

DDR Memories

Comparison and overview

We realize that one of the most important aspects of a computer is its capability to store large amounts of information in what we normally call “memory.” Specifically, it’s random access memory (RAM), and it holds volatile information that can be accessed quickly and directly. And considering the ever growing system need for speed and efficiency, understanding double-data-rate (DDR) memory is important to system developers.

With improvements in processor speeds, RAM memory has evolved into high performance RAM chipsets called DDR synchronous dynamic RAM (SDRAM). It doubles the processing rate by making a data fetch on both the rising and falling-edge of a clock cycle. This is in contrast to the older single-data-rate (SDR) SDRAM that makes a data fetch on only one edge of the clock cycle.

In addition to well-known computer applications, DDR memories are widely used in other high-speed, memory-demanding applications, such as graphic cards, which need to process a large amount of information in a very short time to achieve the best graphics processing efficiency. Blade servers using many blades, or single purpose boards, powered by a single, more efficient power supply also need fast memory access. This allows the blades to quickly transmit reliable information among each other and create greater opportunities to reduce power consumption. Memory devices are also required in networking and communications applications with tasks ranging from simple address lookups to traffic shaping/policing and buffer management.

This article describes the main characteristics of DDR memories as well as the specifications of power supplies required for these types.

DDR Memory Characteristics

DDR memory’s primary advantage is the ability to fetch data on both the rising and falling edge of a clock cycle, doubling the data rate for a given clock frequency. For example, in a DDR200 device the data transfer frequency is 200 MHz, but the bus speed is 100 MHz.

DDR1, DDR2 and DDR3 memories are powered up with 2.5, 1.8 and 1.5V supply voltages respectively, thus producing less heat and providing more efficiency in power management than normal SDRAM chipsets, which use 3.3V.

Temporization is another characteristic of DDR memories. Memory temporization is given through a series of numbers, such as 2-3-2-6-T1, 3-4-4-8 or 2-2-2-5 for DDR1. These numbers indicate the number of clock pulses that it takes the memory to perform a certain operation—the smaller the number, the faster the memory.

The operations that these numbers represent are the following: CL- tRCD – tRP – tRAS - CMD. To understand them, you have to keep in mind that the memory is internally organized as a matrix, where the data is stored at the intersection of the rows and columns.

- CL: Column address strobe (CAS) latency is the time it takes between the processor asking memory for data and memory returning it.
- tRCD: Row address strobe (RAS) to CAS delay is the time it takes between the activation of the row (RAS) and the column (CAS) where data is stored in the matrix.
- tRP: RAS precharge is the time between disabling the access to a row of data and the beginning of the access to another row of data.
- tRAS: Active to precharge delay is how long the memory has to wait until the next access to memory can be initiated.
- CMD: Command rate is the time between the memory chip activation and when the first command may be sent to the memory. Sometimes this value is not informed. It usually is T1 (1 clock speed) or T2 (2 clock speeds).

Table 1 is a comparison of clock and transfer rates for the RAM memory chipsets that can be found in today’s computers, including SDR, DDR, DDR2 and future DDR3 modules.

Table 1: SDRAM Memories Speed Comparison

Memory	Technology	Rated Clock	Real Clock	Maximum Transfer Rate
PC66	SDRAM	66 MHz	66 MHz	533 MB/s
PC100	SDRAM	100 MHz	100 MHz	800 MB/s
PC133	SDRAM	133 MHz	133 MHz	1,066 MB/s
DDR200	DDR-SDRAM	200 MHz	100 MHz	1,600 MB/s
DDR266	DDR-SDRAM	266 MHz	133 MHz	2,100 MB/s
DDR333	DDR-SDRAM	333 MHz	166 MHz	2,700 MB/s
DDR400	DDR-SDRAM	400 MHz	200 MHz	3,200 MB/s
DDR2-400	DDR2-SDRAM	400 MHz	200 MHz	3,200 MB/s
DDR2-533	DDR2-SDRAM	533 MHz	266 MHz	4,264 MB/s
DDR2-667	DDR2-SDRAM	667 MHz	333 MHz	5,336 MB/s
DDR2-800	DDR2-SDRAM	800 MHz	400 MHz	6,400 MB/s
DDR3-800	DDR3-SDRAM	800 MHz	400 MHz	6,400 MB/s
DDR3-1066	DDR3-SDRAM	1066 MHz	533 MHz	8,528 MB/s
DDR3-1333	DDR3-SDRAM	1333 MHz	666 MHz	10,664 MB/s
DDR3-1600	DDR3-SRAM	1600 MHz	800 MHz	12,800 MB/s


