DesignCon 2003

High-Performance Systems Design Conference

Beyond DDR, Signal Integrity and Timing Analysis of Quad Band Memory (QBM) Systems.

Douglas Burns
Barry Katz
Eric Brock
Walter Katz, PhD

Signal Integrity Software, Inc.
WWW.SISOFT.COM

Abstract

As system performance continues to increase the need for higher-bandwidth memory is ever increasing. A novel memory architecture, Quad Band Memory (QBM), that utilizes standard DDR memory devices, a custom PLL, and a "2" to "1" bus switch is introduced. This architecture doubles the available data bandwidth of current DDR SDRAM memory technology. This paper will present an overview of the design, technical challenges, system tradeoffs, as well as the key methodology features required. Signal integrity and timing are explored in depth while considering inter-symbol interference, reflections, termination, interconnect loss, and populations schemes. Feasibility of this technology at 533MHz data-rates will be presented.

Biographies

Douglas J. Burns graduated Magna cum Laude from the University of Massachusetts in 1981 with his BSEE and received his MSEE from Northeastern University in 1986. Between 1981 and 2000, Doug worked for Honeywell, Digital, and Compaq Computer. He has led VLSI design implementation teams, signal integrity teams, ASIC implementation teams, and influenced the system architecture of ALPHA systems. Doug consulted on SI and implementation issues across many groups within Digital/Compaq, and has a successful track record of bringing products to market. Current, Doug holds four patents in computer system design and has three additional patents pending. Doug's expertise spans a wide range of engineering disciplines including: system architecture, VLSI design, package design, timing analysis, interconnect analysis, cross-talk analysis, electromagnetic modeling, IO buffer analysis and selection, simultaneously switching output analysis, termination selection, clock distribution and skew analysis, and high-speed bus design.

Barry Katz has been at the forefront of high-speed design throughout his career, devoting his efforts to solving the problems faced by designers of leading edge high-speed systems. In 1995, Barry founded Signal Integrity Software, Inc. (SiSoft) where he has assembled a team of world-class experts committed to solving the most challenging high-speed design problems facing the industry today by delivering a comprehensive design methodology, software tools, and expert consulting. Barry has been a major influence to the signal integrity methodology utilized by numerous companies and has personally led multiple signal integrity design teams. He has expertise in all aspects of high-speed design including: timing analysis, interconnect analysis, crosstalk analysis, electromagnetic modeling of interconnect, packages and connectors, IO buffer analysis and selection, decoupling analysis, simultaneously switching output analysis, interconnect topology, termination selection, clock distribution and skew analysis, and high-speed bus design. He currently serves as chairman of the IBIS Quality committee. Barry holds a MSEE from Carnegie Mellon and a BSEE from University of Florida.

Eric Brock received his BS in Electrical Engineering as well as a BS in Physics from Portland State University in 1998. After graduating, he joined Compaq Computer as a Signal Integrity Engineer in the Alpha Development Group. During his time at Compaq, Eric was a member of a system-level signal integrity group that was involved in timing analysis, parts selection, clock skew analysis, clock distribution, interconnect analysis, waveform quality, and termination selection, on Compaq's latest Alpha processor systems, and led a group of signal integrity engineers in the creation of design rules, simulation, analysis, and design verification of Compaq's next generation Alpha processor systems. Eric joined Signal Integrity Software Inc. in 2002 as a Senior Signal Integrity Consultant.

Dr. Walter Katz is a pioneer in the development of constraint driven printed circuit board routers. He developed SciCards, the first commercially successful autorouter. Walter founded Layout Concepts and sold routers through Cadence, Zuken, Daisix, Intergraph and Accel. More than 20,000 copies of his tools have been used worldwide. Walter developed the first signal integrity tools for a 17 MHz 32-bit minicomputer in the seventies. In 1991, IBM used his software to design a 1 GHz computer. Walter holds a PhD from University of Rochester and a BS from Polytechnic Institute of Brooklyn.

