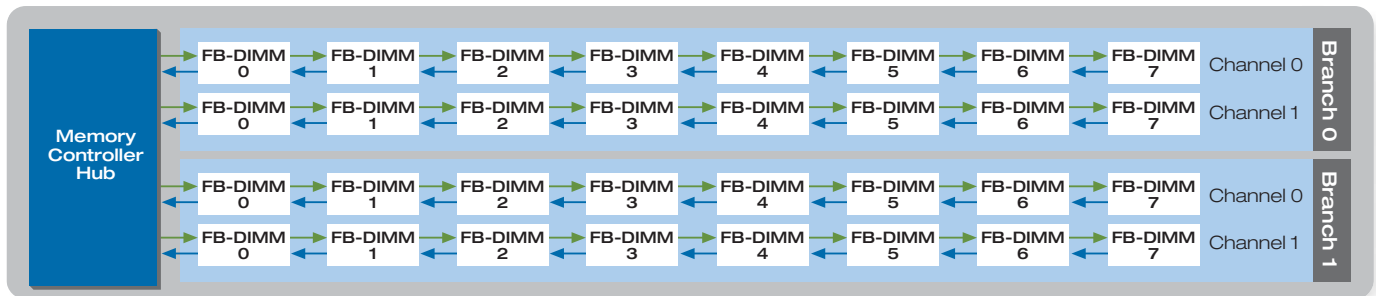


Figure 8. FB-DIMM memory controller hub supports four channels with each channel having up to eight FB-DIMMs. A total of 32GB of system memory is possible using 1GB FB-DIMMs. This can easily grow to 128GB using 4GB FB-DIMMs. This large memory capacity is very attractive for performance servers and high end workstations.

One FB-DIMM channel supports up to eight FB-DIMMs. FB-DIMM motherboard layouts are easier because only 24 high speed signals traces per channel are needed from the memory controller hub and these traces do not have to be the same trace length. UDIMMs or RDIMMs, however, need large numbers of parallel signals (commands, address, data, etc.) from the memory controller hub and these trace lengths must match for proper operation. As a result of the fewer FB-DIMM signals per memory channel and the simpler circuit board trace layouts, the memory controller hub can support more FB-DIMM memory channels (See Figure 8). This means more memory bandwidth, since each memory channel can be active at the same time.

FB-DIMMs are being used in servers, which require very large memory systems. It is expected that high-end workstations will also use FB-DIMMs.

Serial Presence Detect

The Serial Presence Detect (SPD) function is on all computer DIMMs and is used to provide DIMM memory configuration information such as memory size, speed,

latency, timing, manufacturer, etc. to the computer's BIOS during the computer's power-on (See Table 8). At power-on, the BIOS (Basic Input Output Software) reads the configuration information of each DIMM using the SPD function. This information is then used to configure the memory controller hub and the DRAM mode and extended mode registers on each UDIMM and RDIMM. The SPD functions are specified by JEDEC standards.

For UDIMMs and RDIMMs the SPD functions are implemented in a small nonvolatile memory IC with a slow speed I²C interface that is located on each DIMM. The motherboard has an I²C interface with a unique address (0 through 7) for each DIMM slot. At power-on, each DIMM slot is checked using the I²C interface. If a DIMM is present the SPD values are read by the BIOS.

For FB-DIMMs the SPD functions are implemented in the AMB, which has the I²C interface. The FB-DIMM I²C interface is called SMBus (System Management bus). The SMBus is used to configure the AMB in each FB-DIMM.

