SLDRAM Architectural and Functional Overview

by Peter Gillingham, MOSAID Technologies Inc.

SLDRAM Benefits

The primary objective of DRAM is to offer the largest memory capacity at the lowest possible cost. This is achieved by two means; First, by optimizing the process and the design to minimize manufacturing cost, and second, ensuring that the device can serve high volume markets and can be manufactured in the highest possible volume to achieve the greatest economies of scale.

Serving a wide variety of Applications

SLDRAM has been specified as a general purpose high performance DRAM, serving a wide variety of applications. As computer main memory, in desktop, mobile, and high end servers and workstations SLDRAM offers high sustainable bandwidth, low latency, low power, user upgradeability, and support for large hierarchical memory configurations. For video, graphics, and telecommunications applications SLDRAM provides multiple independent banks, fast read/write bus turn-around, and the capability for small fully pipelined bursts. SLDRAM addresses the requirements of all major DRAM applications.

An Evolutionary Solution

SLDRAM is an evolutionary technology representing the next step in DRAM's evolution from EDO to SDRAM to DDR. SDRAM included several important architectural features over standard EDO, including multiple internal banks, a clocked synchronous interface, terminated small swing signaling, and programmable data bursts. These changes decoupled internal DRAM address and control paths from the data interface to achieve higher bandwidth. The emerging standard for Double Data Rate (DDR) DRAM has added data clocking on both clock edges and a return clock, allowing higher speed operation with improved system timing margin. SLDRAM builds on the features of SDRAM and DDR with the addition of a packetized address/control protocol, in-system timing and signaling optimization, and full compatibility from generation to generation.

With an evolutionary approach systems designers do not have to discard existing designs and start from scratch. Over more than 20 years encompassing 9 generations of DRAM, the evolutionary solution has always been embraced by the market.

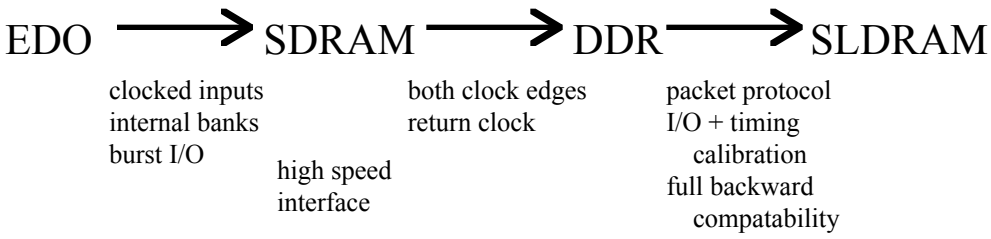

Figure 1. DRAM evolution from EDO to SLDRAM

The SLDRAM protocol will be upward compatible from generation to generation. The command packets include spare bits to accommodate addressing out to the 4G generation. SLDRAM will be sampled first as 64M devices operating at 400Mb/s/pin. With a 16bit wide data interface, these 1st generation SLDRAMs offer a data bandwidth of 800MB/s. Later, 600Mb/s/pin, 800Mb/s/pin, and >1Gb/s/p interfaces will be introduced as they become cost effective. The SLDRAM protocol allows different speed interfaces to be mixed. An 800Mb/s/pin device, for example, can be plugged into a 400Mb/s/pin system and operate correctly at 400Mb/s/pin.

Cost vs. Performance

A new high performance DRAM interface will be adopted once the cost premium over the conventional alternative falls below 5%. This has been the case with SDRAM. High performance is achieved by improving the interface while leaving the DRAM core relatively unchanged. Core performance could easily be improved with reduced bitline, wordline, and datapath loading, but this would lead to significant die area penalty due to increased array fragmentation, grossly exceeding the 5% cost target. SDRAM, DDR, and SLDRAM all make use of a DRAM core having a page mode cycle time of roughly 10ns. To obtain an interface rate higher than the core cycle time, several words must be fetched in parallel. In a x16 200Mb/s/pin DDR device, 2 I/O words must be read or written over a 32 bit internal datapath in one 10ns core cycle.

