COMIT® SYSTEMS
The Contract Engineering Company

Evaluating DDR2 SDRAM Controller Cores

With a Comparison of DDR Memories and an Evaluation Checklist

Double Data Rate SDRAM (DDR1 and DDR2) is a variant of SDRAM where data is transferred twice in one cycle. DDR1 is available at speeds from 200MHz to 400MHz, while DDR2 is available from 400MHz to 800MHz. The theoretical maximum bandwidth for 200MHz DDR is 3.2 GB/sec (200MHz x 2 x 8 bytes).

Other types of available memories include GDDR (optimized for graphics applications), RDRAM (Rambus DRAM, proprietary to Intel and Rambus, capable of running at 800MHz), XDR (Extreme data rate, currently sampling 3.2GHz devices, 4.8GHz expected by end of 2005, 8GHz in the future) and RLDRAM, to name a few. This article will focus on controllers for DDR1 and DDR2 SDRAM.

Some DDR1 and DDR2 related terms and conventions

SDRAM is referred to by the clock speed (PC100, PC133, and so on). DDR and the latest Rambus chips are labeled using the maximum theoretical bandwidth (PC1600, PC2100, and so on) - largely for marketing reasons. Given below are contemporary PC standards, circa 2005

PC2100 (266MHz)
PC2700 (333MHz)
PC3200 (400MHz)
PC3500 (433MHz)
PC3700 (466MHz)
PC4000 (500MHz)
PC4200 (533MHz)
PC4400 (550MHz)

The number in brackets represents the data rate (twice the clock rate), and the number immediately following PC is simply eight times the data rate, rounded up or down to the nearest hundred.

PC memory is available in the following configurations:

•Single Inline Memory Module (SIMM)
•Dual Inline Memory Module (DIMM)
•Buffered/Non-buffered DIMMS
•Registered DIMM
•SO-DIMM for laptops

Selecting a DDR2 Controller

No IP core is ever plug and play. This is especially true for a DDR2 controller, which is always matched carefully to the system in which it will be deployed.

Chips accessing memory have different requirements in terms of bandwidth, latency, read/write request size etc. Usually, there will be many blocks of a chip that will be competing for memory access. Along with providing each of them with the necessary bandwidth, the memory manager needs to maintain a fair balance among their requests. The number of blocks accessing the memory and the bandwidth/latency requirement of each block is different in different chips. All of these combine together to complicate the DDR2 controller.

Here are three examples of factors that differentiate DDR2 controller IP from other types of IP:

1. **Arbitration:** The number of blocks accessing memory vary from chip to chip, or, even within a chip from one functionality to another. So one needs to a have a fair arbitration algorithm that can accommodate all the blocks. It should also be able to keep an equitable balance between accesses from different blocks, while satisfying the bandwidth needs of each block.

2. **Clock crossing per agent:** The memory interfaces from the individual blocks of a chip accessing the memory may run at different speeds. It is the memory controller's job to ensure that all clock crossing issues are taken care of.

3. **Write/Read buffers per agent:** No wait states can be added to the data flow at memory. Once a read request is given to the memory, the requestor should be ready to accept all the data words fetched because of the read request. Similarly for write, once write data starts flowing to memory, no waits can be introduced.

All DDR memory controller IP cores are not created equal, and different memory controllers from different vendors can be vastly different in the features that contribute to effective utilization of DDR bandwidth. At the minimum, it is wise to check for support of the following:

a. **Programmability of DDR timing parameters:** This is extremely important from a product-market standpoint, as one never knows whether a particular manufacturer's memory that one's product has been using will suddenly be in short supply. Experience has shown that entire systems can stop working due to minor differences in memories from different manufacturers. A memory controller that can be dynamically configured to match variations in memory timing can help one's product get to market, since one's chip can be made to coexist with different manufacturers' memories
b. **Configurable address mapping:** This is important if one is to be able to optimize DDR bandwidth. Configurable address mapping enables playing around with various mapping schemes till the one that provides the maximum throughput for a particular application is found
c. **Command pipeline:** Effective command pipelining will also result in optimal use of DDR bandwidth. Command pipeline, although supported by most controllers, differs widely in quality of implementation from vendor to vendor. The best way to determine how efficient a particular implementation is, is to provide potential vendors with a defined sequence of accesses and analyze the waveforms produced by their respective controllers
d. **JEDEC / Custom initialization support:** This may sound both similar and contrary to the programmability advantage at the same time. Having programmability, adherence to JEDEC and custom, together in one controller, offers many advantages. The ability to custom initialize ensures that one can implement all those whiz-bang tricks of one's own, while JEDEC support will ensure that one will be able to easily incorporate support for

incremental advancements in DDR specs. Depending on the level of programmability offered, it will be easy to add support for, say, GDDR memories instead of DDR.