QBM System Definition

QBM is a novel memory architecture that utilizes industry standard DDR memory devices, a custom PLL, and a "2" to "1" Bus switch to effectively double the data bandwidth of current DDR memory technology. A simple block diagram of the data paths and strobe paths is shown in Figure 1 below.

Figure 1 QBM Block Diagram

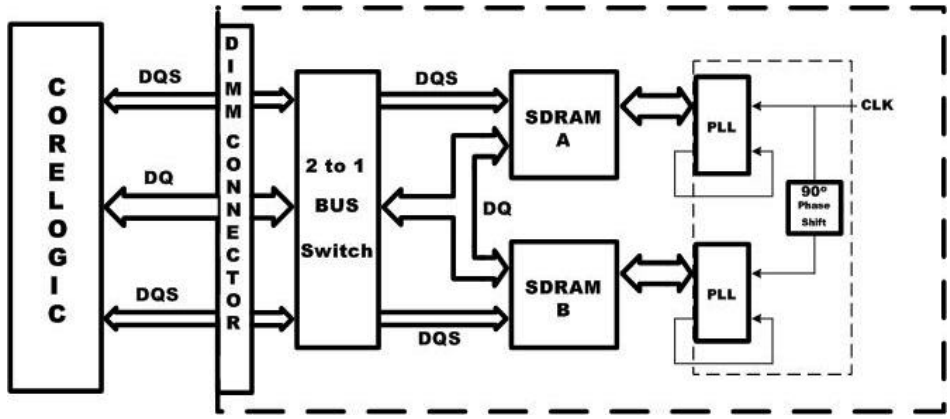

QBM Reads

For read operations, data is driven from bank A and bank B, with bank B clocking 90 degrees out of phase from Bank A. Both devices are running at standard DDR speeds (i.e. 200 MHz data-rate). The data is then multiplexed through the bus switch by a select pin (BE) running at twice the frequency of the DDR SDRAMs (i.e. 400 MHz data-rate). This results in an output data rate of twice the DDR rate. The associated data strobes are generated from the bus switch using BE to toggle between static high and static low inputs. Data is then captured at the corelogic with the associated strobes.

QBM Writes

For write operation, data is generated from the corelogic at twice standard DDR speeds (i.e. 400 MHz) along with the associated strobes. For each 8 bit slice, there are two strobes, which run at the standard DDR frequency (i.e. 200 MHz data-rate), but are phase shifted by 90 degrees with respect to each other. One strobe connects to bank A while the second connects to bank B. The bus switch connects the output from the corelogic to both SDRAMs through a latching structure to provide an extended data eye at each SDRAM. Data is clocked into each respective SDRAM via the 90 degree phase shifted strobes at standard DDR speeds (i.e. 200 MHz).

Analysis Methodology

Using SiSoft's Core-to-Core™ methodology, a rigorous signal integrity and timing analysis was performed over a wide solution space spanning variations in process, temperature, voltage, topology, termination, drive strength, data patterns, lengths, etc. For these analyses, a 3 DIMM configuration was assumed and each DIMM and the CoreLogic was simulated as the bus driver. Population loadings were not performed as part of this feasibility study and will be considered in future work. The objective of this study was to mine this comprehensive solution space to maximize both waveform quality and timing margins. Optimizing a solution for these interfaces proved a challenge because waveform quality often traded off against timing and reads traded off against writes.

Performance metrics used to rate a solution included the following:

Eye Jitter (Vmeas)
Eye Height (inner and outer)
Eye Width (Vmeas +/-)
Setup Margin
Hold Margin
Overshoot
Slew rate
Ring back

Signal Integrity Tool Set

The signal integrity toolset chosen for this analysis was SiSoft's SiAuditor™. SiAuditor utilizes the industry-standard Hspice simulation engine and provides an integrated waveform and timing analysis environment with advanced data management techniques.