The need to widen the internal datapath leads to a die cost penalty. The 16M generation 100Mb/s/pin x16 SDRAM has recently fallen below a 5% die cost adder. For 100Mb/s/pin operation the core cycles at the same rate as the interface. 200Mb/s/pin DDR will become cost effective for 64M devices, where there will be a sufficient number of active memory sub-arrays to support a 32 bit internal datapath without substantial area

penalty. SLDRAM devices will first become available with a 400Mb/s/pin x16 interface employing a 64 bit internal datapath. These devices will be cost effective in 256M density. 800Mb/s/pin SLDRAM devices employing 128 bit internal datapaths will be introduced later. SLDRAM or any other memory solution utilizing a 128 bit internal datapath will not be cost effective until the 1Gbit generation.

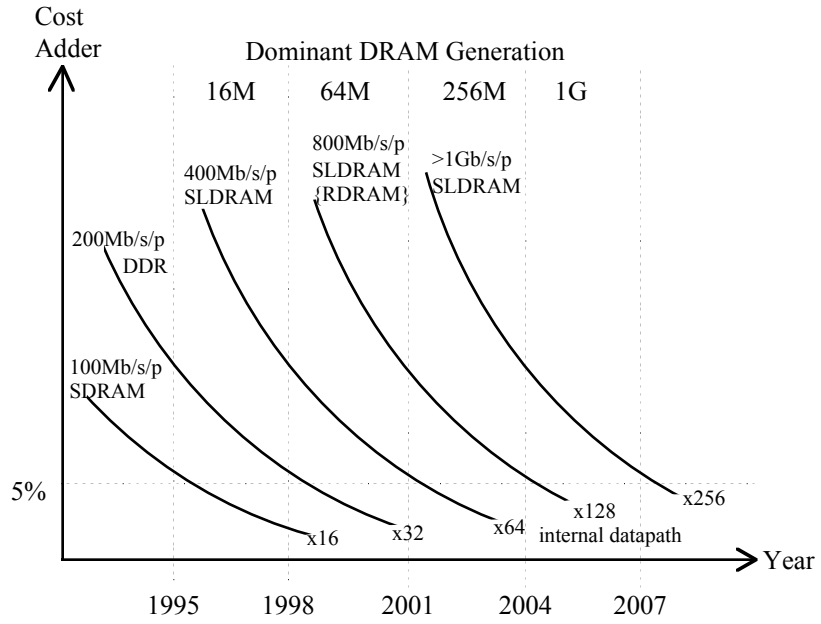

Figure 2. DRAM Bandwidth Cost Penalty for 16 bit I/O

Manufacturability

DRAM's low cost is attributable to high manufacturing yields but the tight timing specifications required for high frequency operation are incompatible with high yield. SLDRAM is defined such that virtually any functional part will meet the specification, in one speed grade or another. The system is able to ascertain on power-up the speed performance capabilities of all SLDRAMs present, and then make the appropriate adjustments. The SLDRAM packet protocol permits the setup and hold time, data delay, and output drive levels of individual memory devices to be adjusted and matched for consistent system operation. Periodically during operation devices are re-calibrated to account for system drift. The flexibility afforded by SLDRAM packet protocol ensures high yield and low cost.

Low System Cost

SLDRAM achieves low system cost through the use of conventional packaging and PCB technology. The SLDRAM devices themselves are packaged in standard 0.5mm pitch TSOP or a 0.8mm staggered pitch vertically surface mounted package (VSMP). With buffered modules an SLDRAM controller is required to support only 33 high speed signals to accommodate Gbyte memory configurations. Wide modules can add additional 16/18 bit data channels without additional control overhead for increased memory bandwidth. Conventional low cost printed circuit board material with 5mil tracks, using 2 of 4 layers for interconnect, is recommended for the SLDRAM interface.

Open Standard

SLDRAM is an open standard to be formalized by IEEE and JEDEC specifications. Open standards allow manufacturers to develop differentiated products to address emerging applications and niche opportunities. Open competition will ensure the continued rapid pace of development of DRAM technology, at the lowest possible cost.

SLDRAM Bus Topology

The SLDRAM bus is a multidrop bus with one memory controller and up to 8 loads. A load can be either a single SLDRAM device or buffered module with many SLDRAM devices. Command, address and control information from the memory controller is sent to the SLDRAMs on the unidirectional CommandLink. Read and write data is exchanged between controller and SLDRAM on the bi-directional DataLink. Both CommandLink and DataLink operate at the same rate, i.e. 400Mb/s/p, 600Mb/s/p, 800Mb/s/p, etc.