e. **Self refresh / Pre-charge:** For low power designs, we will want to shut off the clock to the memory, and sometimes all other signals as well. Without going into the details of how each is implemented, suffice is to say that one needs self-refresh and pre-charge if power is at a premium.
f. **ODT (On-die termination) support:** Termination of signals is necessary in order to eliminate bouncing and ringing that occurs whenever a signal hits an interface in its path. In that case, the signal is reflected into the bus where the reflection will cause interference with any real signal. The proper method to eliminate this kind of ringing is to terminate the reflection by adding termination to the path. ODT implements termination resistors on the memory die, rather than on the board. On board terminations add stubs on board, causing signal integrity problems. This reduces the available margin for the data-valid window. This is why DDR2 memories provide termination on die, eliminating stubs, saving space and enabling dynamic switching of termination (controller side on read, DDR side on write). Having more than one DIMM on board complicates ODT control. The DDR2 controller should be able to use ODT effectively.
g. **2T support:** When 2T is on, address and most control signals are generated one clock in advance (hence, *2*T, or two time periods), providing better signal margins on the board
h. **OCD (Off-Chip Driver calibration):** This is important for some applications. With OCD, both parts of the differential strobes are calibrated against each other and against the DQ signal. Through this sort of calibration, the ramping voltages are optimized for the buffer impedances to reduce over and undershooting at the rising and falling edges. DQS and /DQS are matched so that their cross point coincides with the DQ signal crossing the reference voltage to eliminate DQ-DQS skew. The scheme results in better compatibility between different memory designs, higher signal integrity through minimization of DQ-DQS skew and reduced overshoot / undershoot for better signal quality.

There are a couple more things that have to be carefully considered and analyzed – the **DDR PHY** inside the chip, and the **board design**. Since at the speed of 400MHz (800Mbps data rate), the margins are very little, these two components must be fine tuned to each other to meet timing. Support from IP vendors is a must.

Comit DDR2 SDRAM Controller Core

The Comit DDR2 SDRAM Controller core is, at once, feature-rich and optimized for high bandwidth, providing cutting edge performance, with a low, 2-clock, latency. The core supports operating frequency of up to 333MHz in TSMC0.13G, with 0.2ns clock uncertainty and wire load model. Higher operating frequency is possible at 90 nanometers. FPGA performance figure on the Xilinx Virtex-II Pro is 200MHz, while that for the Virtex-4 device is 266MHz. The core is highly programmable, allowing available features to be modified by the user, simply by changing parameters.

The controller supports both DDR1 and DDR2 memory: buffered/non-buffered, single or dual rank DIMMs with 4 or 8 banks, and configurable data widths, 2X or 4X of application data width, and a generic application interface. Address mapping between application and DDR2 memory is configurable. It supports a 4-deep command pipeline, with the order of transactions being maintained. Timing parameters such as trcd, tras, trfc, trp, twr, twtr etc. are programmable. DQ:DQS ratio is 8:1

The controller provides maximum request size as large as 1024Bytes, and supports address space of 4GB, expandable up to 16GB. Further, it supports auto refresh, self-refresh and pre-charge power down. Initialization can be either JEDEC or programmable or both, and the core supports write byte masking. Read and write CAS latencies are programmable, and the core supports both sequential and interleaved bursts. ECC is supported if desired.

Among the DDR2 special features, ODT is supported. The Comit DDR2 SDRAM Controller also provides a 2T feature for better signals margin on board.

The feature-set to be implemented is mostly customizable, allowing for flexibility of further boosting performance by configuring only the needed features. Several features are changeable at run-time, while others can be configured at initialization.

With high-bandwidth, low-latency and programmable feature set, the Comit DDR2 SDRAM Controller Core offers ultra-fast performance, high configurability and reliable integration.

