Dynamic RAMs

From Asynchrounos to DDR4

Contents

Articles

References

Article Licenses

Dynamic random-access memory

Dynamic random-access memory (**DRAM**) is a type of random-access memory that stores each bit of data in a separate capacitor within an integrated circuit. The capacitor can be either charged or discharged; these two states are taken to represent the two values of a bit, conventionally called 0 and 1. Since capacitors leak charge, the information eventually fades unless the capacitor charge is refreshed periodically. Because of this refresh requirement, it is a *dynamic* memory as opposed to SRAM and other *static* memory.

The main memory (the "RAM") in personal computers is dynamic RAM (DRAM). It is the RAM in laptop and workstation computers as well as some of the RAM of video game consoles.

The advantage of DRAM is its structural simplicity: only one transistor and a capacitor are required per bit, compared to four or six transistors in SRAM. This allows DRAM to reach very high densities. Unlike flash memory, DRAM is volatile memory (cf. non-volatile memory), since it loses its data quickly when power is removed. The transistors and capacitors used are extremely small; billions can fit on a single memory chip.

History

The cryptanalytic machine code-named *"Aquarius"* used at Bletchley Park during World War II incorporated a hard-wired dynamic memory. Paper tape was read and the characters on it "were remembered in a dynamic store. ... The store used a large bank of capacitors, which were either charged or not, a charged capacitor representing cross (1) and an uncharged capacitor dot (0). Since the charge gradually leaked away, a periodic pulse was applied to top up those still charged (hence the term 'dynamic')".[1]

In 1964 Arnold Farber and Eugene Schlig, working for IBM, created a hard-wired memory cell, using a transistor gate and tunnel diode latch. They replaced the latch with two transistors and two resistors, a configuration which became known as the Farber-Schlig cell. In 1965 Benjamin Agusta and his team at IBM created a 16-bit silicon memory chip based on the Farber-Schlig cell, with 80 transistors, 64 resistors, and four diodes. In 1966 DRAM was invented by Dr. Robert Dennard at the IBM Thomas J. Watson Research Center. He was granted U.S. patent number 3,387,286 [2] in 1968. Capacitors had been used for earlier memory schemes such as the drum of the Atanasoff–Berry Computer, the Williams tube and the Selectron tube.

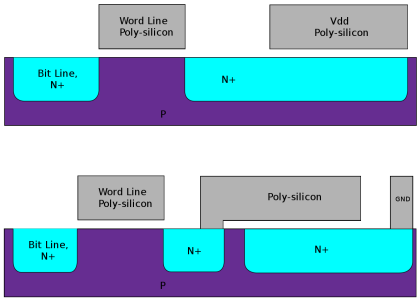

Schematic drawing of original designs of DRAM patented in 1968.

The Toshiba *"Toscal" BC-1411* electronic calculator, which was introduced in November 1966,[3] used a form of dynamic RAM built from discrete components.[4]

In 1969 Honeywell asked Intel to make a DRAM using a 3-transistor cell that they had developed. This became the Intel 1102 (1024x1) in early 1970. However, the 1102 had many problems, prompting Intel to begin work on their own improved design, in secrecy to avoid conflict with Honeywell. This became the first commercially available DRAM memory, the Intel 1103 (1024x1), in October 1970, despite initial problems with low yield until the fifth revision of the masks. The 1103 was designed by Joel Karp and laid out by Barbara Maness.

The first DRAM with multiplexed row and column address lines was the Mostek MK4096 (4096x1) designed by Robert Proebsting and introduced in 1973. This addressing scheme, a radical advance, enabled it to fit into packages with fewer pins, a cost advantage that grew with every jump in memory size. The MK4096 proved to be a very robust design for customer applications. At the 16K density, the cost advantage increased; the Mostek MK4116 16K DRAM, introduced in 1976, achieved greater than 75% worldwide DRAM market share. However, as density increased to 64K in the early 80s, Mostek was overtaken by Japanese DRAM manufacturers selling higher quality

DRAMs using the same multiplexing scheme at below-cost prices. See Japan–United States relations#Frictions in the semiconductor industry sector

Operation principle

DRAM is usually arranged in a rectangular array of charge storage cells consisting of one capacitor and transistor per data bit. The figure to the right shows a simple example with a 4 by 4 cell matrix. Modern DRAM matrices are many thousands of cells in height and width.

The long horizontal lines connecting each row are known as word-lines. Each column of cells is composed of two bit-lines, each connected to every other storage cell in the column (the illustration to the right does not include this important detail). They are generally known as the + and − bit-lines.

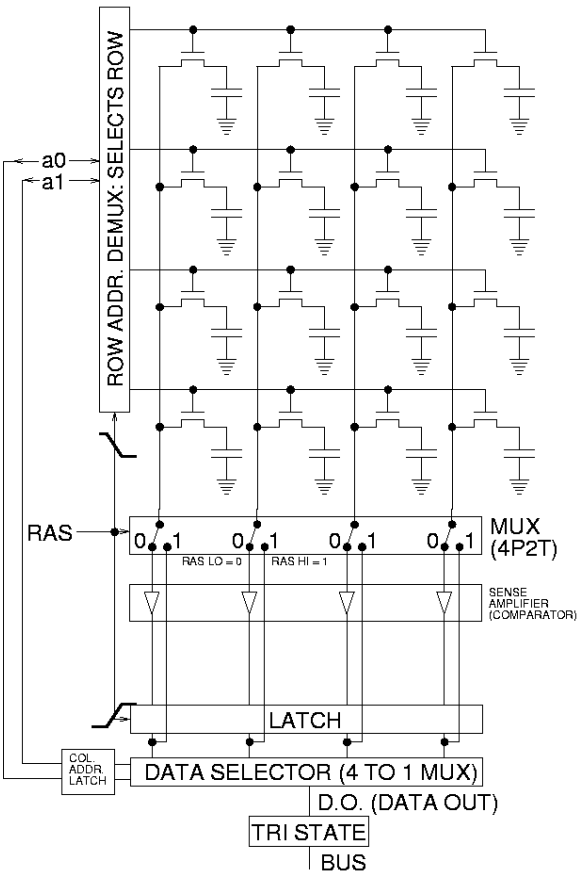

Principle of operation of DRAM read, for simple 4 by 4 array.

Operations to read a data bit from a DRAM storage cell

1. The sense amplifiers are disconnected.
2. The bit-lines are precharged to exactly equal voltages that are in between high and low logic levels. The bit-lines are physically symmetrical to keep the capacitance equal, and therefore the voltages are equal.
3. The precharge circuit is switched off. Because the bit-lines are relatively long, they have enough capacitance to maintain the precharged voltage for a brief time. This is an example of dynamic logic.
4. The desired row's word-line is then driven high to connect a cell's storage capacitor to its bit-line. This causes the transistor to conduct, transferring charge between the storage cell and the connected bit-line. If the storage cell's capacitor is discharged, it will greatly decrease the voltage on the bit-line as the precharge is used to charge the storage capacitor. If the storage cell is charged, the bit-line's voltage only decreases very slightly. This occurs because of the high capacitance of the storage cell capacitor compared to the capacitance of the bit-line, thus allowing the storage cell to determine the charge level on the bit-line.
5. The sense amplifiers are connected to the bit-lines. Positive feedback then occurs from the cross-connected inverters, thereby amplifying the small voltage difference between the odd and even row bit-lines of a particular column until one bit line is fully at the lowest voltage and the other is at the maximum high voltage. Once this has happened, the row is "open" (the desired cell data is available).
6. All storage cells in the open row are sensed simultaneously, and the sense amplifier outputs latched. A column address then selects which latch bit to connect to the external data bus. Reads of different columns in the same row can be performed without a row opening delay because, for the open row, all data has already been sensed and latched.
7. While reading of columns in an open row is occurring, current is flowing back up the bit-lines from the output of the sense amplifiers and recharging the storage cells. This reinforces (i.e. "refreshes") the charge in the storage

cell by increasing the voltage in the storage capacitor if it was charged to begin with, or by keeping it discharged if it was empty. Note that due to the length of the bit-lines there is a fairly long propagation delay for the charge to be transferred back to the cell's capacitor. This takes significant time past the end of sense amplification, and thus overlaps with one or more column reads.

8. When done with reading all the columns in the current open row, the word-line is switched off to disconnect the storage cell capacitors (the row is "closed") from the bit-lines. The sense amplifier is switched off, and the bit lines are precharged again.

To write to memory

To store data, a row is opened and a given column's sense amplifier is temporarily forced to the desired high or low voltage state, thus causing the bit-line to charge or discharge the cell storage capacitor to the desired value. Due to the sense amplifier's positive feedback configuration, it will hold a bit-line at stable voltage even after the forcing voltage is removed. During a write to a particular cell, all the columns in a row are sensed simultaneously just as during reading, so although only a single column's storage-cell capacitor charge is changed, the entire row is refreshed (written back in), as illustrated in the figure to the right.

Refresh rate

Typically, manufacturers specify that each row must have its storage cell capacitors refreshed every 64 ms or less, as defined by the JEDEC (Foundation for developing Semiconductor Standards) standard. Refresh logic is provided in a DRAM controller which automates the periodic refresh, that is no software or other hardware has to perform it. This makes the controller's logic circuit more complicated, but this drawback is outweighed by the fact that DRAM is much cheaper per storage cell and because each storage cell is very simple, DRAM has much greater capacity per unit of surface than SRAM.

Some systems refresh every row in a burst of activity involving all rows every 64 ms. Other systems refresh one row at a time staggered throughout the 64 ms interval. For example, a system with $2^{13} = 8192$ rows would require a staggered refresh rate of one row every 7.8 μs which is 64 ms divided by 8192 rows. A few real-time systems refresh a portion of memory at a time determined by an external timer function that governs the operation of the rest of a system, such as the vertical blanking interval that occurs every 10–20 ms in video equipment. All methods require some sort of counter to keep track of which row is the next to be refreshed. Most DRAM chips include that counter. Older types require external refresh logic to hold the counter.

Under some conditions, most of the data in DRAM can be recovered even if the DRAM has not been refreshed for several minutes.

Memory timing

Many parameters are required to fully describe the timing of DRAM operation. Here are some examples for two timing grades of asynchronous DRAM, from a data sheet published in 1998:[5]

	"50 ns"	"60 ns"	Description
t_{RC}	84 ns	104 ns	Random read or write cycle time (from one full /RAS cycle to another)
t_{RAC}	50 ns	60 ns	Access time: /RAS low to valid data out
t_{RCD}	11 ns	14 ns	/RAS low to /CAS low time
t_{RAS}	50 ns	60 ns	/RAS pulse width (minimum /RAS low time)
t_{RP}	30 ns	40 ns	/RAS precharge time (minimum /RAS high time)
t_{PC}	20 ns	25 ns	Page-mode read or write cycle time (/CAS to /CAS)
t_{AA}	25 ns	30 ns	Access time: Column address valid to valid data out (includes address setup time before /CAS low)
t_{CAC}	13 ns	15 ns	Access time: /CAS low to valid data out
t_{CAS}	8 ns	10 ns	/CAS low pulse width minimum

Thus, the generally quoted number is the /RAS access time. This is the time to read a random bit from a precharged DRAM array. The time to read additional bits from an open page is much less.

When such a RAM is accessed by clocked logic, the times are generally rounded up to the nearest clock cycle. For example, when accessed by a 100 MHz state machine (i.e. a 10 ns clock), the 50 ns DRAM can perform the first read in five clock cycles, and additional reads within the same page every two clock cycles. This was generally described as "5-2-2-2" timing, as bursts of four reads within a page were common.

When describing synchronous memory, timing is described by clock cycle counts separated by hyphens. These numbers represent t_{CL}-t_{RCD}-t_{RP}-t_{RAS} in multiples of the DRAM clock cycle time. Note that this is half of the data transfer rate when double data rate signaling is used. JEDEC standard PC3200 timing is 3-4-4-8[6] with a 200 MHz clock, while premium-priced high performance PC3200 DDR DRAM DIMM might be operated at 2-2-2-5 timing.[7]

	PC-3200 (DDR-400)				PC2-6400 (DDR2-800)				PC3-12800 (DDR3-1600)				Description
	Typical		Fast		Typical		Fast		Typical		Fast		
	cycles	time	cycles	time	cycles	time	cycles	time	cycles	time	cycles	time	
t_{CL}	3	15 ns	2	10 ns	5	12.5 ns	4	10 ns	9	11.25 ns	8	10 ns	/CAS low to valid data out (equivalent to t_{CAC})
t_{RCD}	4	20 ns	2	10 ns	5	12.5 ns	4	10 ns	9	11.25 ns	8	10 ns	/RAS low to /CAS low time
t_{RP}	4	20 ns	2	10 ns	5	12.5 ns	4	10 ns	9	11.25 ns	8	10 ns	/RAS precharge time (minimum precharge to active time)
t_{RAS}	8	40 ns	5	25 ns	16	40 ns	12	30 ns	27	33.75 ns	24	30 ns	Row active time (minimum active to precharge time)

The improvement over 11 years is not that significant. Minimum random access time has improved from t_{RAC} = 50 ns to t_{RCD} + t_{CL} = 22.5 ns, and even the premium 20 ns variety is only 2.5 times better compared to the typical case (~2.22 times better). CAS latency has improved even less, from t_{CAC} = 13 ns to 10 ns. However, the DDR3 memory does achieve 32 times higher bandwidth; due to internal pipelining and wide data paths, it can output two words every 1.25 ns (1600 Mword/s), while the EDO DRAM can output one word per t_{PC} = 20 ns (50 Mword/s).

Timing abbreviations

- t_{CL} – CAS latency
- t_{CR} – Command rate
- t_{PTP} – precharge to precharge delay
- t_{RAS} – RAS active time
- t_{RCD} – RAS to CAS delay
- t_{REF} – Refresh period
- t_{RFC} – Row refresh cycle time
- t_{RP} – RAS precharge
- t_{RRD} – RAS to RAS delay
- t_{RTP} – Read to precharge delay
- t_{RTR} – Read to read delay
- t_{RTW} – Read to write delay
- t_{WR} – Write recovery time
- t_{WTP} – Write to precharge delay
- t_{WTR} – Write to read delay
- t_{WTW} – Write to write delay

Error detection and correction

Electrical or magnetic interference inside a computer system can cause a single bit of DRAM to spontaneously flip to the opposite state. The majority of one-off ("soft") errors in DRAM chips occur as a result of background radiation, chiefly neutrons from cosmic ray secondaries, which may change the contents of one or more memory cells or interfere with the circuitry used to read/write them. Recent studies[8] show that single event upsets due to cosmic radiation have been dropping dramatically with process geometry and previous concerns over increasing bit cell error rates are unfounded.

This problem can be mitigated by using redundant memory bits and memory controllers that exploit these bits, usually implemented within DRAM modules. These extra bits are used to record parity and to enable missing data to be reconstructed by error-correcting code (ECC). Parity allows the detection of all single-bit errors (actually, any odd number of wrong bits). The most common error-correcting code, a SECDED Hamming code, allows a single-bit error to be corrected and, in the usual configuration, with an extra parity bit, double-bit errors to be detected.

An ECC-capable memory controller as used in many modern PCs can typically detect and correct errors of a single bit per 64-bit "word" (the unit of bus transfer), and detect (but not correct) errors of two bits per 64-bit word. Some systems also 'scrub' the errors, by writing the corrected version back to memory. The BIOS in some computers, and operating systems such as Linux, allow counting of detected and corrected memory errors; this allows identification and replacement of failing memory modules.

Recent studies give widely varying error rates with over 7 orders of magnitude difference, ranging from 10^{-10}–10^{-17} error/bit·h, roughly one bit error, per hour, per gigabyte of memory to one bit error, per century, per gigabyte of memory.[8][9][10] The Schroeder et al. 2009 study reported a 32% chance that a given computer in their study would suffer from at least one correctable error per year, and provided evidence that most such errors are intermittent hard rather than soft errors. A 2010 study at Rochester University also gave evident that a substantial fraction of memory errors are intermittent hard errors.[11] Large scale studies on non-ECC RAM in PCs and laptops suggest that undetected memory errors account for a substantial number of system failures: the study reported a 1 in 1700 chance per 1.5% of memory tested (extrapolating to an approximately 26% chance for total memory) that a computer would have a memory error per 8 months. [12]

Packaging

For economic reasons, the large (main) memories found in personal computers, workstations, and non-handheld game-consoles (such as PlayStation and Xbox) normally consist of dynamic RAM (DRAM). Other parts of the computer, such as cache memories and data buffers in hard disks, normally use static RAM (SRAM).

General DRAM formats

Dynamic random access memory is produced as integrated circuits (ICs) bonded and mounted into plastic packages with metal pins for connection to control signals and buses. In early use individual DRAM ICs were usually either installed directly to the motherboard or on ISA expansion cards; later they were assembled into multi-chip plug-in modules (DIMMs, SIMMs, etc.). Some standard module types are:

- **DRAM chip (Integrated Circuit or IC)**
  - Dual in-line Package (DIP)
- **DRAM (memory) modules**
  - Single In-line Pin Package (SIPP)
  - Single In-line Memory Module (SIMM)
  - Dual In-line Memory Module (DIMM)
  - Rambus In-line Memory Module (RIMM), technically DIMMs but called RIMMs due to their proprietary slot.
  - Small outline DIMM (SO-DIMM), about half the size of regular DIMMs, are mostly used in notebooks, small footprint PCs (such as Mini-ITX motherboards), upgradable office printers and networking hardware like routers. Comes in versions with:
    - 72-pin (32-bit)
    - 144-pin (64-bit) used for SDRAM
    - 200-pin (72-bit) used for DDR SDRAM and DDR2 SDRAM
    - 204-pin (64-bit) used for DDR3 SDRAM
  - Small outline RIMM (SO-RIMM). Smaller version of the RIMM, used in laptops. Technically SO-DIMMs but called SO-RIMMs due to their proprietary slot.
- **Stacked vs. non-stacked RAM modules**
  - Stacked RAM modules contain two or more RAM chips stacked on top of each other. This allows large modules to be manufactured using cheaper low density wafers. Stacked chip modules draw more power, and tend to run hotter than non-stacked modules. Stacked modules can be built using the older TSOP or the newer BGA style IC chips.

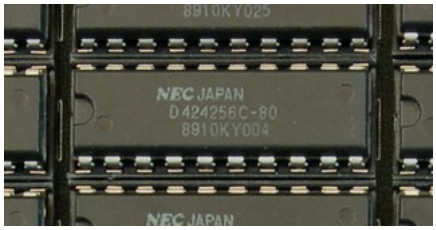

A 256 k x 4 bit 20-pin DIP DRAM on an early PC memory card (k = 1024), usually Industry Standard Architecture

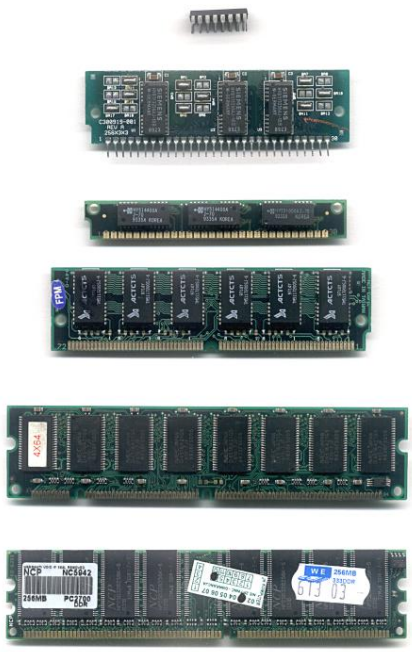

Common DRAM packages. From top to bottom: DIP, SIPP, SIMM (30-pin), SIMM (72-pin), DIMM (168-pin), DDR DIMM (184-pin).

Common DRAM modules

Common DRAM packages as illustrated to the right, from top to bottom:

1. DIP 16-pin (DRAM chip, usually pre-fast page mode DRAM (FPRAM))
2. SIPP 30-pin (usually FPRAM)
3. SIMM 30-pin (usually FPRAM)
4. SIMM 72-pin (often extended data out DRAM (EDO DRAM) but FPRAM is not uncommon)

5. DIMM 168-pin (SDRAM)
6. DIMM 184-pin (DDR SDRAM)
7. RIMM 184-pin (RDRAM) *not pictured*
8. DIMM 240-pin (DDR2 SDRAM and DDR3 SDRAM) *not pictured*

Memory size of a DRAM module

The exact number of bytes in a DRAM module is always an integral power of two. A 512 MB (as marked on a module) SDRAM DIMM, actually contains 512 MiB (mebibytes) = 512×2^{20} bytes = 2^{29} bytes = 536,870,912 bytes exactly), and might be made of 8 or 9 SDRAM chips, each containing exactly 512 Mib (mebibits) of storage, and each one contributing 8 bits to the DIMM's 64- or 72- bit width. For comparison, a 2 GB SDRAM module contains 2 GiB (gibibytes) = 2×2^{30} bytes = 2^{31} bytes = 2,147,483,648 bytes of memory, exactly. The module usually has 8 SDRAM chips of 256 MiB each.

Versions

While the fundamental DRAM cell and array has maintained the same basic structure (and performance) for many years, there have been many different interfaces for communicating with DRAM chips. When one speaks about "DRAM types", one is generally referring to the interface that is used.