Transfer Nets

A *Transfer Net* is an implementation independent network description that defines high-level connectivity of components on a bus and valid transfers that can occur between these components. The QBM memory sub-system is comprised of two main areas: Corelogic/Bus switch interface and the Bus switch/DDR SDRAM interface (Figure 2). The DQ_MUX2SDRAM/DQS_MUX2SDRAM transfer nets define the components, Data Rate, and allowable transfers between the DDR SDRAMs and the Bus Switch. The DQ_CORE2MUX/DQS_CORE2MUX describe the transfers between the Corelogic and the Bus Switch. These transfer nets form the basis for all of the simulations. Figure 3 shows the actual transfer net description for the DQ_MUX2SDRAM.

Figure 2 QBM Transfer Nets

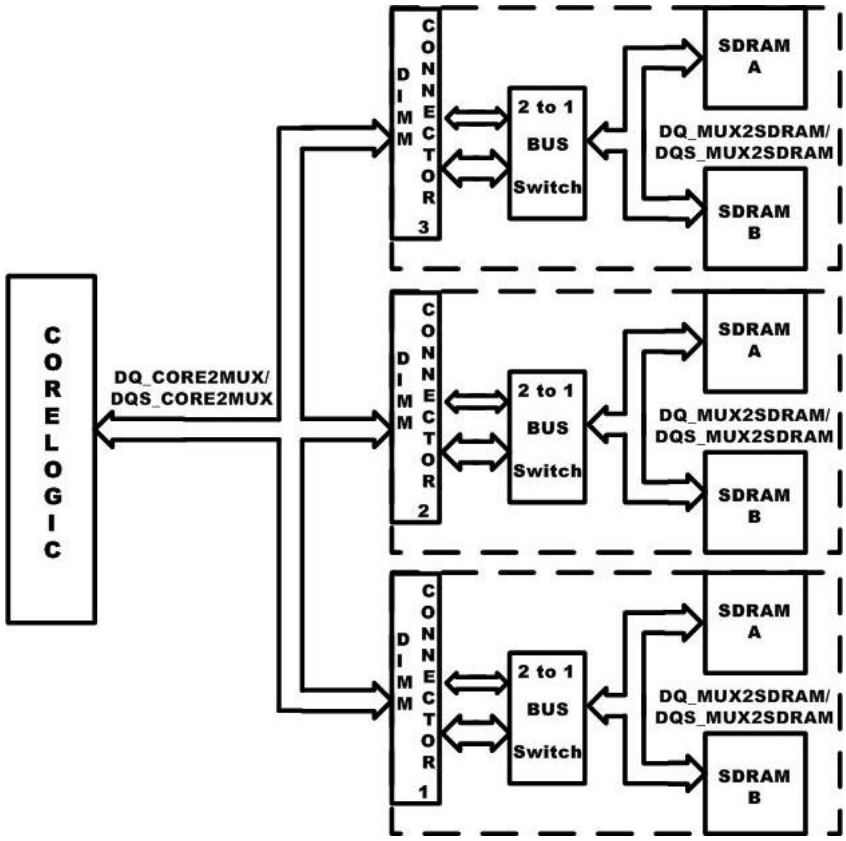

Figure 3 DQ_MUX2SDRAM Transfer Net

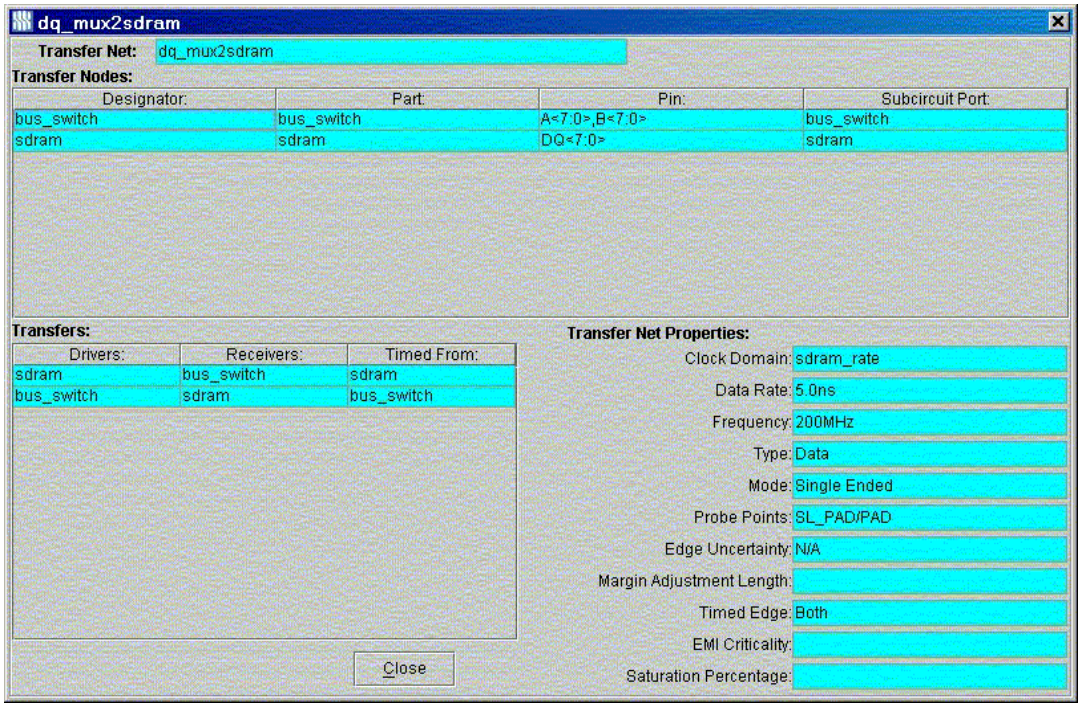

Waveform Analysis

Signaling memory buses at 400MHz data-rates results in networks that do not settle in one bit time, thus it is extremely important to consider inter-symbol interference (ISI), coupling, package variations, and more. Furthermore, exciting the various energy states on a network requires the use of complicated stimulus patterns and requires waveform analysis on all edges of the output signal for quality and interconnect delay. Capturing all of the timing and waveform quality information requires a robust set of measurement criteria. Figure 4 shows the types of voltage levels required to ensure accurate capture of waveform information. The SiAuditor toolset utilizes all of these levels in processing the signal waveforms and applies these levels to every simulation edge. Using the method of Progressive Discovery, the tool presents multiple views of the data, in a clear and concise format, which eases the analysis effort.

