Engineer To Engineer Note EE-126

Technical Notes on using Analog Devices' DSP components and development tools
Phone: (800) ANALOG-D, FAX: (781) 461-3010, EMAIL: dsp.support@analog.com, FTP: ftp.analog.com, WEB: www.analog.com/dsp

Contributed by Robert Hoffmann, European DSP Applications *Rev1 (20-March-02)*

The ABC of SDRAMemory

As new algorithms for digital signal processing are developed, the memory requirements for these applications will continue to grow. Not only while these applications require more memory, but they will require increased efficiency for data storage and retrieval.

The ADSP-21065L, ADSP-21161N and ADSP-TS101S from the floating-point family of DSPs from Analog Devices Inc. have been designed with an on-chip SDRAM interface allowing applications to gluelessly incorporate less expensive SDRAM devices into designs. Also, some new members from the fixed-point family, ADSP-21532 and ADSP-21535, will have an SDRAM interface.

Introduction

This application note will demonstrate the complexity of the SDRAM technology. It will illustrate, that an SDRAM is not "just a memory".

The first part shows the basic internal DRAM circuits with their timing specifications. In the second part, the SDRAM architecture units with their features are discussed. Furthermore, the timing specs and set of commands are illustrated with the help of state diagrams. The last part deals with the different access modes, burst performance and controller's address mapping schemes. Moreover some SDRAM standards are introduced.

1 – DRAM Technology

Two common storage options are static random access memory (SRAM) and dynamic random access memory (DRAM). The functional difference between an SRAM and DRAM is how the device stores its data.

In an SRAM data is stored in up to 6 transistors, which hold their value until you overwrite them with new data. The DRAM stores data in the capacitors, which gradually loose their charge and, without refreshing, will loose data.

The Synchronous DRAM-Technology, originally offered for main storage use in personal computers, is not completely new, but is an advanced development from the DRAM-Technology. The interface works in a synchronous manner, which makes the hardware requirements easier to fulfill.

1.1 – The storage cell

As figure 1 points out, the binary information is stored in a unit consisting of a transistor and a very small capacitor of about 20-40 fF (Femtofarad, 0.020-0.040 pF) for each cell. A charged capacitor has a logical 1, a discharged capacitor a logical 0.

Additionally, the figure shows an example with a 1bit I/O structure. Typical structures are 4, 8, 16 and 32 bit. The 1 bit architecture has disappeared, because of the high density memory requests from the market. Structures of 4, 8, 16 or 32 bit require less hardware intensive solutions.

1.2 – The surrounding circuits

The capacitor storage cell needs surrounding circuits such as precharge circuits, sense amplifiers, I/O gates, word and bit lines.

Idle State

Just by powering up the device brings the memory into an undefined state. Then, a command is required to bring the bank in Idle mode. The memory is now in a defined state and can be accessed properly.

Row Activation

The row address decoder (figure 1) starts accessing the precharge circuit with the word line, when the ~RAS line is asserted. Both inputs (positive and negative bit line) of the sense amp (opamp) are precharged with VDD/2. This sequence is current intensive and requires some time. In the meantime, the row's sense amp starts gating. Both inputs of the sense amp are now the same voltage VDD/2. The word line switch connects the storage cell to the positive bit line. Depending on the capacitor's cell charge, the potential increases (VDD/2 +ΔV) or decreases (VDD/2 -ΔV) on the positive bit line.

Because the capacitance of 40 fF (0.040 pF) is much smaller than the transistor's and the bit line's capacitance, the voltage difference is typically only about ΔV=100 mV. The small sensed voltage is amplified to level VDD or 0. The sense amp acts as a latch to store the sensed value. After the sensing has finished (spec *tRCD*), the device is ready for read or writes operation.

Note: The advantage of the precharge technique is, that only the difference between the positive and negative bit lines must be amplified, not the absolute value of 100 mV, thus increasing reliability.

Figure 1: Basic DRAM storage cell

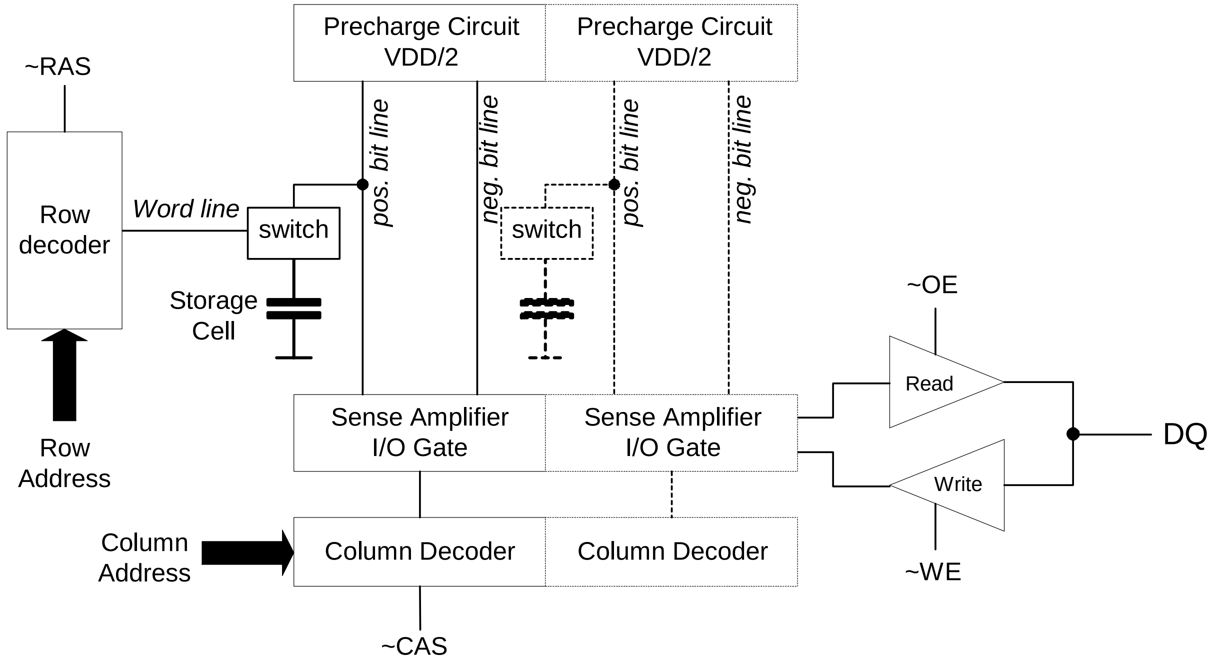

Column Write

After the spec *tRCD* is satisfied, the assertion of ~CAS causes the column decoder to select the dedicated sense amp. In parallel, the ~WE enables the write drivers to feed the data directly into the sense amp.

Column Read

Basically, there is a difference between read and write operations. After the spec *tRCD* is satisfied, the assertion of ~CAS causes the column decoder to select the dedicated sense amp. In parallel, the ~OE starts the read latch to drive the data from the sense amp to the output.

Note: The read operation (unlike write) discharges the capacitor. In order to restore the information in the cell, additional logic performs a precharge sequence.

Row Precharge

If the next write or read access falls in another row, the current row (page) must be closed or “precharged”. Binary zeros or ones stored in the sense amp during the row activation will rewrite the storage cell during precharge (spec *tRP*).

Note: After precharge, the row returns in idle state.

Note: Don’t mix up the precharge circuit with precharge command.

Row Refresh

The refresh is simply a sequence based on activation followed by a precharge (spec *tRC=tRAS+tRP*) but with disabled I/O buffer. The sense amp reads the storage cell during the *tRAS* time, immediately followed by a precharge *tRP* to rewrite the cell.

Note: Refresh must occur periodically for each row after the specified time tREF

1.3 – Timing Issues

The capacitor cell is accessed in a multiplexed manner (figure 2):

The RAS line is asserted through the activate command; all bit lines are now biased to VDD/2. In parallel, all the row’s sense amps (depending on page size) are gated. Finally, the values are stored, requiring the time *tRCD* (RAS to CAS delay).

Now, any column can be opened by a read or write command (CAS line asserted). If you write to the cell, the write command and the data will be sampled in the same clock cycle. The next read access falls in a different row, causing a precharge sequence within the following time frame *tRP* (precharge period).

Figure 2: DRAM's level-controlled Write and Read

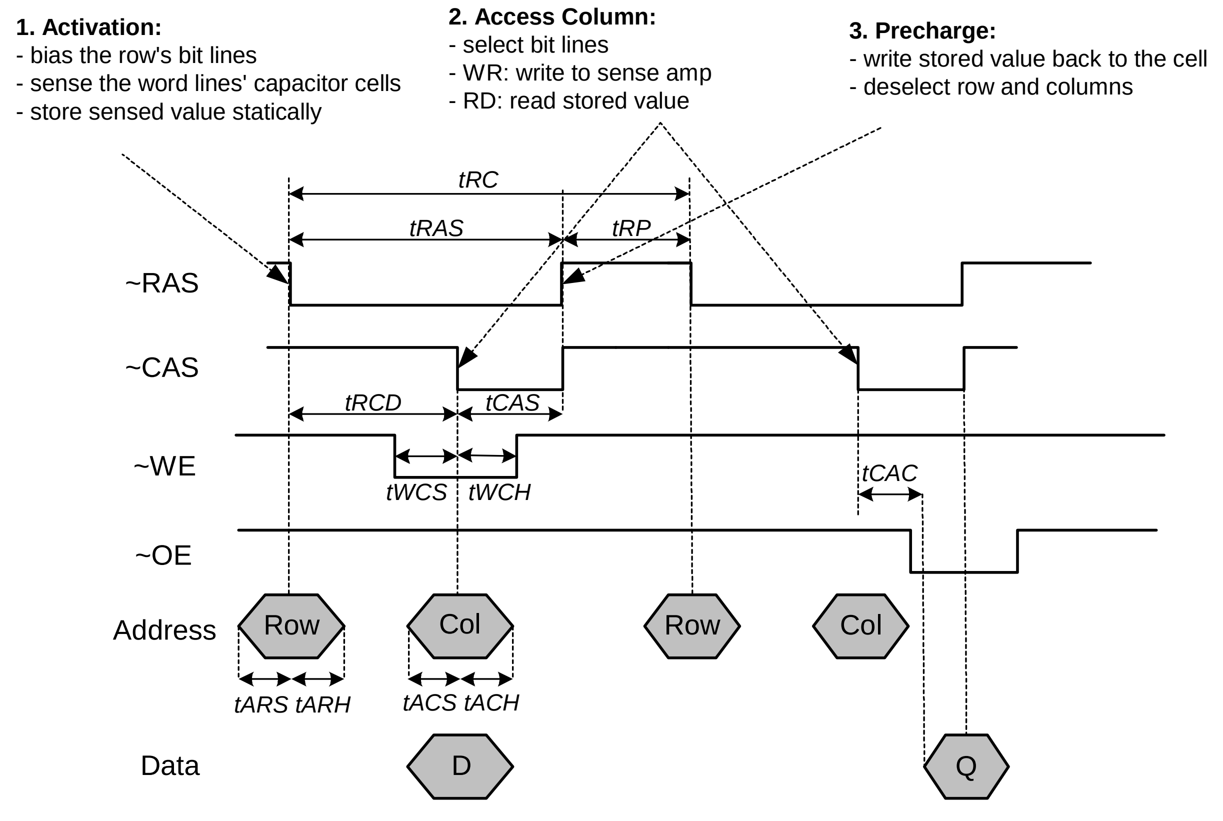

1. Activation:
- bias the row's bit lines
- sense the word lines' capacitor cells
- store sensed value statically

2. Access Column:
- select bit lines
- WR: write to sense amp
- RD: read stored value

3. Precharge:
- write stored value back to the cell
- deselect row and columns

Note: DRAM accesses are multiplexed, first row address followed by column address.