Asynchronous DRAM

This is the basic form from which all others derive. An asynchronous DRAM chip has power connections, some number of address inputs (typically 12), and a few (typically one or four) bidirectional data lines. There are four active-low control signals:

- **/RAS**, the Row Address Strobe. The address inputs are captured on the falling edge of /RAS, and select a row to open. The row is held open as long as /RAS is low.
- **/CAS**, the Column Address Strobe. The address inputs are captured on the falling edge of /CAS, and select a column from the currently open row to read or write.
- **/WE**, Write Enable. This signal determines whether a given falling edge of /CAS is a read (if high) or write (if low). If low, the data inputs are also captured on the falling edge of /CAS.
- **/OE**, Output Enable. This is an additional signal that controls output to the data I/O pins. The data pins are driven by the DRAM chip if /RAS and /CAS are low, /WE is high, and /OE is low. In many applications, /OE can be permanently connected low (output always enabled), but it can be useful when connecting multiple memory chips in parallel.

This interface provides direct control of internal timing. When /RAS is driven low, a /CAS cycle must not be attempted until the sense amplifiers have sensed the memory state, and /RAS must not be returned high until the storage cells have been refreshed. When /RAS is driven high, it must be held high long enough for precharging to complete.

Although the RAM is asynchronous, the signals are typically generated by a clocked memory controller, which limits their timing to multiples of the controller's clock cycle.

RAS Only Refresh (ROR)

Classic asynchronous DRAM is refreshed by opening each row in turn.

The refresh cycles are distributed across the entire refresh interval in such a way that all rows are refreshed within the required interval. To refresh one row of the memory array using /RAS Only Refresh, the following steps must occur:

1. The row address of the row to be refreshed must be applied at the address input pins.
2. /RAS must switch from high to low. /CAS must remain high.
3. At the end of the required amount of time, /RAS must return high.

This can be done by supplying a row address and pulsing /RAS low; it is not necessary to perform any /CAS cycles. An external counter is needed to iterate over the row addresses in turn.[13]

CAS before RAS refresh (CBR)

For convenience, the counter was quickly incorporated into RAM chips themselves. If the /CAS line is driven low before /RAS (normally an illegal operation), then the DRAM ignores the address inputs and uses an internal counter to select the row to open. This is known as /CAS-before-/RAS (CBR) refresh.

This became the standard form of refresh for asynchronous DRAM, and is the only form generally used with SDRAM.

Hidden refresh

Given support of CAS-before-RAS refresh, it is possible to deassert /RAS while holding /CAS low to maintain data output. If /RAS is then asserted again, this performs a CBR refresh cycle while the DRAM outputs remain valid. Because data output is not interrupted, this is known as "hidden refresh".[14]

Video DRAM (VRAM)

VRAM is a dual-ported variant of DRAM that was once commonly used to store the frame-buffer in some graphics adaptors.

Window DRAM (WRAM)

WRAM is a variant of VRAM that was once used in graphics adaptors such as the Matrox Millenium and ATI 3D Rage Pro. WRAM was designed to perform better and cost less than VRAM. WRAM offered up to 25% greater bandwidth than VRAM and accelerated commonly used graphical operations such as text drawing and block fills.[15]

Fast page mode DRAM (FPM DRAM)

Fast page mode DRAM is also called FPM DRAM, FPRAM, Page mode DRAM, Fast page mode memory, or Page mode memory.

In page mode, a row of the DRAM can be kept "open" by holding /RAS low while performing multiple reads or writes with separate pulses of /CAS so that successive reads or writes within the row do not suffer the delay of precharge and accessing the row. This increases the performance of the system when reading or writing bursts of data.

Static column is a variant of page mode in which the column address does not need to be stored in, but rather, the address inputs may be changed with /CAS held low, and the data output will be updated accordingly a few nanoseconds later.

Nibble mode is another variant in which four sequential locations within the row can be accessed with four consecutive pulses of /CAS. The difference from normal page mode is that the address inputs are not used for the second through fourth /CAS edges; they are generated internally starting with the address supplied for the first /CAS

edge.

Extended data out DRAM (EDO DRAM)

EDO DRAM, sometimes referred to as Hyper Page Mode enabled DRAM, is similar to Fast Page Mode DRAM with the additional feature that a new access cycle can be started while keeping the data output of the previous cycle active. This allows a certain amount of overlap in operation (pipelining), allowing somewhat improved performance. It was 5% faster than FPM DRAM, which it began to replace in 1995, when Intel introduced the 430FX chipset that supported EDO DRAM.

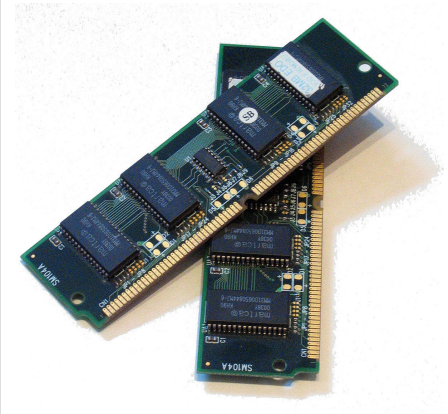

A pair of 32 MB EDO DRAM modules.

To be precise, EDO DRAM begins data output on the falling edge of /CAS, but does not stop the output when /CAS rises again. It holds the output valid (thus extending the data output time) until either /RAS is deasserted, or a new /CAS falling edge selects a different column address.

Single-cycle EDO has the ability to carry out a complete memory transaction in one clock cycle. Otherwise, each sequential RAM access within the same page takes two clock cycles instead of three, once the page has been selected. EDO's performance and capabilities allowed it to somewhat replace the then-slow L2 caches of PCs. It created an opportunity to reduce the immense performance loss associated with a lack of L2 cache, while making systems cheaper to build. This was also good for notebooks due to difficulties with their limited form factor, and battery life limitations. An EDO system with L2 cache was tangibly faster than the older FPM/L2 combination.

Single-cycle EDO DRAM became very popular on video cards towards the end of the 1990s. It was very low cost, yet nearly as efficient for performance as the far more costly VRAM.

Much equipment taking 72-pin SIMMs could use either FPM or EDO. Problems were possible, particularly when mixing FPM and EDO. Early Hewlett-Packard printers had FPM RAM built in; some, but not all, models worked if additional EDO SIMMs were added.[16]

Burst EDO DRAM (BEDO DRAM)

An evolution of EDO DRAM, **Burst EDO DRAM**, could process four memory addresses in one burst, for a maximum of 5-1-1-1, saving an additional three clocks over optimally designed EDO memory. It was done by adding an address counter on the chip to keep track of the next address. BEDO also added a pipelined stage allowing page-access cycle to be divided into two components. During a memory-read operation, the first component accessed the data from the memory array to the output stage (second latch). The second component drove the data bus from this latch at the appropriate logic level. Since the data is already in the output buffer, quicker access time is achieved (up to 50% for large blocks of data) than with traditional EDO.

Although BEDO DRAM showed additional optimization over EDO, by the time it was available the market had made a significant investment towards synchronous DRAM, or SDRAM [17]. Even though BEDO RAM was superior to SDRAM in some ways, the latter technology quickly displaced BEDO.

Multibank DRAM (MDRAM)

Multibank DRAM applies the interleaving technique for main memory to second-level cache memory to provide a cheaper and faster alternative to SRAM. The chip splits its memory capacity into small blocks of 256 kB and allows operations to two different banks in a single clock cycle.

This memory was primarily used in graphic cards with Tseng Labs ET6x00 chipsets, and was made by MoSys. Boards based upon this chipset often used the unusual RAM size configuration of 2.25 MB, owing to MDRAM's ability to be implemented in various sizes more easily. This size of 2.25 MB allowed 24-bit color at a resolution of 1024×768, a very popular display setting in the card's time.

Synchronous graphics RAM (SGRAM)

SGRAM is a specialized form of SDRAM for graphics adaptors. It adds functions such as bit masking (writing to a specified bit plane without affecting the others) and block write (filling a block of memory with a single colour). Unlike VRAM and WRAM, SGRAM is single-ported. However, it can open two memory pages at once, which simulates the dual-port nature of other VRAM technologies.

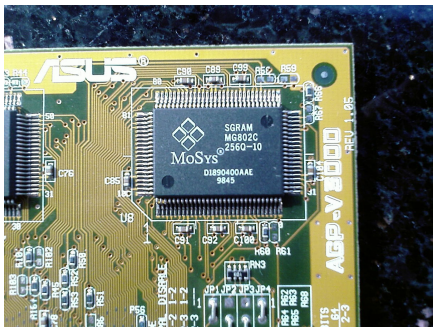

MoSys SGRAM

Synchronous dynamic RAM (SDRAM)

SDRAM significantly revises the asynchronous memory interface, adding a clock (and a clock enable) line. All other signals are received on the rising edge of the clock.

The /RAS and /CAS inputs no longer act as strobes, but are instead, along with /WE, part of a 3-bit command:

SDRAM Command summary

/CS	/RAS	/CAS	/WE	Address	Command
H	x	x	x	x	Command inhibit (No operation)
L	H	H	H	x	No operation
L	H	H	L	x	Burst Terminate: stop a read or write burst in progress
L	H	L	H	column	Read from currently active row
L	H	L	L	column	Write to currently active row
L	L	H	H	row	Activate a row for read and write
L	L	H	L	x	Precharge (deactivate) the current row
L	L	L	H	x	Auto refresh: Refresh one row of each bank, using an internal counter
L	L	L	L	mode	Load mode register: Address bus specifies DRAM operation mode.

The /OE line's function is extended to a per-byte "DQM" signal, which controls data input (writes) in addition to data output (reads). This allows DRAM chips to be wider than 8 bits while still supporting byte-granularity writes.

Many timing parameters remain under the control of the DRAM controller. For example, a minimum time must elapse between a row being activated and a read or write command. One important parameter must be programmed into the SDRAM chip itself, namely the CAS latency. This is the number of clock cycles allowed for internal operations between a read command and the first data word appearing on the data bus. The "Load mode register" command is used to transfer this value to the SDRAM chip. Other configurable parameters include the length of read and write bursts, i.e. the number of words transferred per read or write command.

The most significant change, and the primary reason that SDRAM has supplanted asynchronous RAM, is the support for multiple internal banks inside the DRAM chip. Using a few bits of "bank address" which accompany each command, a second bank can be activated and begin reading data *while a read from the first bank is in progress*. By alternating banks, an SDRAM device can keep the data bus continuously busy, in a way that asynchronous DRAM cannot.

Single data rate (SDR)

Single data rate SDRAM (sometimes known as **SDR**) is a synchronous form of DRAM.

Double data rate (DDR)

Double data rate SDRAM (**DDR**) was a later development of SDRAM, used in PC memory beginning in 2000. Subsequent versions are numbered sequentially (**DDR2**, **DDR3**, etc.). DDR SDRAM internally performs double-width accesses at the clock rate, and uses a double data rate interface to transfer one half on each clock edge. DDR2 and DDR3 increased this factor to 4× and 8×, respectively, delivering 4-word and 8-word bursts over 2 and 4 clock cycles, respectively. The internal access rate is mostly unchanged (200 million per second for DDR-400, DDR2-800 and DDR3-1600 memory), but each access transfers more data.

Direct Rambus DRAM (DRDRAM)

Direct RAMBUS DRAM (**DRDRAM**) was developed by Rambus.

Pseudostatic RAM (PSRAM)

PSRAM or **PSDRAM** is dynamic RAM with built-in refresh and address-control circuitry to make it behave similarly to static RAM (SRAM). It combines the high density of DRAM with the ease of use of true SRAM. PSRAM (made by Numonyx) is used in the Apple iPhone and other embedded systems as XFlar Platform.[18]

Some DRAM components have a "self-refresh mode". While this involves much of the same logic that is needed for pseudo-static operation, this mode is often equivalent to a standby mode. It is provided primarily to allow a system to suspend operation of its DRAM controller to save power without losing data stored in DRAM, not to allow operation without a separate DRAM controller as is the case with PSRAM.

An embedded variant of PSRAM is sold by MoSys under the name 1T-SRAM. It is technically DRAM, but behaves much like SRAM. It is used in Nintendo Gamecube and Wii consoles.

Reduced Latency DRAM (RLDRAM)

Reduced Latency DRAM is a high performance double data rate (DDR) SDRAM that combines fast, random access with high bandwidth, mainly intended for networking and caching applications.

1T DRAM

Unlike all of the other variants described in this section of this article, **1T DRAM** is a different way of constructing the basic DRAM bit cell. 1T DRAM is a "capacitorless" bit cell design that stores data in the parasitic body capacitor that is an inherent part of silicon on insulator (SOI) transistors. Considered a nuisance in logic design, this floating body effect can be used for data storage. Although refresh is still required, reads are non-destructive; the stored charge causes a detectable shift in the threshold voltage of the transistor.[19]

There are several types of 1T DRAM memories: the commercialized Z-RAM from Innovative Silicon, the TTRAM from Renesas and the A-RAM from the UGR/CNRS consortium.

The classic one-transistor/one-capacitor (1T/1C) DRAM cell is also sometimes referred to as "1T DRAM", particularly in comparison to 3T and 4T DRAM which it replaced in the 1970s.

Security

Although dynamic memory is only specified and *guaranteed* to retain its contents when supplied with power and refreshed every 64 ms, the memory cell capacitors often retain their values for significantly longer, particularly at low temperatures.[20] Under some conditions most of the data in DRAM can be recovered even if it has not been refreshed for several minutes.[21]

This property can be used to circumvent security and recover data stored in memory and assumed to be destroyed at power-down by quickly rebooting the computer and dumping the contents of the RAM, or by cooling the chips and transferring them to a different computer. Such an attack was demonstrated to circumvent popular disk encryption systems, such as the open source TrueCrypt, Microsoft's BitLocker Drive Encryption, and Apple's FileVault.[20] This type of attack against a computer is often called a cold boot attack.

References

[1] Copeland B. Jack, and others (2006) *Colossus: The Secrets of Bletchley Park's Codebreaking Computers* Oxford: Oxford University Press, p301.

[2] http://patft1.uspto.gov/netacgi/nph-Parser?patentnumber=3387286

[3] Spec Sheet for Toshiba "TOSCAL" BC-1411 (http://www.oldcalculatormuseum.com/s-toshbc1411.html)

[4] Toshiba "Toscal" BC-1411 Desktop Calculator (http://www.oldcalculatormuseum.com/toshbc1411.html) (The introduction date is listed here as November 1965, but this is a year too early and appears to be a typographical error.)

[5] Micron 4 Meg x 4 EDO DRAM data sheet (http://download.micron.com/pdf/datasheets/dram/d47b.pdf)

[6] cmx1024-3200.ai (http://www.corsairmemory.com/corsair/products/specs/cmx1024-3200.pdf)

[7] http://www.corsairmemory.com/corsair/products/specs/twinx1024-3200xl.pdf

[8] Borucki, "Comparison of Accelerated DRAM Soft Error Rates Measured at Component and System Level", 46th Annual International Reliability Physics Symposium, Phoenix, 2008, pp. 482–487

[9] http://www.cs.toronto.edu/~bianca/papers/sigmetrics09.pdf

[10] http://www.ece.rochester.edu/~xinli/usenix07/

[11] Li, Huang, Shen, Chu (2010). ""A Realistic Evaluation of Memory Hardware Errors and Software System Susceptibility". Usenix Annual Tech Conference 2010" (http://www.cs.rochester.edu/~kshen/papers/usenix2010-li.pdf). .

[12] "Cycles, cells and platters: an empirical analysis of hardware failures on a million consumer PCs. Proceedings of the sixth conference on Computer systems (EuroSys '11). pp 343-356" (http://research.microsoft.com/pubs/144888/eurosys84-nightingale.pdf). 2011. .

[13] (http://www.ece.cmu.edu/~ece548/localcpy/dramop.pdf)

[14] Various Methods of DRAM Refresh (http://www.downloads.reactivemicro.com/Public/Electronics/DRAM/DRAM Refresh.pdf) Micron Technical Note TN-04-30

[15] The PC Guide (http://www.pcguide.com/ref/video/techWRAM-c.html), definition of WRAM.

[16] Page on memory upgrades for HP printers (http://www.hpprintermemory.com/index.html#y)

[17] http://www.tomshardware.com/1998/10/24/ram_guide/page7.html

[18] EE Times teardown of iPhone 3G (http://www.eetimes.com/showArticle.jhtml?articleID=209000014)

[19] Sallese, Jean-Michel (2002-06-20). "Principles of the 1T Dynamic Access Memory Concept on SOI" (http://legwww.epfl.ch/ekv/mos-ak/wroclaw/MOS-AK_JMS.pdf). *MOS Modeling and Parameter Extraction Group Meeting*. Wroclaw, Poland. . Retrieved 2007-10-07.

[20] "Center for Information Technology Policy » Lest We Remember: Cold Boot Attacks on Encryption Keys" (http://web.archive.org/web/20110722182409/http://citp.princeton.edu/memory/). . 080222 citp.princeton.edu

[21] Scheick, Leif Z.; Guertin, Steven M.; Swift, Gary M. (December 2000). "Analysis of radiation effects on individual DRAM cells" (http://parts.jpl.nasa.gov/docs/DRAM_Indiv-00.pdf) (PDF). *IEEE Trans. on Nuclear Science* **47** (6): 2534–2538. doi:10.1109/23.903804. ISSN 0018-9499. . Retrieved 2008-11-03.

External links

- Modern DRAM Memory Systems:Performance analysis and a high performance, power-constrained DRAM scheduling algorithm (http://www.ece.umd.edu/~blj/papers/thesis-PhD-wang--DRAM.pdf) – PhD dissertation by David Tawei Wang, has very well written and detailed discussion on how DRAM works.
- DRAM density and speed trends (http://www.eecs.berkeley.edu/~culler/courses/cs252-s05/lectures/cs252s05-lec01-intro.ppt#359,15,Memory Capacity (Single Chip DRAM)) has some interesting historical trend charts of DRAM density and speed from 1980.
- Back to Basics—Memory, part 3 (http://www.computerwriter.com/archives/1998/cw052198.htm)
- Benefits of Chipkill-Correct ECC for PC Server Main Memory (http://www-1.ibm.com/servers/eserver/pseries/campaigns/chipkill.pdf) — A 1997 discussion of SDRAM reliability—some interesting information on "soft errors" from cosmic rays, especially with respect to Error-correcting code schemes
- [a] Tezzaron Semiconductor Soft Error White Paper (http://www.tezzaron.com/about/papers/soft_errors_1_1_secure.pdf) 1994 literature review of memory error rate measurements.
- Soft errors' impact on system reliability (http://www.edn.com/article/CA454636.html) — Ritesh Mastipuram and Edwin C Wee, Cypress Semiconductor, 2004
- Scaling and Technology Issues for Soft Error Rates (http://www.nepp.nasa.gov/DocUploads/40D7D6C9-D5AA-40FC-829DC2F6A71B02E9/Scal-00.pdf) A Johnston—4th Annual Research Conference on Reliability Stanford University, October 2000
- Challenges and future directions for the scaling of dynamic random-access memory (DRAM) (http://www.research.ibm.com/journal/rd/462/mandelman.html) — J. A. Mandelman, R. H. Dennard, G. B. Bronner, J. K. DeBrosse, R. Divakaruni, Y. Li, and C. J. Radens, IBM 2002
- Ars Technica: RAM Guide (http://arstechnica.com/paedia/r/ram_guide/ram_guide.part1-2.html)
- Versatile DRAM interface for the 6502 CPU (http://www.howell1964.freeserve.co.uk/projects/DRAM_6502.htm)
- David Tawei Wang (2005). *Modern DRAM Memory Systems: Performance Analysis and a High Performance, Power-Constrained DRAM-Scheduling Algorithm* (http://www.ece.umd.edu/~blj/papers/thesis-PhD-wang--DRAM.pdf). PhD thesis, University of Maryland, College Park. Retrieved 2007-03-10. A detailed description of current DRAM technology.
- The Toshiba *"Toscal" BC-1411* Desktop Calculator (http://www.oldcalculatormuseum.com/toshbc1411.html) — An early electronic calculator that uses a form of dynamic RAM built from discrete components.
- Mitsubishi's 3D-RAM And Cache DRAM (http://findarticles.com/p/articles/mi_m0EIN/is_1998_July_21/ai_50179297) incorporate high performance, on-board SRAM cache
- Multi-port Cache DRAM — **MP-RAM** (http://www.cs.berkeley.edu/~pattrsn/294)
- How to install DRAM installation guide (http://www.memorydepot.com/guide_install.asp) for desktop and laptop computers.
- (http://www.cs.toronto.edu/~bianca/papers/sigmetrics09.pdf) – DRAM Errors in the Wild: A Large-Scale Field Study.
- What every programmer should know about memory (https://lwn.net/Articles/250967/) by Ulrich Drepper, continued with: CPU Caches (http://lwn.net/Articles/252125/), Virtual memory (http://lwn.net/Articles/253361/), NUMA systems (http://lwn.net/Articles/254445/), What programmers can do - cache optimization (http://lwn.net/Articles/255364/), What programmers can do - multi-threaded optimizations (http://lwn.net/Articles/256433/), Memory performance tools (http://lwn.net/Articles/257209/), Future technologies (https://lwn.net/Articles/258154/)

Synchronous dynamic random-access memory

Synchronous dynamic random access memory (**SDRAM**) is dynamic random access memory (DRAM) that is synchronized with the system bus. Classic DRAM has an asynchronous interface, which means that it responds as quickly as possible to changes in control inputs. SDRAM has a synchronous interface, meaning that it waits for a clock signal before responding to control inputs and is therefore synchronized with the computer's system bus. The clock is used to drive an internal finite state machine that pipelines incoming commands. The data storage area is divided into several *banks*, allowing the chip to work on several memory access commands at a time, interleaved among the separate banks. This allows higher data access rates than an asynchronous DRAM.