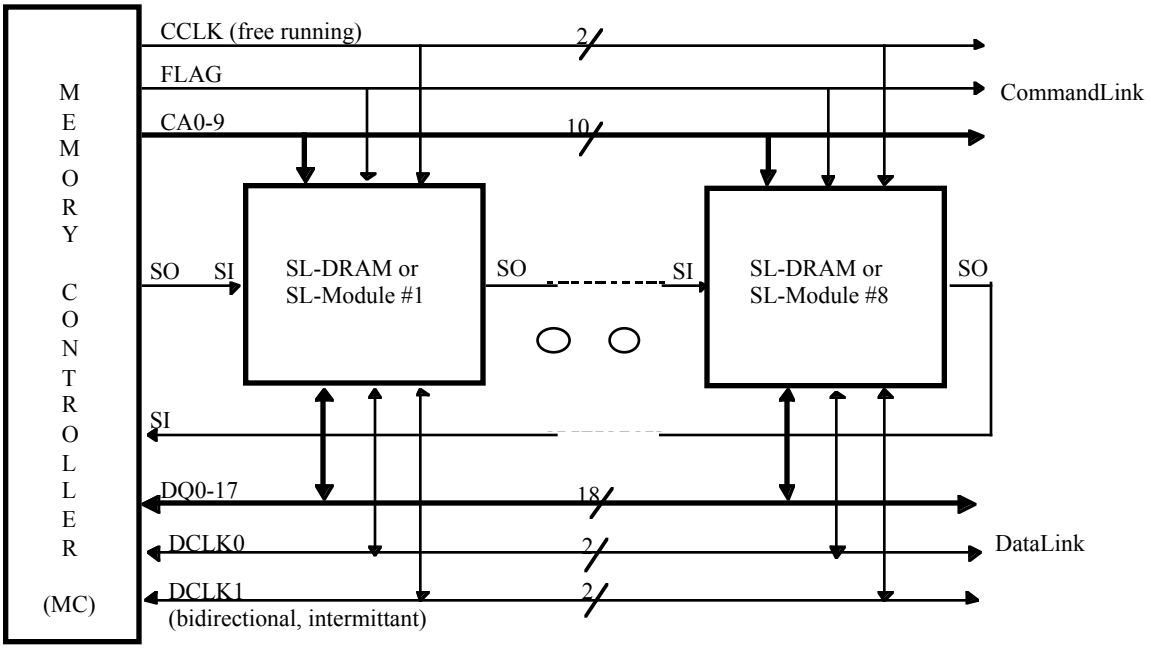

Figure 3. SLDRAM Bus Topology

The CommandLink consists of the signals CCLK, CCLK*, FLAG, CA[9:0], LISTEN, LINKON, and RESET. Commands consist of 4 consecutive 10 bit words on CA[9:0]. The first word of a command is indicated by a '1' on the FLAG bit. Both edges of the differential free running clock CCLK/CCLK* are used by the SLDRAMs to latch command words. While the LISTEN pin is high the SLDRAMs monitor the CommandLink for commands. When LISTEN is low there can be no commands on the CommandLink so SLDRAMs can enter a power saving standby mode. The SLDRAM can exit standby mode and accept a command within 2 clock cycles of LISTEN going high. When LINKON is brought low the SLDRAMs enter shutdown mode in which CCLK can be turned off to achieve zero power on the Link. A RESET signal puts the SLDRAMs in a known state on power-up.

The DataLink is a bi-directional bus for the transmission of write data from controller to SLDRAMs and read data from SLDRAMs back to the controller. It consists of DQ[17:0],

DCLK0, DCLK0*, DCLK1, and DCLK1*. Read and write data packets having a minimum burst length of 4 are accompanied by either differential clock DCLK0/DCLK0* or DCLK1/DCLK1*. The 2 sets of DCLKs allow control of the DataLink to be passed from one device to another with minimum gap.

A daisy chained serial bus with input SI and output SO on each device is used on power up to synchronize the SLDRAMs and assign unique IDs to each one.