1.4 – Refresh Issues

The DRAM must refresh the row each time the spec *tREF* is elapsed. The row refresh pattern is free until the time *tREF* is satisfied for each row. 3 different refresh modes are available:
Note: The Refresh uses internal read during tRAS and write during tRP.
Note: The Refresh row cycle tRC=tRAS+tRP.

RAS only Refresh

The external row address during the falling edge of the ~RAS pin starts a refresh each time it is required.
Note: The RAS only refresh requires an external address counter.

Hidden Refresh

This mode is similar to the CBR refresh. The external row address (falling edge of ~RAS) and column address (falling edge of the ~CAS) starts an internal hidden refresh using the internal refresh counter, each time it is required.
Note: The Hidden Refresh can only be used for continuous access of DRAM.

CAS before RAS Refresh

The CBR- or auto refresh is started by deassertion of ~CAS followed by the deassertion of ~RAS, that means in reversed order. Hereby, the device requires no external address to full fill a refresh. The

internal refresh counter will handle this job. The time gap between refreshing two successive rows in a classical DRAM is 15,625µs. The refresh period adds up to Rows/*tREF* in spec terms. In this particular mode, the data transactions are periodically interrupted by auto refresh commands.
Note: The CBR refresh is comfortable and reduces the power dissipation.

Figure 3: The Capacitors Refresh

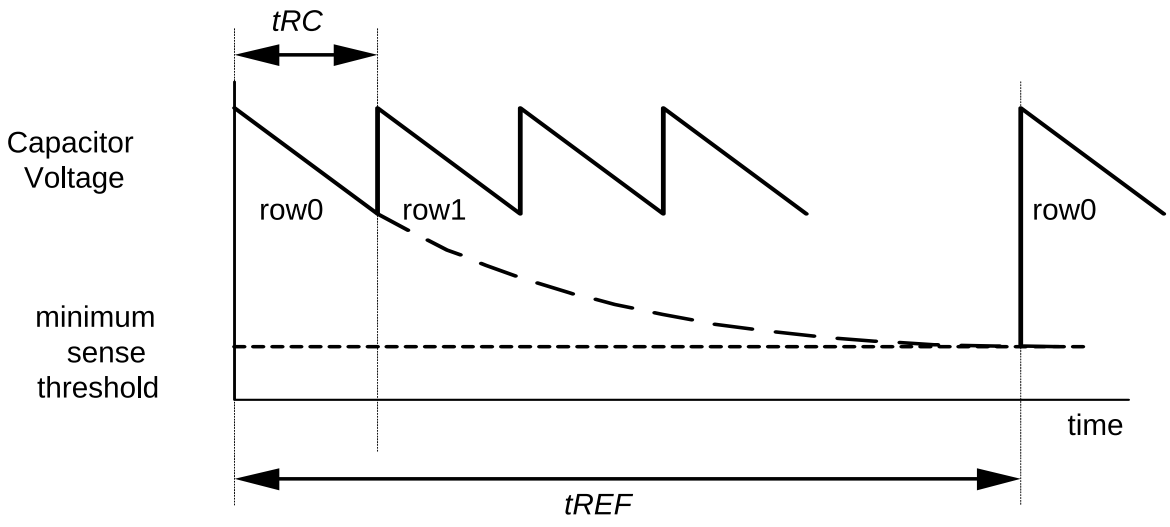

1.5 – SRAM vs. DRAM

SRAMs are generally simple from a hardware and software perspective. Every read or write instruction is a single access, and wait states can be programmed to access slower memories if desired. The disadvantage of SRAMs is that large memories and fast memories, for system that desire zero wait states, are expensive. DRAMs have the advantage of address multiplexing, thus needing less address lines. Additionally, they are available in larger capacities than SRAMs because of the high-density cell. The main disadvantage is the need for refresh and precharge operations.

2 – Architecture SDRAM

As the speed of processors continues to increase, the speed of standard DRAM's becomes inadequate. In order to improve the overall system performance, the operations have to be synchronized with the system clock cycles. Toshiba's Tecra 700 was the first computer to use SDRAM for main memory, and Kingston Technology has supported the Tecra since its initial release in November 1995. Figure 4 demonstrates the simplified pipelined architecture of an SDRAM.
When synchronous memories use a pipelined architecture (registers for in- and output signals) they produce additional performance gains. In a pipelined device, the internal memory array needs only to present its data to an internal register to be latched rather than pushing the data off the chip to the rest of the system. Because the array only sees the internal delays, it presents data to the latch faster than it would if it had to drive off chip. Further, once the latch captures the array's data, the array can start preparing for the next memory cycle while the latch drives the rest of the system.

Figure 4: Simplified Pipelined Architecture of a 4M x 4bit x 4banks

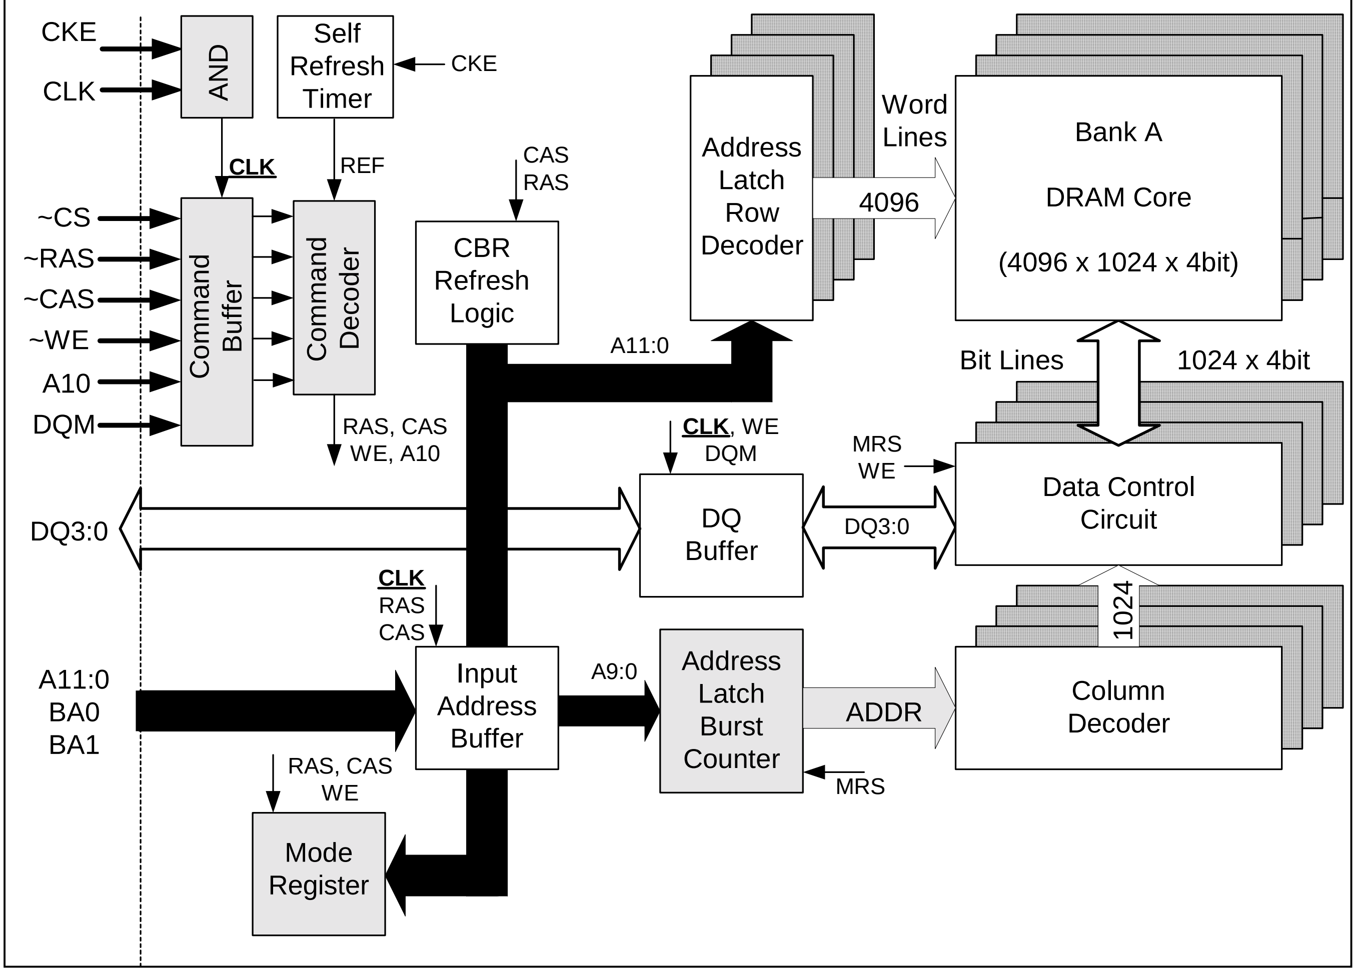

2.1 – Command Units

Relevant units: Control buffer, Command decoder, Mode register

Command buffer

All input control signals are sampled with the positive edge of the CLK, making the timing requirements (setup- and hold times) much easier to meet. The CKE pin is used to enable the CLK operation of the SDRAM.

Note: The pulsed external SDRAM timing uses internal DRAM timing.

Command Decoder

This unit is the heart of the memory device: The inputs trigger a state machine, which is part of the command logic. During the rising CLK edge, the command logic decodes the lines ~RAS, ~CAS, ~WE, ~CS, A10 and executes the command.

Note: The command decoder is enabled with ~CS low

Mode Register

The mode register stores the data for controlling the various operation modes of SDRAM. The current mode is determined by the address lines values.

During mode register setup, all the addresses and bank select pins are used to configure the SDRAM.

2.2 - Data Units

Relevant units: address buffer, address latches for row- and column, decoder for row- and column, refresh logic, burst counter, data control circuit and DQ buffer

Address Buffer

The input address buffer latches the current address of the specific command. The RAS and CAS strobes from the command decoder indicate whether the row or the column address latch is selected. The buffer is used for address pipelining, which means that during reads more than one address (depending from read latency) can be latched until data is available.
Note: The address pipeline is an important performance benefit vs. asynchronous memories.

Address Decoder

The row decoder drives the selected word lines of the array. To access i.e. 4096 rows, you need 12 address lines. The column decoder drives the selected bit lines; its length represents the page size. Typical I/O-structures are:

4 bit => 4096 words page size
4 bit => 2048 words page size
4 bit => 1024 words page size
8 bit => 512 words page size
16 bit => 256 words page size
32 bit => 256 words page size

Note: The bigger the I/O-structure the smaller the page size.

Decoding 1024 words takes 10 address lines. The matrix is called a memory array or memory bank. The matrix size is 4096 x 1024 x 4bit = 4M x 4bit each bank. You can find 2 or 4 independent banks; this value depends typical on the SDRAM size:

- 16 Mbit => 2 banks
- >16 Mbit => 4 banks

Refresh Logic

SDRAMs use the CBR refresh to benefit from the internal refresh counter. All rows must be refreshed during the specified maximum refresh time *tREF* in order to avoid data loss. The refresh counter starts addressing the rows in all banks simultaneously each time requests arrive from the external controller or the internal timer (self refresh) by asserting CAS before RAS line. The pointer increments automatically to the next address after each refresh and wraps around after a full period is over.