Types of DDR Memories

There are presently three generations of DDR memories:

1. DDR1 memory, with a maximum rated clock of 400 MHz and a 64-bit (8 bytes) data bus is now becoming obsolete and is not being produced in massive quantities. Technology is adopting new ways to achieve faster speeds/data rates for RAM memories.
2. DDR2 technology is replacing DDR with data rates from 400 MHz to 800 MHz and a data bus of 64 bits (8 bytes). Widely produced by RAM manufacturers, DDR2 memory is physically incompatible with the previous generation of DDR memories.
3. DDR3 technology picks up where DDR2 left off (800 Mbps bandwidth) and brings the speed up to 1.6 Gbps. One of the chips already announced by ELPIDA contains up to 512 megabits of DDR3 SDRAM, with a column access time of 8.75 ns (CL7 latency) and data transfer rate of 1.6 Gbps at 1.6 GHz. The 1.5V DDR3 voltage level also saves some power compared to DDR2 memory. What is more interesting is that at an even lower 1.36V, the DDR3 RAM runs fine at 1.333 GHz (DDR3-1333) with a CL6 latency (8.4 ns total CAS time), which matches the CAS time of the fastest current DDR2 memory.

Figure 2.0.1 shows the on die termination (ODT) for DDR2/DDR3 memory types compared to the motherboard termination of DDR1, which is the primary reason why the first two types are physically incompatible with DDR1 devices.

Figure 2.0.1: Termination Differences between DDR-I and DDR-II



DDR2 versus DDR3

The primary differences between the DDR2 and DDR3 modules are:

- DDR2 memories include 400 MHz, 533 MHz, 667 MHz and 800 MHz versions, while DDR3 memories include 800 MHz, 1066 MHz, 1333 MHz and 1600 MHz versions. Both types double the data rate for a given clock frequency. Therefore, the listed clocks are nominal clocks, not real ones. To get the real clock divide the nominal clock by two. For example, DDR2-667 memory in fact works at 333 MHz.
- Besides the enhanced bandwidth, DDR3 also uses less power than DDR2 by operating on 1.5V—a 16.3 percent reduction compared to DDR2 (1.8V). Both DDR2 and DDR3 memories have power saving features, such as smaller page sizes and an active power down mode. These power consumption advantages make DDR3 memory especially suitable for notebook computers, servers and low power mobile applications.
- A newly introduced automatic calibration feature for the output data buffer enhances the ability to control the system timing budget during variations in voltage and temperature. This feature helps enable robust, high-performance operation, one of the key benefits of the DDR3 architecture.
- DDR3 devices introduce an interrupt reset for system flexibility
- In DDR2 memories, the CL parameter, which is the time the memory delays delivering requested data, can be three to five clock cycles, while on DDR3 memories CL can be of five to ten clock cycles.
- In DDR2 memories, depending on the chip, there is an additional latency (AL) of zero to five clock cycles. So in a DDR2 memory with CL4 the AL1 latency is five.
- DDR2 memories have a write latency equal to the read latency (CL + AL) minus one.
- Internally, the controller in DDR2 memories works by preloading 4 data bits from the storage area (a task known as prefetch) while the controller inside DDR3 memories works by loading 8 bits in advance.

Typical DDR Power Requirements

VTT and VREF Considerations for DDRx SDRAM Devices

The termination voltage (VTT) supply must sink and source current at one-half output voltage ($\frac{1}{2} V_{DDQ}$). This means that a standard switching power supply cannot be used without a shunt to allow for the supply to sink current. Since each data line is connected to VTT with relatively low impedance, this supply must be extremely stable. Any noise on this supply goes directly to the data lines.

Sufficient bulk and bypass capacitance must be provided to keep this supply at $\frac{1}{2} V_{DDQ}$. This same noise issue requires that the voltage reference (VREF) signal cannot be derived from VTT, but instead must be derived from VDDQ with a one percent or better resistor divider. Do not try to generate VREF with one divider and route it from the controller to the memory devices. Generate a local VREF at every device. Use discrete resistors to generate VREF. Do not use resistor packs.