DDR Application Module (DAM)

The DAM is an add-on to the Comit DDR2 controller. It is an interface between the application side and the controller, enabling additional feature enhancements without having to modify the controller core. For example, the DAM allows command queue depth to be easily increased, read/write buffers to be added, and read-modify-write to be supported. It also provides a facility to pump in many write commands along with write data, optimizing bandwidth. One or more bits inside a word can be modified, the application and DDR controller can run at different clocks, the address from the application can be non 64Byte aligned, more agents and arbitration can be configured, and BIST incorporated.

The DAM enables flexibility and enhancement while retaining the reliability of integrating a proven controller core.

Comparing SDRAM, DDR1 & DDR2 memories

SDRAM Vs DDR1

	Clock	Data Transfer	Source synchronous	I/O
SDRAM	Single ended	Single edge of clock	No	
DDR1	Differential	Both edges of clock	Yes. Data strobe signals added. Data has same phase relationship with the strobe across the system	SSTL-2 (2.5V)

DDR implementation also provides for more data pre-fetch from core for more speed

DDR1 Vs DDR2

	DDR1	DDR2
Frequency		
Data Rate	200/266/333/400 Mbps*	400/533/667/800 Mbps*
Bus Frequency	100/133/166/200 MHz	200/266/333/400 MHz
DRAM Core Frequency	100/133/166/200 MHz	100/133/(166) MHz
Prefetch Size	2 bit	4 bit
Burst Length	2/4/2008	4/8**
Data Strobe	Single DQS	Differential Strobe: DQS, /DQS***
CAS Latency	1.5, 2, 2.5	3+, 4, 5
Write Latency	1T	Read Latency-1
Power		
Core Voltage (VDD)	2.5V. DDR400 uses 2.6V	1.8V
I/O Voltage (VDDQ)	SSTL_2 (2.5V)	SSTL_1.8 (1.8V)
Footprint		
Packaging	TSOP (II), TBGA	FBGA
Backward Compatibility		
Command Set		Same as DDR 1
Timing Parameters		Same as DDR I
Bus Utilization and Signal Integrity		
Special Features		ODT OCD-calibration Posted CAS Additive Latency+++

* Megabit/pin/sec
** The original burst length was defined as 4 QW, however, a burst of 8 QW has been added.
*** DDR1 only uses a single DQS, using the cross point with the reference voltage. DDR2 uses a differential DQS. Optionally, single ended DQS may be enabled using MRS settings
+ CAS-3 is possible using a 533 MHz CAS-4 speed bin in DDR400 mode.
++ All current DDR1 components are running at 1.8V internally, using voltage regulators to reduce VDD from 2.5V to 1.8V.
+++ Additive Latency can be 0,1,2,3,4T, the actual Read Latency is the sum of CAS latency and Additive Latency, e.g. CAS-4 + AL-4 =8T Read Latency.

(Adapted from *Lostcircuits*)

Some recent (circa 1Q05) updates to DDR2

50 ohm ODT (previously 75/150 ohms)
OCD being de-emphasized
Edge rate de-rating
SSTL_18 signaling
Setup/Hold numbers in DDR are not from mid point

The setup/hold numbers has to be de-rated as slew rate changes
Set-up = tDS + deltaDS (derating) + tTS (transition from Vref)
Hold = tDH + deltaDH (derating) +tTH (transition from Vref)

Important DDR2 Timing Parameters

There are several timing parameters that can be important to a DDR2 controller selection, depending on the particular implementation planned. A comprehensive list can be found in DDR2 specs, and usually on vendor data sheets. Here are a few important ones

tRCD (Row Address Strobe to Column Address Strobe Delay): As the name suggests this is the delay between RAS to CAS. This is generally 3,4 or 5 clock cycles.

tCL (Cas Latency): This is the time interval between issuance of a CAS command and the instant in time that the DDR2 data is available at the DDR2 pins. This parameter is frequency dependent. For DDR2, tCL is 3, 4 or 5 clocks. No half clock cycle latencies are supported in DDR2, but, in DDR1, 2.5 clock latency is also available.

tRP (Precharge time): This is the time taken to precharge an open page. The controller has to wait for this time before it can open any page in the same bank after the precharge command. For DDR2, this is generally 4 or 5 clocks.