Note: The auto refresh (CBR refresh) is the refresh mode for SDRAM used in standard data transactions.

List for refresh values:

Size	*Row*	*tREF*	*Refresh Rate*	*tRC/Row*
16Mbit	2k	32 ms	64 kHz	15,625 µs
64Mbit	2k	32 ms	64 kHz	15,625 µs
64Mbit	4k	64 ms	64 kHz	15,625 µs
128Mbit	4k	64 ms	64 kHz	15,625 µs
256Mbit	8k	64 ms	128 kHz	7,812 µs
512Mbit	8k	64 ms	128 kHz	7,812 µs

DQ Buffer

The DQs buffer register the data on the rising edge of clock. The DQM pin (mask pin) controls the data buffer. In read mode, DQM controls the output buffers like the conventional ~OE pin on DRAMs. DQM=high and DQM=low switch the output buffer off and on.

In write mode, ~WE is asserted, DQM controls the word mask. Input data is written to the cell if DQM is low but not if DQM is high.

The fixed DQM latency is:

- 2 clock cycles for reads
- no latency for writes

Vendors offer independent DQM[x] pins depending on the I/O structure. It’s featured to control the data nibble- or byte wise to allow for instance byte write accesses. If not desired, the DQM[x] pins must be interconnected.

Note: The SDRAM controller controls the DQ buffer by the state of DQM pin.

I/O size	*number of DQMs*	*masked word size*
4 bit	1	1 nibble
8 bit	2	1 nibble
16 bit	2	1 byte
32 bit	4	1 byte

Additional, masking is used to block SDRAM’s data buffer during precharge, while invalid data may be written at the same clock cycle as the precharge command. To prevent this from happening, the DQM pin is tied high at the same clock as the precharge, this blocks the data of the burst operation.

Moreover, masking during read to write transitions is useful to avoid data contention caused by different latencies.

Note: The DQM pin is used to optimize read to write transitions.

Burst Counter

Basically, the burst counter acts as column address latch. The number of columns defines the SDRAM's page size. For instance, the burst counter provided with 10 multiplexed address lines can address up to 1024 locations in the same row. It is used to improve the reliability during high-speed transfer over the PCB. Only the start address is driven to the SDRAM. Next sequential- or interleaved addresses are incremented by the SDRAM's burst counter. Therefore, no activities are registered on the address bus during bursting. This improves the reliability and reduces the power consumption.
Note: The CLK signal is used to increment the burst counter during the bursts. The SDRAM ignores all next addresses while the input address buffer is not sampled during the burst.
Note: Bursting in conjunction with address pipeline during reads optimize the throughput enormously.

Data Control Circuit

This unit is based on sense amplifiers, I/O gating like read latches and write driver. The WE is used to select between read and write gates. Moreover, the Read Latency takes affect here.

2.3 – Memory Unit

Relevant units: Precharge circuits, memory core

Precharge Circuits

These circuits are used to equalize the sense amps with VDD/2 during row's activation.

DRAM Core

The memory array with its storage cell is addressed over the bit- and the word lines. The use of multiple arrays guarantees better throughput during off page accesses. While one bank precharges, the other bank drives data.

2.4 - Additional Units

Relevant units: Power-down, Suspend, Self refresh

Self Refresh Timer

The Self-refresh mode can be used to retain data in the SDRAM, even if the rest of the system is powered down. An internal timer meets all the refresh requirements. This feature provides the capability to reduce power consumption during refresh operations.
Note: In self-refresh mode, the address buffer and command interface are disabled.

Power-Down Mode

In order to reduce standby power consumption, a power down mode is available.

Suspend Mode

This mode freezes the internal clock and extends data read and write operations.

3 - Timing Issues

3.1 – AC Parameters

Asynchronous memories depend on properly timed and shaped pulses on their control lines. With total cycle times approaching 10 ns (for 100-MHz systems), the pulse shape becomes increasingly intolerant of error, and therefore, harder to design.

Synchronous memories avoid the need for critical pulse shapes, depending only on the placement of clock edges relative to the other data, often using the same clock as the rest of the system.

SDRAM: All signals are now sampled in the command decoder with the rising clock edge (figure 5). This simplifies the timings enormously, as the timing is pulsed and the commands are normalized to the system clock in order to fulfill the demands. The design must tolerate the setup and hold times, which are very short related to the clock. Therefore, all asynchronous timings are self timed or done internally with the effect of an increased reliability.

Figure 5: SDRAM's pulse-controlled Write and Read

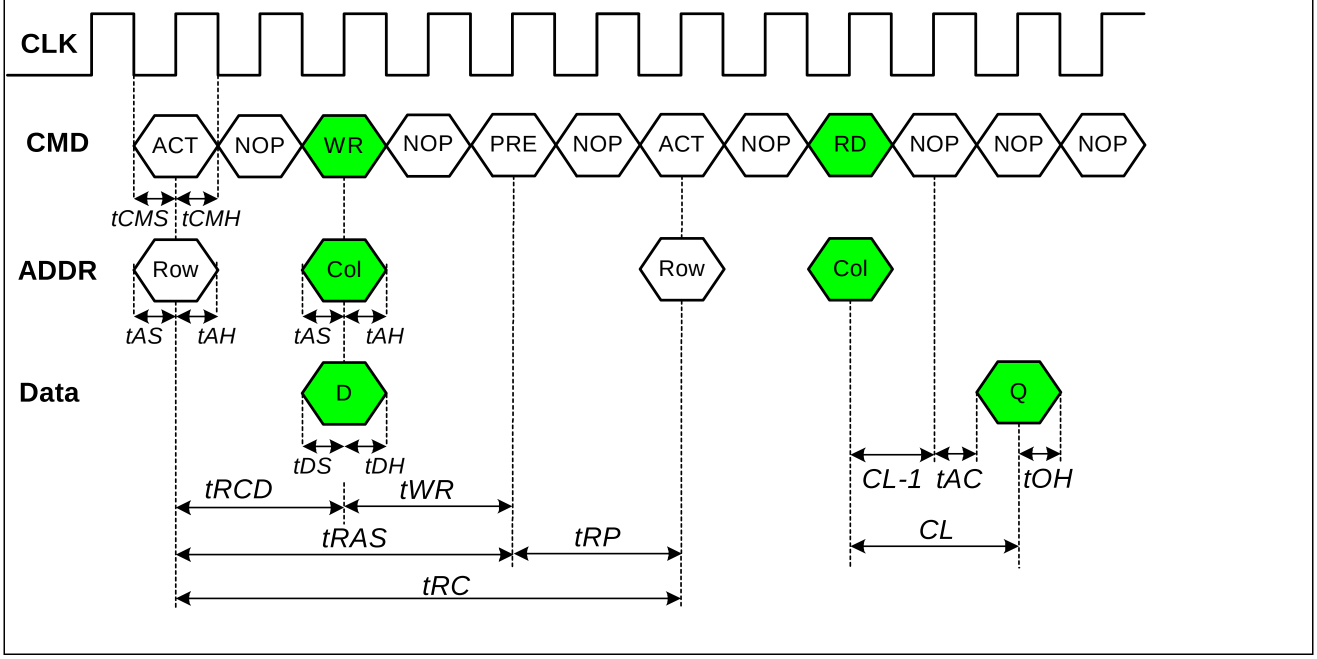

Next list of timing specs are the most important in SDRAM designs. The vendor's datasheets use generally these abbreviations:

Timing spec	*Description*	*Note*
tCK	clock cycle time	reference
CL	Read Latency	normalized to tCK
tRAS	activate period	internal implicit refresh read

tRP	precharge period	internal implicit refresh write
tRCD	RAS to CAS delay	first read or write
tWR (tDPL)	write recovery	tRAS=tRCD+tWR
tRC	activate A to activate A	tRC=tRAS+tRC (row refresh)
tRRD	activate A to activate B	tRRD=1/2*tRC (banking)
tCCD	CAS to CAS delay	1 cycle for pipeline architecture
tMRD (tRSC)	mode register to command	
tXSR	self refresh to activate	
tAC	access time from clock edge	
tREF	row refresh period	tREF=15.625µs or 7.81µs *rows

3.2 – DC Parameters

Most of the SDRAMs are designed LVTTL-compatible for input- and output levels. The vendor's datasheets use generally these abbreviations:

Timing spec	***Value***	***Description***
VCC	3,3V±0,3V	Voltage supply for DQ buffer
VSS		Ground for core
VCCQ	3,3V±0,3V	Voltage supply for DQ buffer
VSSQ		Ground for DQ buffers
VIL	max. 0,8V	Input low voltage
VIH	min. 2,0V	Input high voltage
VOL	max. 0,4V	Output low voltage
VOH	min. 2,4V	Output high voltage
TA	-40° C to 85° C	Extended operating ambient temperature
CAx	typ. 4pF	Input capacitance for address and control
CIx	typ. 6pF	Input capacitance for data

Note: The separate VDDQ VSSQ pins are used to improve the SNR of the device.

3.3 - DRAM vs. SDRAM

The next table summarizes most differences:

DRAM	*SDRAM*
No system clock	Runs off system clock
Level control	Pulsed clock control, internally self timed
No address pipeline	Address pipeline
1 bank operation	2 or 4 banks for on-chip interleaving
1 transfer per column access	burst of 1, 2, 4, 8 or full page transfer per column
No programmable Read latency	programmable Read latency for different speed

3.4 – SDRAM's Evaluation

Type	*Description*	*Performance 4 words*
DRAM	asynchronous	5-5-5-5
FPM DRAM	fast page mode	5-3-3-3
EDO DRAM	extended data out	5-2-2-2
BEDO DRAM	burst extended data out	5-1-1-1
SDRAM	synchronous	5-1-1-1

4 – Memory Organization

The JEDEC standard offers sizes depending on the word- and bit lines length (memory core). Moreover, the multiple bank support in relationship to the I/O structure shows the vendor's program:

16 Mbits / 2 MBytes – 2 banks

Type	I/O	Row x Page size	RAS	CAS
1M x 16	16bit	(2048 x 256)	A[10:0]	A[7:0]
2M x 8	8bit	(2048 x 512)	A[10:0]	A[8:0]
4M x 4	4bit	(2048 x 1024)	A[10:0]	A[9:0]

64 Mbits / 4 MBytes – 4 banks

Type	I/O	Row x Page size	RAS	CAS
2M x 32	32bit	(2048 x 256)	A[10:0]	A[7:0]
4M x 16	16bit	(4096 x 256)	A[11:0]	A[7:0]
8M x 8	8bit	(4096 x 512)	A[11:0]	A[8:0]
16M x 4	4bit	(4096 x 1024)	A[11:0]	A[9:0]

128 Mbits / 8 MBytes – 4 banks

Type	I/O	Row x Page size	RAS	CAS
4M x 32	32bit	(4096 x 256)	A[11:0]	A[7:0]
8M x 16	16bit	(4096 x 512)	A[11:0]	A[8:0]
16M x 8	8bit	(4096 x 1024)	A[11:0]	A[9:0]
32M x 4	4bit	(4096 x 2048)	A[11:0]	A[9:0], A[11]

256 Mbits / 16 MBytes – 4 banks

Type	I/O	Row x Page size	RAS	CAS
8M x 32	32bit	(8192 x 256)	A[12:0]	A[7:0]
16M x 16	16bit	(8192 x 512)	A[12:0]	A[8:0]
32M x 8	8bit	(8192 x 1024)	A[12:0]	A[9:0]
64M x 4	4bit	(8192 x 2048)	A[12:0]	A[9:0], A[11]

512 Mbits / 32 MBytes – 4 banks

Type	I/O	Row x Page size	RAS	CAS
16M x 32	32bit	TBD		
32M x 16	16bit	(8192 x 1024)	A[12:0]	A[9:0]
64M x 8	8bit	(8192 x 2048)	A[12:0]	A[9:0], A[11]
128M x 4	4bit	(8192 x 4096)	A[12:0]	A[9:0], A[11:12]

Note: For column accesses, address 10 acts as command pin. For page sizes ≥ 2048 words, the next valid address is pin 11.