Byte Number	Function Described	SPD HEX Value
0	Defines number of bytes written into serial memory at module manufacturer	80
1	Total number of SPD memory bytes	08
2	Fundamental memory type (FPM or EDO)	01 or 02
3	Number of row addresses on the memory device	0C
4	Number of column addresses on memory device	0A
5	Number of physical banks on this memory module	01
6	Data width of this module	40
7	Data width (continued)	00
8	Module voltage interface levels	01
9	RAS access time of this assembly	3C
10	CAS access time of this assembly	0F
11	Module error correction configuration type (non-parity, parity, ECC)	00 or 01 or 02
12	Refresh rate/type	00 or 83
13	Primary DRAM width	10
14	Error checking DRAM data width	00
15 - 61	Reserved	00
62	SPD revision	00
63	Checksum for bytes 0-62	Calculated
64	Manufacturer JEDEC ID code	2C
65 - 71	Manufacturer JEDEC ID code (continued)	00
72	Manufacturing location	01 - 08
73 - 90	Manufacturer's part number	Variable
91	PCB identification code	01 - 09
92	PCB identification code (continued)	00
93	Year of manufacture	Variable
94	Week of manufacture	Variable
95 - 98	Module serial number	Variable
99 - 125	Manufacturer-specific data	Variable
126 - 127	Reserved	00
128 - 255	Open user free-form area not defined	FF

Table 8. The computer BIOS reads the DIMM configuration using the Serial Presence Detect (SPD) interface. The SPD data is specified by JEDEC standards.

Fundamentals of SDRAM Memory

► Primer

Memory System Design

The first few steps of product design are product requirements, product architecture design and subsystem design. One of the subsystem designs is the memory system. Memory system design is dependant on memory size, speed, power, existing standards, new developing standards, reuse of existing designs and other requirements.

The computer chipset manufacturers heavily influence memory system designs for computers. Some of these computer chipset manufacturers have their own testing procedures, validation processes and workshops to test products. Typically, these computer chipset manufacturers' web sites list memory products that passed their compatibility testing.

Design Simulation

A key part of memory system design is design simulation. The importance of comprehensive memory system design simulation cannot be understated. Experience has shown that a resistor value change of only a few ohms can have a significant impact on having a reliable operating memory system.

Memory system design simulation should include the effects of probing loading caused by any instrument when it is connected to the prototype memory system. The verification and debugging process will be very difficult if the

prototype stops operating because of probe loading. Also, simulation should analyze the signals at probe test points with the loading of the instrument probe. The data valid window will change along the signal trace from the memory controller hub driver to the SDRAM pins.

Probing test points should be as close as possible to the receiver pins so that the instrument shows the signal that the receiver is seeing. Sometimes this is not possible and interposers, test adapter boards and other special probing fixtures and aids are used to retrieve difficult to access signals. These probing aids should also be included in design simulations to understand their effect on the SDRAM signals and the measurement of the signals.

Design Verification

Using new DRAM features in designs requires new design methods and techniques, which range from new techniques in design simulation to new BIOS operation. As a result, DRAM design implementations require complete verification and testing, ranging from circuit board construction to software operation to ensure reliable memory operation. Product reliability will suffer if a memory system has infrequent random errors because of a design implementation that has not been fully verified. In addition, the customer may require a product to satisfy various compliance testing requirements that have been defined by JEDEC or by other manufacturers.

Verification Strategy

It is important to have a strategy for effectively and quickly debugging design problems in any design implementation. Quick time-to-market product development requires verification/debugging planning early in the design. This plan should identify the following requirements:

- what are new design elements and what are reused design elements,
- what to avoid and change based on past designs,
- what level of validation and testing is needed, does the testing require special operating modes or signal patterns,
- what special design-in features are needed (e.g., probing test points or test fixtures), has simulation analysis accounted for probing the prototype, are signal stimuli needed, is special software needed to exercise the hardware,
- what environmental tests are needed (e.g., temperature, humidity, etc.),
- what circuit operational visibility do you have in order to debug it,
- what regulatory compliance testing is required, will the validation/debug test points be used to test the product in manufacturing, will the validation/debug test points be used to repair the product in service, and how do you manage the risk of what you do not know today.

For example, some verification strategies include building a validation prototype with numerous probing test points to verify the new system architecture with new ASICs/FPGAs. It is best that the validation prototype operates at full speed to verify at-speed operation and performance. Complex designs require more comprehensive visibility of their real-time operation in order to pin-point problems quickly. Once the validation prototype is running correctly and has completed validation, the final prototype is implemented with reduced test points.