DDR Memories Power Management

DDR1 memory has a push-pull output buffer, while the input receiver is a differential stage requiring a reference bias midpoint, VREF. Therefore, it requires an input voltage termination capable of sourcing as well as sinking current. This last feature differentiates the DDR VTT from other terminations present on the computer motherboard. The termination for the front system bus (FSB), connecting the CPU to the memory channel hub (MCH), requires only sink capability due to its termination to the positive rail. Hence, such DDR VTT termination can't re-use or adapt previous VTT termination architectures and requires a new design. Figure 3.2.1 shows the typical power management configuration for a DDR1 memory.

Between any output buffer from the driving chipset and the corresponding input receiver on the memory module, you must terminate a routing trace or stub with resistors R_T and R_S (Figure 3.2.1). Accounting for all the impedances, including the output buffers, each terminated line can sink or source ± 16.2 mA (this is more than the older value of ± 15.2 mA, per the June 2000 Revision of JESD79). For systems with longer trace lengths, between transmitter and receiver, it may be necessary to terminate the line at both ends, which doubles the current.

Peak and average current consumption for VTT and VDDQ are two parameters for the correct sizing of our power supply system. To find the peak power requirements for the termination voltage, we must determine the total lines in the memory system.

Figure 3.2.1: DDR Power Management




Table 2: SDRAM Devices Comparison

Items	DDR3 SDRAM	DDR2 SDRAM	DDR SDRAM
Clock frequency	400/533/667/ 800 MHz	200/266/333/400 MHz	100/133/166/200 MHz
Transfer data rate	800/1066/1333/ 1600 Mbps	400/533/667/800 Mbps	200/266/333/400 Mbps
I/O width	x4/x8/x16	x4/x8/x16	x4/x8/x16/x32
Prefetch bit width	8-bit	4-bit	2-bit
Clock input	Differential clock	Differential clock	Differential clock
Burst length	8, 4 (Burst chop)	4, 8	2, 4, 8
Data strobe	Differential data strobe	Differential data strobe	Single data strobe
Supply voltage	1.5V	1.8V	2.5V
Interface	SSTL_15	SSTL_18	SSTL_2
/CAS latency (CL)	5, 6, 7, 8, 9, 10 clock	3, 4, 5 clock	2, 2.5, 3 clock
On die termination (ODT)	Supported	Supported	Unsupported
Component package	FBGA	FBGA	TSOP(II) / FBGA / LQFP

Power Sequencing Violations

DDR memories must be powered up in a specific manner. Any violation of the power up sequencing will result in undefined operations. Also, power down sequencing serves as a power saving tool when the DDR device needs to be shut down without all other devices connected to the same power supply. Violating these sequences will present poor power saving results.

Typical Applications of DDRx Memories

Market analyses indicate that DDR is currently utilized in over 50 percent of all electronic systems, and usage is expected to increase to 80 percent over the next several years. DDR is not, and will never be, an “all things to all designs” technology. DDR memory is well suited for those designs that have a high read to write ratio. Quad-data-rate memory, for example, is designed for applications that require a 50 percent read/write ratio.

Today, DDR memories are used mainly for computer applications in dual in-line memory module (DIMM) RAM devices. DDR memories can be used when interfacing with 32-bit microcontrollers and DSPs. Since many memories work with a 64-bit data bus, microcontrollers have to make a double data acquisition to get the less significant 32 bits first and then the more significant 32 bits. DDR memories have also been used to interface with FPGAs, giving those devices wide programming flexibility. FPGAs are often used to customize applications which combine many digital modules, such as a microcontroller with specific application hardware, USB controllers and printing modules. It is also common to find FPGAs used specifically as memory controllers to interface with

other devices. DDR memories are suitable for these types of devices because they are fast and inexpensive compared with other RAM memories.

Conclusions

To remain competitive, you have to create new designs that take advantage of today's advanced technology. To understand the best solution for a specific design challenge, it is important to research well, giving you the opportunity to weigh the advantages of technologies, such as DDR memory, to incorporate them in new product developments. This article gives you a brief and useful overview of DDR memories, describing how they are different from other memories to give you a good starting point in choosing the best solution for your specific needs.

References

- JEDEC STANDARD JESD79, June 2000 and JESD8-9 of Sep.1998.
- Understanding Ram Timings (Article)
www.hardwaresecrets.com/article/26/1
- DDR Memories require efficient power management. (Article)
powerelectronics.com/mag/power_ddr_memories_require/
- ELPIDA 1 Gb DDR2 SDRAM memory datasheet

Joaquin Romo is an application engineer working in the Power Management and Motor Control Applications organization. Romo is in charge of characterizing new power management products and developing application notes and evaluation boards focused on the consumer electronics market.