

Design Optimization Techniques for Double Data Rate SDRAM Modules

Abstract

Just as PC133 SDRAM has gained widespread acceptance in high performance computing systems, Double-Data-Rate (DDR) SDRAM is fast becoming the memory technology of choice in computing systems from Internet servers to personal notebooks. This new memory standard has moved to a source-synchronous architecture to increase memory system bandwidth. In a source-synchronous design both data and clock are sent from a transmitter to a receiver. The self-timed nature of this type of architecture allows for higher data-rates and bandwidth to be achieved across the memory interface.

This paper will discuss the timing specifications and supporting ICs that comprise the DDR standard. In addition, the paper will give system designers insight into possible design enhancement techniques that result in improved module timing and reliability.

Introduction

Over the past decade, system memory bandwidth has continued a steady climb, achieving over 1GByte/sec with the introduction of PC133 memory in 1998. Processors in the not to distant future will easily surpass 1GHz operation. This type of processing power has spurred revolutionary developments that will continue to push memory systems faster and faster.

Even today's developments in chipset and motherboard design have pushed beyond the bandwidth of conventional (PC100/PC133) SDRAM; the next stage of evolutionary migration for standard DRAMs is double data rate (DDR) SDRAM. By taking an evolutionary approach to doubling memory bandwidth incremental product costs have been well contained. This sets the stage for widespread market acceptance of products with the right cost vs. performance tradeoff for main memory in platforms such as PC desktops, servers, and workstations.

Double Data Rate SDRAM

Many similarities exist between DDR SDRAMs and standard SDRAMs, these include packaging, command / address protocols, and connector design. These similarities leverage off existing infrastructure and allow for the same test equipment, handlers, and back-end support.

Figure 1 is a block diagram of a typical PC266 main memory system, suitable for PC desktops, workstations, and servers. There are variations to this diagram, based

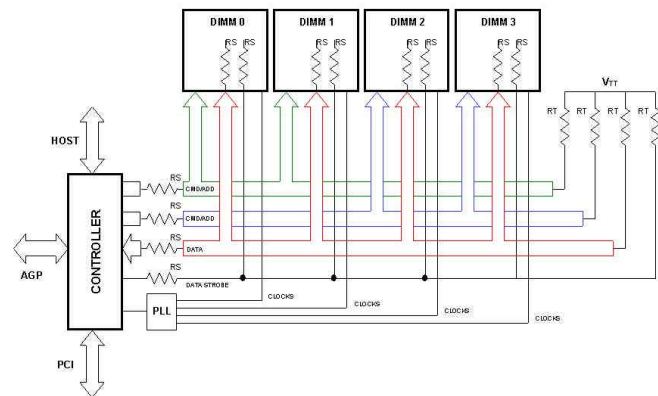


Figure 1: DDR PC266 Memory Subsystem

on the number of modules supported by the system, the type of module system (non-buffered, registered, or both), and the designer's choice of clock topologies. For the majority of desktop systems only one command / address bus is required. As the memory size is increased, multiple address / command buses are required or a single one needs to be multiplexed, in order to meet system timing requirements. The controller chip is the PC North Bridge, and it provides communication and control to the PCI adapter bus, the AGP bus, the microprocessor host bus, and to the main memory.

Command, address, and data signals are routed to the dual in-line memory modules (DIMM) from the motherboard chipset. The data signals include series resistors on the DIMMs and on the motherboard near the first DIMM socket, and are parallel terminated to V_{TT} after the last socket. The data signals are double data rate and transfer on both edges of the clock signals, while the command/address signals will only switch on the positive edge of the clock signal. The DDR SDRAM input specifications allow command / address signals to be received as either SSTL-2 or 2.5V LVCMOS signals.

SSTL-2 stands for Series Stub Terminated Logic and has been defined and standardized within JEDEC. Although it is applicable for many different applications, SSTL-2 has been optimized for the main memory environment, which has long stubs off the motherboard bus due to the DIMM routing traces.