SDRAM Verification

DRAM verification and testing techniques depend upon what is being designed. DRAM designs are grouped into the following types: computer memory controller hub ICs, memory ICs, AMB ICs, DIMMs, computer motherboards and embedded systems. Each of these products requires different validation strategies, different validation tests and different test equipment. For example, a memory IC designer will not be verifying circuit board construction whereas the DIMM designer will be verifying the DIMM circuit board construction.

The memory controller is typically designed by the embedded systems designer because of its unique requirements to work with a specific processor and unique embedded system input/output configuration. As a result, a significant part of the design work is designing the memory controller and designing the circuit board layout between the memory controller and the memory ICs. Verifying this part of the design is critical for reliable operation.

Fundamentals of SDRAM Memory

► Primer

Verification	Tasks	Instruments
Circuit Board Construction	Single-ended trace impedances	Sampling Oscilloscope with TDR
	Differential trace impedances	Sampling Oscilloscope with TDR
	Trace lengths	Sampling Oscilloscope with TDR
	Crosstalk	Sampling Oscilloscope with TDT
Electrical Power & Signals	Power supply quality, noise, glitches & ground bounce	Oscilloscope
	Clock signal quality, rise & fall times/slew rates, spread spectrum clocking profile	Oscilloscope with jitter analysis software
	Command, address & data valid windows, clock, strobes & data signal skew	Oscilloscope with jitter analysis software
	FB-DIMM serial signals data valid windows	Oscilloscope with serial data compliance and analysis software, Signal Sources & FB-DIMM fixtures
Protocols Sequences & Timing	Memory system power up initialization protocols & timing	Logic analyzer with SDRAM support packages
	SDRAM mode register operation	Logic analyzer with SDRAM support packages
	SDRAM command protocols & timing	Logic analyzer with SDRAM support packages
	Read/Write data valid windows	Logic analyzer with SDRAM support packages
	Refresh operations	Logic analyzer with SDRAM support packages
	Memory channel traffic	Logic analyzer with FB-DIMM support packages

Table 9. Verification tasks with associated test equipment.

DRAM verification and testing techniques require a range of test and measurement equipment such as sampling oscilloscopes, oscilloscopes, logic analyzers, probes, test fixtures, analysis software, compliance software, etc. (See Table 9). Test equipment needs to provide unobtrusive probing, precise acquisition and complete system visibility of electrical signals and protocol layers. To help designers quickly verify memory operation, powerful analysis capability is also necessary.

Monitoring a computer system or embedded system with a logic analyzer creates a powerful verification and debugging development environment. The logic analyzer is used to trace and correlate the processor bus activity, the memory activity and the input/output operations. Complete system visibility on logic analyzer display provides critical design

insight to real-time system operation. In addition, using integrated oscilloscope and logic analyzer probing, triggering and display provides complete design visibility from the software listing, the protocol listing, the digital waveforms and the analog waveforms on the same display. The result is a powerful, comprehensive and efficient analysis of the prototype.

Tektronix offers a comprehensive tool set including industry-leading oscilloscopes, true differential TDRs, and logic analyzers with Nexus Technology memory supports to enable embedded and computer designers to perform quick and accurate electrical testing and operational validation of their memory designs. Collectively, this tool set provides superior performance with unparalleled ease-of-use, making it an ideal solution for embedded systems and computer memory systems verification and debugging.

Glossary

For easy reference, the glossary also includes common terms not used in this document.

Advanced Memory Buffer (AMB): Provides intelligent Southbound and Northbound channel initialization to align the high-speed serial clocks, locate frame boundaries and verify channel connectivity.

Amplitude: The magnitude of a quantity or strength of a signal. In electronics, amplitude usually refers to either voltage or power.

Analog Signal: A signal with continuously variable voltages.

Analog-to-Digital Converter (ADC): A digital electronic component that converts an electrical signal into discrete binary values.

Asynchronous: Not synchronized. The logic analyzer runs its own sampling clock. The clock is independent and unaware of the timing on the device under test. This is the basis of the “timing” acquisition mode.

Attenuation: A decrease in signal amplitude during its transmission from one point to another.