	Signal Name	Description	Direction	Level
CommandLink	CCLK/CCLK*	Command Clock	MC --> SLDRAM	SLIO
	CA[9:0]	Command Address Bus	MC --> SLDRAM	SLIO
	LISTEN	Standby Mode	MC --> SLDRAM	SLIO
	LINKON	Shutdown Mode	MC --> SLDRAM	LVCMOS
	RESET*	Hard Reset	MC --> SLDRAM	LVCMOS
DataLink	DCLK0/DCLK0*	Data Clock 0	MC <-> SLDRAM	SLIO
	DCLK1/DCLK1*	Data Clock 1	MC <-> SLDRAM	SLIO
	DQ[17:0]	Data Bus	MC <-> SLDRAM	SLIO
Serial	SI	Serial Input	MC --> SLDRAM, SLDRAM --> SLDRAM	LVCMOS
	SO	Serial Output	SLDRAM --> SLDRAM, SLDRAM --> MC	LVCMOS

Figure 4. SLDRAM Signal Names and Description

Signaling

SLDRAM uses a subset of SSTL_2 for high speed signaling. Output drive levels are more tightly specified to achieve faster bus settling and improved noise margin. The Controller and SLDRAM devices employ a 2.5v Vddq supply to power their output drivers. CommandLink and DataLink signals are single-end terminated to a midpoint reference level of 1.25v at the far end of the bus. Single end termination saves power over double-ended termination. Series stub resistors of approximately 20Ω isolate the SLDRAM module stubs from the main bus. Output drive levels are calibrated on power-up by the memory controller to achieve 0.9v and 1.6v drive levels on the main bus.

Pipelined Transactions

The following timing diagram shows a series of Page Read and Page Write commands issued by the memory controller to the SLDRAMs. For purposes of illustration all burst lengths are 4N although both 4N and 8N bursts can be dynamically mixed. Longer bursts can be achieved by concatenating 4N and/or 8N bursts. The read access time to an open

bank, also known as Page Read Latency, is shown here as 12N (30ns). The first two commands are Page Reads to different banks in SLDRAM #0. The read data appears on the databus along with DCLK0, which provides the memory controller the necessary edges to strobe in read data. Since the first two Page Read commands are for the same SLDRAM, it is not necessary to insert a gap between the two 4N data bursts. The SLDRAM itself ensures that DCLK0 is driven continuously without any glitch.

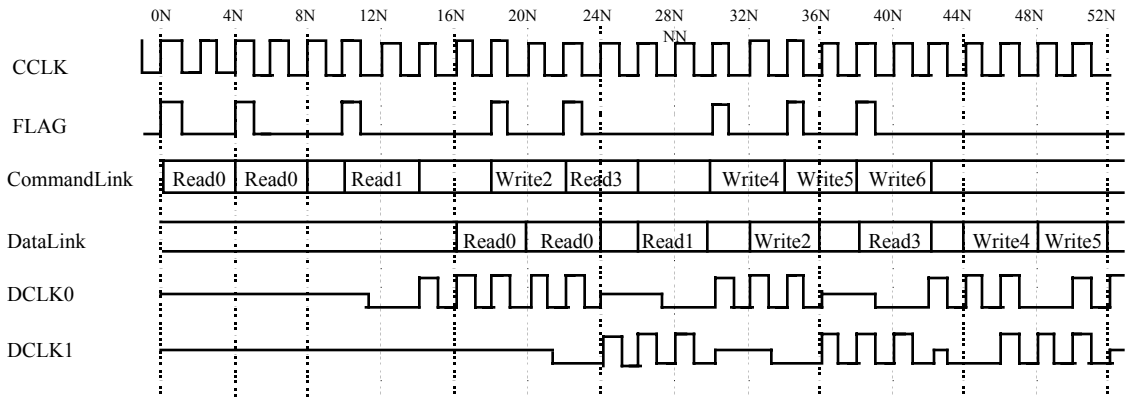

Figure 5. SLDRAM Bus Transactions

The data burst for the following Page Read, to SLDRAM #1, must be separated by a 2N gap to allow settling of the DataLink bus and for timing uncertainty between SLDRAM #0 and SLDRAM #1. A 2N gap is required any time control of the DataLink passes from one device to another, as in reads to different SLDRAMs or read to write and write to read transtitions between SLDRAMs and the memory controller. The memory controller creates the 2N gap between data by inserting a 2N gap between commands. Data for the Read1 command is accompanied by DCLK1, allowing SLDRAM #1 to begin driving the DCLK lines well in advance of the actual data burst.

The next command is a write command using DCLK0 to strobe write data into SLDRAM #2. The Page Write Latency of the SLDRAM is programmed to Page Read Latency - 2N. In order to create a 2N gap between Read1 and Write2 data on the DataLink the Write2 command must be delayed 4N after the Read1 command. Programming write latency in this manner creates an open 4N command slot on the CommandLink, which could be used for non-data commands such as row open or close, register write or refresh, without affecting DataLink utilization. The following read command to SLDRAM #3 does not require any additional delay to achieve the 2N gap on the DataLink.