The SSTL-2 specification requires adequate output drive so that the parallel termination schemes can be implemented effectively. This is important for high-speed signaling, since it allows proper termination of the bus transmission lines, reducing signal reflections. The result is improved signal quality, higher clock frequencies, and lower EMI emissions. Series resistors are incorporated in the SSTL-2 signaling technology for main memory applications. The data lines on the DIMMs use 22Ω stub resistors near the connector contacts to improve the signal integrity of the system. These resistors can be very effective in dissipating any reflected wave energy traveling along the module traces and isolating the module stubs from the main memory bus.

SSTL-2 inputs are typically a differential pair common source amplifier with one input tied to the V_{TT} reference. This type of circuit provides excellent gain and bandwidth. The topology also works to minimize variation in the threshold voltage over process, voltage, and temperature variations. The result is that smaller input voltage swings can be used reliably.

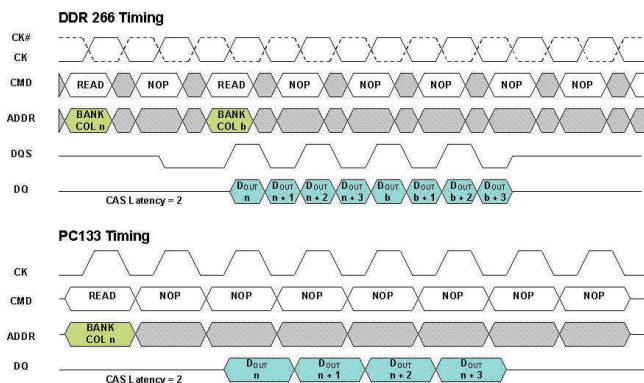


Figure 2: Source Synchronous Advantage

The advantages of source synchronous timing can be seen in the waveform diagrams in Figure 2 which show basic "read" operations from memory. The use of a data strobe signal to self time data arriving back at the memory controller tightens the timing equations sufficiently to effectively send or receive data on both edges of the differential clock. Even though the DDR memory chips both use a CAS latency of 2 cycles, the DDR memory system, with the same basic clock rate,

can offer up to twice the effective throughput of the synchronous system as shown in the diagram.

The JEDEC DDR specification defines the electrical and mechanical requirements for 184-pin, 2.5 Volt, PC200/PC266, 72-bit wide, Registered Double Data Rate Synchronous DRAM Dual In-Line Memory Modules (DDR SDRAM DIMMs). These SDRAM DIMMs are intended for use as main memory when installed in systems such as servers and workstations.

Support chip specifications included in this standard provide an initial basis for Registered DIMM designs. Enhancements or modifications to the initial set of reference designs may be required to meet all system timing, signal integrity and thermal requirements for PC200/PC266 support. All registered DIMM implementations must use simulations and lab verification to ensure proper timing requirements and signal integrity in the design.

Fairchild Semiconductor supplies SSTV16857 Registers and FM34W02 Serial Presence Detect EEPROM Devices which meet the performance attributes set out in the JEDEC standard.

Clocking

The clock to SDRAM delay is intended to be optimized for high speed operation, while permitting a variety of component layout options. As with Registered SDRAM DIMMs, the entire clock delay is present between the clock tab pin and the PLL input, and is adjusted by balancing the feedback network with respect to the actual output net. The clock proposed 'Reference net' is provided for use during module simulation to ensure an accurate clock delay, since measurement of the delay is impractical due to the reflections at the clock tab pin.

The clock delay from the input of the PLL to the input of any SDRAM is designed to be 0ns (nominal). The clock arrival time at the PLL input should not be adjusted, although sources of timing variation include PLL input capacitance, padding capacitor and series resistor tolerances, and DIMM impedance variations. Due to these variations, it is possible that there will be a difference between the input of the PLL and the input of the SDRAM. A reasonable target for this variation is ± 100 ps.