Ball Grid Array (BGA): An integrated circuit package.

Bandwidth: A frequency range, usually limited by -3 dB.

Bit: a binary character whose state may be either 1 or 0.

Byte: a unit of digital information usually consisting of eight bits.

Chip Enable (CE#): Activates the device.

Chip Select (CS#): Selects the device.

Clock Rate: Fundamental rate in cycles per second at which a device performs its most basic operation.

Column Address Select (CAS#): Selects the address column of interest within the device.

Cursor: An on-screen marker that you can align with a waveform to make more accurate measurements.

Cyclic Redundancy Code (CRC): A number derived from, and stored or transmitted with, a block of data in order to detect corruption. By recalculating the CRC and comparing it to the value originally transmitted, the receiver can detect some types of transmission errors.

Decibel (dB): Unit used to express the relative difference in power between two electric signals, equal to ten times the common logarithm of the ratio of the two levels.

Device Under Test (DUT): The device being tested by the measurement instrument.

Digital Oscilloscope: A type of oscilloscope that uses an analog-to-digital converter (ADC) to convert the measured voltage into digital information. Three types: digital storage, digital phosphor, and digital sampling oscilloscopes.

Digital Phosphor Oscilloscope (DPO): A type of digital oscilloscope that closely models the display characteristics of an analog oscilloscope while providing traditional digital oscilloscope benefits (waveform storage, automated measurements, etc.) The DPO uses a parallel-processing architecture to pass the signal to the raster-type display, which provides intensity-graded viewing of signal characteristics in real time. The DPO displays signals in three dimensions: amplitude, time and the distribution of amplitude over time.

Digital Sampling Oscilloscope: A type of digital oscilloscope that employs equivalent-time sampling method to capture and display samples of a signal, ideal for accurately capturing signals whose frequency components are much higher than the oscilloscope's sample rate.

Digital Signal: A signal whose voltage samples are represented by discrete binary numbers.

Fundamentals of SDRAM Memory

► Primer

Digital Storage Oscilloscope (DSO): A digital oscilloscope that acquires signals via digital sampling (using an analog-to-digital converter). It uses a serial-processing architecture to control acquisition, user interface, and the raster display.

Digitize: The process by which an analog-to-digital converter (ADC) in the horizontal system samples a signal at discrete points in time and converts the signal's voltage at these points into digital values called sample points.

Double Data Rate (DDR): The peak data rate is twice the rate at which commands can be clocked into the device.

Dual Inline Memory Module (DIMM): The prevalent packaging scheme for dynamic random access memory components in PC platforms.

Dynamic Random Access Memory (DRAM): A type of memory that stores each bit of data in a separate capacitor.

Error Correction Code (ECC): Eight check bits used for error detection and correction.

Field Effect Transistor (FET): A transistor in which the output current is controlled by a variable electric field.

Fine-pitch Ball Grid Array (FBGA): An integrated circuit package.

Frequency: The number of times a signal repeats in one second, measured in Hertz (cycles per second). The frequency equals 1/period.

Fully Buffered Dual Inline Memory Module (FB-DIMM): A next generation memory architecture.

Gigabit (Gb): 1 billion bits of information.

Gigabyte (GB): 1 billion bytes of information.

Gigahertz (GHz): 1 billion Hertz.

Gigatransfers per Second (GT/s): One billion data transfers per second.

Glitch: An intermittent, high-speed error in a circuit.

Hertz (Hz): One cycle per second. The unit of frequency.

iCapture™ Multiplexing: provides simultaneous digital and analog acquisition through a single logic analyzer probe.

iLink™ Toolset: consists of several elements designed to speed problem detection and troubleshooting, including: iCapture™ Multiplexing, iView™ Display, and iVerify™ Analysis.

Input/Output (I/O): Typically referring to signals going into or out of a device.

Integrated Circuit (IC): A set of components and their interconnections etched or imprinted on a chip.

Interleave: To intersperse or place at regular intervals.

iVerify™ Analysis: offers multi-channel bus analysis and validation testing using oscilloscope-generated eye diagrams.

iView™ Display: delivers time-correlated, integrated logic analyzer and oscilloscope measurements on the logic analyzer display.