Note: The x4, x8, x16 parts are available in 54-pin TSOP, x32 is available in 86-pin TSOP.

5 – Initialization and Power up Mode

The SDRAM internal condition after power-up is undefined. It is required to follow this sequence:

5.1 – Hardware Initialization

Required steps:

- 1. Apply power and start clock, attempt to deselect the command decoder with ~CS high
- 2. Maintain power stable (VDD and VDDQ simultaneously), stable clock (CLK) for a minimum of 200 µs

5.2 – Software Power up mode

Required steps:

- 3. Precharge all banks to run into defined Idle state (PREA)
- 4. Assert 8 auto refresh cycles (REF)
- 5. Program the mode register (MRS)
- 6. SDRAM in Idle mode, ready for normal operation (ACT)

Note: Some vendors swap step 4 and 5 in the power-up sequence. For other vendors, the order of steps 4 and 5 are meaningless.

Note: During power-up mode, 8 refresh cycles are used charging internal nodes.

6 – Commands

6.1 – Pulse-Controlled Commands

There are 4 pulsed signals (figure 6): ~CS, ~RAS, ~CAS and ~WE to issue up to 8 commands during the rising clock edge of the SDRAM. Using the A10 signal, the number of commands increases to 12.

6.2 – Level-Controlled Commands

The level-controlled CKE pin is defined as a clock-enable signal and it is also responsible for putting the SDRAM into low power state.

Figure 6: SDRAM Commands

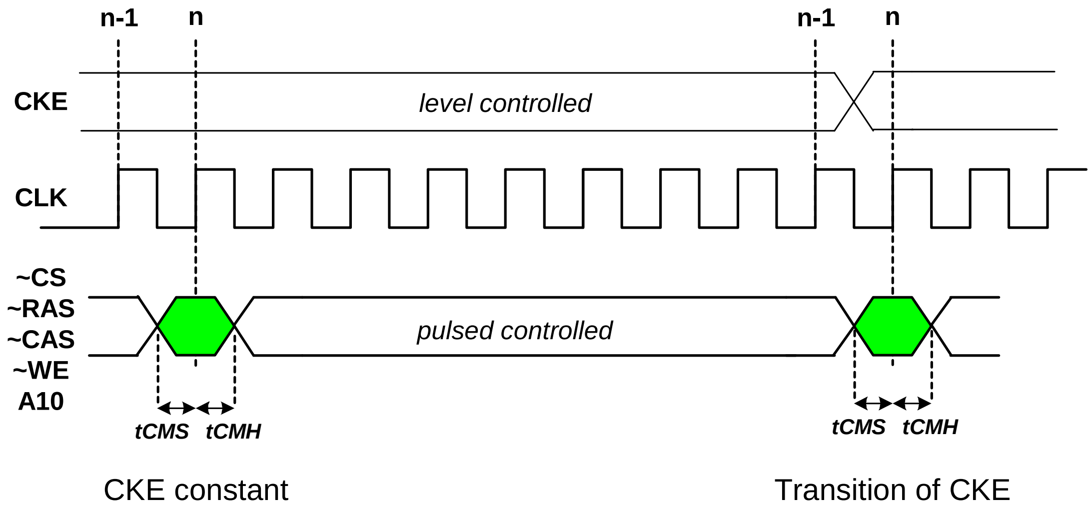

6.2 – Setup and Hold Times

Synchronous operations use the clock as the reference. The commands are latched on the rising edge of clock. The valid time margins around the rising edge are defined as setup time *tCMS* (time before rising edge) and hold time *tCMH* (time after rising edge) to guarantee that both the master and slave are working reliably together. The slew rate (typ. ≥ 1V/ns, time between the input low and high level) of clock-, command- and address signal, propagation delays (PCB) and capacitive loads (devices) influence these parameters and should be taken into consideration.

6.3 – Commands with CKE High

Note: CKE(n-1) is the logic state of the previous clock edge; CKE(n) is the logic state of the current clock edge.

Mode Register Set (MRS)

This command specifies the Read Latency and the working mode of the burst counter.

Once a mode register is programmed, the contents of the register will be held until re-programmed by another MRS command.

Figure 7: Data for Mode Register Set

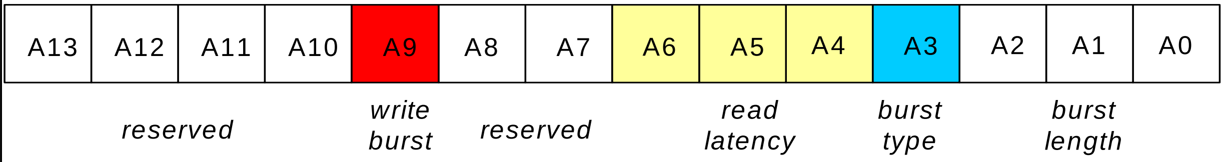

A13	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
reserved				*write burst*	*reserved*		*read latency*			*burst type*	*burst length*		

During MRS, the following address lines (figure 7) are used:

A[0:2]	Burst length (fixed or full page)
A[3]	Burst type (sequential or interleaved address counting)
A[4:6]	Read Latency (CAS latency)
A[7:8]	reserved zeros
A[9]	Write burst mode (programmable burst length or single write access)
A[10:13]	reserved zeros

Note: Only during mode register set, the address lines are used to transfer the specific data for the setup.

Note: The default value of the mode register is not defined, therefore the MRS must be written to during power up sequence.

Fixed Burst Length (1-8)

In this mode, the burst counter is configured to a length of 1, 2, 4 or 8. That means, only the address for the first access is issued, the second and third one are internally incremented by the burst counter. After the last transfer, the operation is automatically stopped. For a burst length of 1, the controller must provide the address for every access, so it is a no burst mode.

Full Page Burst

In this other mode, after each access the counter increments the address by 1. The main difference with fixed length bursts: the counter doesn’t stop automatically, but wraps around the page like a pointer in a circular buffer and starts over again. This procedure can be interrupted with another burst or stopped with burst stop command.

Note: The length full page is an optional feature for SDRAMs.

Sequential Counting

Hereby, the burst counter increments the address by simply adding 1 to the current address.

Note: The full page burst requires sequential counting.

Interleaved Counting

This counting mode is especially used in Intel based systems. Hereby, the burst counter increments interleaved (scrambled) depending on the starting address like for example:

0-1-2-3 or 1-0-3-2 or 2-3-0-1 or 3-2-1-0.

Note: Interleaved burst counting goes hand in hand with burst length 4.

Read Latency

The pipelined architecture causes a delay between the read command and the first data output (CAS-, read latency or pipeline deep) for different direction of address and data (write: address and data same direction, no latency required). The first data output cycle can be programmed as a fraction of clock cycles. Minimum read latency, (spec value C*L)* is based on the maximum clock frequency and depends also on the used speed grade. The spec *CL-1+tAC* describes when the data is driven first. For instance, if CL is set to one, the first data is driven after the time *tAC* is elapsed.

The following example illustrates the influence of the different operating frequencies on the latencies:

Speed grade		*-10 (100MHz)*	*-12 (83MHz)*
CL (Cycles)	*CL-1+tAC*	*Speed (MHz)*	*Speed (MHz)*
1 (2 or 3)	tAC	≤ 33	≤ 33
2 (3)	1+tAC	≤ 66	≤ 66
3	2+tAC	≤ 100	≤ 83

The values in brackets can also be set, but this will cause decreased performance.

Note: The higher the normalized clock for a speed grade, the higher the read latency.

Note: The programmable read latency feature is provided to allow efficient use of the SDRAM over a wide range of clock frequencies.

Extended Mode Register Set (MRS)

The extended mode register controls the function beyond those controlled by the mode register set. These additional functions are special features for mobile SDRAMs. They include temperature compensated self-refresh and partial self-refresh.

Figure 7a: Data for extended Mode Register Set

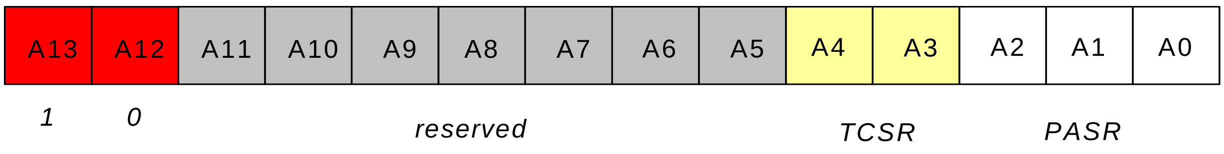

Note: Mobile SDRAMs have an additional extended mode register.

During extended MRS, the following address lines (figure 7a) are used:

A[0:2]	PASR (Partial array self-refresh)
A[3:4]	TCSR (Temperature compensated self-refresh)
A[5:11]	reserved (zeros)
A[12]	low for extended MRS
A[13]	high for extended MRS

Note: If during standard MRS the bits A[12:13] are all zero, the extended MRS is not initiated.

Activate (ACT)

Bank activate is used to select a random row in an idle bank. The SDRAM has 2 or 4 internal banks in the same chip and shares parts of the internal circuitry to reduce chip area. But the noise restricts the activation off all banks simultaneously.

Read (RD)

This command is used to perform data reads in the activated row. The first output appears in read (CAS) latency numbers of clock cycles after execution of read command. The output goes into high-impedance at the end of the operation. Execution of single- or burst reads is possible. For burst, the first read command is followed by multiple NOPs. The read can be initiated on any column address of the active row. Furthermore, every burst operation can be interrupted by another read, write, burst stop or precharge.

Note: Read cannot be interrupted directly with write, because of the read latency.

Write (WR)

The write command is similar to read. The data inputs are provided for the initial address in the same clock cycle as the write command. The write command can be initiated on any column address of the active row. Execution of single- or burst writes is possible. For burst, the first write command is followed by multiple NOPs. The write can be initiated on any column address of the active row. Furthermore, every burst can be interrupted by another write, read, burst stop or precharge.

Note: Interrupting write with precharge (spec tWR) must satisfy the tRAS spec.

Read with Autoprecharge (RDA)

This mode executes the earliest possible precharge automatically after fixed burst end. Unlike the standard read command, the A10 pin must be sampled high, while latching the column address.

Note: The Intel based multiple bank activation uses Autoprecharge.

Write with Autoprecharge (WRA)

Just like the RDA.

Note: The Autoprecharge is illegal in full page burst length.