The most important factor in clock measurements is to ensure consistent clock arrival times at the SDRAM. The Registered DIMM clock reference board provides a standard way to measure and evaluate PLLs for this measurement. The DIMM suppliers must adjust the value of the PLL feedback capacitor to place the clock arrival at the SDRAM on the DIMM within ± 100 ps of the clock arrival at the Registered DIMM clock reference net. This measurement should always be taken as a 'mean', since the PLL and source jitter makes this measurement difficult. The PLL manufacturer's recommended feedback capacitor value should be used, and should include worst case contributions of PLL skew, PLL phase error, feedback capacitor variations, and DIMM variations.

On DDR DIMMs, the clock to the register should be in phase with the SDRAM clock. The target range for variation is ± 100 ps. The actual delay may vary as a result of the input capacitance of the SDRAMs, the register clock input capacitance, PLL output skew variations of the PCB parasitics, and measurement skew. In the case when the register clocks arrive prior to the

SDRAM input clock, this relationship may be adjusted by increasing the value of the register clock padding capacitor.

It is important to understand the timing relationship between the PLL feedback capacitor and the Register padding capacitor as it may be necessary to adjust both capacitors to optimize both Register and DDR memory clock timing.

Clocks are the life-blood of any synchronous system and doubly so for a source synchronous environment. In a source synchronous system multiple clock domains are required to implement the memory interface and retime the data once it reaches the memory controller. Ideally the DDR memories will receive a differential clock signal with positive and negative going edges that cross at or around $V_{DDQ} * 0.5$, typically 1.25V.

Post Register Timing

The post register timing for DDR modules is critical for proper and reliable module operation. The table in Figure 3 describes a method by which worst case conditions may be analyzed for proper module operation. Results of this analysis indicate a positive timing margin at the limits of register and PLL operation. This method is “Time to V_M ” or “Time to the measurement Voltage”. This method uses the register timing into its specified test load (instead of T_{CO} , open circuit) and adds or subtracts the timing into the SDRAM net and loads.

Symbol	Parameter	PC200 Setup	PC200 Hold	PC266 Setup	PC266 Hold
t_{CLK}	Clock cycle time.	10.0	NA	7.5	NA
t_{PD}	Maximum time for the signal to exit the register. This is measured into a standard test load of 30pF to GND, 50 Ohms to V_{DD} .	2.8	1.1	2.8	1.1
$t_{PROPDELAY}$	Maximum time for the signal to propagate from the register to any SDRAM. The switching point for a rising edge is $V_{REF} + 100mV$.	2.0	0.5	2.0	0.5
t_{REG}	Shift in the register clock in relationship to the SDRAM. The input clocks to the registers and the SDRAM are intended to track.	0.10	0.10	0.10	0.10
t_{SET}	Setup time required for the SDRAM inputs.	1.1	NA	1.1	NA
t_{HOLD}	Hold time required for the SDRAM inputs.	NA	1.1	NA	1.1
t_{SKEW}	Clock jitter and skew on the DIMM.	0.275	0.2	0.275	0.2
t_{SSO}	Simultaneous Switch effect addition to propagation delay.	0.35	NA	0.35	NA
$t_{CROSSTALK}$	Crosstalk Adder.	0.10	NA	0.10	NA
MARGIN		3.475	0.2	0.975	0.2

Figure 3: Register Timing Table

Taking the data in the Post-Register timing analysis and putting it into a graphical form makes it easier to “see” the available margin (Figure 4).