Joint Electron Device Engineering Council (JEDEC): The semiconductor engineering standardization body of the Electronic Industries Alliance (EIA), a trade association that represents all areas of the electronics industry.
www.jedec.org

Kilohertz (kHz): 1 thousand Hertz.

Latency: The time that elapses between a stimulus and the response. For example, the time from the first step to the third step of the read cycle when data is available on DQ pins.

Loading: The unintentional interaction of the probe and oscilloscope with the circuit being tested which distorts a signal.

Logic Analyzer: An instrument used to make the logic states of many digital signals visible over time. It analyzes the digital data and can represent the data as real-time software execution, data flow values, state sequences, etc.

MagniVu™ Acquisition: a unique high-resolution sampling architecture at the heart of every TLA Series logic analyzer. MagniVu acquisition provides a dynamic record of signal activity surrounding the trigger point with higher resolution.

Megabit (Mb): One million bits of information.

Megabyte (MB): One million bytes of information.

Megahertz (MHz): One million Hertz.

Megasamples per second (MS/s): A sample rate unit equal to one million samples per second.

Megatransfers per second (MT/s): One million data transfers per second.

Memory Cycle Time: The time from the first step to completion of the last step in a read cycle.

Microsecond (μs): A unit of time equivalent to 0.000001 seconds.

Millisecond (ms): A unit of time equivalent to 0.001 seconds.

Miniature Dual Inline Memory Module (Mini-DIMM): Smaller than SO-DIMMs and typically used in single board computers.

Motherboard: A computer's main system circuit board containing processor, memory controller, hard disk controller, input/output interface chipset, etc. Other circuit boards such as DIMMs and video cards are plugged into the motherboard.

Nanosecond (ns): A unit of time equivalent to 0.000000001 seconds.

Noise: An unwanted voltage or current in an electrical circuit.

Oscilloscope: An instrument used to make voltage changes visible over time. The word oscilloscope comes from "oscillate," since oscilloscopes are often used to measure oscillating voltages.

Output Enable (OE#): Activates the device output.

Period: The amount of time it takes a wave to complete one cycle. The period equals 1/frequency.

Pre-trigger Viewing: The ability of a digital instrument to capture what a signal did before a trigger event. Determines the length of viewable signal both preceding and following a trigger point.

Precharge: The phase in the access cycle of DRAM during which the storage capacitors are charged to the appropriate value.

Probe: A measurement instrument input device, usually having a pointed metal tip for making electrical contact with a circuit element, a lead to connect to the circuit's ground reference, and a flexible cable for transmitting the signal and ground to the instrument.

Pulse: A common waveform shape that has a fast rising edge, a width, and a fast falling edge.

Pulse Train: A collection of pulses traveling together.

Pulse Width: The amount of time the pulse takes to go from low to high and back to low again, conventionally measured at 50% of full voltage.

Ramps: Transitions between voltage levels of sine waves that change at a constant rate.

Random Access Memory (RAM): A memory device in which information can be accessed in any order.

Read Cycle: Periodically repeated sequence of events used to read from a device.

Fundamentals of SDRAM Memory

► Primer

Record Length: The number of waveform points used to create a record of a signal.

Refresh: To maintain by sending a new electric pulse to recharge the chips.

Registered Dual Inline Memory Module (RDIMM): Reduces some of the problems of the tree stub architecture by buffering the RDIMM SDRAMs clock, command signals and address signals on the RDIMM.

Rise Time: The time taken for the leading edge of a pulse to rise from its low to its high values, typically measured from 10% to 90%.

Row Address Select (RAS#): Selects the address row of interest within the device.

Sample Point: The raw data from an ADC used to calculate waveform points.

Sample Rate: Refers to how frequently a digital measurement instrument takes a sample of the signal, specified in samples per second (S/s).

Sampling: The conversion of a portion of an input signal into a number of discrete electrical values for the purpose of storage, processing and/or display by an instrument.