Figure 4 Waveform Measurements

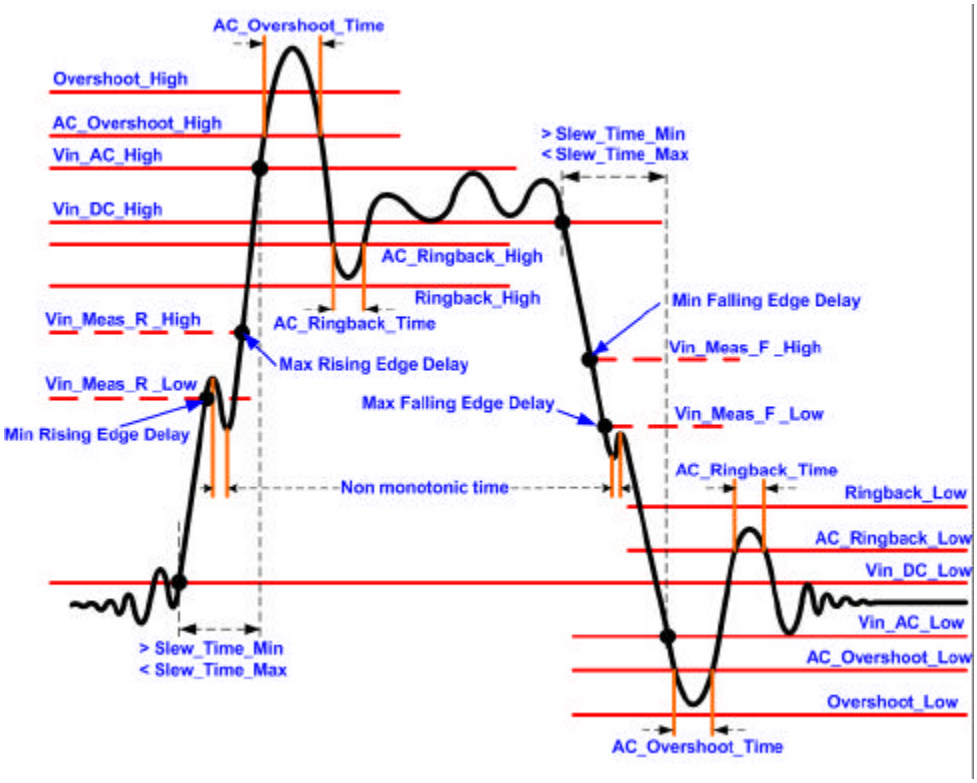

For waveform analysis, every edge of every waveform is checked against these levels at the pad of the receivers and at the drivers.

Timing Analysis

The QBM interface is comprised of both synchronous and source-synchronous transfers. Setup and Hold margins are determined by incorporating AC specs of source and target devices, interconnect delays for data and strobes, clock distribution skew and jitter. It is very important to consider both correlated and uncorrelated effects when adding up skews to determine timing margins.

Calculation of the wire delay required simulations of each device into its standard load as well as into the actual network topologies. The standard load delay is subtracted from the actual network delay to generate a network wire delay that is used by timing analysis. Delay measurements are made to the VMEAS points defined in the IBIS model. Delay measurements are made for each edge and the resulting min/max extremes used to compute the timing.

Design Focus

Many systems signal memory up to 200MHz, but getting designs operating in the 400MHz to 533MHz range is a challenging task. Conventional designs have required specialized memory devices to address these increased performance needs. In reality, the problem is not the need for better memory devices, but in better IO interfaces. The QBM switch isolates the memory interface from the system interface, thus allowing standard memory devices to support higher operating rates.

The design of the QBM memory required analysis in 2 distinct domains:

SDRAM to Bus Switch: This involves the analysis of the SDRAM data valid windows for reads and positioning the bus switch selector. For writes, the source synchronous transfer from the switch to the SDRAM must be analyzed.

Corelogic to Bus Switch: Both the read and write transfers are source synchronous and require that multiple loading configurations (populations) be analyzed.

Simulation background

Analysis Scope

Utilizing a 3 DIMM implementation, a comprehensive solution space analysis was explored by varying the following parameters:

- AC Termination (values and position)
- Parallel Termination (Rterm and Vterm and position)
- Series Termination (value and position)
- Etch impedances and process variations
- Mother board topology (balanced star, daisy chain, unbalanced star)
- Network lengths
- Corelogic drive strength
- SDRAM Drive strength
- SS/FF Process
- Voltage Variation
- Temperature Variation

Assumptions

Design goals were focused upon design feasibility and not on design optimization. The system topology chosen was extracted from an existing PC motherboard and was used as the baseline for all simulations. Only minor modifications to the etch characteristics were allowed. The IO interfaces on the Bus Switch were identical to those used by the SDRAM.

Results

The DQ_MUX2SDRAM network is point-to-point as shown in Figure 5 (L_DIMM). Simulations of the SDRAM driving to the Bus Switch (READ's) were performed across a range of etch lengths as well as process corner. The timing goal was to place the clock pulse generated by the 2x PLL output inside the data valid window provided by the SDRAM. For writes, the switch drives a source synchronous transfer, involving both the DQ_MUX2SDRAM and DQS_MUX2SDRAM transfer nets, to the SDRAM.