Deselect (DESL)

The Deselect function prevents new commands from being executed by the SDRAM, regardless of the CLK signal being enabled. The device is effectively deselected by the ~CS pin. Operations already in progress are not affected.

Note: ~CS low enables the command decoder. If ~CS is high, all commands will be masked, including the refresh command. Internal operation such as burst cycles will not be suspended.

No Operation (NOP)

No operation is used to perform an access to the SDRAM selected by ~CS pin. NOP does not initiate any new operation, yet it is needed to complete operations requiring more than a single clock cycle, like bank activate, burst read and write, and auto refresh. NOP and DESL commands have the same effect on the devices.

Note: The SDRAM's burst counter will recognize bursting every time a single read or write is followed by multiple NOPs or DESLs. The external addresses will be ignored.

Burst Stop (BST)

The burst stop is used to truncate either fixed length or full page bursts. In fixed length bursts, it will stop automatically after its length.
Note: Burst length 1 requires no burst stop command.
Note: In full page burst, it is used to generate an arbitrary burst length.

Precharge Single Bank (PRE)

Precharging is used to close the accessed bank, the *tRAS* spec must be satisfied. Care should be taken to make sure that burst write is completed or the DQM pin is used to inhibit writing before asserting the precharge command.
Note: write followed by precharge (spec tWR) must satisfy the spec tRAS.

Precharge All Banks (PREA)

Just like PRE, but all banks are simultaneously closed due to A10 set to high. The precharge all is useful in systems, where more than 1 bank is open at a time (interleaved or banking mode). Otherwise, it precharges the active bank only.
Note: First command issued to the SDRAM during MRS.

Auto Refresh (REF)

Auto refresh is used during normal SDRAM operation and it is the same as CBR Refresh. This command is non persistent and is with every external refresh. The addresses are still handled by the internal refresh counter.
Note: No addresses are latched during NOP, BST, PRE, PREA commands.

6.4 – Commands with Transition of CKE

While the CKE line toggles its value asynchronously, the commands are always registered synchronously to the CLK signal. There are 3 modes of CKE: enter – maintain - exit
Note: The CKE signal toggles asynchronously in power-down, self refresh and suspend mode

Read Suspend (RDS)

The clock suspend mode occurs when a read burst is in progress and CKE is registered low. In the clock suspend mode, the internal clock is deactivated and the synchronous logic is frozen. For each rising clock edge with CKE sampled low, the access is extended and the burst counter will not be incremented.

Write Suspend (WRS)

Just like RDS.

Read Autoprecharge Suspend (RDAS)

The clock suspend mode occurs when an autoprecharge read burst is in progress and CKE is registered low. In the clock suspend mode, the internal clock is deactivated, the synchronous logic is frozen For each rising clock edge with CKE sampled low, the access is extended and the burst counter will not be incremented.

Write Autoprecharge Suspend (WRAS)

Just like RDAS.

Self Refresh (SREF)

When in self-refresh mode, the SDRAM retains data without external clocking. The command is initiated like an auto refresh except the fact that CKE is low. Once the self refresh is registered, all inputs become don’t care, with the exception of CKE, which must remain low. A 128 Mbit part needs for a 15 µs CBR refresh about 3 mA. In self-refresh mode, it takes 2 mA.

Power Down (PD)

This mode is initiated when a NOP or DESL command occurs while CKE is low. During PD, only the clock and CKE pins are active.
Note: This mode does not perform any refresh operation; therefore the device cannot remain in PD any longer than the refresh period tREF of the device.

7 – Command Coding

The truth table gives an overview of all commands issued to the SDRAM.

7.1 – Pin Description

Pin	*Type*	*Signal*	*Description*
CLK	(I)	pulse	master clock input
CKE:	(I)	level	command input
~CS:	(I)	pulse	command input
~RAS:	(I)	pulse	command input
~CAS:	(I)	pulse	command input
~WE:	(I)	pulse	command input
DQM	(I)	pulse	DQ-buffer control
A[10]	(I)	level/pulse	address/command input
A[X:0]	(I)	level	address
DQ[X:0]	(I/O)	level	data
BA0 (BS0)	(I)	level	bank select lines (2banks: BA, 4banks: BA0, BA1)
BA1 (BS1)	(I)	level	bank select lines (4banks: BA0, BA1)

I=Input, O=output

7.2 – Truth Table with CKE high

Command	*CKE(n-1)*	*CKE(n)*	*~CS*	*~RAS*	*~CAS*	*~WE*	*A10*	*ADDR*
MRS	1	1	0	0	0	0	v	v
ACT	1	1	0	0	1	1	v	v
RD	1	1	0	1	0	1	0	v
WR	1	1	0	1	0	0	0	v
RDA	1	1	0	1	0	1	1	v
WRA	1	1	0	1	0	0	1	v
Command without validity of address								
DESL	1	1	1	x	x	x	x	x
NOP	1	1	0	1	1	1	x	x
BST	1	1	0	1	1	0	x	x
PRE	1	1	0	0	1	0	0	x
PREA	1	1	0	0	1	0	1	x
REF	1	1	0	0	0	1	x	x

7.3 – Truth Table with Transition of CKE

Command	*CKE(n-1)*	*CKE(n)*	*~CS*	*~RAS*	*~CAS*	*~WE*	*A10*	*ADDR*
RDS En	1	0	0	1	0	1	0	v
RDS Ma	0	0	0	1	0	1	0	v
RDS Ex	0	1	0	1	0	1	0	v
WRS En	1	0	0	1	0	0	0	v
WRS Ma	0	0	0	1	0	0	0	v
WRS Ex	0	1	0	1	0	0	0	v
RDAS En	1	0	0	1	0	1	1	v
RDAS Ma	0	0	0	1	0	1	1	v
RDAS Ex	0	1	0	1	0	1	1	v
WRAS En	1	0	0	1	0	0	1	v
WRAS Ma	0	0	0	1	0	0	1	v
WRAS Ex	0	1	0	1	0	0	1	v
Command without validity of address								
SREF En	1	0	0	0	0	1	x	x
SREF Ma	0	0	x	x	x	x	x	x
SREF Ex	0	1	1	x	x	x	x	x
PD En	1	0	x	x	x	x	x	x
PD Ma	0	0	x	x	x	x	x	x
PD Ex	0	1	x	x	x	x	x	x

x=”don’t care”, v=valid data input, 0=logic 0, 1=logic 1, En=entry, Ma=maintain, Ex=exit

7.4 – Truth Table Address 10

This pin is used for several tasks as listed in the following table:

Command	***A10***
ACT	address 10
PRE	0
PREA	1
RD	0
WR	0
RDA	1
WRA	1

Note: A10 pin has multifunctional character.

7.5 – Truth Table Bank Access

The BA0/BA1 lines are decoded to access the corresponding bank in the SDRAM

Accessing 2 banks:

Access	***BA***	***A10***
Bank_A	0	0
Bank_B	1	0
All_Banks	x	1

Accessing 4 banks:

Access	***BA0***	***BA1***	***A10***
Bank_A	0	0	0
Bank_B	0	1	0
Bank_C	1	0	0
Bank_D	1	1	0
All_Banks	x	x	1

x=”don’t care”, 0=logic 0, 1=logic 1

Note: Since the A10 has a higher priority, all banks are accessed independent from the BA0/BA1 lines.

8 – On Page Column Access

The column access mode is allowed within the following modes:

- sequential counting, depending on MRS
- interleaved counting, depending on MRS
- random access (non burst)
- transitions (read to write or write to read)

8.1 – Sequential Column Counting

An on page access falls in a random column address in a dedicated row. The burst length defines, how many sequential addresses are incremented automatically in a row (page).
Figure 8 demonstrates the burst effect: If no burst mode (burst length 1) is selected, the controller must issue a column address for each operation. In the second case, the burst is set to four. After the first column address, the controller issues a NOP, which signals the SDRAM to start the internal burst counter. After the length has elapsed, the procedure restarts. The full page starts the burst with the first column address. Following sequential addresses are incremented endlessly in the page, the counter wraps around the page and restarts or it is interrupted with a burst stop command.
Note: The Burst length is independent from the performance throughput.

8.2 – Interleaved Column Counting

Just like sequential, but MRS configured to interleaved burst.

Figure 8: 6 sequential or interleaved reads using address pipeline, different burst lengths