The final burst of three consecutive write commands shows that the 2N gap between data bursts is not required when writing to different SLDRAM devices. Different DCLKs must be used so that each SLDRAM can identify the start of it’s write data burst. Since all write data originates from the memory controller there will be no glitchs on the DCLKs.

Data Clocks

When control of the DataLink is passed from one device to another the bus will remain at a midpoint level for nominally 2N. This will result in indeterminate logic levels and possibly multiple transitions at the input buffers. This is acceptable for the data lines themselves, but not the data clocks which are used to strobe data. To solve this problem the data clocks have a 00010 preamble before the transition associated with the first bit of data occurs. The device receiving data can enable it's DCLK input buffer anytime during the first 000 period. A dummy transition 10 is included in the preamble to remove pulsewidth dependent skew from the DCLK signal. The receiving device ignores the first rising and falling edge of DCLK and begins clocking data on the 2nd rising edge. Two data clocks are provided so that gapless 4N write bursts to different SLDRAMs and 4N read bursts from different SLDRAMs can be accommodated. The controller indicates in each command packed which DCLK is to be used.

The controller transmits CCLK edges coincident with edges of CA[9:0] and FLAG data. DCLK edges originating from the controller are also coincident with DQ[17:0] data. The SLDRAMs add fractional delay to incoming CCLK and DCLKs to sample commands and write data at the optimum time. SLDRAMs are programmed by the controller to add fractional delay to DCLKs, allowing the controller read data input registers to directly strobe in read data using the received DCLK without the need for any internal delay adjustments.

Timing Adjustment

The controller programs each SLDRAM with 4 timing parameters; Page Read Latency, Page Write Latency, Bank Read Latency, and Bank Write Latency. Latency is defined as the time between the command burst and the start of the associated data burst. For consistent operation of the memory subsystem, it is recommended that each SLDRAM should be programmed to obtain equal values for the 4 different latencies viewed from the controller pins. Due to different round-trip bus delay through SLDRAMs at different bus positions, the presence or absence of buffers, and varying performance of individual SLDRAM devices, the actual values programmed in each SLDRAM may differ considerably. On power up reset the SLDRAM latency registers are set to their minimum values. The controller can observe the response of each SLDRAM and then make appropriate adjustments to obtain consistant operation.

Read latency is adjustable in coarse increments of unit bit intervals and fine increments of fractional bit intervals. The controller will program the coarse and fine read latency of each SLDRAM so that read data bursts from different devices, at different electrical distances from the controller, all arrive back at the controller with equal delay from the command packet.

Write latency is only adjustable in coarse increments. The write latency values determine when the SLDRAM begins looking for transitions on DCLK in order to strobe in write data.

SLDRAM Command Format

The SLDRAM command packet is organized as four 10bit words. The following command packet is for a 64M SLDRAM with 8 banks, 1024 row addresses, and 128 column addresses. There are 3 bits for bank address, BNK[2:0], 10 bits for row address, ROW[9:0], and 7 bits for column address, COL[6:0]. Many other organizations and densities can be accommodated within the same 40 bits. On power up the memory controller polls the SLDRAMs to determine how many banks, rows and columns each device has. The controller can then include the appropriate number of address bits in the command packet for each individual SLDRAM.

FLAG	CA9	CA8	CA7	CA6	CA5	CA4	CA3	CA2	CA1	CA0
1	ID8	ID7	ID6	ID5	ID4	ID3	ID2	ID1	ID0	CMD5
0	CMD4	CMD3	CMD2	CMD1	CMD0	BNK2	BNK1	BNK0	ROW9	ROW8
0	ROW7	ROW6	ROW5	ROW4	ROW3	ROW2	ROW1	ROW0	0	0
0	0	0	0	COL6	COL5	COL4	COL3	COL2	COL1	COL0

Figure 6. SLDRAM Command Format

Chip ID and Multicasting

The first word of the command packet contains the chip ID bits. An SLDRAM will ignore any command that does not match the local ID. Chip ID is assigned by the controller on power-up using the SI/SO signals. This allows the controller to uniquely address every SLDRAM in the system without the need for separate chip enable signals or glue logic. The first 8 bits of chip ID allow addressing of up to 256 SLDRAMs on a single hierarchical DataLink. The 9th bit of the Chip ID field is used for multicasting, which permits any group of 2, 4, 8, 16, 32, 64, 128 or all 256 SLDRAMs to be addressed with a single command. This is useful for initialization, refresh, and multiple DataLink configurations.