Pipelining means that the chip can accept a new command before it has finished processing the previous one. In a pipelined write, the write command can be immediately followed by another command, without waiting for the data to be written to the memory array. In a pipelined read, the requested data appears after a fixed number of clock cycles after the read command (latency), clock cycles during which additional commands can be sent. (This delay is called the *latency* and is an important performance parameter to consider when purchasing SDRAM for a computer.)

SDRAM is widely used in computers; from the original SDRAM, further generations of DDR (or *DDR1*) and then DDR2 and DDR3 have entered the mass market, with DDR4 currently being designed and anticipated to be available in 2013.

SDRAM history

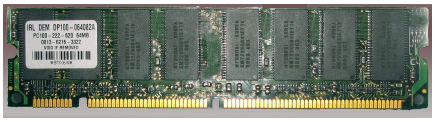

Eight SDRAM ICs on a PC100 DIMM package.

Although the concept of synchronous DRAM has been known since at least the 1970s and was used with early Intel processors, it was only in 1993 that SDRAM began its path to universal acceptance in the electronics industry. In 1993, Samsung introduced its KM48SL2000 synchronous DRAM, and by 2000, SDRAM had replaced virtually all other types of DRAM in modern computers, because of its greater performance.

SDRAM latency is not inherently lower (faster) than asynchronous DRAM. Indeed, early SDRAM was somewhat slower than contemporaneous burst EDO DRAM due to the additional logic. The benefits of SDRAM's internal buffering come from its ability to interleave operations to multiple banks of memory, thereby increasing effective bandwidth.

Today, virtually all SDRAM is manufactured in compliance with standards established by JEDEC, an electronics industry association that adopts open standards to facilitate interoperability of electronic components. JEDEC formally adopted its first SDRAM standard in 1993 and subsequently adopted other SDRAM standards, including those for DDR, DDR2 and DDR3 SDRAM.

SDRAM is also available in registered varieties, for systems that require greater scalability such as servers and workstations.

As of 2007, 168-pin SDRAM DIMMs are not used in new PC systems, and 184-pin DDR memory has been mostly superseded. DDR2 SDRAM is the most common type used with new PCs, and DDR3 motherboards and memory are widely available, and less expensive than still-popular DDR2 products.

Today, the world's largest manufacturers of SDRAM include: Samsung Electronics, Panasonic, Micron Technology, and Hynix.

SDRAM timing

There are several limits on DRAM performance. Most noted is the read cycle time, the time between successive read operations to an open row. This time decreased from 10 ns for 100 MHz SDRAM to 5 ns for DDR-400, but has remained relatively unchanged through DDR2-800 and DDR3-1600 generations. However, by operating the interface circuitry at increasingly higher multiples of the fundamental read rate, the achievable bandwidth has increased rapidly.

Another limit is the CAS latency, the time between supplying a column address and receiving the corresponding data. Again, this has remained relatively constant at 10–15 ns through the last few generations of DDR SDRAM.

In operation, CAS latency is a specific number of clock cycles programmed into the SDRAM's mode register and expected by the DRAM controller. Any value may be programmed, but the SDRAM will not operate correctly if it is too low. At higher clock rates, the useful CAS latency in clock cycles naturally increases. 10–15 ns is 2–3 cycles (CL2–3) of the 200 MHz clock of DDR-400 SDRAM, CL4-6 for DDR2-800, and CL8-12 for DDR3-1600. Slower clock cycles will naturally allow lower numbers of CAS latency cycles.

SDRAM modules have their own timing specifications, which may be slower than those of the chips on the module. When 100 MHz SDRAM chips first appeared, some manufacturers sold "100 MHz" modules that could not reliably operate at that clock rate. In response, Intel published the PC100 standard, which outlines requirements and guidelines for producing a memory module that can operate reliably at 100 MHz. This standard was widely influential, and the term "PC100" quickly became a common identifier for 100 MHz SDRAM modules, and modules are now commonly designated with "PC"-prefixed numbers (PC66, PC100 or PC133 - although the actual meaning of the numbers has changed).

SDR SDRAM

Originally simply known as *SDRAM*, **single data rate** SDRAM can accept one command and transfer one word of data per clock cycle. Typical clock frequencies are 100 and 133 MHz. Chips are made with a variety of data bus sizes (most commonly 4, 8 or 16 bits), but chips are generally assembled into 168-pin DIMMs that read or write 64 (non-ECC) or 72 (ECC) bits at a time.

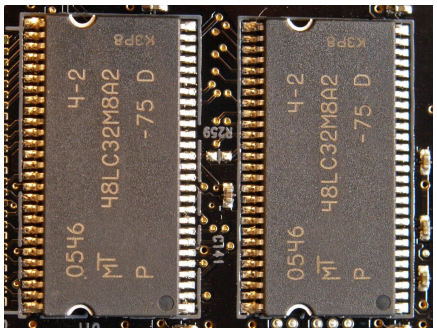

64 MB sound memory of Sound Blaster X-Fi Fatal1ty Pro uses two Micron 48LC32M8A2-75 C SDRAM chips working at 133 MHz (7.5 ns) 8-bit wide[1]

Use of the data bus is intricate and thus requires a complex DRAM controller circuit. This is because data written to the DRAM must be presented in the same cycle as the write command, but reads produce output 2 or 3 cycles after the read command. The DRAM controller must ensure that the data bus is never required for a read and a write at the same time.

Typical SDR SDRAM clock rates are 66, 100, and 133 MHz (periods of 15, 10, and 7.5 ns). Clock rates up to 150 MHz were available for performance enthusiasts.

SDRAM control signals

All commands are timed relative to the rising edge of a clock signal. In addition to the clock, there are 6 control signals, mostly active low, which are sampled on the rising edge of the clock:

- **CKE** Clock Enable. When this signal is low, the chip behaves as if the clock has stopped. No commands are interpreted and command latency times do not elapse. The state of other control lines is not relevant. The effect of this signal is actually delayed by one clock cycle. That is, the current clock cycle proceeds as usual, but the following clock cycle is ignored, except for testing the CKE input again. Normal operations resume on the rising edge of the clock after the one where CKE is sampled high.
  Put another way, all other chip operations are timed relative to the rising edge of a masked clock. The masked clock is the logical AND of the input clock and the state of the CKE signal during the previous rising edge of the input clock.
- **/CS** Chip Select. When this signal is high, the chip ignores all other inputs (except for CKE), and acts as if a NOP command is received.
- **DQM** Data Mask. (The letter Q appears because, following digital logic conventions, the data lines are known as "DQ" lines.) When high, these signals suppress data I/O. When accompanying write data, the data is not actually written to the DRAM. When asserted high two cycles before a read cycle, the read data is not output from the chip. There is one DQM line per 8 bits on a x16 memory chip or DIMM.

Command signals

- **/RAS** Row Address Strobe. Despite the name, this is *not* a strobe, but rather simply a command bit. Along with /CAS and /WE, this selects one of 8 commands.
- **/CAS** Column Address Strobe. Despite the name, this is *not* a strobe, but rather simply a command bit. Along with /RAS and /WE, this selects one of 8 commands.
- **/WE** Write enable. Along with /RAS and /CAS, this selects one of 8 commands. This generally distinguishes read-like commands from write-like commands.

Bank Selection (BAn)

SDRAM devices are internally divided into either 2, 4 or 8 independent internal data banks. One to three Bank Address inputs (BA0, BA1 and BA2) are used to select which bank a command is directed toward.

Addressing (A10/An)

Many commands also use an address presented on the address input pins. Some commands, which either do not use an address, or present a column address, also use A10 to select variants.

Commands

The various DDRx SDRAM standards use essentially the same commands, with minor additions. The selection between additional mode registers (called Extended Mode registers) as added by later standards are distinguished using the Bank Address bits.

The commands are defined as follows:

/CS	/RAS	/CAS	/WE	BA*n*	A10	A*n*	Command
H	x	x	x	x	x	x	Command inhibit (No operation)
L	H	H	H	x	x	x	No operation
L	H	H	L	x	x	x	Burst Terminate: stop a burst read or burst write in progress.
L	H	L	H	bank	L	column	Read: Read a burst of data from the currently active row.
L	H	L	H	bank	H	column	Read with auto precharge: As above, and precharge (close row) when done.
L	H	L	L	bank	L	column	Write: Write a burst of data to the currently active row.
L	H	L	L	bank	H	column	Write with auto precharge: As above, and precharge (close row) when done.
L	L	H	H	bank	row		Active (activate): open a row for Read and Write commands.
L	L	H	L	bank	L	x	Precharge: Deactivate (close) the current row of selected bank.
L	L	H	L	x	H	x	Precharge all: Deactivate (close) the current row of all banks.
L	L	L	H	x	x	x	Auto refresh: Refresh one row of each bank, using an internal counter. All banks must be precharged.
L	L	L	L	0 0	mode		Load mode register: A0 through A9 are loaded to configure the DRAM chip. The most significant settings are CAS latency (2 or 3 cycles) and burst length (1, 2, 4 or 8 cycles)

SDRAM construction and operation

A 512 MB SDRAM DIMM (which contains 512 MiB (mebibytes) = 512 × 2^{20} bytes = 536,870,912 bytes exactly), might be made of 8 or 9 SDRAM chips, each containing 512 Mbit of storage, and each one contributing 8 bits to the DIMM's 64- or 72-bit width. A typical 512 Mbit SDRAM chip internally contains 4 independent 16 MB (MiB) memory banks. Each bank is an array of 8,192 rows of 16,384 bits each. A bank is either idle, active, or changing from one to the other.

The Active command activates an idle bank. It presents a 2-bit bank address (BA0 BA1) and a 13-bit row address (A0 A12), and causes a read of that row into the bank's array of all 16,384 column sense amplifiers. This is also known as "opening" the row. This operation has the side effect of refreshing the dynamic (capacitive) memory storage cells of that row.

Once the row has been activated or "opened", Read and Write commands are possible to that row. Activation requires a minimum amount of time, called the row-to-column delay, or t_{RCD} before reads or writes to it may occur. This time, rounded up to the next multiple of the clock period, specifies the minimum number of wait cycles between an Active command, and a Read or Write command. During these wait cycles, additional commands may be sent to other banks; because each bank operates completely independently.

Both Read and Write commands require a column address. Because each chip accesses 8 bits of data at a time, there are 2048 possible column addresses thus requiring only 11 address lines (A0 A9, A11).

When a Read command is issued, the SDRAM will produce the corresponding output data on the DQ lines in time for the rising edge of the clock 2 or 3 clock cycles later (depending on the configured CAS latency). Subsequent words of the burst will be produced in time for subsequent rising clock edges.

A Write command is accompanied by the data to be written driven on to the DQ lines during the same rising clock edge. It is the duty of the memory controller to ensure that the SDRAM is not driving read data on to the DQ lines at the same time that it needs to drive write data on to those lines. This can be done by waiting until a read burst has finished, by terminating a read burst, or by using the DQM control line.

When the memory controller needs to access a different row, it must first return that bank's sense amplifiers to an idle state, ready to sense the next row. This is known as a "precharge" operation, or "closing" the row. A precharge may be commanded explicitly, or it may be performed automatically at the conclusion of a read or write operation.

Again, there is a minimum time, the row precharge delay, t_{RP}, which must elapse before that bank is fully idle and it may receive another activate command.

Although refreshing a row is an automatic side effect of activating it, there is a minimum time for this to happen, which requires a minimum row access time t_{RAS} delay between an Active command opening a row, and the corresponding precharge command closing it. This limit is usually dwarfed by desired read and write commands to the row, so its value has little effect on typical performance.

Command interactions

The no operation command is always permitted.

The load mode register command requires that all banks be idle, and a delay afterward for the changes to take effect.

The auto refresh command also requires that all banks be idle, and takes a refresh cycle time t_{RFC} to return the chip to the idle state. (This time is usually equal to $t_{RCD}+t_{RP}$.)

The only other command that is permitted on an idle bank is the active command. This takes, as mentioned above, t_{RCD} before the row is fully open and can accept read and write commands.

When a bank is open, there are four commands permitted: read, write, burst terminate, and precharge. Read and write commands begin bursts, which can be interrupted by following commands.

Interrupting a read burst

A read, burst terminate, or precharge command may be issued at any time after a read command, and will interrupt the read burst after the configured CAS latency. So if a read command is issued on cycle 0, another read command is issued on cycle 2, and the CAS latency is 3, then the first read command will begin bursting data out during cycles 3 and 4, then the results from the second read command will appear beginning with cycle 5.

If the command issued on cycle 2 were burst terminate, or a precharge of the active bank, then no output would be generated during cycle 5.

Although the interrupting read may be to any active bank, a precharge command will only interrupt the read burst if it is to the same bank or all banks; a precharge command to a different bank will not interrupt a read burst.

To interrupt a read burst by a write command is possible, but more difficult. It can be done, if the DQM signal is used to suppress output from the SDRAM so that the memory controller may drive data over the DQ lines to the SDRAM in time for the write operation. Because the effects of DQM on read data are delayed by 2 cycles, but the effects of DQM on write data are immediate, DQM must be raised (to mask the read data) beginning at least two cycles before write command, but must be lowered for the cycle of the write command (assuming you want the write command to have an effect).

Doing this in only two clock cycles requires careful coordination between the time the SDRAM takes to turn off its output on a clock edge and the time the data must be supplied as input to the SDRAM for the write on the following clock edge. If the clock frequency is too high to allow sufficient time, three cycles may be required.

If the read command includes auto-precharge, the precharge begins the same cycle as the interrupting command.

SDRAM burst ordering

A modern microprocessor with a cache will generally access memory in units of cache lines. To transfer a 64-byte cache line requires 8 consecutive accesses to a 64-bit DIMM, which can all be triggered by a single read or write command by configuring the SDRAM chips, using the mode register, to perform 8-word bursts.

A cache line fetch is typically triggered by a read from a particular address, and SDRAM allows the "critical word" of the cache line to be transferred first. ("Word" here refers to the width of the SDRAM chip or DIMM, which is 64 bits for a typical DIMM.) SDRAM chips support two possible conventions for the ordering of the remaining words in the cache line.

Bursts always access an aligned block of BL consecutive words beginning on a multiple of BL. So, for example, a 4-word burst access to any column address from 4 to 7 will return words 4 7. The ordering, however, depends on the requested address, and the configured burst type option: sequential or interleaved. Typically, a memory controller will require one or the other.

When the burst length is 1 or 2, the burst type does not matter. For a burst length of 1, the requested word is the only word accessed. For a burst length of 2, the requested word is accessed first, and the other word in the aligned block is accessed second. This is the following word if an even address was specified, and the previous word if an odd address was specified.

For the sequential burst mode, later words are accessed in increasing address order, wrapping back to the start of the block when the end is reached. So, for example, for a burst length of 4, and a requested column address of 5, the words would be accessed in the order 5-6-7-4. If the burst length were 8, the access order would be 5-6-7-0-1-2-3-4. This is done by adding a counter to the column address, and ignoring carries past the burst length.

The interleaved burst mode computes the address using an exclusive or operation between the counter and the address. Using the same starting address of 5, a 4-word burst would return words in the order 5-4-7-6. An 8-word burst would be 5-4-7-6-1-0-3-2. Although more confusing to humans, this can be easier to implement in hardware, and is preferred by Intel microprocessors.

If the requested column address is at the start of a block, both burst modes return data in the same sequential sequence 0-1-2-3-4-5-6-7. The difference only matters if fetching a cache line from memory in critical-word-first order.

SDRAM mode register

Single data rate SDRAM has a single 10-bit programmable mode register. Later double-data-rate SDRAM standards add additional mode registers, addressed using the bank address pins. For SDR SDRAM, the bank address pins and address lines A10 and above are ignored, but should be zero during a mode register write.

The bits are M9 through M0, presented on address lines A9 through A0 during a load mode register cycle.

1. M9: Write burst mode. If 0, writes use the read burst length and mode. If 1, all writes are non-burst (single location).
2. M8, M7: Operating mode. Reserved, and must be 00.
3. M6, M5, M4: CAS latency. Generally only 010 (CL2) and 011 (CL3) are legal. Specifies the number of cycles between a read command and data output from the chip. The chip has a fundamental limit on this value in nanoseconds; during initialization, the memory controller must use its knowledge of the clock frequency to translate that limit into cycles.
4. M3: Burst type. 0 - requests sequential burst ordering, while 1 requests interleaved burst ordering.
5. M2, M1, M0: Burst length. Values of 000, 001, 010 and 011 specify a burst size of 1, 2, 4 or 8 words, respectively. Each read (and write, if M9 is 0) will perform that many accesses, unless interrupted by a burst stop or other command. A value of 111 specifies a full-row burst. The burst will continue until interrupted. Full-row bursts are only permitted with the sequential burst type.

Later (double data rate) SDRAM standards use more mode register bits, and provide additional mode registers called Extended Mode registers. The register number is encoded on the bank address pins during the Load Mode Register command. For example, DDR2 SDRAM has a 13-bit Mode Register, a 13-bit Extended Mode Register #1 (EMR1), and a 5-bit Extended Mode Register #2 (EMR2).

Auto refresh

It is possible to refresh a RAM chip by opening and closing (activating and precharging) each row in each bank. However, to simplify the memory controller, SDRAM chips support an "auto refresh" command, which performs these operations to one row in each bank simultaneously. The SDRAM also maintains an internal counter, which iterates over all possible rows. The memory controller must simply issue a sufficient number of auto refresh commands (one per row, 4096 in the example we have been using) every refresh interval (t_{REF} = 64 ms is a common value). All banks must be idle (closed, precharged) when this command is issued.

Low power modes

As mentioned, the clock enable (CKE) input can be used to effectively stop the clock to an SDRAM. The CKE input is sampled each rising edge of the clock, and if it is low, the following rising edge of the clock is ignored for all purposes other than checking CKE. As long as CKE is low, it is permissible to change the clock rate, or even stop the clock entirely.

If CKE is lowered while the SDRAM is performing operations, it simply "freezes" in place until CKE is raised again.

If the SDRAM is idle (all banks precharged, no commands in progress) when CKE is lowered, the SDRAM automatically enters power-down mode, consuming minimal power until CKE is raised again. This must not last longer than the maximum refresh interval t_{REF}, or memory contents may be lost. It is legal to stop the clock entirely during this time for additional power savings.

Finally, if CKE is lowered at the same time as an auto-refresh command is sent to the SDRAM, the SDRAM enters self-refresh mode. This is like power down, but the SDRAM uses an on-chip timer to generate internal refresh cycles as necessary. The clock may be stopped during this time. While self-refresh mode consumes slightly more power than power-down mode, it allows the memory controller to be disabled entirely, which commonly more than makes up the difference.

SDRAM designed for battery-powered devices offers some additional power-saving options. One is temperature-dependent refresh; an on-chip temperature sensor reduces the refresh rate at lower temperatures, rather than always running it at the worst-case rate. Another is selective refresh, which limits self-refresh to a portion of the DRAM array. The fraction which is refreshed is configured using an extended mode register. The third, implemented in Mobile DDR (LPDDR) and LPDDR2 is "deep power down" mode, which invalidates the memory and requires a full reinitialization to exit from. This is activated by sending a "burst terminate" command while lowering CKE.

Generations of SDRAM

SDR SDRAM (Single Data Rate synchronous DRAM)

This type of SDRAM is slower than the DDR variants, because only one word of data is transmitted per clock cycle (single data rate). But this type is also faster than its predecessors EDO-RAM and FPM-RAM which took typically 2 or 3 clocks to transfer one word of data.

DDR(1) SDRAM

While the access latency of DRAM is fundamentally limited by the DRAM array, DRAM has very high potential bandwidth because each internal read is actually a row of many thousands of bits. To make more of this bandwidth available to users, a double data rate interface was developed. This uses the same commands, accepted once per cycle, but reads or writes two words of data per clock cycle. The DDR interface accomplishes this by reading and writing data on both the rising and falling edges of the clock signal. In addition, some minor changes to the SDR interface timing were made in hindsight, and the supply voltage was reduced from 3.3 to 2.5 V. As a result, DDR SDRAM is not backwards compatible with SDR SDRAM.

DDR SDRAM (sometimes called *DDR1* for greater clarity) doubles the minimum read or write unit; every access refers to at least two consecutive words.

Typical DDR SDRAM clock rates are 133, 166 and 200 MHz (7.5, 6, and 5 ns/cycle), generally described as DDR-266, DDR-333 and DDR-400 (3.75, 3, and 2.5 ns per beat). Corresponding 184-pin DIMMs are known as PC-2100, PC-2700 and PC-3200. Performance up to DDR-550 (PC-4400) is available for a price.

DDR2 SDRAM

DDR2 SDRAM is very similar to DDR SDRAM, but doubles the minimum read or write unit again, to 4 consecutive words. The bus protocol was also simplified to allow higher performance operation. (In particular, the "burst terminate" command is deleted.) This allows the bus rate of the SDRAM to be doubled without increasing the clock rate of internal RAM operations; instead, internal operations are performed in units 4 times as wide as SDRAM. Also, an extra bank address pin (BA2) was added to allow 8 banks on large RAM chips.