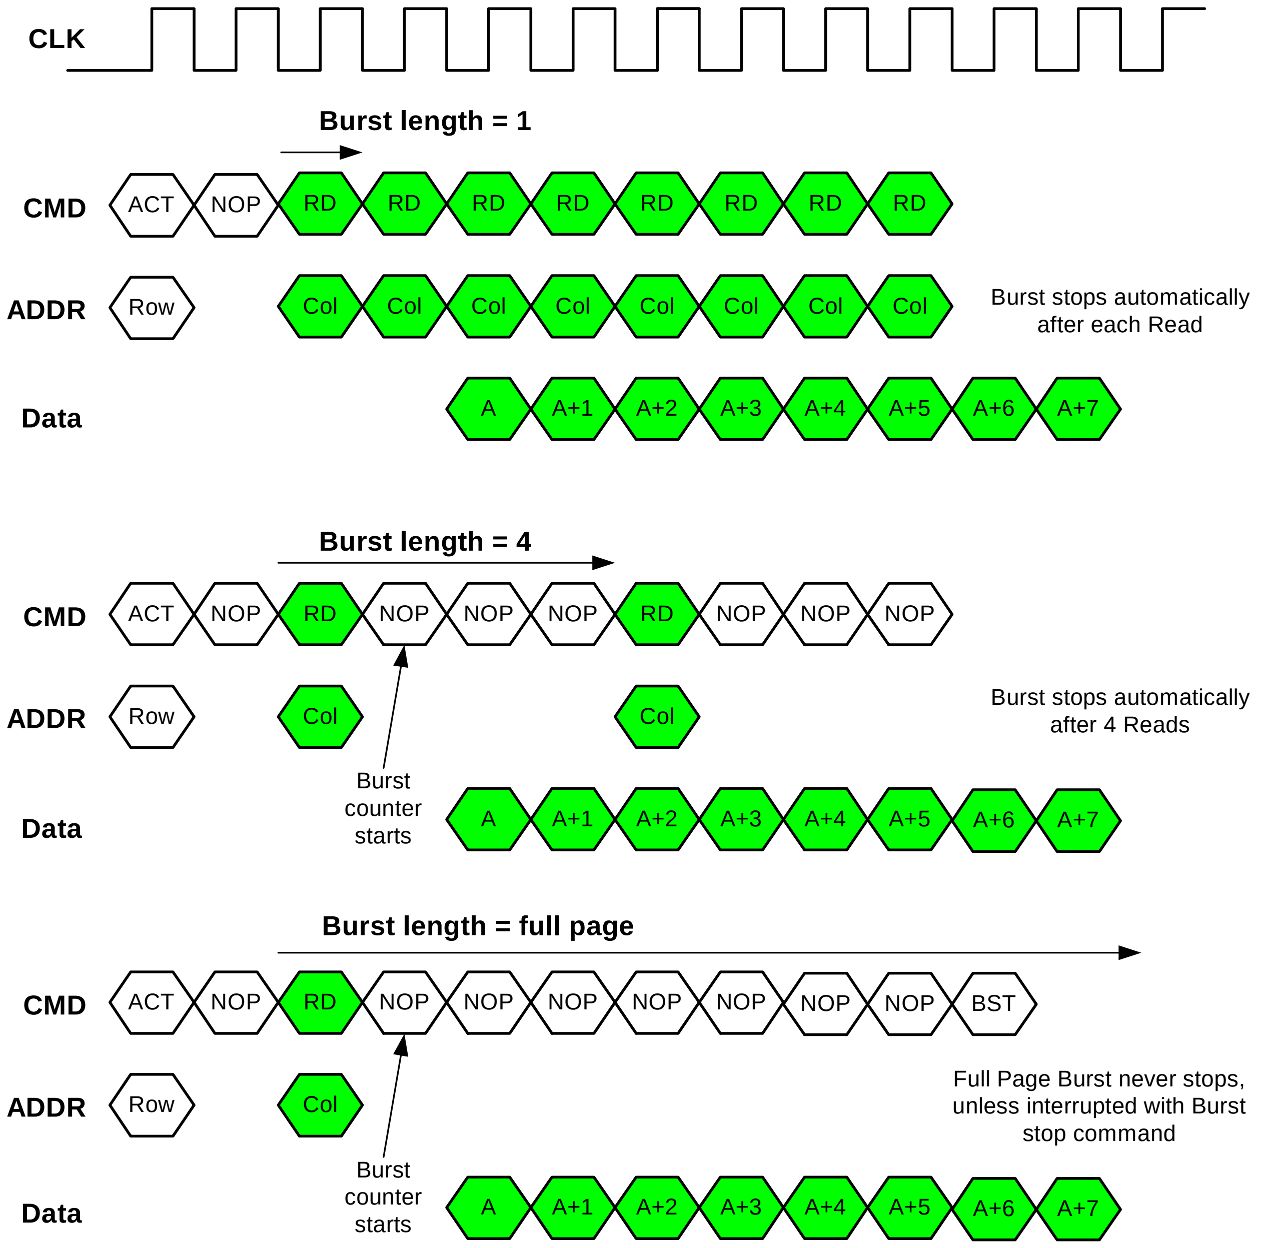

Note: Burst length full page allows sequential counting only.

Figure 9: Random Reads, CL=2

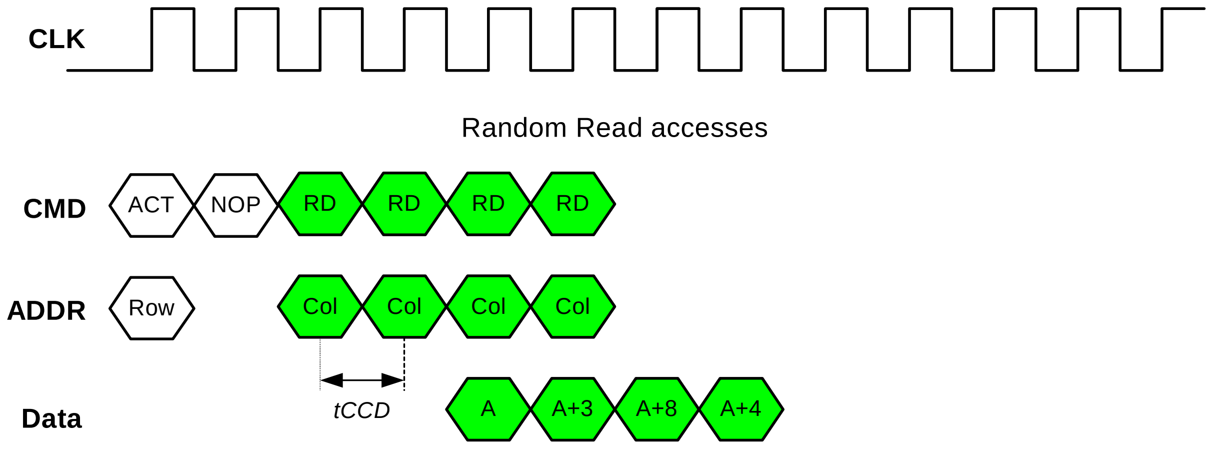

Note: Random Write and Read accesses are similar to burst length 1, each random column address must be issued to the memory to reach the same throughput like bursting.

8.3 – Random Column Access

However, not every application allows sequential accesses. SDRAM’s architecture is pipelined to fulfill the high-speed demands. It allows the randomly column access (figure 9) in each cycle by frequently interrupting the current burst.
Note: This access is independent from mode register set.

8.4 – Read to Write Column Access

The (figure 10) demonstrates a typical read to write operation using the DQM pin to mask data avoiding conflicts caused by different data latencies between read and write. The assertion of the mask pin will take effect on the third data access due to the 2 cycles latency. During this cycle, the Read buffer is disabled. In the next cycle the DQM pin will be deasserted driving data for a write.
Note: Write to Read transitions cause no conflict situation.

Figure 10: Read interrupted by Write using DQ Mask function (CL=2cycles)

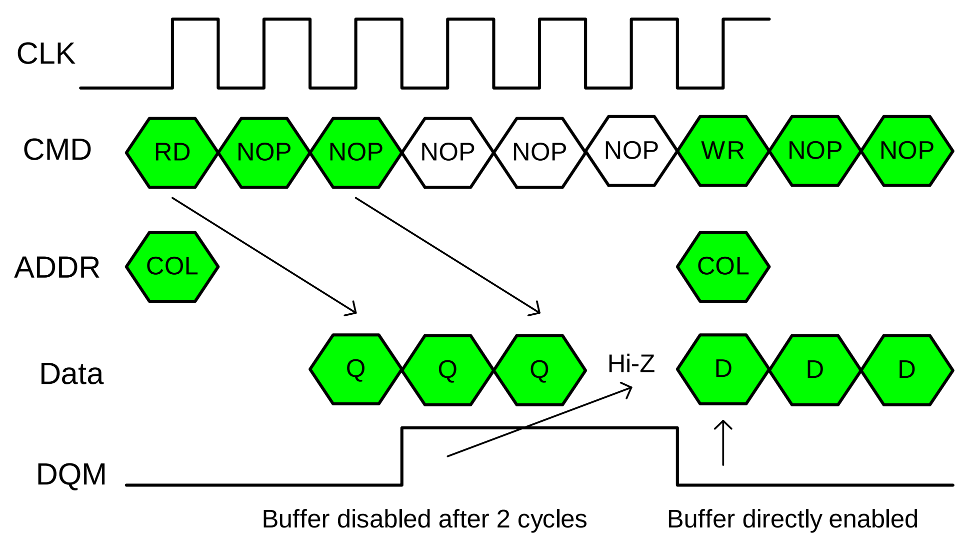

9 – Off Page Column Access

If the next column address falls in another row, the current row must be precharged before the new row is activated.

Two off bank modes are used:

- single bank activation
- multiple bank activation (interleaving or banking mode)

9.1 – Single Bank Activation

The current bank must be precharged (figure 11) and the next bank must be activated before accessing its column. This is also valid for different bank accesses in the same page, because only one bank at the time is active.

Note: Off page accesses need additional overhead caused by burst stop, precharge and activation.

Figure 11: Off page write access for single bank activation (BL=4)

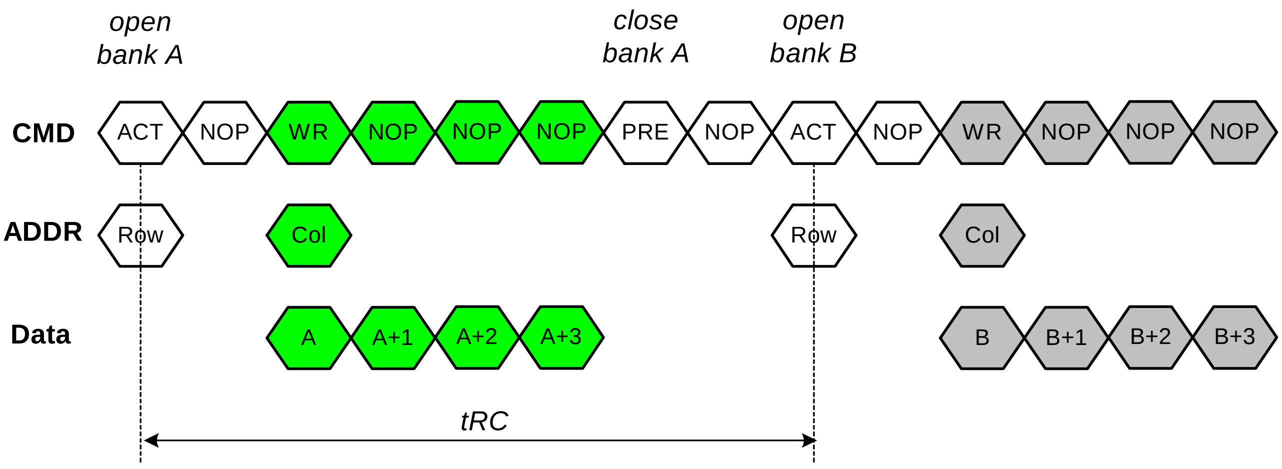

Note: The off page or off bank access puts additional overhead with precharge and activation of the page into account.

9.2 – Multiple Bank Activation

The disadvantage of single bank access results in overhead. A possible solution is the use of multiple bank activation or interleaved banking. This mode (figure 12) is used to ensure Intel compatibility. Designers soon recognized the bottleneck during a page miss. Solution: activation of another bank for bursting data while the current bank is autoprecharged.

Basically, all banks cannot be opened at once; noise generated during the sensing of each bank is high, requiring some time for power supplies to recover before another bank can be sensed reliably. *tRRD* specifies the minimum time required for interleaving between different banks. Comparing both modes, the interleaved mode pushes a higher throughput during off page. After a typical time *tRRD =1/2 tRC*, another bank can be opened.

Note: The banking mode (interleaved) is used for PC technology.

Figure 12: Off page write access for multiple bank activation (BL=4)

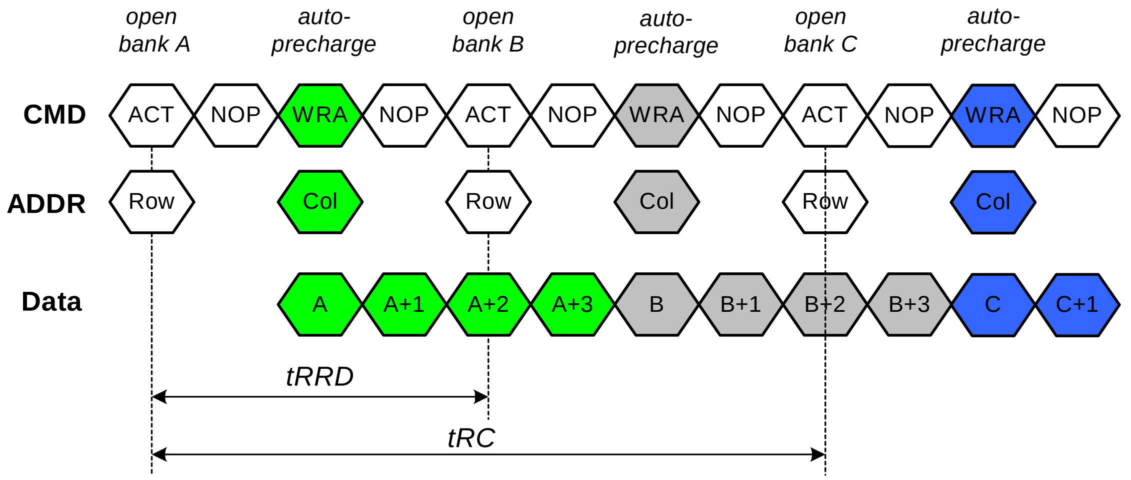

Note: This mode benefits from autoprecharge, because it doesn't require an explicit command.