SLDRAM Commands

The command field consists of 6 bits as shown below. When the most significant bit CMD5 = 0, normal read or write commands are executed. The selection of Page or Bank access, Burst Length, Read or Write, Autoprecharge, and DCLK are independent. Page access requires an open bank. When CMD5=1, row operations, register accesses, events, or special synchronization commands are selected.

CMD5	CMD4	CMD3	CMD2	CMD1	CMD0	Command
0	0	x	x	x	x	Page Access
0	1	x	x	x	x	Bank Access
0	x	0	x	x	x	Burst Length = 4
0	x	1	x	x	x	Burst Length =8
0	x	x	0	x	x	Read
0	x	x	1	x	x	Write
0	x	x	x	0	x	Leave Page Open
0	x	x	x	1	x	Autoprecharge
0	x	x	x	x	0	Use DCLK0
0	x	x	x	x	1	Use DCLK1
1	x	x	x	x	x	Special Commands

Figure 7. SLDRAM Command Format

SLDRAM Calibration

System level calibration of individual SLDRAM timing and output drive levels is key the high manufacturing yield and low cost of SLDRAM components. Individual devices are not required to meet tight AC and DC parametric specifications. Rather, these are calibrated at the system level to compensate for wide variation in individual device parameters.

Initialization and Synchronization

When the SLDRAM memory subsystem is powered up, the controller must take the following steps before normal memory operations can begin;

1. Reset - The RESET* pin on each SLDRAM is held low. This clears the SLDRAM's internal synchronization indication, programs internal device ID=255, sets SO=0, and sets all read and write latency values to their minimum values.

2. Command and Write Data Synchronization - CCLK and DCLK transitions from the controller are coincident with edges of the data on CA[9:0] and DQ[17:0]. Each SLDRAM must internally adjust the phase of the input clocks to correctly sample

incoming data. To begin this process, the controller begins transmitting continuous CCLK and DCLKs transitions, and sets SO=1. On DQ[17:0], CA[9:0], and FLAG the controller transmits inverted and non-inverted versions of the 15 bit repeating psuedo-random SYNC sequence “111101011001000”. The SLDRAMs recognize this sequence by 2 consecutive 1’s on the FLAG bit. The SLDRAMs then determine an optimum internal delay for CCLK and both DCLKs to optimally sample the known pattern. Once the SLDRAM has synchronized and the appropriate delays for CCLK, DCLK0 and DCLK1 have been programmed, the SLDRAM will set SO=SI. The controller stops sending the SYNC pattern when SI=1. It then resets SO=0 which ripples through all SLDRAMs.

3. ID Assignment - The controller sets SO=1 once again and sends an ID Register Write command with data 0. Only the SLDRAM with SI=1 and ID=255 will respond to this command. This SLDRAM overwrites it’s ID Register with the value 0 and sets SO=1. The controller repeats the “ID Register Write” with write data 1 and so on until it observes a high logic level on SI.

4. Voh, Vol Calibration - The controller calibrates SLDRAM I/O levels by instructing each SLDRAM to drive the DataLink with DC ‘0’ and ‘1’ levels, and then issuing appropriate Increment and Decrement Voh/Vol commands until the levels match the controller’s own reference.

5. Read Synchronization - The controller instructs each SLDRAM to transmit the 15 bit psuedo-random synchronization pattern on DQ[17:0] and continuous transitions on DCLK. The controller adjusts the position of the received bits to achieve the desired bit alignment by incrementing the SLDRAM read data vernier.

6. Read Latency Calibration - The controller issues a read command to each SLDRAM and then monitors the specified DCLK for the transitions which accompany the data burst to measure read latency. Once the minimum latencies have been measured for each SLDRAM an appropriate read latency value for the memory system (typically the minimum latency of the slowest device) can be programmed into each device by writing to the Read Delay Registers.