Typical DDR2 SDRAM clock rates are 200, 266, 333 or 400 MHz (periods of 5, 3.75, 3 and 2.5 ns), generally described as DDR2-400, DDR2-533, DDR2-667 and DDR2-800 (periods of 2.5, 1.875, 1.5 and 1.25 ns). Corresponding 240-pin DIMMS are known as PC2-3200 through PC2-6400. DDR2 SDRAM is now available at a clock rate of 533 MHz generally described as DDR2-1066 and the corresponding DIMMs are known as PC2-8500 (also named PC2-8600 depending on the manufacturer). Performance up to DDR2-1250 (PC2-10000) is available for a price.

Note that because internal operations are at 1/2 the clock rate, DDR2-400 memory (internal clock rate 100 MHz) has somewhat higher latency than DDR-400 (internal clock rate 200 MHz).

DDR3 SDRAM

DDR3 continues the trend, doubling the minimum read or write unit to 8 consecutive words. This allows another doubling of bandwidth and external bus rate without having to change the clock rate of internal operations, just the width. To maintain 800–1600 M transfers/s (both edges of a 400–800 MHz clock), the internal RAM array has to perform 100–200 M fetches per second.

Again, with every doubling, the downside is the increased latency. As with all DDR SDRAM generations, commands are still restricted to one clock edge and command latencies are given in terms of clock cycles, which are half the speed of the usually quoted transfer rate (a CAS latency of 8 with DDR3-800 is 8/(400 MHz) = 20 ns, exactly the same latency of CAS2 on PC100 SDR SDRAM).

DDR3 memory chips are being made commercially,[2] and computer systems using them were available from the second half of 2007,[3] with significant usage from 2008 onwards.[4] Initial clock rates were 400 and 533 MHz, which are described as DDR3-800 and DDR3-1066 (PC3-6400 and PC3-8500 modules), but 667 and 800 MHz, described as DDR3-1333 and DDR3-1600 (PC3-10600 and PC3-12800 modules) are now common.[5] Performance up to DDR3-2800 (PC3 22400 modules) are available for a price.[6]

DDR4 SDRAM

DDR4 SDRAM is the successor to DDR3 SDRAM. It was revealed at the Intel Developer Forum in San Francisco in 2008, and is due to be released to market during 2011. The timing has varied considerably during its development - it was originally expected to be released in 2012,[7] and later (during 2010) expected to be released in 2015,[8] before samples were announced in early 2011 and manufacturers began to announce that commercial production and release to market was anticipated in 2012. DDR4 is expected to reach mass market adoption around 2015, which is comparable with the approximately 5 years taken for DDR3 to achieve mass market transition over DDR2.

The new chips are expected to run at 1.2 V or less,[9][10] versus the 1.5 V of DDR3 chips, and have in excess of 2 billion data transfers per second. They are expected to be introduced at frequency rates of 2133 MHz, estimated to rise to a potential 4266 MHz[11] and lowered voltage of 1.05 V[12] by 2013.

DDR4 will *not* double the internal prefetch width again, but will use the same 8*n* prefetch as DDR3.[13] Thus, it will be necessary to interleave reads from several banks to keep the data bus busy.

In February 2009, Samsung validated 40 nm DRAM chips, considered a "significant step" towards DDR4 development[14] since as of 2009, current DRAM chips were only beginning to migrate to a 50 nm process.[15] In January 2011, Samsung announced the completion and release for testing of a 30 nm 2 GB DDR4 DRAM module. It has a maximum bandwidth of 2.13 Gbit/s at 1.2 V, uses pseudo open drain technology and draws 40% less power than an equivalent DDR3 module.[16][17]

Feature map

Type	Feature changes
SDRAM	V_{cc} = 3.3 V Signal: LVTTL
DDR1	Access is ≥2 words Double clocked V_{cc} = 2.5 V 2.5 - 7.5 ns per cycle Signal: SSTL_2 (2.5V)[18]
DDR2	Access is ≥4 words "Burst terminate" removed 4 units used in parallel 1.25 - 5 ns per cycle Internal operations are at 1/2 the clock rate. Signal: SSTL_18 (1.8V)[18]
DDR3	Access is ≥8 words Signal: SSTL_15 (1.5V)[18] Much longer CAS latencies
DDR4	V_{cc} ≤ 1.2 V point-to-point (single module per channel)

Failed successors

In addition to DDR, there were several other proposed memory technologies to succeed SDR SDRAM.

Rambus DRAM (RDRAM)

RDRAM was a proprietary technology that competed against DDR. Its relatively high price and disappointing performance (resulting from high latencies and a narrow 16-bit data channel versus DDR's 64 bit channel) caused it to lose the race to succeed SDR DRAM.

Synchronous-Link DRAM (SLDRAM)

SLDRAM boasted higher performance and competed against RDRAM. It was developed during the late 1990s by the SLDRAM Consortium, which consisted of about 20 major computer industry manufacturers. It is an open standard and does not require licensing fees. The specifications called for a 64-bit bus running at a 200, 300 or 400 MHz clock frequency. This is achieved by all signals being on the same line and thereby avoiding the synchronization time of multiple lines. Like DDR SDRAM, SLDRAM uses a double-pumped bus, giving it an effective speed of 400,[19] 600,[20] or 800 MT/s.

SLDRAM used an 11-bit command bus (10 command bits CA9:0 plus one start-of-command FLAG line) to transmit 40-bit command packets on 4 consecutive edges of a differential command clock (CCLK/CCLK#). Unlike SDRAM, there were no per-chip select signals; each chip was assigned an ID when reset, and the command contained the ID of the chip that should process it. Data was transferred in 4- or 8-word bursts across an 18-bit (per chip) data bus, using one of two differential data clocks (DCLK0/DCLK0# and DCLK1/DCLK1#). Unlike standard SDRAM, the clock was generated by the data source (the SLDRAM chip in the case of a read operation) and transmitted in the same direction as the data, greatly reducing data skew. To avoid the need for a pause when the source of the DCLK changes, each command specified which DCLK pair it would use.[21]

The basic read/write command consisted of (beginning with CA9 of the first word):

SLDRAM Read, write or row op request packet

FLAG	CA9	CA8	CA7	CA6	CA5	CA4	CA3	CA2	CA1	CA0
1	ID8	Device ID							ID0	CMD5
0	Command code				CMD0	Bank			Row	
0	Row (continued)									0
0	0	0	0	Column						

- 9 bits of device ID
- 6 bits of command
- 3 bits of bank address
- 10 or 11 bits of row address
- 5 or 4 bits spare for row or column expansion
- 7 bits of column address

Individual devices had 8-bit IDs. The 9th bit of the ID sent in commands was used to address multiple devices. Any aligned power-of-2 sized group could be addressed. If the transmitted msbit was set, all least-significant bits up to and including the least-significant 0 bit of the transmitted address were ignored for "is this addressed to me?" purposes. (If the ID8 bit is actually considered less significant than ID0, the unicast address matching becomes a special case of this pattern.)

A read/write command had the msbit clear:

- CMD5=0
- CMD4=1 to open (activate) the specified row; CMD4=0 to use the currently open row
- CMD3=1 to transfer an 8-word burst; CMD3=0 for a 4-word burst
- CMD2=1 for a write, CMD2=0 for a read
- CMD1=1 to close the row after this access; CMD1=0 to leave it open
- CMD0 selects the DCLK pair to use (DCLK1 or DCLK0)

A notable omission from the specification was per-byte write enables; it was designed for systems with caches and ECC memory, which always write in multiples of a cache line.

Additional commands (with CMD5 set) opened and closed rows without a data transfer, performed refresh operations, read or wrote configuration registers, and performed other maintenance operations. Most of these commands supported an additional 4-bit sub-ID (sent as 5 bits, using the same multiple-destination encoding as the primary ID) which could be used to distinguish devices that were assigned the same primary ID because they were connected in parallel and always read/written at the same time.

There were a number of 8-bit control registers and 32-bit status registers to control various device timing parameters.

Virtual Channel Memory (VCM) SDRAM

VCM was a proprietary type of SDRAM that was designed by NEC, but released as an open standard with no licensing fees. It is pin-compatible with standard SDRAM, but the commands are different.

The technology was a potential competitor of RDRAM because VCM was not nearly as expensive as RDRAM was. A Virtual Channel Memory (VCM) module is mechanically and electrically compatible with standard SDRAM, so support for both is possible if only the memory controller is capable of it. In the late 1990s, a number of PC northbridge chipsets (such as the popular VIA KX133 and KT133) included VCSDRAM support.

VCM inserts an SRAM cache of 16 "channel" buffers, each 1/4 row "segment" in size, between DRAM banks' sense amplifier rows and the data I/O pins. "Prefetch" and "Restore" command, unique to VCSDRAM, copy data between the DRAM's sense amplifier row and the channel buffers, while the equivalent of SDRAM's Read and Write

commands specify a channel number to access.

Reads and writes may thus be performed independent of the currently active state of the DRAM array, with the equivalent of 4 full DRAM rows being "open" for access at a time. This is an improvement over the 2 open rows possible in a standard 2-bank SDRAM. (There is actually a 17th "dummy channel" used for some operations.)

To read from VCSDRAM, after the Active command, a "Prefetch" command is required to copy data from the sense amplifier array to the channel SDRAM. This command specifies a bank, 2 bits of column address (to select the segment of the row), and 4 bits of channel number. Once this is performed, the DRAM array may be precharged while read commands to the channel buffer continue.

To write, first the data is written to a channel buffer (typically previous initialized using a Prefetch command), then a Restore command, with the same parameters as the Prefetch command, copies a segment of data from the channel to the sense amplifier array.

Unlike a normal SDRAM write, which must be performed to an active (open) row, the VCSDRAM bank must be precharged (closed) when the Restore command is issued. An Active command immediately after the Restore command specifies the DRAM row completes the write to the DRAM array.

There is, in addition, a 17th "dummy channel" which allows writes to the currently open row. It may not be Read from, but may be Prefetched to, Written to, and Restored to the sense amplifier array.[22][23]

Although normally a segment is Restored to the same memory address as it was Prefetched from, the channel buffers may also be used for very efficient copying or clearing of large, aligned memory blocks.

(The use of quarter-row segments is driven by the fact that DRAM cells are narrower than SRAM cells. The SRAM bits are designed to be 4 DRAM bits wide, and are conveniently connected to one of the 4 DRAM bits they straddle.)

Additional commands prefetch a pair of segments to a pair of channels, and an optional command combines prefetch, read, and precharge to reduce the overhead of random reads.

Virtual Channel SDRAM commands[24]

/CS	/RAS	/CAS	/WE	BA	A12–11	A10	A9	A8	A7	A6	A5	A4–2	A1–0	Command
H	x	x	x	x	x	x	x	x	x	x	x	x	x	Command inhibit (No operation)
L	H	H	H	x	x	x	x	x	x	x	x	x	x	No operation
L	H	H	L	bank	channel	AP	channel		L	L	L	x	segment	Prefetch (auto-precharge if AP=H)
L	H	H	L	bank	x	x	x		L	L	H	x	segment	Prefetch to dummy
L	H	H	L	bank	channel	AP	channel		L	H	x	x	segment	Pair prefetch (auto-precharge if AP=H)
L	H	H	L	bank	channel	AP	channel		H	L	x	x	segment	Restore (auto-precharge if AP=H)
L	H	H	L	bank	x	L	x		H	H	x	x	x	Precharge bank
L	H	H	L	x	x	H	x		H	H	x	x	x	Precharge all banks
L	H	L	H	x	channel	x	channel		column					Read channel
L	H	L	L	L	channel	x	channel		column					Write channel
L	H	L	L	H	x	x	x	AR	column					Write dummy channel (auto-restore if AR=H)
L	L	H	H	bank	row									Bank activate
L	L	H	L	seg	channel	seg	channel		column					Prefetch read with auto-precharge (optional)
L	L	L	H	x	x	x	x	x	x	x	x	x	x	Auto refresh
L	L	L	H	reg	mode register data									Mode register set

The above are the JEDEC-standardized commands. Earlier chips did not support the dummy channel or pair prefetch, and used a different encoding for precharge.

A 13-bit address bus, as illustrated here, is suitable for a device up to 128 Mbit. It would have two banks, each containing 8192 rows and 8192 columns. Thus, row addresses are 13 bits, segment addresses are 2 bits, and 8 column address bits are required to select one byte from the 2048 bits (256 bytes) in a segment.

References

[1] "SDRAM Part Catalog" (http://www.micron.com/products/dram/sdram/partlist). . 070928 micron.com

[2] "What is DDR memory?" (http://www.simmtester.com/page/news/showpubnews.asp?num=145). .

[3] Thomas Soderstrom (June 5, 2007). "Pipe Dreams: Six P35-DDR3 Motherboards Compared" (http://www.tomshardware.com/2007/06/05/pipe_dreams_six_p35-ddr3_motherboards_compared/). Tom's Hardware. .

[4] "AMD to Adopt DDR3 in Three Years" (http://news.softpedia.com/news/AMD-to-Adopt-DDR3-in-Three-Years-13486.shtml). .

[5] Wesly Fink (July 20, 2007). "Super Talent & TEAM: DDR3-1600 Is Here!" (http://www.anandtech.com/printarticle.aspx?i=3045). Anandtech. .

[6] Jennifer Johnson (24 April 2012). "G.SKILL Announces DDR3 Memory Kit For Ivy Bridge" (http://hothardware.com/News/GSKILL-Announces-DDR3-Memory-Kit-For-Ivy-Bridge/). .

[7] DDR4 PDF page 23 (http://intel.wingateweb.com/US08/published/sessions/MASS006/SF08_MASS006_100s.pdf)

[8] http://www.semiaccurate.com/2010/08/16/ddr4-not-expected-until-2015/

[9] Looking forward to DDR4 (http://www.pcpro.co.uk/news/220257/idf-ddr3-wont-catch-up-with-ddr2-during-2009.html)

[10] DDR3 successor (http://www.heise-online.co.uk/news/IDF-DDR4-the-successor-to-DDR3-memory--/111367)

[11] "Next-Generation DDR4 Memory to Reach 4.266GHz - Report" (http://www.xbitlabs.com/news/memory/display/20100816124343_Next_Generation_DDR4_Memory_to_Reach_4_266GHz_Report.html). Xbitlabs.com. August 16, 2010. . Retrieved 2011-01-03.

[12] "IDF: DDR4 memory targeted for 2012" (http://www.hardware-infos.com/news.php?news=2332) (in German). hardware-infos.com. . Retrieved 2009-06-16. *English translation* (http://translate.google.com/translate?hl=en&sl=de&u=http://www.hardware-infos.com/news.php?news=2332&ei=bi44Sv_wBouZjAfVzYyjDQ&sa=X&oi=translate&resnum=1&ct=result&prev=/search?q=http://www.hardware-infos.com/news.php%3Fnews%3D2332&hl=en&safe=off&num=100)

[13] "JEDEC Announces Key Attributes of Upcoming DDR4 Standard" (http://www.jedec.org/news/pressreleases/jedec-announces-key-attributes-upcoming-ddr4-standard) (Press release). JEDEC. 2011-08-22. . Retrieved 2011-01-06.

[14] Gruener, Wolfgang (February 4, 2009). "Samsung hints to DDR4 with first validated 40 nm DRAM" (http://www.tgdaily.com/content/view/41316/139/). tgdaily.com. . Retrieved 2009-06-16.

[15] Jansen, Ng (January 20, 2009). "DDR3 Will be Cheaper, Faster in 2009" (http://www.dailytech.com/DDR3+Will+be+Cheaper+Faster+in+2009/article13977.htm). dailytech.com. . Retrieved 2009-06-17.

[16] "Samsung Develops Industry's First DDR4 DRAM, Using 30nm Class Technology" (http://www.samsung.com/us/business/semiconductor/newsView.do?news_id=1202). Samsung. 2011-01-04. . Retrieved 2011-03-13.

[17] http://www.techspot.com/news/41818-samsung-develops-ddr4-memory-up-to-40-more-efficient.html

[18] "EDA DesignLine, januari 12, 2007, The outlook for DRAMs in consumer electronics" (http://www.edadesignline.com/196900432?printableArticle=true). . 100622 edadesignline.com

[19] Dean Kent (1998-10-24), *RAM Guide: SLDRAM* (http://www.tomshardware.com/reviews/ram-guide,89-15.html), Tom's Hardware, , retrieved 2011-01-01

[20] Hyundai Electronics (1997-12-20), *HYSL8M18D600A 600 Mb/s/pin 8M x 18 SLDRAM* (http://icwic.cn/icwic/data/pdf/cd/cd011/12452.pdf) (data sheet), , retrieved 2011-12-27

[21] SLDRAM Inc. (1998-07-09), *SLD4M18DR400 400 Mb/s/pin 4M x 18 SLDRAM* (http://icwic.cn/icwic/data/pdf/cd/cd011/12407.pdf) (data sheet), pp. 32–33, , retrieved 2011-12-27

[22] Siemens Semiconductor Group, *HYB39V64x0yT 64MBit Virtual Channel SDRAM* (http://radiosxema.ru/index.php/datasheet?do=download&file=H|HYB|HYB+39V64x0yT.pdf), , retrieved 2012-07-17

[23] NEC (1999), *128M-BIT VirtualChannel™ SDRAM preliminary datasheet* (http://www.ic72.com/pdf_file/u/32271.pdf), , retrieved 2012-07-17

[24] JEDEC Standard 21-C section 3.11.5.4 (http://www.jedec.org/sites/default/files/docs/3_11_05_04R9.pdf) "VCSDRAM Specific SDRAM Functions"

DDR SDRAM

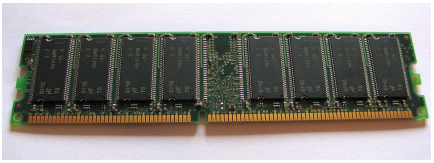

Generic DDR-266 memory in the 184-pin DIMM form

Double data rate synchronous dynamic random-access memory (DDR SDRAM) is a class of memory integrated circuits used in computers. DDR SDRAM (sometimes referred to as **DDR1 SDRAM**) has been superseded by DDR2 SDRAM and DDR3 SDRAM, neither of which is either forward or backward compatible with DDR SDRAM, meaning that DDR2 or DDR3 memory modules will not work in DDR-equipped motherboards, and vice versa.

Corsair DDR-400 memory with heat spreaders

Compared to single data rate (SDR) SDRAM, the DDR SDRAM interface makes higher transfer rates possible by more strict control of the timing of the electrical data and clock signals. Implementations often have to use schemes such as phase-locked loops and self-calibration to reach the required timing accuracy.[1][2] The interface uses double pumping (transferring data on both the rising and falling edges of the clock signal) to lower the clock frequency. One advantage of keeping the clock frequency down is that it reduces the signal integrity requirements on the circuit board connecting the memory to the controller. The name "double data rate" refers to the fact that a DDR SDRAM with a certain clock frequency achieves nearly twice the bandwidth of a SDR SDRAM running at the same clock frequency, due to this double pumping.

With data being transferred 64 bits at a time, DDR SDRAM gives a transfer rate of (memory bus clock rate) × 2 (for dual rate) × 64 (number of bits transferred) / 8 (number of bits/byte). Thus, with a bus frequency of 100 MHz, DDR SDRAM gives a maximum transfer rate of 1600 MB/s.

"Beginning in 1996 and concluding in June 2000, JEDEC developed the DDR (Double Data Rate) SDRAM specification (JESD79)."[3] JEDEC has set standards for data rates of DDR SDRAM, divided into two parts. The first specification is for memory chips, and the second is for memory modules.

Specification standards

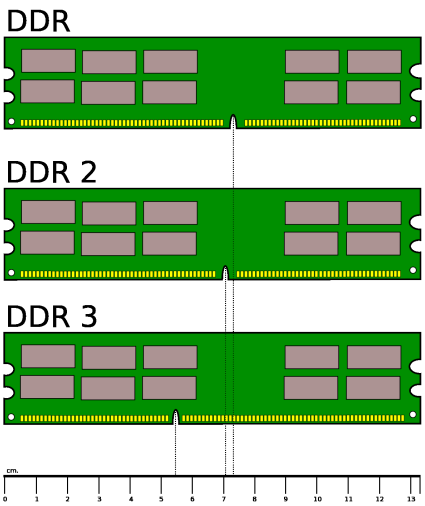

Comparison of memory modules for desktop PCs (DIMM).

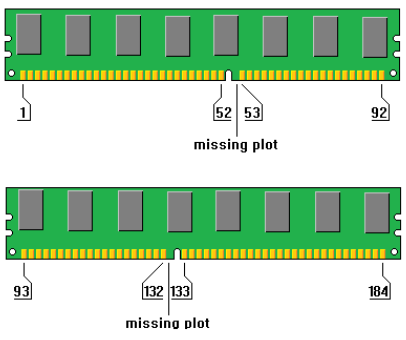

Physical DDR layout

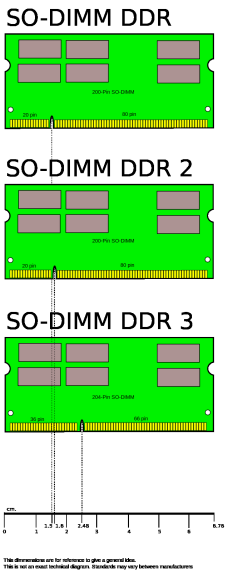

Comparison of memory modules for portable/mobile PCs (SO-DIMM).