10 – State Diagram of Command Decoder

The SDRAM's command decoder is designed as a state machine. The pulsed inputs trigger, synchronous with the clock, the different states. The state diagram illustrates the following:

- Manual and automatic commands
- Level and pulsed commands
- Allows development and timing analyses for SDRAM controller

10.1 – Burst length fixed 1-8

In Figure 13, the SDRAM is configured during the MRS command to a fixed length burst.

Characteristics:

- Burst length of 1, 2, 4, or 8 addresses
- Auto precharge capability
- Burst address sequential or interleaved counting
- Block of columns equal the burst length is effectively selected
- Will wrap around if a boundary is reached
- Burst stop command used to truncate the burst

Note: The interleaved burst counting is especially popular with Intel-based systems.

10.2 – Burst length full page

In Figure 14, the SDRAM is configured during the MRS command to a full page burst. Therefore, the autoprecharge capability is not allowed and reduces the state diagram.

Characteristics:

- Burst lengths of 256, 512, 1024, 2048, 4096 addresses
- Auto precharge illegal (burst will never end)
- Burst address counter: sequential only
- Block of columns equal the full page burst length is effectively selected
- Will wrap around if a boundary is reached
- Burst stop command used to generate arbitrary burst lengths

Note: All the different burst modes feature the same speed performance, but since the address lines are not driven during burst sequences, there is a noticeable improvement of reliability since the address location is generated inside the SDRAM.

Figure 13 : SDRAM configured to fixed length burst

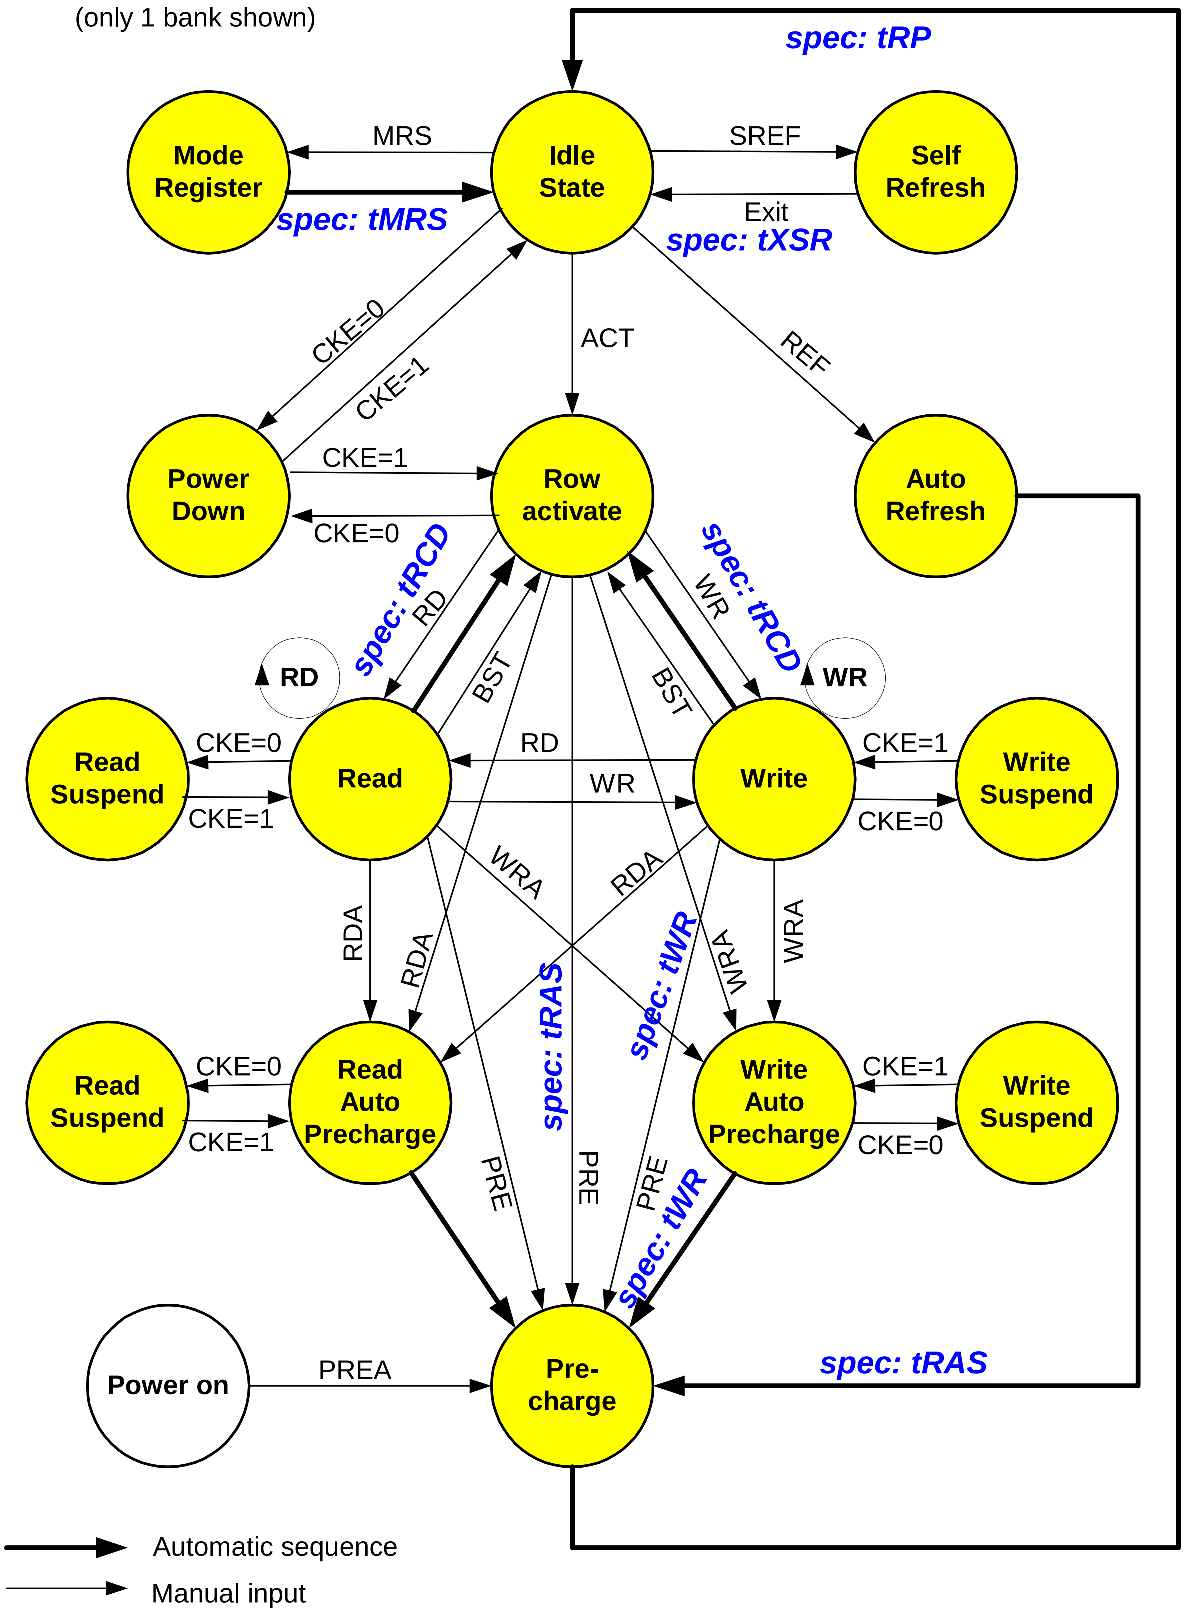

Figure 14: SDRAM configured to full page burst (optional)