7. Write Latency Programming - The controller issues to each SLDRAM a write command with zero write latency and an extended 32N extended burst of data. The data contained in each successive word of the write burst is incremented 0, 1, 2, 3, ... 31. After the write command is completed the controller reads the location addressed by the write command. The first word contains the minimum write latency. With this information the controller can set the write latencies of all SLDRAMs to a common value by writing to the Write Delay Registers. The memory subsystem is now fully calibrated and ready for normal operation.

Recalibration During Operation

During normal operation calibrated timing and voltage level parameters may drift due to changes in temperature and supply voltage. SLDRAMs must be recalibrated periodically to ensure robust operation. Recalibration operations can be hidden during Auto Refresh and Self Refresh periods so as not to affect system performance.

Buffered Modules

The SLDRAM protocol allows the addressing of hierarchical memory subsystems having more devices than a single electrical interface can support. Buffered modules permit deeper and/or wider memory configurations than the basic unbuffered configuration discussed earlier. Buffered and unbuffered modules can be mixed in a system. The initialization and synchronization sequence will accommodate the additional delays added by the buffers.

Buffered CommandLink

A wide module with multiple DataLinks providing higher memory bandwidth is based on a buffered CommandLink. CommandLink overhead, including pins on the controller, motherboard tracks and termination networks, bus termination power dissipation, and module connectors, need not be duplicated for each DataLink.

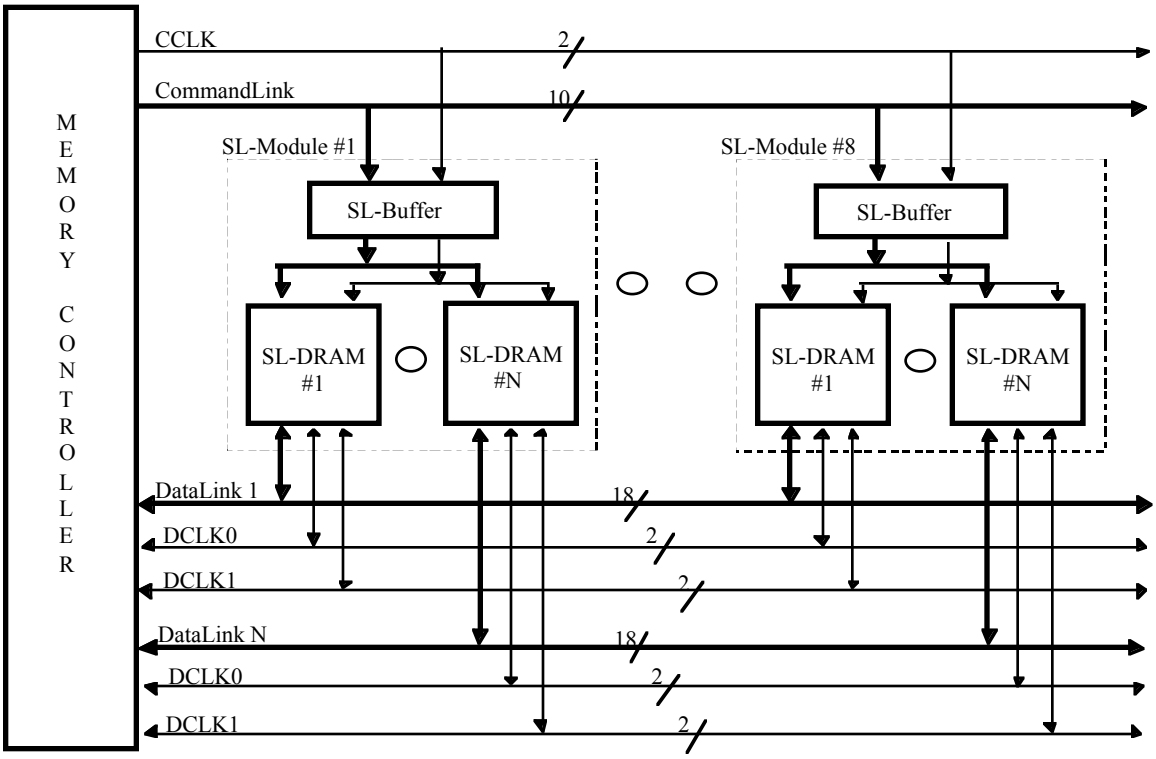

Buffered CommandLink and DataLink

This configuration allows a controller with a single SLDRAM interface to address a larger memory depth than 8 individual unbuffered SLDRAM devices would allow. The CommandLink buffer and the DataLink buffer may be combined into a single chip.