Standard name	Memory clock (MHz)	Cycle time[4] (ns)	I/O bus clock (MHz)	Data rate (MT/s)	V_{DDQ} (V)	Module name	Peak transfer rate (MB/s)	Timings (CL-tRCD-tRP)
DDR-200	100	10	100	200	2.5±0.2	PC-1600	1600	
DDR-266	133⅓	7.5	133⅓	266⅔		PC-2100	2133⅓	2.5-3-3
DDR-333	166⅔	6	166⅔	333⅓		PC-2700	2666⅔	
DDR-400A DDR-400B DDR-400C	200	5	200	400	2.6±0.1	PC-3200	3200	2.5-3-3 3-3-3 3-4-4

Note: All above listed are specified by JEDEC as JESD79F.[5] All RAM data rates in-between or above these listed specifications are not standardized by JEDEC—often they are simply manufacturer optimizations using tighter-tolerance or overvolted chips.

The package sizes in which DDR SDRAM is manufactured are also standardized by JEDEC.

There is no architectural difference between DDR SDRAM designed for different clock frequencies, for example, PC-1600, designed to run at 100 MHz, and PC-2100, designed to run at 133 MHz. The number simply designates the data rate at which the chip is guaranteed to perform, hence DDR SDRAM is guaranteed to run at lower (*underclocking*) and can possibly run at higher (*overclocking*) clock rates than those for which it was made.[6]

DDR SDRAM modules for desktop computers, commonly called DIMMs, have 184 pins (as opposed to 168 pins on SDRAM, or 240 pins on DDR2 SDRAM), and can be differentiated from SDRAM DIMMs by the number of notches (DDR SDRAM has one, SDRAM has two). DDR SDRAM for notebook computers, SO-DIMMs, have 200 pins, which is the same number of pins as DDR2 SO-DIMMs. These two specifications are notched very similarly and care must be taken during insertion if unsure of a correct match. DDR SDRAM operates at a voltage of 2.5 V, compared to 3.3 V for SDRAM. This can significantly reduce power consumption. Chips and modules with DDR-400/PC-3200 standard have a nominal voltage of 2.6 V.

Increasing operating voltage slightly can increase maximum speed, at the cost of higher power dissipation and heating, and at the risk of malfunctioning or damage.

Many new chipsets use these memory types in multi-channel configurations.

Chip characteristics

DRAM density

Size of the chip are measured in megabits. Nearly all motherboards only recognize 1 GB modules if they contain *64M×8* chips *(low density).* If *128M×4 (high density)* 1 GB modules are used, they most likely will not work. The JEDEC standard allows *128M×4* only for slower buffered/registered modules designed specifically for some servers, but some generic manufacturers do not comply.[7]

Organization

The notation like *64M×4* means that the memory matrix has 64 million (the product of *banks* x *rows* x *columns*) 4-bit storage locations. There are *×4, ×8,* and *×16* DDR chips. The *×4* chips allow the use of advanced error correction features like Chipkill, memory scrubbing and Intel SDDC in server environments, while the *×8* and *×16* chips are somewhat less expensive. *x8* chips are mainly used in desktops/notebooks but are making entry into the server market. There are normally 4 banks and only one row can be active in each bank.

Module characteristics

Ranks

To increase memory capacity and bandwidth, chips are combined on a module. For instance, the 64-bit data bus for DIMM requires eight 8-bit chips, addressed in parallel. Multiple chips with the common address lines are called a memory rank. The term was introduced to avoid confusion with chip internal **rows** and **banks**. A memory module may bear more than one rank. The term **sides** would also be confusing because it incorrectly suggests the physical placement of chips on the module.

All ranks are connected to the same memory bus (address+data). The Chip Select signal is used to issue commands to specific rank.

Adding modules to the single memory bus creates additional electrical load on its drivers. To mitigate the resulting bus signaling rate drop and overcome the memory bottleneck, new chipsets employ the multi-channel architecture.

Capacity

Number of DRAM Devices

The number of chips is a multiple of 8 for non-ECC modules and a multiple of 9 for ECC modules. Chips can occupy one side (*single sided*) or both sides (*dual sided*) of the module. The maximum number of chips per DDR module is 36 (9×4) for ECC and 32 (8x4) for non-ECC.

ECC vs non-ECC

Modules that have error correcting code are labeled as ECC. Modules without error correcting code are labeled **non-ECC**.

Timings

CAS latency (CL), clock cycle time (t_{CK}), row cycle time (t_{RC}), refresh row cycle time (t_{RFC}), row active time ($_{tRAS}$).

Buffering

registered (or buffered) vs unbuffered

Packaging

Typically DIMM or SO-DIMM

Power consumption

A test with DDR and DDR2 RAM in 2005 found that average power consumption appeared to be of the order of 1-3W per 512MB module. Increases with clock rate, and when in use rather than idling.[8] A manufacturer has produced calculators to estimate the power used by various types of RAM.[9]

Module and chip characteristics are inherently linked.

Total module capacity is a product of one chip's capacity by the number of chips. ECC modules multiply it by 8/9 because they use one bit per byte for error correction. A module of any particular size can therefore be assembled either from 32 small chips (36 for ECC memory), or 16(18) or 8(9) bigger ones.

DDR memory bus width per channel is 64 bits (72 for ECC memory). Total module bit width is a product of bits per chip by number of chips. It also equals number of ranks (rows) multiplied by DDR memory bus width. Consequently a module with greater amount of chips or using ×8 chips instead of ×4 will have more ranks.

Example: Variations of 1 GB PC2100 Registered DDR SDRAM module with ECC

Module size (GB)	Number of chips	Chip size (Mbit)	Chip organization	Number of ranks
1	36	256	64M×4	2
1	18	512	64M×8	1
1	18	512	128M×4	+

This example compares different real-world server memory modules with a common size of 1 GB. One should definitely be careful buying 1 GB memory modules, because all these variations can be sold under one price position without stating whether they are ×4 or ×8, single or dual ranked.

There is a common belief that number of module ranks equals number of sides. As above data shows, this is not true. One can find 2-side/1-rank or 2-side/4-rank modules. One can even think of a 1-side/2-rank memory module having 16(18) chips on single side ×8 each, but it's unlikely such a module was ever produced.

History

Double data rate (DDR) SDRAM specification

From JEDEC Board Ballot JCB-99-70, and modified by numerous other Board Ballots, formulated under the cognizance of Committee JC-42.3 on DRAM Parametrics.

Standard No. 79 Revision Log:

- Release 1, June 2000
- Release 2, May 2002
- Release C, March 2003 – JEDEC Standard No. 79C.[10]

"This comprehensive standard defines all required aspects of 64Mb through 1Gb DDR SDRAMs with X4/X8/X16 data interfaces, including features, functionality, ac and dc parametrics, packages and pin assignments. This scope will subsequently be expanded to formally apply to x32 devices, and higher density devices as well."

High density vs low density

High density memory here means non-ECC 184 pin SDRAM memory.

Organization

PC3200 is DDR SDRAM designed to operate at 200 MHz using DDR-400 chips with a bandwidth of 3,200 MB/s. Because PC3200 memory transfers data on both the rising and falling clock edges, its effective clock rate is 400 MHz.

1 GB PC3200 non-ECC modules are usually made with sixteen 512 Mbit chips, 8 down each side (512 Mbits × 16 chips) / (8 bits (per byte)) = 1,024 MB. The individual chips making up a 1 GB memory module are usually organized with 64 Mbits and a data width of 8 bits for each chip, commonly expressed as 64M×8. Memory manufactured in this way is low density RAM and will usually be compatible with any motherboard specifying PC3200 DDR-400 memory.

High density RAM

In the context of the 1 GB non-ECC PC3200 SDRAM module, there is very little visually to differentiate low density from high density RAM. High density DDR RAM modules will, like their low density counterparts, usually be double-sided with eight 512 Mbit chips per side. The difference is that for each chip, instead of being organized in a 64M×8 configuration, it is organized with 128 Mbits and a data width of 4 bits, or 128M×4.

High density memory modules are assembled using chips from multiple manufacturers. These chips come in both the familiar 22 × 10 mm (approx.) TSOP2 and smaller squarer 12 × 9 mm (approx.) FBGA package sizes. High density chips can be identified by the numbers on each chip.

High density RAM devices were designed to be used in registered memory modules for servers. JEDEC standards do not apply to high-density DDR RAM in desktop implementations. JEDEC's technical documentation, however, supports 128M×4 semiconductors as such that contradicts 128×4 being classified as high density. As such, *high density* is a relative term, which can be used to describe memory which is not supported by a particular motherboard's memory controller.

Variations

DDR SDRAM Standard	Bus clock (MHz)	Internal rate (MHz)	Prefetch (min burst)	Transfer Rate (MT/s)	Voltage	DIMM pins	SO-DIMM pins	MicroDIMM pins
DDR	100–200	100–200	2n	200–400	2.5/2.6	184	200	172
DDR2	200–533	100–266	4n	400–1066	1.8	240	200	214
DDR3	400–1066	100–266	8n	800–2133	1.5	240	204	214

DDR (DDR1) was superseded by DDR2 SDRAM, which had modifications for higher clock frequency and again doubled throughput, but operates on the same principle as DDR. Competing with DDR2 was Rambus XDR DRAM. DDR2 dominated due to cost and support factors. DDR2 was in turn superseded by DDR3 SDRAM which offered higher performance for increased bus speeds and new features. DDR3 will likely be superseded by DDR4 SDRAM, which was first produced in 2011 and whose standards are still in flux (2012) with significant architectural changes.

DDR's prefetch buffer depth is 2(bits), while DDR2 uses 4. Although the effective clock rates of DDR2 are higher than DDR, the overall performance was no greater in the early implementations, primarily due to the high latencies of the first DDR2 modules. DDR2 started to be effective by the end of 2004, as modules with lower latencies became available.[11]

Memory manufacturers stated that it was impractical to mass-produce DDR1 memory with effective transfer rates in excess of 400 MHz (i.e. 400MT/s and 200 MHz external clock) due to internal speed limitations. DDR2 picks up where DDR1 leaves off, utilizing internal clock rates similar to DDR1, but is available at effective transfer rates of 400 MHz and higher. DDR3 advances extended the ability to preserve internal clock rates while providing higher effective transfer rates by again doubling the prefetch depth.

RDRAM was a particularly expensive alternative to DDR SDRAM, and most manufacturers dropped its support from their chipsets. DDR1 memory's prices substantially increased since Q2 2008 while DDR2 prices declined. In January 2009, 1 GB DDR1 was 2–3 times more expensive than 1 GB DDR2. High density DDR RAM will suit about 10% of PC motherboards on the market while low density will suit almost all motherboards on the PC Desktop market.

MDDR

MDDR is an acronym that some enterprises use for Mobile DDR SDRAM, a type of memory used in some portable electronic devices, like mobile phones, handhelds, and digital audio players. Through techniques including reduced voltage supply and advanced refresh options, Mobile DDR can achieve greater power efficiency.

References

[1] Northwest Logic DDR Phy datasheet (http://www.nwlogic.com/docs/ASIC_DDR_PHY.pdf)
[2] Memory Interfaces Data Capture Using Direct Clocking Technique (Xilinx application note) (http://www.xilinx.com/support/documentation/application_notes/xapp701.pdf)
[3] "The Love/Hate Relationship with DDR SDRAM Controllers" (http://www.design-reuse.com/articles/13805/the-love-hate-relationship-with-ddr-sdram-controllers.html). .
[4] Cycle time is the inverse of the I/O bus clock frequency; e.g., 1/(100 MHz) = 10 ns per clock cycle.
[5] DOUBLE DATA RATE (DDR) SDRAM STANDARD (http://www.jedec.org/standards-documents/docs/jesd-79f)
[6] "What is the difference between PC-2100 (DDR-266), PC-2700 (DDR-333), and PC-3200 (DDR-400)?" (http://www.crucial.com/support/memory_speeds.aspx). Micron Technology, Inc.. .
[7] Low Density vs High Density memory modules (http://reviews.ebay.com/Myth-Low-Density-vs-High-Density-memory-modules_W0QQugidZ10000000001236178)
[8] Mike Chin: Power Distribution within Six PCs (http://www.silentpcreview.com/article265-page4.html)
[9] Micron: System Power Calculators (http://www.micron.com/products/support/power-calc)
[10] http://www.jedec.org/download/search/JESD79F.pdf DOUBLE DATA RATE (DDR) SDRAM SPECIFICATION (Release F)
[11] DDR2 vs. DDR: Revenge Gained (http://www.xbitlabs.com/articles/memory/display/ddr2-ddr.html)

External links

- Official JEDEC website (http://www.jedec.org/)

DDR2 SDRAM

DDR2 SDRAM is a double data rate synchronous dynamic random-access memory interface. It supersedes the original DDR SDRAM specification and has itself been superseded by DDR3 SDRAM. DDR2 is neither forward nor backward compatible with either DDR or DDR3.

Two PC2-6400 modules encased in heat spreaders

In addition to double pumping the data bus as in DDR SDRAM (transferring data on the rising and falling edges of the bus clock signal), DDR2 allows higher bus speed and requires lower power by running the internal clock at half the speed of the data bus. The two factors combine to require a total of four data transfers per internal clock cycle. With data being transferred 64 bits at a time, DDR2 SDRAM gives a transfer rate of (memory clock rate) × 2 (for bus clock multiplier) × 2 (for dual rate) × 64 (number of bits transferred) / 8 (number of bits/byte). Thus with a memory clock frequency of 100 MHz, DDR2 SDRAM gives a maximum transfer rate of 3200 MB/s.

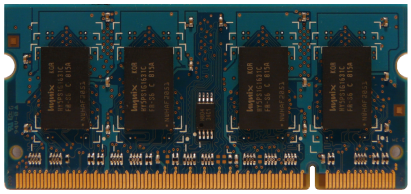

PC2-6400 DDR2 SO-DIMM (for notebooks)

Since the DDR2 internal clock runs at half the DDR external clock rate, DDR2 memory operating at the same external data bus clock rate as DDR results in DDR2 being able to provide the same bandwidth but with higher latency. Alternatively, DDR2 memory operating at twice the external data bus clock rate as DDR may provide twice the bandwidth with the same latency. The best-rated DDR2 memory modules are at least twice as fast as the best-rated DDR memory modules.

Overview

Like all SDRAM implementations, DDR2 stores data in memory cells that are activated with the use of a clock signal to synchronize their operation with an external data bus. Like DDR before it, the DDR2 I/O buffer transfers data both on the rising and falling edges of the clock signal (a technique called "double pumping"). The key difference between DDR and DDR2 is that for DDR2 the memory cells are clocked at 1 quarter (rather than half) the rate of the bus. This requires a 4-bit-deep prefetch queue, but, without changing the memory cells themselves, DDR2 can effectively operate at twice the bus speed of DDR.

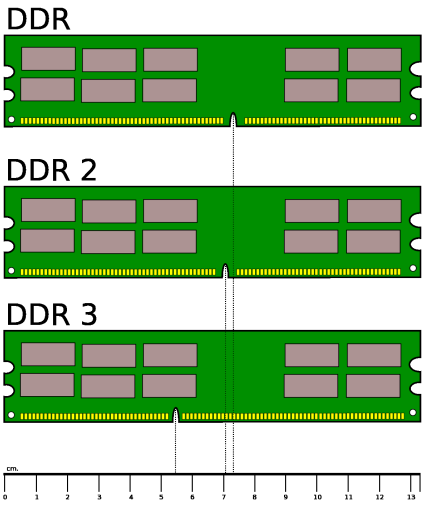

Comparison of memory modules for desktop PCs (DIMM).

DDR2's bus frequency is boosted by electrical interface improvements, on-die termination, prefetch buffers and off-chip drivers. However, latency is greatly increased as a trade-off. The DDR2 prefetch buffer is 4 bits deep, whereas it is two bits deep for DDR and eight bits deep for DDR3. While DDR SDRAM has typical read latencies of between 2 and 3 bus cycles, DDR2 may have read latencies between 4 and 6 cycles. Thus, DDR2 memory must be operated at twice the data rate to achieve the same latency.

Another cost of the increased bandwidth is the requirement that the chips are packaged in a more expensive and more difficult to assemble BGA package as compared to the TSSOP package of the previous memory generations such as DDR SDRAM and SDR SDRAM. This packaging change was necessary to maintain signal integrity at higher bus speeds.

Power savings are achieved primarily due to an improved manufacturing process through die shrinkage, resulting in a drop in operating voltage (1.8 V compared to DDR's 2.5 V). The lower memory clock frequency may also enable power reductions in applications that do not require the highest available data rates.

According to JEDEC[1] the maximum recommended voltage is 1.9 volts and should be considered the absolute maximum when memory stability is an issue (such as in servers or other mission critical devices). In addition, JEDEC states that memory modules must withstand up to 2.3 volts before incurring permanent damage (although they may not actually function correctly at that level).

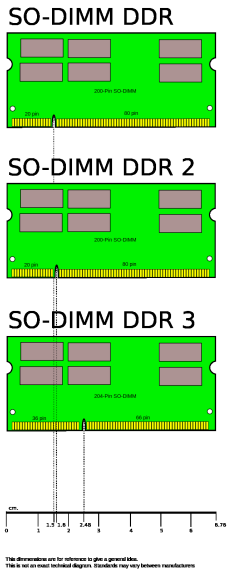

Comparison of memory modules for portable/mobile PCs (SO-DIMM).

Specification standards

Chips and modules

For use in computers, DDR2 SDRAM is supplied in DIMMs with 240 pins and a single locating notch. Laptop DDR2 SO-DIMMs have 200 pins and often come identified by an additional **S** in their designation. DIMMs are identified by their peak transfer capacity (often called bandwidth).

Standard name	Memory clock (MHz)	Cycle time (ns)	I/O bus clock (MHz)	Data rate (MT/s)	Module name	Peak transfer rate (MB/s)	Timings[2][3] (CL-tRCD-tRP)	CAS latency (ns)
DDR2-400B DDR2-400C	100	10	200	400	PC2-3200	3200	3-3-3 4-4-4	15 20
DDR2-533B DDR2-533C	133⅓	7½	266⅔	533⅓	PC2-4200*	4266⅔	3-3-3 4-4-4	11¼ 15
DDR2-667C DDR2-667D	166⅔	6	333⅓	666⅔	PC2-5300*	5333⅓	4-4-4 5-5-5	12 15
DDR2-800C DDR2-800D DDR2-800E	200	5	400	800	PC2-6400	6400	4-4-4 5-5-5 6-6-6	10 12½ 15
DDR2-1066E DDR2-1066F	266⅔	3¾	533⅓	1066⅔	PC2-8500*	8533⅓	6-6-6 7-7-7	11¼ 13⅛

** Some manufacturers label their DDR2 modules as PC2-4300, PC2-5400 or PC2-8600 instead of the respective names suggested by JEDEC. At least one manufacturer has reported this reflects successful testing at a higher-than standard data rate[4] whilst others simply round up for the name.*

Note: DDR2-xxx denotes data transfer rate, and describes raw DDR chips, whereas PC2-xxxx denotes theoretical bandwidth (with the last two digits truncated), and is used to describe assembled DIMMs. Bandwidth is calculated by taking transfers per second and multiplying by eight. This is because DDR2 memory modules transfer data on a

bus that is 64 data bits wide, and since a byte comprises 8 bits, this equates to 8 bytes of data per transfer.

In addition to bandwidth and capacity variants, modules can

1. Optionally implement ECC, which is an extra data byte lane used for correcting minor errors and detecting major errors for better reliability. Modules with ECC are identified by an additional ECC in their designation. PC2-4200 ECC is a PC2-4200 module with ECC.
2. Be "registered" ("buffered"), which improves signal integrity (and hence potentially clock rates and physical slot capacity) by electrically buffering the signals at a cost of an extra clock of increased latency. Those modules are identified by an additional **R** in their designation, whereas non-registered (a.k.a. "unbuffered") RAM *may be* identified by an additional **U** in the designation. PC2-4200R is a registered PC2-4200 module, PC2-4200R ECC is the same module but with additional ECC.
3. Be fully buffered modules, which are designated by **F** or **FB** and do not have the same notch position as other classes. Fully buffered modules cannot be used with motherboards that are made for registered modules, and the different notch position physically prevents their insertion.

Note: registered and un-buffered SDRAM generally cannot be mixed on the same channel.

Note that the highest-rated DDR2 modules in 2009 operate at 533 MHz (1066 MT/s), compared to the highest-rated DDR modules operating at 200 MHz (400 MT/s). At the same time, the CAS latency of 11.2 ns = 6 / (Bus clock rate) for the best PC2-8500 modules is comparable to that of 10 ns = 4 / (Bus clock rate) for the best PC-3200 modules.

Debut

DDR2 was introduced in the second quarter of 2003 at two initial clock rates: 200 MHz (referred to as PC2-3200) and 266 MHz (PC2-4200). Both performed worse than the original DDR specification due to higher latency, which made total access times longer. However, the original DDR technology tops out at a clock rate around 200 MHz (400 MT/s). Higher performance DDR chips exist, but JEDEC has stated that they will not be standardized. These modules are mostly manufacturer optimizations of highest-yielding chips, drawing significantly more power than slower-clocked modules, and usually do not offer much, if any, greater real-world performance.

DDR2 started to become competitive with the older DDR standard by the end of 2004, as modules with lower latencies became available.[5]

Backward compatibility

DDR2 DIMMs are not designed to be backward compatible with DDR DIMMs. The notch on DDR2 DIMMs is in a different position from DDR DIMMs, and the pin density is higher than DDR DIMMs in desktops. DDR2 is a 240-pin module, DDR is a 184-pin module. Notebooks have 200-pin modules for DDR and DDR2, however the notch on DDR modules is in a slightly different position than that on DDR2 modules.