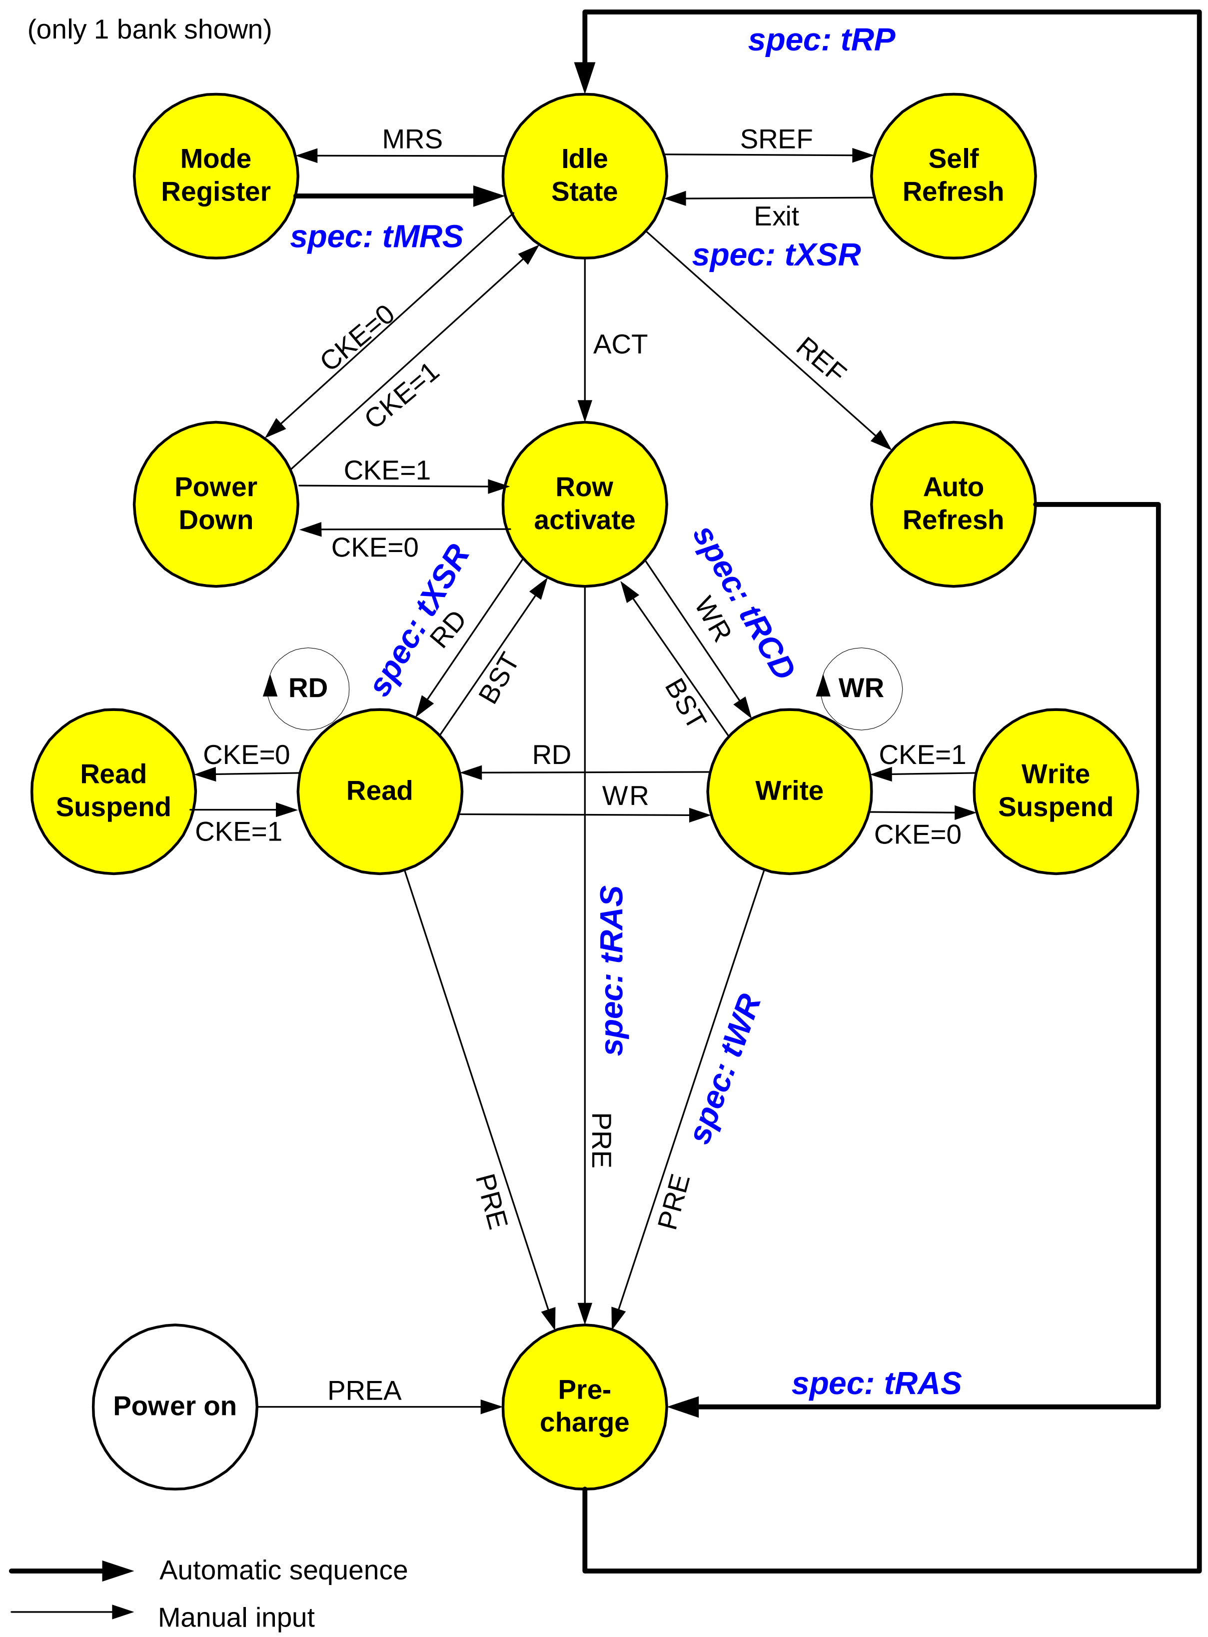

11 – SDRAM Controller

11.1 – Architecture

Your DSP or CPU controlled system just wants to see the SDRAM (figure 15) like a normal asynchronous memory. This is possible with the help of an SDRAM controller. In order to access an SDRAM, the controller must duplicate the SDRAM's timing protocol. For instance, Analog Devices Inc. uses on some DSP family members SDRAM controller, which are designed for glue less interface usage. The controller handling is therefore transparent to the user, who simply reads from or writes to the SDRAM memory. Before starting operation, the controller is configured with the necessary settings for the state machine. The CPU, supporting address pipelining, requests the controller with multiple read addresses, which are multiplexed to the SDRAM. In the meantime, the control interface gets activated providing pulsed commands to SDRAM. The busy line is asserted to prevent the CPU from accessing the controller during overhead cycles like refresh, precharge and activation.

Figure 15: Signal flow between CPU and SDRAM

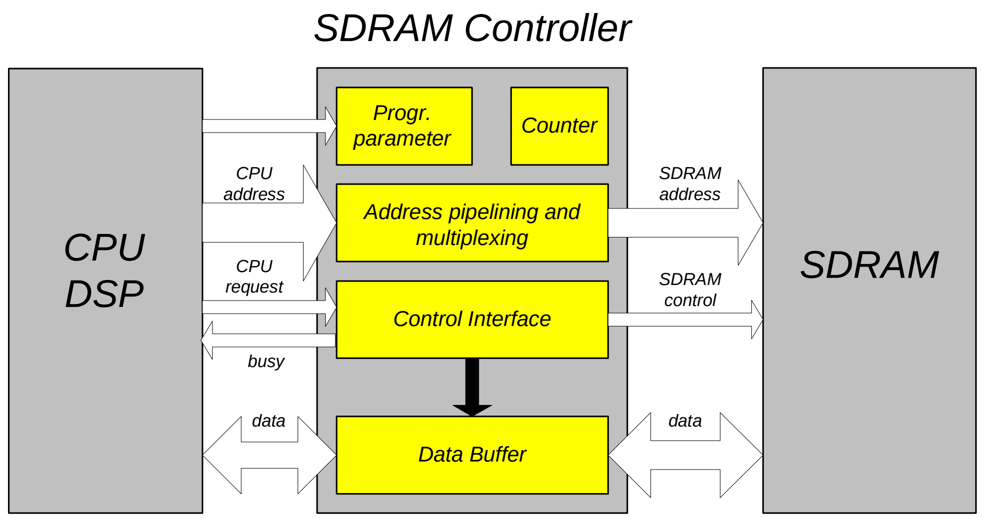

11.2 – State Machine

A state machine must be designed for a fixed length burst or a full page burst operation. For full page burst, the state machine is reduced, because the autoprecharge mode is illegal. Furthermore, the specs should be used to configure the timing gaps between the states in working mode to comply for all speed situations.

General decisions must be made if the memory is working in fixed or full page burst in conjunction with the read latency. A timer must be implemented to generate periodic refresh requests depending on the speed. Common issues to keep in mind are burst interruption and precharge termination.

Features orientated decisions like suspend, power-down modes, and use of multi bank activations can be additionally implemented. E.g. for multi bank operation, the controller needs command logic to precharge each bank and also to support the autoprecharge functionality. Furthermore, the spec *tRRD* (instead of *tRC)* comes into play and adds on to the state diagram. In single bank mode, only precharge-all is required. An important performance issue is based if the DSP or CPU supports address pipelining. For instance, a read latency of 2 requires latching of 2 addresses before the data are driven off chip. Otherwise, the throughput for non-sequential reads will slow down.

11.3 – Address Mapping Scheme

The mapping scheme describes, how the controller's input address is placed into the SDRAM in a multiplexed manner. Here, different mapping schemes could be used, depending on the system's performance requirements. For instance, in Intel based systems, a byte address scheme is used with the help of the DQM pins to mask the I/O size during writes. For DSPs, a non-byte controlled address mapping is most useful in order to use the full bandwidth.

Figure 16 shows two different physical usages of SDRAMs. For sequential mapping, the controller addresses each bank in a row before incrementing the row address. For row mapping, it addresses all rows before switching to the next bank.

Figure 16: SDRAM controller's address mapping schemes

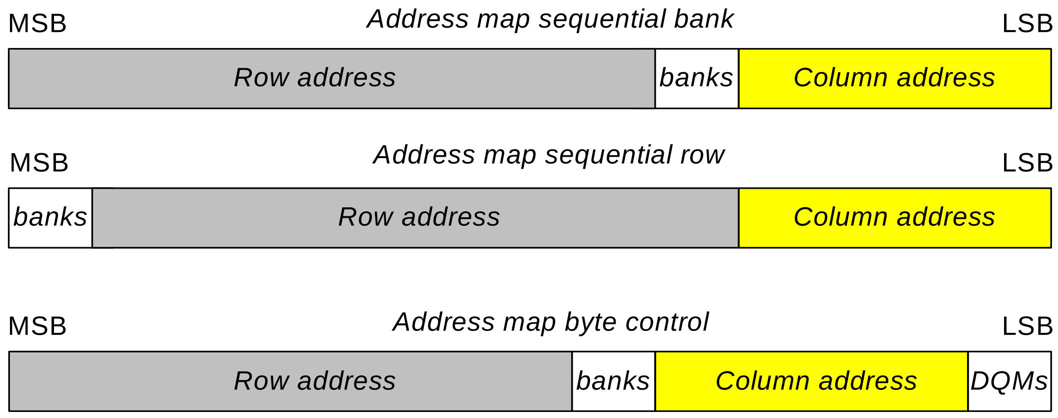

12 – SDRAM Standards

12.1 – JEDEC Standard

Some major SDRAM suppliers have invested resources to develop and standardize the device at JEDEC (Joint Electron Device Engineering Council) standard for memories No.21-C (SDRAM superset) and have educated the user community in how these devices operate.

Unlike SBSRAMs, SDRAMs have a JEDEC standard to follow. However, not all of them do. The major difference between standard and nonstandard SDRAMs occurs in the use of ~RAS line. In JEDEC-standard SDRAMs, ~RAS and other pulsed control lines form a command, after execution, these signals are no longer needed.

Experience with conventional DRAMs suggests that SDRAMs conforming to the JEDEC standard would be compatible. Unfortunately, there are still architectural variations within the standard. Some devices have dual internal memory arrays (prefetch architecture). This architecture allows them to fetch data from two consecutive addresses (even and odd) simultaneously, than present them in succession to the output buffers. An alternative design uses internal pipelining to retrieve the next data word from the internal array while the device presents the first word to the output. The pipelined architecture only offers a single memory array.

The difference between the two architectures is transparent to the user during normal burst reads. If the burst gets interrupted, however, and memory access resumes at a new column address, the prefetch design requires a 1-clock delay (2-N rule) before changing column addresses. On the other side, the pipelined design can respond immediately (1-N rule). To accommodate both architectures, the JEDEC standard requires that designers allow 2 clocks between address cycles when changing column addresses within the same row. Following this rule ensures that both architectures behave in the same way in the design, but it sacrifices the performance advantages of the pipelined memory.

12.2 – Intel Standard

Because the SDRAM market is based on the personal computer (PC) market, Intel has defined standards depending on different speeds to ensure compliance. The PC66, PC100 and PC133 standards describe the requirements of the DIMM modules, which are available in 32, 64 or 72 bit I/0 sizes for parity. Furthermore, the modules are registered or unbuffered and containing SPD-EEPROMs (serial presence detect) used for PC BIOS setup.

For instance, the PCxxx standard defines the dimension of PCB, pin layout, number of layers, capacitance and terminations.

In order to define a general module performance, it is characterized with Read Latency (*CL*)- RAS to CAS Delay (*tRCD*) - Row Precharge (*tRP*), all normalized to the systems clock.

For instance, the PC100 spec requires a minimum of (2-2-2) and a maximum of (3-3-3) cycles.

Links and References

http://www.jedec.org/download/pub21/default.htm
http://developer.intel.com/technology/memory/pcsdram/spec/index.htm
http://www.kingston.com
http://www.infineon.com
http://www.micron.com
http://www.necel.com
http://www.chips.ibm.com/products/memory
http://www.mitsubishichips.com/data/datasheets/memory/index.html
http://www.hea.com/product/dram.htm
http://www.samsungelectronics.com/semiconductors/products/products_index.html
http://www.toshiba.com
http://www.fujitsu.com
http://www.edram.com
http://www.xilinx.com
http://www.altera.com
http://www.celoxica.com