tRAS (Active to Precharge command delay): This is the time interval between issuance of Active command to the issuance of Precharge command to the same bank. This has minimum and maximum values set for it. The controller cannot precharge the bank before the minimum value has elapsed, and must do so within the maximum time limit. The minimum value is 45ns, typically, although it may vary. The maximum is 70,000ns.

tRFC (Refresh to Refresh/Active command delay): This is the time interval between Refresh to the next Refresh, or Active command. In short, the controller must wait for tRFC time interval after the Refresh command before it can issue the next command to the DDR memory. The value for tRFC in JEDEC standard is defined as follows:

For DDR2:

256Mb : 75ns
512Mb : 105ns
1Gb : 127.5ns
2Gb : 195ns
4Gb : TBD

For DDR1: 70ns to 120ns.

tWR (Write Recovery time): The time needed for the actual write to happen to sense amplifier after the end of a write burst. The value is 15ns for DDR2. For DDR1, this time is measured from the last intended write data.

tWTR (Write to READ command time): The time between end of a WRITE burst to the READ command. For DDR1, this time is measured from the last intended write data.

tRRD (RAS to RAS delay): This is the minimum time between two RAS commands. Typically, 2-3 clocks.

tCCD (CAS to CAS delay): This is the minimum time between two CAS commands. For DDR2, this is 2 clocks.

tRTP (Read to Precharge delay): This is the minimum time between the end of a READ burst to Precharge (to the same bank) command. This is only applicable to DDR2.

A Checklist for DDR2 Controller Evaluation

Given below is a checklist that can be used to tabulate the results of an evaluation of DDR2 controllers. This is based on the Comit DDR2 Controller, but can be modified to suit particular application evaluation needs. The format easily lends itself to setting up as a spreadsheet, with columns added to the right for multiple vendor offerings

Features	Must Have	Desirable	Nice to Have
DDR Data Width (bits)	Insert Value	Insert Value	Insert Value
Application data width (Enter **2X** or **4X** of DDR width)			
Burst length (Fixed. 8 DDR words)	8	8	8
Minimum application request size (Bytes)			
Maximum application request size(Bytes)	1024	1024	1024
Command queue depth (Fixed.)	4	4	4
Desired number of Ranks support (Enter **1** or **2**)	Insert Value	Insert Value	Insert Value
DQ:DQS ratio (Fixed.)	8:1	8:1	8:1
Max addressable Memory needed (Enter one of **4**GB, **8**GB or **16**GB)			
DDR Features:			
-Self Refresh (Enter **Y**ES or **N**O)			
-Auto Refresh (Enter **Y**ES or **N**O)			
-Auto precharge (Enter **Y**ES or **N**O)			
-Precharge power down (Enter **Y**ES or **N**O)			
-Programmable CAS Latency (Enter **Y**ES or **N**O)			
-Power up initialization sequence (Enter **J**edec or **P**rogrammable or **B**oth)			
-Word mask desired on write (Enter **Y**ES or **X** = Don't care)			
ECC desired (Enter **Y**ES or **N**O: 1Bit correct & 2 bit detect hamming code)			
ODT desired (Enter **Y**ES or **N**O)			
2T feature desired (Enter **Y**ES or **N**O)			
DDR1/DDR2 (Enter **1** or **2** or **B**oth)			
DIMM Types (Enter Bu**F**fered, **N**on-buffered or **B**oth)			
Optional Extras: (Leave blank where not desired)			
-Read Modify Write desired (Enter **Y**ES or **N**O)			
-Write Byte Mask with ECC desired (Enter **Y**ES or **N**O)			
-Application clock frequency different from DDR (Enter **Y**ES or **N**O)			
-BIST desired (Enter **Y**ES or **N**O)			
-Variable command queue depth	Insert Value	Insert Value	Insert Value
-Write Buffer depth	Insert Value	Insert Value	Insert Value
-Read Buffer depth	Insert Value	Insert Value	Insert Value
-No of application agents (Enter a value from **1 thru 16**)	Insert Value	Insert Value	Insert Value

Evaluating DDR2 SDRAM Controller Cores

3375 Scott Blvd. Suite 139
Santa Clara, CA 95054 USA
(408) 988-2988

www.comit.com