Serial Presence Detect (SPD): Uses a separate, electronically erasable/programmable, read-only memory (EEPROM) device to hold module density, timing, and performance parameters.

Signal Integrity: The accurate reconstruction of a signal, determined by the systems and performance considerations of an instrument, in addition to the probe used to acquire the signal.

Signal Source: A test device used to inject a signal into a circuit input; the circuit's output is then read by a measurement instrument. Also known as a signal generator.

Small Outline Dual Inline Memory Module (SO-DIMM): Small size DIMMs used in laptops and other space constant implementations.

Synchronous: Synchronized. A logic analyzer state acquisition is said to be synchronous because the logic analyzer receives its clock information from an external source, usually the DUT. This causes the two systems to be synchronized, and the logic analyzer acquires data only when the DUT is active. This is known as the "state" acquisition mode.

Synchronous Dynamic Random Access Memory (SDRAM): Designed to synchronize the DRAM operation to the rest of the computer system and to eliminate defining all the different modes of memory operations based on the sequence of CE#, RAS#, CAS# and WE# edge transitions.

System Under Test (SUT): The system being tested by the measurement instrument.

Time Domain Reflectometry (TDR): A convenient way to evaluate impedance values and variations along a transmission line such as cables, connectors or a microstrip on a PC board.

Trigger: The circuit that references a horizontal sweep on a measurement instrument.

Trigger Holdoff: A control that allows you to adjust the period of time after a valid trigger during which the instrument cannot trigger.

Trigger Level: The voltage level that a trigger source signal

must reach before the trigger circuit initiates a sweep.

Unregistered Dual Inline Memory Module (UDIMM):

UDIMMs were the first implementation of DIMMs. UDIMM has no buffering of the DDR, DDR2 and DDR3 SDRAMs signals on the DIMM.

Very Low Profile Dual Inline Memory Module

(VLP-DIMM): DIMMS that are shorter in height and often used in blade servers.

Volt (V): The unit of electric potential difference.

Voltage: The difference in electric potential, expressed in volts, between two points.

Wave: The generic term for a pattern that repeats over time. Common types include: sine, square, rectangular, saw-tooth, triangle, step, pulse, periodic, non-periodic, synchronous, asynchronous.

Write Enable (WE#): Activates the write ability to the device.

Contact Tektronix:

- ASEAN / Australasia / Pakistan** (65) 6356 3900
- Austria** +41 52 675 3777
- Balkan, Israel, South Africa and other ISE Countries** +41 52 675 3777
- Belgium** 07 81 60166
- Brazil & South America** 55 (11) 3741-8360
- Canada** 1 (800) 661-5625
- Central East Europe, Ukraine and Baltics** +41 52 675 3777
- Central Europe & Greece** +41 52 675 3777
- Denmark** +45 80 88 1401
- Finland** +41 52 675 3777
- France & North Africa** +33 (0) 1 69 86 81 81
- Germany** +49 (221) 94 77 400
- Hong Kong** (852) 2585-6688
- India** (91) 80-22275577
- Italy** +39 (02) 25086 1
- Japan** 81 (3) 6714-3010
- Luxembourg** +44 (0) 1344 392400
- Mexico, Central America & Caribbean** 52 (55) 56666-333
- Middle East, Asia and North Africa** +41 52 675 3777
- The Netherlands** 090 02 021797
- Norway** 800 16098
- People's Republic of China** 86 (10) 6235 1230
- Poland** +41 52 675 3777
- Portugal** 80 08 12370
- Republic of Korea** 82 (2) 528-5299
- Russia & CIS** 7 095 775 1064
- South Africa** +27 11 254 8360
- Spain** (+34) 901 988 054
- Sweden** 020 08 80371
- Switzerland** +41 52 675 3777
- Taiwan** 886 (2) 2722-9622
- United Kingdom & Eire** +44 (0) 1344 392400
- USA** 1 (800) 426-2200

For other areas contact Tektronix, Inc. at: 1 (503) 627-7111

Last Updated June 15 2005

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com



Copyright © 2005, Tektronix. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

11/05 FLG/WOW

52W-19200-0