Higher-speed DDR2 DIMMs are compatible with lower-speed DDR2 DIMMs although the motherboard or CPU memory controller will be bound to the limits of the lower-performance modules.

Relation to GDDR memory

The first commercial product to claim using the "DDR2" technology was the NVIDIA GeForce FX 5800 graphics card. However, it is important to note that this GDDR2 memory used on graphics cards is not DDR2 per se, but rather an early midpoint between DDR and DDR2 technologies. Using "DDR2" to refer to GDDR2 is a colloquial misnomer. In particular, the performance-enhancing doubling of the I/O clock rate is missing. It had severe overheating issues due to the nominal DDR voltages. ATI has since designed the GDDR technology further into GDDR3, which is based on DDR2-SDRAM, though with several additions suited for graphics cards.

GDDR3 is now commonly used in modern graphics cards and some tablet PCs. However, further confusion has been added to the mix with the appearance of budget and mid-range graphics cards which claim to use "GDDR2". These cards actually use standard DDR2 chips designed for use as main system memory although operating with higher latencies to achive higher clockrates. These chips cannot achieve the clock rates of GDDR3 but are inexpensive and fast enough to be used as memory on mid-range cards.

References

[1] JEDEC JESD 208 (http://www.jedec.org/download/search/JESD208.pdf) (section 5, tables 15 and 16)

[2] *DDR2 SDRAM SPECIFICATION* (http://www.jedec.org/download/search/JESD79-2E.pdf). JESD79-2E. JEDEC. April 2008. pp. 78. . Retrieved 2009-03-14.

[3] *SPECIALITY DDR2-1066 SDRAM* (http://www.jedec.org/download/search/JESD208.pdf). JEDEC. November 2007. pp. 70. . Retrieved 2009-03-14.

[4] Mushkin PC2-5300 vs. Corsair PC2-5400 (http://www.metku.net/index.html?path=reviews/ddr2_1/index_eng)

[5] Ilya Gavrichenkov. "DDR2 vs. DDR: Revenge gained" (http://www.xbitlabs.com/articles/memory/display/ddr2-ddr.html). X-bit Laboratories. .

Further reading

- JEDEC standard: DDR2 SDRAM Specification (http://www.jedec.org/download/search/JESD79-2F.pdf) (JESD79-2F, November 2009)
- JEDEC standard: DDR2-1066 (http://www.jedec.org/download/search/JESD208.pdf)
- "JEDEC Standard No. 21C: 4.20.13 240-Pin PC2-5300/PC2-6400 DDR2 SDRAM Unbuffered DIMM Design Specification" (http://www.jedec.org/download/search/4_20_13R18.pdf) (PDF). JEDEC Solid State Technology Association. 2008-10. Retrieved 2008-12-26.
- Razak Mohammed Ali. "DDR2 SDRAM interfaces for next-gen systems" (http://www.eetasia.com/ARTICLES/2006OCT/PDF/EEOL_2006OCT16_INTD_STOR_TA.pdf) (PDF). Electronic Engineering Times.

External links

- JEDEC website (http://www.jedec.org)
- Overview of DDR-II technology (http://www.lostcircuits.com/mambo//index.php?option=com_content&task=view&id=35&Itemid=60)
- DDR2 low latency vs high bandwidth, Core 2 Duo (Conroe) performance (http://www.xbitlabs.com/articles/memory/display/core2duo-memory-guide.html)

DDR3 SDRAM

PC3-10600 DDR3 SO-DIMM (204 pins)

In computing, **DDR3 SDRAM**, an abbreviation for **double data rate type three synchronous dynamic random access memory,** is a modern kind of dynamic random access memory (DRAM) with a high bandwidth interface, and has been in use since 2007. DDR3 SDRAM is neither forward nor backward compatible with any earlier type of random access memory (RAM) due to different signaling voltages, timings, and other factors.

DDR3 is a DRAM interface specification. The actual DRAM arrays that store the data are similar to earlier types, with similar performance.

The primary benefit of DDR3 SDRAM over its immediate predecessor, DDR2 SDRAM, is its ability to transfer data at twice the rate (eight times the speed of its internal memory arrays), enabling higher bandwidth or peak data rates. With two transfers per cycle of a quadrupled clock signal, a 64-bit wide DDR3 module may achieve a transfer rate of up to 64 times the memory clock speed in megabytes per second (MB/s). With data being transferred 64 bits at a time per memory module, DDR3 SDRAM gives a transfer rate of (memory clock rate) × 4 (for bus clock multiplier) × 2 (for data rate) × 64 (number of bits transferred) / 8 (number of bits/byte). Thus with a memory clock frequency of 100 MHz, DDR3 SDRAM gives a maximum transfer rate of 6400 MB/s. In addition, the DDR3 standard permits chip capacities of up to 8 gigabytes.

Overview

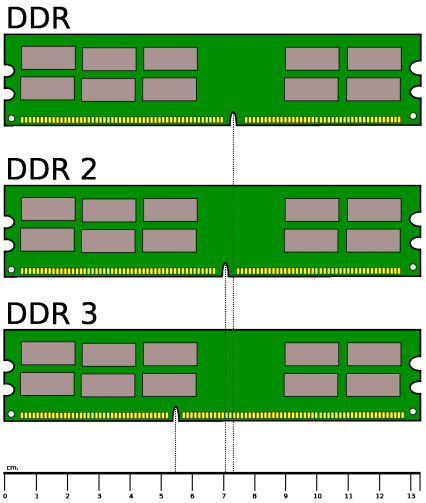

Comparison of memory modules for desktop PCs (DIMM).

Compared to DDR2 memory, DDR3 memory uses 30% less power. This reduction comes from the difference in supply voltages: 1.8 V or 2.5 V for DDR2, and 1.5 V for DDR3. The 1.5 V supply voltage works well with the 90 nanometer fabrication technology used in the original DDR3 chips. Some manufacturers further propose using "dual-gate" transistors to reduce leakage of current.[1]

According to JEDEC,[2] 1.575 volts should be considered the absolute maximum when memory stability is the foremost consideration, such as in servers or other mission-critical devices. In addition, JEDEC states that memory modules must withstand up to 1.975 volts before incurring permanent damage, although they are not required to function correctly at that level.

Another benefit is its prefetch buffer, which is 8-burst-deep. In contrast, the prefetch buffer of DDR2 is 4-burst-deep, and the prefetch buffer of DDR is 2-burst-deep. This advantage is an enabling technology in DDR3's transfer speed.

DDR3 modules can transfer data at a rate of 800–2133 MT/s using both rising and falling edges of a 400–1066 MHz I/O clock. Sometimes, a vendor may misleadingly advertise the I/O clock rate by labeling the MT/s as MHz. The MT/s is normally twice that of MHz by double sampling, one on the rising clock edge, and the other, on the

falling. In comparison, DDR2's current range of data transfer rates is 400–1066 MT/s using a 200–533 MHz I/O clock, and DDR's range is 200–400 MT/s based on a 100–200 MHz I/O clock. High-performance graphics was an initial driver of such bandwidth requirements, where high bandwidth data transfer between framebuffers is required.

DDR3 does use the same electric signaling standard as DDR and DDR2, Stub Series Terminated Logic, albeit at different timings and voltages. Specifically, DDR3 uses SSTL_15.[3]

DDR3 prototypes were announced in early 2005. Products in the form of motherboards appeared on the market in June 2007[4] based on Intel's P35 "Bearlake" chipset with DIMMs at bandwidths up to DDR3-1600 (PC3-12800).[5] The Intel Core i7, released in November 2008, connects directly to memory rather than via a chipset. The Core i7 supports only DDR3. AMD's first socket AM3 Phenom II X4 processors, released in February 2009, were their first to support DDR3.

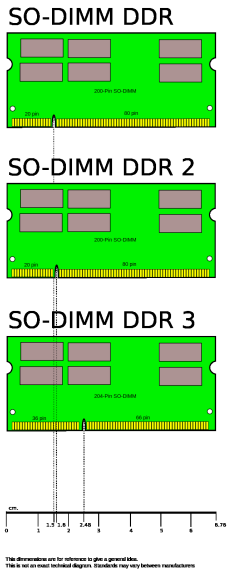

Comparison of memory modules for portable/mobile PCs (SO-DIMM).

DDR3 DIMMs have 240 pins and are electrically incompatible with DDR2. The two are prevented from being accidentally interchanged by different key notch positions on the DIMMs. Not only are DDR2 and DDR3 keyed differently, DDR2 has rounded notches on the side and the DDR3 modules have square notches on the side. [6] DDR3 SO-DIMMs have 204 pins.[7]

GDDR3 memory, sometimes incorrectly referred to as "DDR3" due to its similar name, is an entirely different technology, as it is designed for use in graphics cards and is based on DDR2 SDRAM.

DDR3L

The "L" in DDR3L stands for low-voltage. JEDEC introduced two low-voltage standards. The DDR3L standard is 1.35V and has the label "PC3L" for its modules. Examples include DDR3L-800, DDR3L-1066, DDR3L-1333, and DDR3L-1600. The DDR3U standard is 1.25V and has the label "PC3U" for its modules.

Latencies

While the typical latencies for a JEDEC DDR2 device were 5-5-5-15, some standard latencies for JEDEC DDR3 devices include 7-7-7-20 for DDR3-1066 and 8-8-8-24 for DDR3-1333.

DDR3 latencies are numerically higher because the I/O bus clock cycles by which they are measured are shorter; the actual time interval is similar to DDR2 latencies (around 10 ns). There is some improvement because DDR3 generally uses more recent manufacturing processes, but this is not directly caused by the change to DDR3.

As with earlier memory generations, faster DDR3 memory became available after the release of the initial versions. DDR3-2000 memory with 9-9-9-28 latency (9 ns) was available in time to coincide with the Intel Core i7 release.[8] CAS latency of 9 at 1000 MHz (DDR3-2000) is 9 ns, while CAS latency of 7 at 667 MHz (DDR3-1333) is 10.5 ns.

(CAS / Frequency (MHz)) × 1000 = X ns

Example:

(7 / 667) × 1000 = 10.4948 ns

Power consumption

Power consumption of individual SDRAM chips (or, by extension, DIMMs) varies based on many factors, including speed, type of usage, voltage, etc. Dell's Power Advisor calculates that 4 GB ECC DDR1333 RDIMMs use about 4 W each.[9] By contrast, a more modern mainstream desktop-oriented part 8 GB, DDR3/1600 DIMM, is rated at 2.58 W, despite being significantly faster.[10]

Extensions

Intel Corporation officially introduced the eXtreme Memory Profile (XMP) Specification on March 23, 2007 to enable enthusiast performance extensions to the traditional JEDEC SPD specifications for DDR3 SDRAM.[11]

Modules

JEDEC standard modules

Standard name	Memory clock (MHz)	Cycle time (ns)	I/O bus clock (MHz)	Data rate (MT/s)	Module name	Peak transfer rate (MB/s)	Timings (CL-tRCD-tRP)	CAS latency (ns)
DDR3-800D DDR3-800E	100	10	400	800	PC3-6400	6400	5-5-5 6-6-6	12 1/2 15
DDR3-1066E DDR3-1066F DDR3-1066G	133⅓	7 1/2	533⅓	1066⅔	PC3-8500	8533⅓	6-6-6 7-7-7 8-8-8	11 1/4 13 1/8 15
DDR3-1333F* DDR3-1333G DDR3-1333H DDR3-1333J*	166⅔	6	666⅔	1333⅓	PC3-10600	10666⅔	7-7-7 8-8-8 9-9-9 10-10-10	10 1/2 12 13 1/2 15
DDR3-1600G* DDR3-1600H DDR3-1600J DDR3-1600K	200	5	800	1600	PC3-12800	12800	8-8-8 9-9-9 10-10-10 11-11-11	10 11 1/4 12 1/2 13 3/4
DDR3-1866J* DDR3-1866K DDR3-1866L DDR3-1866M*	233⅓	4 2/7	933⅓	1866⅔	PC3-14900	14933⅓	10-10-10 11-11-11 12-12-12 13-13-13	10 5/7 11 11/14 12 6/7 13 13/14
DDR3-2133K* DDR3-2133L DDR3-2133M DDR3-2133N*	266⅔	3 3/4	1066⅔	2133⅓	PC3-17000	17066⅔	11-11-11 12-12-12 13-13-13 14-14-14	10 5/16 11 1/4 12 3/16 13 1/8

* optional

CL - Clock cycles between sending a column address to the memory and the beginning of the data in response

tRCD - Clock cycles between row activate and reads/writes

tRP - Clock cycles between row precharge and activate

Fractional frequencies are normally rounded down, but rounding up to -667 is common due to the exact number being -666⅔ and rounding to the nearest whole number. Some manufacturers also round to a certain precision or round up instead. For example, PC3-10666 memory could be listed as PC3-10600 or PC3-10700.[12]

Note: All items listed above are specified by JEDEC as JESD79-3D.[13] All RAM data rates in-between or above these listed specifications are not standardized by JEDEC—often they are simply manufacturer optimizations using higher-tolerance or overvolted chips. Of these non-standard specifications, the highest reported speed reached was equivalent to DDR3-2544, as of May 2010.[14]

DDR3-xxx denotes data transfer rate, and describes raw DDR chips, whereas PC3-xxxx denotes theoretical bandwidth (with the last two digits truncated), and is used to describe assembled DIMMs. Bandwidth is calculated by taking transfers per second and multiplying by eight. This is because DDR3 memory modules transfer data on a bus that is 64 data bits wide, and since a byte comprises 8 bits, this equates to 8 bytes of data per transfer.

In addition to bandwidth and capacity variants, modules can:

1. Optionally implement ECC, which is an extra data byte lane used for correcting minor errors and detecting major errors for better reliability. Modules with ECC are identified by an additional **ECC** or **E** in their designation. For example: "PC3-6400 ECC", or PC3-8500E.[15]
2. Be "registered", which improves signal integrity (and hence potentially clock rates and physical slot capacity) by electrically buffering the signals with a register, at a cost of an extra clock of increased latency. Those modules are identified by an additional **R** in their designation, whereas non-registered (a.k.a. "unbuffered") RAM *may be* identified by an additional **U** in the designation. PC3-6400R is a registered PC3-6400 module, and PC3-6400R ECC is the same module with ECC.
3. Be fully buffered modules, which are designated by **F** or **FB** and do not have the same notch position as other classes. Fully buffered modules cannot be used with motherboards that are made for registered modules, and the different notch position physically prevents their insertion.

Feature summary

DDR3 SDRAM components

- Introduction of asynchronous RESET pin
- Support of system-level flight-time compensation
- On-DIMM mirror-friendly DRAM pinout
- Introduction of CWL (CAS write latency) per clock bin
- On-die I/O calibration engine
- READ and WRITE calibration

DDR3 modules

- Fly-by command/address/control bus with on-DIMM termination
- High-precision calibration resistors
- Are **not** backwards compatible—DDR3 modules do not fit into DDR2 sockets; forcing them can damage the DIMM and/or the motherboard[16]

Technological advantages compared to DDR2

- Higher bandwidth performance, up to 2133 MT/s standardized
- Slightly improved latencies, as measured in nanoseconds
- Higher performance at low power (longer battery life in laptops)
- Enhanced low-power features

Development and market penetration

In May 2005, Desi Rhoden, chairman of the JEDEC committee responsible for creating the DDR3 standard, stated that DDR3 had been under development for "about 3 years".[17] DDR3 was launched in 2007, but sales were not expected to overtake DDR2 until the end of 2009, or possibly early 2010, according to Intel strategist Carlos Weissenberg, speaking during the early part of their roll-out in August 2008.[18] (The same timescale for market penetration had been stated by market intelligence company DRAMeXchange over a year earlier in April 2007,[19] and by Desi Rhoden in 2005.[17]) The primary driving force behind the increased usage of DDR3 has been new Core i7 processors from Intel and Phenom II processors from AMD, both of which have internal memory controllers: the latter recommends DDR3, the former requires it. IDC stated in January 2009 that DDR3 sales will account for 29 percent of the total DRAM units sold in 2009, rising to 72% by 2011.[20]

Successor

JEDEC's planned successor to DDR3 is DDR4, whose standard is currently in development.[21] The primary benefits of DDR4 compared to DDR3 include a higher range of clock frequencies and data transfer rates[22] and significantly lower voltage. Some manufacturers have already demonstrated DDR4 chips for testing purposes.[23]

References

[1] McCloskey, Alan, *Research: DDR FAQ* (http://www.ocmodshop.com/ocmodshop.aspx?a=868), , retrieved 2007-10-18

[2] JEDEC JESD 79-3B (http://www.jedec.org/download/search/JESD79-3B.pdf) (section 6, table 21 and section 7, table 23)

[3] Jaci Chang *Design Considerations for the DDR3 Memory Sub-system*. Jedex, 2004, p. 4. http://www.jedex.org/images/pdf/samsung%20-%20jaci_chang.pdf

[4] Soderstrom, Thomas (2007-06-05). "Pipe Dreams: Six P35-DDR3 Motherboards Compared" (http://www.tomshardware.com/2007/06/05/pipe_dreams_six_p35-ddr3_motherboards_compared/). Tom's Hardware. .

[5] Fink, Wesley (2007-07-20). "Super Talent & TEAM: DDR3-1600 Is Here!" (http://www.anandtech.com/printarticle.aspx?i=3045). AnandTech. .

[6] "DocMemory" (2007-02-21). "Memory Module Picture 2007" (http://www.simmtester.com/page/news/showpubnews.asp?title=Memory+Module+Picture+2007&num=150). .

[7] "JEDEC" (2010-12-01). "204-Pin DDR3 SDRAM SO-DIMM Specification" (http://www.jedec.org/download/search/4_20_18R20A.pdf). .

[8] Shilov, Anton (2008-10-29). "Kingston Rolls Out Industry's First 2GHz Memory Modules for Intel Core i7 Platforms" (http://www.xbitlabs.com/news/memory/display/20081029141143_Kingston_Rolls_Out_Industry_s_First_2GHz_Memory_Modules_for_Intel_Core_i7_Platforms.html). Xbit Laboratories. . Retrieved 2008-11-02.

[9] http://essa.us.dell.com/DellStarOnline/DCCP.aspx?c=us&l=en&s=corp&Template=6945c07e-3be7-47aa-b318-18f9052df893

[10] http://www.kingston.com/dataSheets/KVR16N11_8.pdf

[11] "Intel Extreme memory Profile (Intel XMP) DDR3 Technology" (http://www.intel.com/assets/pdf/whitepaper/319124.pdf). . Retrieved 2009-05-29.

[12] *Pc3 10600 vs. pc3 10666 What's the difference - New-System-Build* (http://www.tomshardware.com/forum/274587-31-10600-10666-what-difference#t2045244), Tomshardware.com, , retrieved 2012-01-23

[13] *DDR3 SDRAM STANDARD* (http://www.jedec.org/standards-documents/docs/jesd-79-3d), Jedec.org, , retrieved 2012-01-23

[14] *Kingston's 2,544 MHz DDR3 On Show at Computex* (http://news.softpedia.com/news/Kingston-s-2-544-MHz-DDR3-On-Show-at-Computex-143379.shtml), News.softpedia.com, 2010-05-31, , retrieved 2012-01-23

[15] (PDF) *Memory technology evolution: an overview of system memory technologies* (http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00256987/c00256987.pdf), Hewlett-Packard, p. 18,

[16] "DDR3: Frequently Asked Questions" (http://www.kingston.com/channelmarketingcenter/hyperx/literature/MKF_1223_DDR3_FAQ.pdf). . Retrieved 2009-08-18.

[17] Sobolev, Vyacheslav (2005-05-31). "JEDEC: Memory standards on the way" (http://www.digitimes.com/news/a20050530PR201.html). *digitimes.com*. . Retrieved 2011-04-28. "*JEDEC is already well along in the development of the DDR3 standard, and we have been working on it for about three years now.... Following historical models, you could reasonably expect the same three-year transition to a new technology that you have seen for the last several generations of standard memory*"

[18] "IDF: "DDR3 won't catch up with DDR2 during 2009"" (http://www.pcpro.co.uk/news/220257/idf-ddr3-wont-catch-up-with-ddr2-during-2009.html). pcpro.co.uk. 19 August 2008. . Retrieved 2009-06-17.

[19] Bryan, Gardiner (April 17, 2007). "DDR3 Memory Won't Be Mainstream Until 2009" (http://www.extremetech.com/article2/0,2845,2115031,00.asp). extremetech.com. . Retrieved 2009-06-17.

[20] Salisbury, Andy (2009-01-20). "New 50nm Process Will Make DDR3 Faster and Cheaper This Year" (http://www.maximumpc.com/article/news/new_50nm_process_will_make_ddr3_faster_and_cheaper_this_year). maximumpc.com. . Retrieved 2009-06-17.

[21] "KH Kim Receives 2011 JEDEC Technical Recognition Award" (http://www.jedec.org/news/jedec-awards-program/kh-kim-2011-tr-award). jedec.org. . Retrieved 2011-07-31.

[22] Shilov, Anton (August 16, 2010). "Next-Generation DDR4 Memory to Reach 4.266GHz – Report" (http://www.xbitlabs.com/news/memory/display/20100816124343_Next_Generation_DDR4_Memory_to_Reach_4_266GHz_Report.html). Xbitlabs.com. . Retrieved 2011-01-03.

[23] "Samsung develops DDR4 memory with up to 40 percent better energy efficiency than DDR3" (http://www.engadget.com/2011/01/04/samsung-develops-ddr4-memory-with-up-to-40-percent-better-energy/). Engadget.com. January 4, 2011. . Retrieved 2011-07-31.