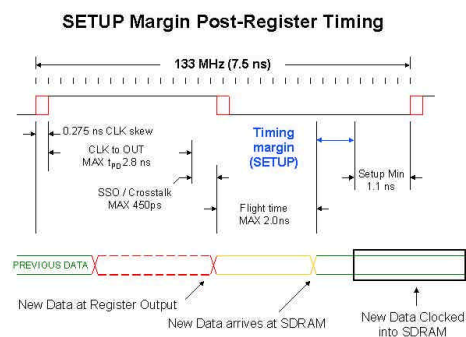


Figure 4: Graphical Timing Analysis

No high-speed analysis would be complete without simulating and observing the signal integrity of the system. Poor signal quality can lead to system problems with EMI and difficult to trace field failures. The SSTV16857 register is responsible for receiving, storing, and redriving command and address information to all the DDR memories on the module. Failure to do this job reliably results in immediate system failure.

Layout

As stated earlier, the reference designs are meant to be a starting point for Registered DIMM designs. Making small, but important improvements to design schematics and layout will add

margin to the overall design. As always, when dealing with high-performance systems care must be taken to simulate and then empirically evaluate all changes thoroughly in order to quantify the benefits of any changes that are made.

In looking at the register itself, we've touched on the fact that this register will be driving several types of transmission lines with distributed SDRAM loads. The basic DDR transmission line topologies lie well within the physical constraints of a DIMM outline, memory IC footprints, and IC to IC spacing requirements for manufacturability. As with register propagation delays, the layout of DIMM transmission lines should minimize the maximum flight time while maintaining the tightest possible flight time window between signals. This is one case where individual timing improvements may be small, but taken together can reduce the flight time window a couple hundred picoseconds.

Another factor effecting signal flight time is the load capacitance of all SDRAM IC input pins. Memory ICs exhibiting lower input capacitance will reduce the filtering effects of a distributed load on register output transitions. Lower input capacitance results in a higher effective transmission line impedance, lowering flight times, and reducing any register propagation delay degradation due to multiple outputs switching simultaneously. Propagation delay degradation can be reduced further by balancing the number of outputs used on each register IC.

Net structures and lengths must satisfy signal quality and setup/hold time requirements for the memory interface. Net structure diagrams for each signal group and raw card formats have been developed in the JEDEC standard. All signal groups are accompanied by a trace length information that lists the minimum and maximum allowable lengths for each trace segment and/or net. To verify DIMM functionality, a full simulation of all signal integrity and timing is required.

The general routing requirements as defined in the JEDEC specification are as follows: Route all signal traces using 4 mil traces and a 6 mil minimum spacing to adjacent nets. Clocks must be routed using 4 mil lines and 6 mil spacing between the differential signal pairs. In addition 90% of the length of the clock signal needs to be routed on the inner layers of the module. These requirements are critical in achieving nets with the appropriate impedance and good signal integrity.

Relying on equal length clock traces to result in nominal register and DDR memory clock timing is like running with a set of blinders on. It is impossible to always run at the nominal clock timing, however back annotation of the clock nets may show the designers that individual traces may be electrically slow or fast. Because some DDR memory devices are physically closer to the clock distribution PLL, some clock lines are forced to meander and change layers more than others. By making adjustments before the actual board is produced, clock timing inside the target range of $\pm 100\text{ps}$ will be far easier to achieve across the entire module.

Because of increases in component power consumption, operating frequency, faster edge rates, and lower supply voltages the power planes or power distribution system now plays a critical role in the signal integrity of a system. For simulations to accurately predict the characteristics of real boards, care must be taken to implement a power plane decoupling strategy that provides a minimal impedance over at least several harmonics of the circuits fundamental frequency range of operation.

One of the most critical areas is that of the decoupling capacitor. Choosing the right style of surface mount capacitor and using a low inductance mounting methodology will dramatically lower the overall inductance and series resistance. This results in a power distribution system that exhibits constant low impedance to many hundreds of megahertz.

Conclusion

As memory standards evolve, higher frequencies, reduced voltages, differential signals and advanced architectures will be used to provide new generations of processors and computing systems the volume of information needed for efficient operation. To overcome the engineering obstacles of these new standards, IC designers and system architects must work together to optimize small pieces of the design. Finished pieces are assembled and simulated as a whole, furthering the optimization process.