DDR4 SDRAM

DDR4 SDRAM

Type	synchronous dynamic random-access memory
Release date	September 2012[1]
Predecessor	DDR3 SDRAM
Website	JEDEC [2]

In computing, **DDR4 SDRAM**, an abbreviation for **double data rate (fourth generation) synchronous dynamic random-access memory**, is a type of dynamic random-access memory (DRAM) with a high bandwidth interface expected to be released to market in 2013.[3][4] It is one of several variants of DRAM, some of which have been in use since the early 1970s[5] and is not compatible with any earlier type of random access memory (RAM) due to different signaling voltages, physical interface and other factors.

Benefits

Its primary benefits compared to DDR3 include a higher range of clock frequencies and data transfer rates (2133–4266 MT/s compared to DDR3's 800 and higher[6][7][8]) and lower voltage (1.05–1.2 V for DDR4,[7] compared to 1.2–1.65 V for DDR3) with current remaining the same.[9] DDR4 also anticipates a change in topology. It discards the multiple DIMMs per channel approach in favor of a point-to-point topology where each channel in the memory controller is connected to a single DIMM.[7][10] Switched memory banks are also an anticipated option for servers.[7]

Development and market history

Standards body JEDEC began working on a successor to DDR3 around 2005,[11] about 2 years before the launch of DDR3 in 2007.[12][13] The high-level architecture of DDR4 was planned for completion in 2008.[14]

Some advance information was published in 2007,[15] and a guest speaker from Qimonda provided further public details in a presentation at the August 2008 San Francisco Intel Developer Forum (IDF).[15][16][17][18] DDR4 was described as involving a 30 nm process at 1.2 volts, with bus frequencies of 2133 MT/s "regular" speed and 3200 MT/s "enthusiast" speed, and reaching market in 2012, before transitioning to 1 volt in 2013.[16][18]

Subsequently, further details were revealed at MemCon 2010, Tokyo (a computer memory industry event), at which a presentation by a JEDEC director titled "Time to rethink DDR4" [19] with a slide titled "New roadmap: More realistic roadmap is 2015" led some websites to report that the introduction of DDR4 was probably[20] or definitely[21][22] delayed until 2015. However, DDR4 test samples were announced in line with the original schedule in early 2011 at which time manufacturers began to advise that large scale commercial production and release to market was scheduled for 2012.[3]

DDR4 is expected to represent 5% of the DRAM market in 2013,[3] and to reach mass market adoption and 50% market penetration around 2015;[3] the latter is comparable with the approximately 5 years taken for DDR3 to achieve mass market transition over DDR2.[7] In part, this is because changes required to other components would impact all other parts of computer systems, which would need to be updated to work with DDR4.[6]

In February 2009, Samsung validated 40 nm DRAM chips, considered a "significant step" towards DDR4 development[23] since in 2009, DRAM chips were only beginning to migrate to a 50 nm process.[24] In January 2011, Samsung announced the completion and release for testing of a 2 GB DDR4 DRAM module based on a process

between 30 and 39 nm.[25] It has a maximum data transfer rate of 2133 Mb/s at 1.2 V, uses pseudo open drain technology (adapted from graphics DDR memory[26]) and draws 40% less power than an equivalent DDR3 module.[25][27][28]

Three months later in April 2011, Hynix announced the production of 2 GB DDR4 modules at 2400 MT/s, also running at 1.2 V on a process between 30 and 39 nm (exact process unspecified),[3] adding that it anticipated commencing high volume production in the second half of 2012.[3] Semiconductor processes for DDR4 are expected to transition to sub-30 nm at some point between late 2012 and 2014.[7][29]

In May 2012 Micron announced[4] it is aiming at starting production in late 2012 of 30 nm modules.

In July 2012, Samsung Electronics Co., Ltd., announced that it has begun sampling the industry's first 16GB registered dual inline memory modules (RDIMMs) using DDR4 SDRAM for enterprise server systems.[30][31]

In September 2012 JEDEC released the final specification of DDR4.[1]

DIMM differences

DDR4 memory comes in a 284-pin DIMM, similar to a 240-pin DDR-2/DDR-3 DIMM.[32]:11 The pins are spaced more closely (0.85 mm instead of 1.0) to fit more within the standard $5\frac{1}{4}$-inch (133.35 mm) DIMM width, the height is increased slightly (31.25 mm/1.23 in instead of 30.35 mm/1.2 in) to make signal routing easier, and the thickness is increased (to 1.2 mm from 1.0) to accommodate more signal layers.

DDR4 SO-DIMMs have 256 pins (rather than 204), spaced closer (0.5 mm rather than 0.6), and are 1.0 mm wider (68.6 mm rather than 67.6), but retain the same 30 mm height.[32]:11

Technical description

The new chips will use a 1.2 V supply[32]:16[33][34] with a 2.5 V auxiliary supply V_{PP},[32]:16, versus the standard 1.5 V of DDR3 chips, with lower voltage variants at 1.05 V appearing in 2013. They are expected to be introduced at transfer rates of 2133 MT/s,[32]:18 estimated to rise to a potential 4266 MT/s[6] by 2013. The minimum transfer rate of 2133 MT/s was said to be due to progress made in DDR3 speeds which, being likely to reach 2133 MT/s, left little commercial benefit to specifying DDR4 below this speed.[6][7] Techgage interpreted Samsung's January 2011 engineering sample as having CAS latency of 13 clock cycles, described as being comparable to the move from DDR2 to DDR3.[26]

Internal banks are increased to 16 (4 bank select bits), with up to 8 ranks per DIMM.[32]:16

Protocol changes include:[32]:20

- Parity on the command/address bus
- Data bus inversion (like GDDR4)
- CRC on the data bus
- Independent programming of individual DRAMs on a DIMM, to allow better control of on-die termination.

Increased memory density is anticipated, possibly using TSV ("through-silicon via") or other 3D stacking processes.[6][7][10][35] The DDR4 specification will include standardized 3D stacking "from the start" according to JEDEC,[35] with provision for up to 8 stacked dies.[32]:12 X-bit Labs predicted that "as a result DDR4 memory chips with very high density will become relatively inexpensive".[6] Prefetch remains at 8n[32]:16 with bank groups, including the use of two or four selectable bank groups.[36]

DDR4 also anticipates a change in topology. It discards the multi-drop bus approach used in previous generations in favor of point-to-point where each channel in the memory controller is connected to a single module.[7][10] This mirrors the trend also seen in the earlier transition from PCI to PCI Express, where parallelism was moved from the interface to the controller,[10] and is likely to simplify timing in modern high-speed data buses.[10] Switched memory banks are also an anticipated option for servers.[7][10]

In 2008, concerns were raised in the book *Wafer Level 3-D ICs Process Technology* that non-scaling analog elements such as charge pumps and voltage regulators, and additional circuitry "have allowed significant increases in bandwidth but they consume much more die area". Examples include CRC error-detection, on-die termination, burst hardware, programmable pipelines, low impedance, and increasing need for sense amps (attributed to a decline in bits per bitline due to low voltage). The authors noted that as a result, the amount of die used for the memory array itself has declined over time from 70–78% with SDRAM and DDR1, to 47% for DDR2, to 38% for DDR3 and potentially to less than 30% for DDR4.[37]

Command encoding

Although it still operates in fundamentally the same way, DDR4 makes one major change to the command formats used by previous SDRAM generations. A new command signal **/ACT** is low to indicate the activate (open row) command.

The activate command requires more address bits than any other (18 row address bits in an 8 Gb part), so the standard **/RAS**, **/CAS** and **/WE** signals are shared with high-order address bits that are not used when /ACT is high. The combination of /RAS, /CAS and /WE that previously encoded an activate command is unused.

As in previous SDRAM encodings, A10 is used to select command variants: auto-precharge on read and write commands, and one bank vs all banks for the precharge command. It also selects two variants of the ZQ calibration command.

In addition, A12 is used to request *burst chop*: truncation of an 8-transfer burst after 4 transfers. Although the bank is still busy and unavailable for other commands until 8 transfer times have elapsed, a different bank can be accessed.

Also, the number of bank addresses has been increased greatly. There are 4 bank select bits to select up to 16 banks within each DRAM: 2 bank address bits (BA0, BA1), and 2 bank group bits (BG0, BG1). There are additional timing restrictions when accessing banks within the same bank group; it is faster to access a bank in a different bank group.

In addition, there are 3 chip select signals (C0, C1, C2), allowing up to 8 stacked chips to be placed inside a single DRAM package. These effectively act as three more bank select bits, bringing the total to 7 (128 possible banks).

DDR4 command encoding[38]

/CS	BG*n*, BA*n*	/ACT	A17	A16 /RAS	A15 /CAS	A14 /WE	A13	A12	A11	A10	A9–0	Command
H	— x —											Deselect (No operation)
L	bank	L	Row address									Active (activate): open a row
L	x	H	x	H	H	H	— x —					No operation
L	x	H	x	H	H	L	x			long	x	ZQ Calibration
L	bank	H	x	H	L	H	x	BC	x	AP	Column	Read (BC=burst chop)
L	bank	H	x	H	L	L	x	BC	x	AP	Column	Write (AP=auto-precharge)
L	x	H	x	L	H	H	— x —					(Unassigned, reserved)
L	x	H	x	L	H	L	x			H	x	Precharge all banks
L	bank	H	x	L	H	L	x			L	x	Precharge one bank
L	x	H	x	L	L	H	— x —					Refresh
L	register	H	0	L	L	L	0	data				Mode register set (MR0–MR6)

Note: x bits are "don't care", but must be at a valid voltage level, either 0 or 1.

Standard transfer rates are 1600, 1866, 2133 and 2400 MT/s.[38] (12/15, 14/15, 16/15 and 18/15 GHz clock speeds, double data rate.) 2666 and 3200 MT/s (20/15 and 24/15 GHz clock speeds) are provided for, but the specifications are not yet complete.

References

[1] JEDEC Announces Publication of DDR4 Standard | JEDEC (http://www.jedec.org/news/pressreleases/jedec-announces-publication-ddr4-standard)

[2] http://www.jedec.org/standards-documents/results/jesd79-4%20ddr4

[3] "Hynix produces its first DDR4 modules" (http://www.behardware.com/news/11425/hynix-produces-its-first-ddr4-modules.html). *behardware.com*. 2011-04-05. . Retrieved 2012-04-14.

[4] *Micron teases working DDR4 RAM* (http://www.engadget.com/2012/05/08/micron-teases-working-ddr4-ram-module), engadget.com, 2012-05-08, , retrieved 2012-05-08

[5] *The DRAM Story* (http://www.ieee.org/portal/cms_docs_societies/sscs/PrintEditions/200801.pdf), www.ieee.org, 2008, p. 10, , retrieved 2012-01-23

[6] Shilov, Anton (2010-08-16), *Next-Generation DDR4 Memory to Reach 4.266 GHz* (http://www.xbitlabs.com/news/memory/display/20100816124343_Next_Generation_DDR4_Memory_to_Reach_4_266GHz_Report.html), Xbitlabs.com, , retrieved 2011-01-03

[7] 後藤弘茂 ("Gotou Shigehiro"). "メモリ4Gbps時代へと向かう次世代メモリDDR4 ("Towards Next-Generation 4Gbps DDR4 Memory")" (http://pc.watch.impress.co.jp/docs/column/kaigai/20100816_387444.html). *2010-08-16*. PC Watch (Japan). . Retrieved 2011-04-25. (English translation (http://translate.google.com/translate?js=y&prev=_t&hl=en&ie=UTF-8&layout=1&eotf=1&u=http://pc.watch.impress.co.jp/docs/column/kaigai/20100816_387444.html&sl=ja&tl=en))

[8] At the time of publication DDR3 was described as reaching a limit at 2133 MHz; since then a wide range of DDR3 memory with speeds of 2400 or higher have become available, so the cited upper limit is now outdated.

[9] Samsung announces 16GB DDR4 DIMM to be released in 2014 :: TweakTown USA Edition (http://www.tweaktown.com/news/24799/samsung_announces_16gb_ddr4_dimm_to_be_released_in_2014/index.html)

[10] Swinburne, Richard (2010-08-26). "DDR4: What we can Expect" (http://www.bit-tech.net/hardware/memory/2010/08/26/ddr4-what-we-can-expect/1). *bit-tech.net*. . Retrieved 2011-04-28. – Page 1 (http://www.bit-tech.net/hardware/memory/2010/08/26/ddr4-what-we-can-expect/2) – – Page 3 (http://www.bit-tech.net/hardware/memory/2010/08/26/ddr4-what-we-can-expect/3)

[11] Sobolev, Vyacheslav (2005-05-31). "JEDEC: Memory standards on the way" (http://www.digitimes.com/news/a20050530PR201.html). *digitimes.com*. . Retrieved 2011-04-28. "*Initial investigations have already started on memory technology beyond DDR3. JEDEC always has about three generations of memory in various stages of the standardization process: current generation, next generation, and future.*"

[12] "DDR3: Frequently asked questions" (http://www.kingston.com/channelmarketingcenter/hyperx/literature/MKF_1223-1_DDR3_FAQ.pdf). Kingston Technology. . Retrieved 2011-04-28. "*"DDR3 memory launched in June 2007""*

[13] Valich, Theo (2007-05-02). "DDR3 launch set for May 9th" (http://www.theinquirer.net/inquirer/news/1016272/ddr3-launch-set-may-9th). *The Inquirer*. . Retrieved 2011-04-28.

[14] Hammerschmidt, Christoph (2007-08-29). "Non-volatile memory is the secret star at JEDEC meeting" (http://www.eetimes.com/electronics-news/4189215/Non-volatile-memory-is-the-secret-star-at-JEDEC-meeting?pageNumber=1). *eetimes.com*. . Retrieved 2011-04-28.

[15] "DDR4 – the successor to DDR3 memory" (http://www.h-online.com/newsticker/news/item/IDF-DDR4-the-successor-to-DDR3-memory-736983.html). *The "H" (h-online.com)*. 2008-08-21. . Retrieved 2011-04-28. "*The JEDEC standardisation committee cited similar figures around one year ago*"

[16] Graham-Smith, Darien (2008-08-19). "IDF: DDR3 won't catch up with DDR2 during 2009" (http://www.pcpro.co.uk/news/220257/idf-ddr3-wont-catch-up-with-ddr2-during-2009). *PC Pro*. . Retrieved 2011-04-28.

[17] Volker Risska (Volker Rißka) (2008-08-21). "IDF: DDR4 als Hauptspeicher ab 2012 ["Intel Developer Forum: DDR4 as the main memory from 2012" (http://www.computerbase.de/news/hardware/arbeitsspeicher/2008/august/idf-ddr4-als-hauptspeicher-ab-2012/)"]. *computerbase.de*. . Retrieved 2011-04-28. (English (http://translate.google.com/translate?hl=en&sl=de&tl=en&u=http://www.computerbase.de/news/hardware/arbeitsspeicher/2008/august/idf-ddr4-als-hauptspeicher-ab-2012/))

[18] Novakovic, Nebojsa (2008-08-19). "Qimonda: ddr3 moving forward" (http://www.theinquirer.net/inquirer/news/1012591/qimonda-ddr3-moving-forward). *The Inquirer*. . Retrieved 2011-04-28.

[19] Gervasi, Bill. "Time to rethink DDR4" (http://discobolusdesigns.com/personal/20100721a_gervasi_rethinking_ddr4.pdf). *July 2010*. Discobolus Designs. . Retrieved 2011-04-29.

[20] "DDR4-Speicher kommt wohl später als bisher geplant ("DDR4 memory is probably later than previously planned")" (http://www.heise.de/newsticker/meldung/DDR4-Speicher-kommt-wohl-spaeter-als-bisher-geplant-1060545.html). *heise.de*. 2010-08-17. . Retrieved 2011-04-29. (English (http://translate.google.com/translate?hl=en&sl=de&tl=en&u=http://www.heise.de/newsticker/meldung/DDR4-Speicher-kommt-wohl-spaeter-als-bisher-geplant-1060545.html))

[21] Nilsson, Lars-Göran (2010-08-16). "DDR4 not expected until 2015" (http://semiaccurate.com/2010/08/16/ddr4-not-expected-until-2015/). *semiaccurate.com*. . Retrieved 2011-04-29.

[22] By 'annihilator' (2010-08-18). "DDR4 memory in Works, Will reach 4.266GHz" (http://wccftech.com/2010/08/18/ddr4-memory-works-reach-4266ghz/). *wccftech.com*. . Retrieved 2011-04-29.

[23] Gruener, Wolfgang (February 4, 2009). "Samsung hints to DDR4 with first validated 40 nm DRAM" (http://www.tgdaily.com/content/view/41316/139/). tgdaily.com. . Retrieved 2009-06-16.

[24] Jansen, Ng (January 20, 2009). "DDR3 Will be Cheaper, Faster in 2009" (http://www.dailytech.com/DDR3+Will+be+Cheaper+Faster+in+2009/article13977.htm). dailytech.com. . Retrieved 2009-06-17.

[25] "Samsung Develops Industry's First DDR4 DRAM, Using 30nm Class Technology" (http://www.samsung.com/us/business/semiconductor/newsView.do?news_id=1202). *Samsung*. 2011-04-11. . Retrieved 26 April 2011.

[26] Perry, Ryan (2011-01-06). "Samsung Develops the First 30nm DDR4 DRAM" (http://techgage.com/news/samsung_develops_the_first_30nm_ddr4_dram/). *techgage.com*. . Retrieved 2011-04-29.

[27] "Samsung Develops Industry's First DDR4 DRAM, Using 30nm Class Technology" (http://www.samsung.com/us/business/semiconductor/newsView.do?news_id=1202). Samsung. 2011-01-04. . Retrieved 2011-03-13.

[28] Protalinski, Emil (2011-01-04), *Samsung develops DDR4 memory, up to 40% more efficient* (http://www.techspot.com/news/41818-samsung-develops-ddr4-memory-up-to-40-more-efficient.html), Techspot.com, , retrieved 2012-01-23

[29] "Diagram: Anticipated DDR4 timeline" (http://translate.googleusercontent.com/translate_c?hl=en&ie=UTF-8&sl=ja&tl=en&u=http://pc.watch.impress.co.jp/img/pcw/docs/387/444/html/kaigai-09.jpg.html&prev=_t&rurl=translate.google.com&twu=1&usg=ALkJrhiDMjWRSrDnnnL4mro50bRDlxQlVw). *2010-08-146*. PC Watch (Japan). . Retrieved 2011-04-25. (Linked from cited PC Watch article dated August 16, 2010)

[30] Samsung Samples Industry's First DDR4 Memory Modules for Servers. (http://www.xbitlabs.com/news/memory/display/20120702221021_Samsung_Samples_Industry_s_First_DDR4_Memory_Modules_for_Servers.html)

[31] Samsung Samples Industry's First 16-Gigabyte Server Modules Based on DDR4 Memory technology (http://www.samsung.com/global/business/semiconductor/news-events/press-releases/detail?newsId=11701)

[32] Jung, J.Y. (2012-09-11), "How DRAM Advancements are Impacting Server Infrastructure" (https://intel.activeevents.com/sf12/scheduler/catalog.do), *Intel Developer Forum 2012*, Samsung, , retrieved 2012-09-15

[33] *Looking forward to DDR4* (http://www.pcpro.co.uk/news/220257/idf-ddr3-wont-catchup-with-ddr2-during-2009.html), Pcpro.co.uk, 2008-08-19, , retrieved 2012-01-23

[34] *IDF: DDR4 - the successor to DDR3 memory* (http://www.heise-online.co.uk/news/IDF-DDR4-the-successor-to-DDR3-memory--/111367), Heise-online.co.uk, 2008-08-21, , retrieved 2012-01-23

[35] "JEDEC Announces Broad Spectrum of 3D-IC Standards Development" (http://www.jedec.org/news/pressreleases/jedec-announces-broad-spectrum-3d-ic-standards-development). *JEDEC*. 2011-03-17. . Retrieved 26 April 2011.

[36] "Main Memory: DDR3 & DDR4 SDRAM" (http://www.jedec.org/category/technology-focus-area/main-memory-ddr3-ddr4-sdram). *jedec.org*. . Retrieved 2012-04-14.

[37] Tan, Gutmann and Reif (2008). *Wafer Level 3-D ICs Process Technology* (http://books.google.com/books?id=fhen8HeoC1AC&pg=PA278&dq=ddr4+sdram&hl=en&ei=iLa2Ta6LAY_Cswa1m9DNDQ&sa=X&oi=book_result&ct=result&resnum=13&ved=0CIcBEOgBMAw#v=onepage&q=ddr4 sdram&f=false). Springer. pp. 278 (sections 12.3.4–12.3.5). .

[38] *JEDEC Standard JESD79-4: DDR4 SDRAM* (http://www.jedec.org/standards-documents/docs/jesd79-4), JEDEC Solid State Technology Association, September 2012, , retrieved 2012-10-11. Username "cypherpunks" and password "cypherpunks" will allow download.

Article Sources and Contributors

Dynamic random-access memory *Source*: http://en.wikipedia.org/w/index.php?oldid=536477746 *Contributors*: 2mcm, 7265, A5b, Aapo Laitinen, Adambro, Aij, Ajbuckle1, Aksi great, Alecv, Ali@gwc.org.uk, Ambulnick, Anastrophe, Animum, Anonywiki, Ariel., ArnoldReinhold, Arteum, Atreyu42, Atropos235, Avenged Eightfold, B jonas, B.Zsolt, B4hand, Bakuryuu, Baxeel, Belovedfreak, Big Brother 1984, Blainster, Bloodshedder, Borgx, Brighterorange, Bryan Derksen, Brycen, Burntsauce, C+C, Cab88, Caliston, CapitalR, Ceros, Chrihern, Christian75, Colin Douglas Howell, Colin Marquardt, Colonel Lord 80, Cyferz, DHR, Daio, Damian Yerrick, Daniel.West, Darin-0, Darth Panda, DaveJB, David Latapie, DavidCary, Dead3y3, Denniss, Dewritech, Discospinster, Djtexho, Dogcow, Douglas Whitaker, Download, DrSeehas, Drakferion, Dreddnott, Drewster1829, ELLinng, ENeville, ESkog, Eagleal, Edward, Electron9, Emijrp, Engineerism, Espoo, Etienne.navarro, Exonie, Face, Falcon9x5, Feneeth of Borg, Flickboy, Fram, Frap, Fuzzbox, Fyrael, Gaius Cornelius, Geoffsmithuk, Glenn, Glogger, Graham87, Gretchenpatti, GroundHog, Gurch, Guruleninn, Gypsum Fantastic, Haiviet, Ham Pastrami, Hanifbbz, Harry Q. Hammer, Hatster301, Hdante, HeirloomGardener, HenkeB, Heron, Hitchhiker89, Hmains, Hooperbloob, Hydrogen Iodide, Hyins, ICE77, Icarus3, Imperator3733, Intgr, Isnow, JTorkyt, Jamesontai, Jason One, Jbanes, Jcarroll, Jce, Jdwinx, Jeblad, Jez9999, Jhsounds, Jim Grisham, Jni, Johan Jönsson, John of Reading, Johnnyspastic, Jwortzel, Karlhendrikse, Kbh3rd, Kenyon, Kha0sK1d, Kipholbeck, Kledsky, Klower, Kostmo, Kwamikagami, LOL, Lakche, Lavenderbunny, Lehk, Leszek Jańczuk, Liftarn, LilHelpa, LindsayH, MER-C, Ma Baker, Madhero88, Magus732, Mako098765, MaltaGC, Mamyles, Mani kunnathettu, Mariushm, Matt Britt, Maury Markowitz, MaverickSolutions, Maxdlink, Mboverload, Mbz113, Megaloman, Megapixie, Mirror Vax, Moscito, Moxfyre, MrOllie, Nancy, Nicolas de Marqué, Nile, Nrnkpeukdzr, Nturton, Nwndarkness, Olof nord, Omnipedian, Onkelringelhuth, Orderud, Oregonerik, Otaviogood, Owenmann, Oxymoron83, PartierSP, Patrick, Pengo, Permacultura, PerryTachett, Peruvianllama, Pgan002, Phsource, Plasticup, Plugwash, Pmj, Pnm, Pol098, Polytope4d, Prithvi Ravichandran, Qbgeekjtw, Quibik, R'n'B, R. Matick, RSRScrooge, RTC, Raven4x4x, Raymondwinn, Rchandra, Reaper Eternal, Redgrittybrick, RenamedUser01302013, ReneS, RevRagnarok, Rigadoun, Robertvan1, Rwwww, Sagaciousuk, Salam32, Salamurai, Sam Hocevar, Sarenne, Scientus, Scuarty, Sfahey, Skarebo, Slice4e, Smartguy777, Solinym, Some jerk on the Internet, Sophus Bie, SoulCity17, SpaceFlight89, SpeedyGonsales, Stefan Heinzmann, Stratocracy, Subspace1250, Susan123fg, Swaaye, Syrthiss, Tabletop, TakuyaMurata, Tarquin, Tempshill, The Anome, TheFeds, Thecheesykid, Theo Pardilla, Thirteenangrymen, Thumperward, Tide rolls, Tim Starling, Timo Honkasalo, Tobyc75, Tobycat, Tom harrison, Tonyle, Trckq5229, Trekphiler, Tritochuk, Tsman, Tsuna20, UkPaolo, VampWillow, Velella, Vladmy, Voidxor, Vrenator, WLior, Wavelength, Wbm1058, WeißNix, Wiki alf, Wikipedeur, William Avery, Wizard IT, Wizardist, Wtshymanski, Yonaa, Ysangkok, Yuyudevil, 493 anonymous edits

Synchronous dynamic random-access memory *Source*: http://en.wikipedia.org/w/index.php?oldid=528801175 *Contributors*: Alereon, ArchStanton69, ArielGold, Ash73, AxelBoldt, B4hand, Biker Biker, Blound, Borgx, Brintoneng, C 1, CAkira, CCFreak2K, Cacadril, Capsitan, Ceros, Chepry, Chetvorno, Cogiati, CosineKitty, Curtains01, Dabomb87, DannyKitty, Darin-0, Ddrane, Deanpcmad, Debresser, Denniss, Dgessner, DmitryKo, Donhoraldo, Dozens, DragonHawk, E2eamon, ELCleanup, Edokter, Electron9, Emijrp, Engineerism, Eranhi, Eternal nirvana21, FT2, Frap, GreyCat, Haham hanuka, Ham Pastrami, Havianas123, Hdante, Hecline3rd, Hellbus, Heresiarch, Hlinasp, Hmrox, Intgr, JCDenton2052, Jamelan, Jamesjv, Jared Preston, Javalenok, Jeffq, Jezaraknid, Jhw38, Jim1138, Jimmi Hugh, JohnOwens, Jonkerz, Jshayden, Jts888, Karam.Anthony.K, Kigali1, Krawi, Krazymike, Kveri, Kyleess123, Landofjacks, Lanny.ch, LordBlizzard, LorenzoB, ManiacK, Marek69, Matejhowell, Materialscientist, Matt Britt, Mindmatrix, Monty845, Mr Stephen, Music Sorter, NewEnglandYankee, Niceguyedc, Nikpapag, Oscillon, Oxymoron83, Pbb, Peter17, Pitel, Polytope4d, Rastapoulous, Rich Farmbrough, Rilak, Seaphoto, Slakr, Smb1001, SpartanPhalanx, Squantmuts, Sriharsh1234, Sterremix, Stillwaterising, Syp, Tcnuk, Tehom2000, Terryeo, Tgr, The toiletbrush, Thewikigod12345, Thunderbird2, Tom strom, Tony Fox, Trolleytimes, Tuli, Wadsworth, Why Not A Duck, WillMak050389, WillWare, WorkingBeaver, Wtshymanski, ZZninepluralZalpha, Zhangxc, 282 anonymous edits

DDR SDRAM *Source*: http://en.wikipedia.org/w/index.php?oldid=533526037 *Contributors*: 194.81.54.xxx, A876, ALEF7, AS, Aasmith, Abbas mansoor, Abdull, Abrahami, Adam Zivner, Aegith, Ahodacsek, Ahoerstemeier, Alansohn, Aleksey Gerasimov, Alereon, Alexander Abramov, Amin10, Anon lynx, Anonymous Dissident, Aruffell, Aryaniae, Ayecee, BenFrantzDale, Bento00, Beyond silence, Bhyst, Blac (VS), Blehfu, Bobet, Bobo192, Bookofjude, Bovineone, Brian Patrie, Bryan Derksen, Bunnyhop11, Burt777, Can't sleep, clown will eat me, Canadacow, CaptainVindaloo, Carmichael, Ceecookie, Ceros, ChangChienFu, Chawmp, Cloudrider7, Cobaltbluetony, Conversion script, CourtPoet, Crispmuncher, Cxz111, DHR, DVdm, Daio, Damian Yerrick, Dan2paul, Dance-a-day, Darin-0, DavidCary, Dawnseeker2000, Denniss, DmitryKo, DocWatson42, DoomBringer, Dratman, Drfrogsplat, Duckbill, Ed g2s, Eliz81, Elvey, Engineerism, Enochlau, Epbr123, Excirial, FT2, Facts707, Fernvale, Fleminra, Fnagaton, Frap, Fred Bradstadt, Fumitol, Glen, GratefulGeek, Gugustiuci, Gurch, Gyrferret, Ham Pastrami, Hawksj, Heron, Hinrik, Holek, Humayun me, Hut 8.5, Hyungjin Ahn, Iammaxus, Inter, Intgr, J.delanoy, JB001, Jaakobou, James086, Javalenok, Javed.siddiki, Jdelvare, Jeffq, Jpo, Jts888, JulesH, Justinc, KGyST, Kanhef, Katieh5584, Kevin Steinhardt, Kingpin13, Kraftlos, Krash, Kurykh, Kyng, Lanny.ch, LarryLACa, Liftarn, Lotu, Luigiacruz, Malerin, Martiniturbide, Martyx, MatthewWilcox, Mboverload, Mdotley, Mediran, MementoVivere, Metasquares, Michaeloqu, Mike Rosoft, Mikeatwwu, Mild Bill Hiccup, Mix Bouda-Lycaon, Mnts, Mrwojo, Mtmsmile, NawlinWiki, Ncjones, Neelix, Npc, OlivierMiR, Orborde, Orderud, OrgasGirl, Pablo Mayrgundter, Paranoidmage, Pbb, Peplm, Pion, Pol098, Poppafuze, Procedure, Promotie, Quibik, Qviri, Qxz, Rajat Bhargav, Rebroad, Rexcor, Richwales, Robertwalczyk, Royan, Sander123, Sarenne, Savvo, Saxbryn, ScwB, Secretlondon, Seipher, Seksek, Shawnc, Sibi antony, Sir Lothar, Smappy, Sophus Bie, Spinningspark, Struhs, StuffOfInterest, Tarikash, Taw, Tedickey, Tehnimax, The High Fin Sperm Whale, Thomas Willerich, Thumperward, Thunderbird2, Tide rolls, Tobias Hoevekamp, Tooki, Useight, VampWillow, Voidxor, W.F.Galway, Warren314, Wasabimaro, Wayne Hardman, Wernher, WhiteDragon, Wjejskenewr, Woz2, Zodon, 491 anonymous edits

DDR2 SDRAM *Source*: http://en.wikipedia.org/w/index.php?oldid=537524584 *Contributors*: A876, Aaronmz, Adrianwn, Algotr, Ali@gwc.org.uk, Alonduro, Although, Altonbr, Alveo, Amog, AndrewRH, Angelic Wraith, Angelpeream, AngryParsley, Annirak, Arny, BRUTE, Barneyboo, Basilicofresco, Benklop, Benstown, Bob Wiyadabebe-Iytsaboi, Bodhitha, Boffy b, Br77rino, BraneJ, Bryant1410, CTF83!, Caiaffa, Can't sleep, clown will eat me, Canadian-Bacon, Canaima, Catbar, Ceros, CesarB, Chasingsol, Chris S, Chrumps, Cootiequits, Courcelles, Crazytales, Crazyvas, Cristan, Cuvtixo, Cypherpunks, DARTH SIDIOUS 2, Damian Yerrick, Dancraggs, Darin-0, DaveJB, Davhorn, David Latapie, David Rolek, David.Monniaux, Davidfraser, Denniss, Desktiger, Devon nekechuk, Diablo-D3, Digisus, Dilane, DmitryKo, Drego5, Drongou, Dtcdthingy, Duckbill, Dysprosia, Długosz, ESkog, Eggman64, Electron9, Engineerism, Equinox76, F Tech, Facts707, Faisal imran, Favonian, Fnagaton, Fæ, Ganimoth, Gionnico, Greverod, Gth-au, Haham hanuka, Ham Pastrami, Hearingpocket, Hyperflux, Idamlaj, Intgr, JForget, JHunterJ, Jaakobou, Jaanus.kalde, Jacob.jose, Javalenok, Jdpipe, Jebus989, Jeffq, Jorunn, Jrdoane, Jts888, JulesH, Kadoo, Kar.ma, Kelledin, KennethJ, Kenyon, Klossner, Krash, Kyro, LAk loho, Lbecque, Leszek Jańczuk, Levander, Luigiacruz, M1ss1ontomars2k4, MSTCrow, Machinephilosophy, Makelelecba, Manning Bartlett, Manuishino, Martiniturbide, Martyx, Mattbrundage, MatthewWilcox, Meaghan, Mild Bill Hiccup, Mor, MorbiousStone, Mortense, Mr. PIM, Msikma, Mushroom, N2e, NJA, Nburden, Nczempin, Nikpapag, Nixdorf, Ohthelameness, Optikos, Orderud, Oxymoron83, PBurns, Parikhjigish, Pax85, Pearle, Peripitus, Phantomsteve, Pwnvds, Qatter, Qxl32, Qxz, R Man, RadetzkyVonRadetz, Randomtime, Reddyuday, Redline, Reinderien, Renaissancee, Rjwilmsi, RunOrDie, S.Örvarr.S, Sarenne, Sasquatch, Saviourmachine, Saxbryn, ScootrNova, Seksek, Shnout, SirRedK, Sirex98, SlackerMom, Spunch, Square87, Starwiz, Stassats, StephenFerg, SteveSims, Surgeprotector, Tansm, Tarikash, Tarquin, TastyPoutine, Terryeo, The Thing That Should Not Be, Themindset, Tide rolls, Tigga en, Twile, Ufundo, Victorrocha, Voidxor, Voltron, WarBaCoN, Warren314, Weedwhacker128, WhiteDragon, Xmath, Yagibear, Yamla, Yggdræsil, 519 anonymous edits

DDR3 SDRAM *Source*: http://en.wikipedia.org/w/index.php?oldid=536887397 *Contributors*: ACookr, AOK30, Adhirk, Airplaneman, Ajs1984, Alansohn, Ali@gwc.org.uk, Allstarecho, Alonduro, Angrytoast, Arch dude, Areseepee, AySz88, Bender235, Benjamin albert, Bettymnz4, Binba, BlindWanderer, Bob A, Bobblewik, Boemanneke, Boing! said Zebedee, Bowmanjj, C. Foultz, Can't sleep, clown will eat me, Cesiumfrog, Colin Douglas Howell, CombatWombat42, CommonsDelinker, Corvus cornix, Csigabi, D. Recorder, DVdm, Damian Yerrick, Danpritts, Darin-0, Dave laird, DaveJB, Defender of torch, Denniss, Diablo-D3, Digital infinity, DivineLight, Djgandy, DocWatson42, Drjt87, Drmies, Dual Freq, ELCleanup, Edetic, Ee79, Epbr123, Eugene-elgato, FT2, Facts707, Fatsamsgrandslam, Feneeth of Borg, Fernvale, FerrousTigrus, Finell, Fiskars007, FlyingPenguins, Fnagaton, Frap, GOV, Gameking123, Gavin86, Gcprakashh, Geekosaurus, Gildos, Giseong, Gogo Dodo, GregorB, GreyCat, Guy Harris, HallucigeniaUK, Ham Pastrami, Harizotoh9, Hdante, Hellcat fighter, Henk.muller, Hugowolf, IT DoNT, Iggee85, Iiaiiappa, Immunize, Isilanes, Itinerant1, Jaakobou, JavierMC, Jeffq, Jesse Viviano, Joey-das-WBF, Jts888, JustinRossi, KDesk, Kar.ma, Karam.Anthony.K, Kigali1, Kyng, LOL, Laptcd, Leandrod, Lezitor, Llloic, Locos epraix, Locriani, Ltwizard, Luigiacruz, Maarten Rail, Maghnus, Majidjanvaljan, Makelelecba, Markhahn, Marky540, Martiniturbide, Master-andra, Materialscientist, Mboehn, Megahmad, Mike.lifeguard, MilerWhite, Mindmatrix, MrESaulved!, MureninC, NJA, Nigeldh, Nikpapag, Ninjagecko, Noahspurrier, Nullzero, Olin Coles, Optikos, Otus, P2501, Paranoidmage, Pboyd04, Pedz221, Pfhorapedia, Pgk1, Picomp314, PluniAlmoni, QTCaptain, Quelrod, Ravagewing, Reddevil0728, Reddyuday, Rich Farmbrough, Rjwilmsi, Satusguy, Saxbryn, Sergei, SeriousWorm, Shadowjams, Shandris, Snoofer, Srikeit, Sterremix, Steveprutz, Superchad, TBrandley, Tarikash, Tax24cat, Therealdp, Thinpig, Thue, Timl2k4, Tomaxer, Tratten, Treedee, Versus22, ViperSnake151, Voidxor, W.F.Galway, Warren314, WhosAsking, Wikiwill, Winston Chuen-Shih Yang, XUnlimited, Z hosen, Zalgo, Zangoo, קיפודנחש, 461 anonymous edits

DDR4 SDRAM *Source*: http://en.wikipedia.org/w/index.php?oldid=534908959 *Contributors*: Akanago, Alereon, Anarij, Andros 1337, Arsenikk, Artem-S-Tashkinov, Attila Oláh, BD2412, Bender235, Bennettp123, Bobby122, Chris81w, Christian75, Colinstu, Cyberkagami, Damian Yerrick, Dark Bane, Darkwind, Dougweller, Drmies, Ee79, Electron9, FT2, Fram, Frap, Harizotoh9, Hello71, Hmains, HumphreyW, InternetMeme, Iqinn, J04n, Jadine, JamesAndersonJr, Jimmi Hugh, Jonjames1986, KDesk, Kigali1, Kyng, Lanny.ch, Lord Aro, Magioladitis, Malleus Fatuorum, Muleattack, Mydmdm, Nikpapag, Pabouk, Pfhorapedia, Phy1729, Rangeenbasu, Reify-tech, Rilak, Robbiemuffin, Rrohbeck, Ryan Postlethwaite, Ryrynz, Sivar, Telanis, Theopolisme, Thue, Walter Görlitz, XP1, 87 anonymous edits

Image Sources, Licenses and Contributors

Image:Original 1T1C DRAM design.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Original_1T1C_DRAM_design.svg *License*: GNU Free Documentation License *Contributors*: Cyferz

Image:square array of mosfet cells read.png *Source*: http://en.wikipedia.org/w/index.php?title=File:Square_array_of_mosfet_cells_read.png *License*: GNU Free Documentation License *Contributors*: Glogger at en.wikipedia Later version(s) were uploaded by Tene at en.wikipedia.

Image:256Kx4 DRAM.JPG *Source*: http://en.wikipedia.org/w/index.php?title=File:256Kx4_DRAM.JPG *License*: Creative Commons Attribution-ShareAlike 3.0 Unported *Contributors*: Original uploader was Redgrittybrick at en.wikipedia

Image:RAM n.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:RAM_n.jpg *License*: GNU Free Documentation License *Contributors*: Afrank99, Amada44, ArachanoxReal, Grendelkhan, Jpk, Mindmatrix, Museo8bits, NaSH, TOR, Tothwolf, Wagner51, WikipediaMaster, Wouterhagens, 10 anonymous edits

Image:Pair32mbEDO-DRAMdimms.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Pair32mbEDO-DRAMdimms.jpg *License*: Creative Commons Attribution-Sharealike 2.5 *Contributors*: User:Douglas Whitaker

File:MoSYS SGRAM.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:MoSYS_SGRAM.jpg *License*: Public Domain *Contributors*: Hyins

Image:SDR SDRAM-1.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:SDR_SDRAM-1.jpg *License*: Creative Commons Attribution 2.5 *Contributors*: Royan

Image:Micron 48LC32M8A2-AB.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Micron_48LC32M8A2-AB.jpg *License*: GNU Free Documentation License *Contributors*: Chepry, Jpk, MMuzammils, Solomon203

Image:Generic DDR Memory (Xytram).jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Generic_DDR_Memory_(Xytram).jpg *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: Martyn M aka Martyx

Image:Corsair DDR PC-3200 Memory (Xytram).jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:Corsair_DDR_PC-3200_Memory_(Xytram).jpg *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: Martyn M aka Martyx

Image:Desktop DDR Memory Comparison.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Desktop_DDR_Memory_Comparison.svg *License*: Public Domain *Contributors*: Martini

Image:DDR layout sketch.png *Source*: http://en.wikipedia.org/w/index.php?title=File:DDR_layout_sketch.png *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: Tehnimax

Image:Laptop SODIMM DDR Memory Comparison V2.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Laptop_SODIMM_DDR_Memory_Comparison_V2.svg *License*: Public Domain *Contributors*: Martini, Tothwolf

Image:High Performance RAM.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:High_Performance_RAM.jpg *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: Victorrocha

Image:1GB DDR2 SO-DIMM.png *Source*: http://en.wikipedia.org/w/index.php?title=File:1GB_DDR2_SO-DIMM.png *License*: Creative Commons Attribution-Sharealike 3.0 *Contributors*: 1029man, Afrank99, Denniss, Kyro

Image:Desktop_DDR_Memory_Comparison.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Desktop_DDR_Memory_Comparison.svg *License*: Public Domain *Contributors*: Martini

Image:Laptop_SODIMM_DDR_Memory_Comparison_V2.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Laptop_SODIMM_DDR_Memory_Comparison_V2.svg *License*: Public Domain *Contributors*: Martini, Tothwolf

File:4GB DDR3 SO-DIMM.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:4GB_DDR3_SO-DIMM.jpg *License*: Creative Commons Attribution-Sharealike 3.0,2.5,2.0,1.0 *Contributors*: Tobias b köhler

File:Desktop DDR Memory Comparison.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Desktop_DDR_Memory_Comparison.svg *License*: Public Domain *Contributors*: Martini

File:Laptop SODIMM DDR Memory Comparison V2.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:Laptop_SODIMM_DDR_Memory_Comparison_V2.svg *License*: Public Domain *Contributors*: Martini, Tothwolf

License