Figure 5 DQ_MUX2SDRAM

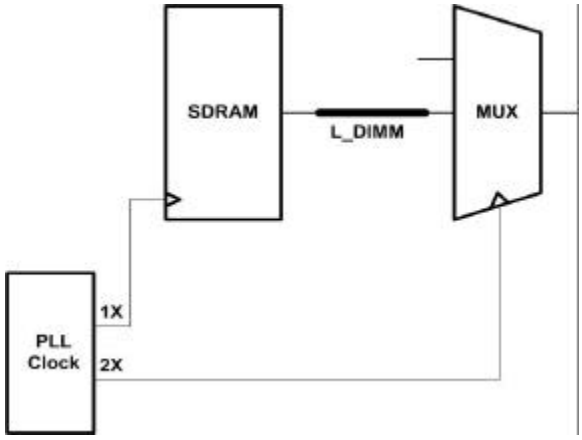

The system topology for the DQ_CORE2MUX and DQS_CORE2MUX transfer nets is shown in Figure 6. This topology is one of a number of topologies investigated and it reduced the number of simulations needed since all 3 Dimm's drive into an electrically identical environment. Variables of simulation included etch impedance, termination values (R_Series, Rup, Rdown), Voltage, Temperature, and Process corner. All simulations involved a complex stimulus that would excite inter-symbol interference (ISI). Both coupled and uncoupled environments were examined.

Figure 6 System Topology

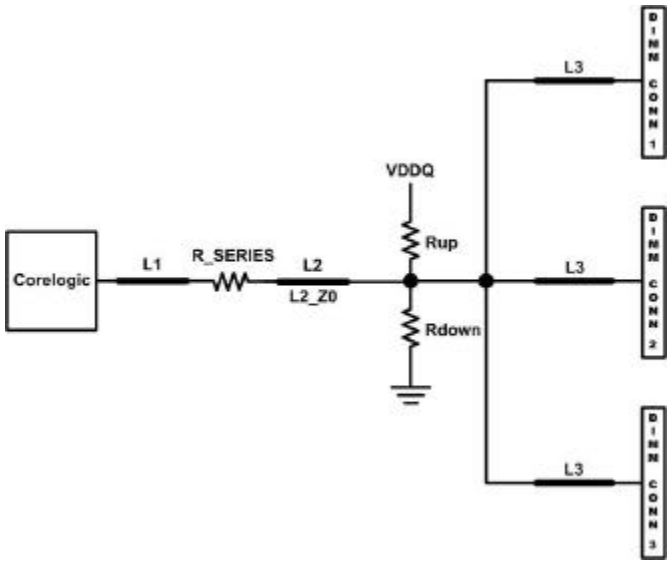

Waveform Quality

DIMM Read Transfers

The DQ_MUX2SDRAM and DQS_MUX2SDRAM transfers are point to point and operate at 200MHz data rate. These nets exist entirely on the DIMM. Figure 7 and Figure 8 shows the eye diagrams for the DQ_MUX2SDRAM path during memory READ's for the SS and FF process. Each has ample eye height and very little jitter. While there is signal overshoot and ringback, signaling margins and device limits are not exceeded.

Figure 7 DQ_MUX2SDRAM, READ, SS Process Corner

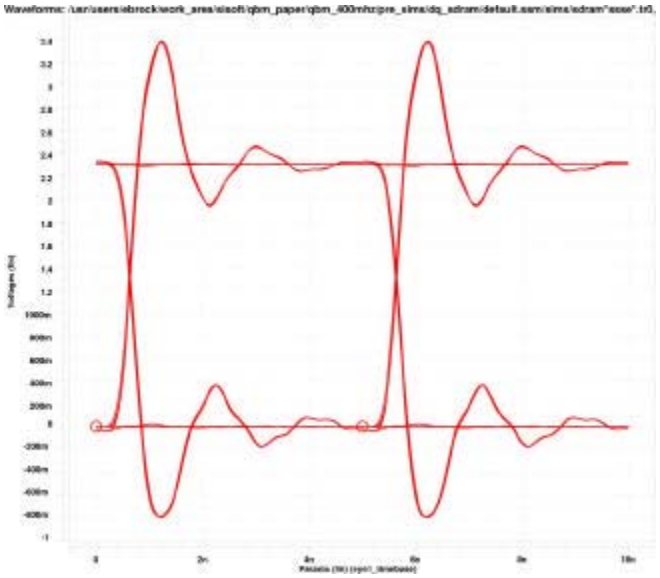

Figure 8 DQ_MUX2SDRAM, READ, FF Process Corner

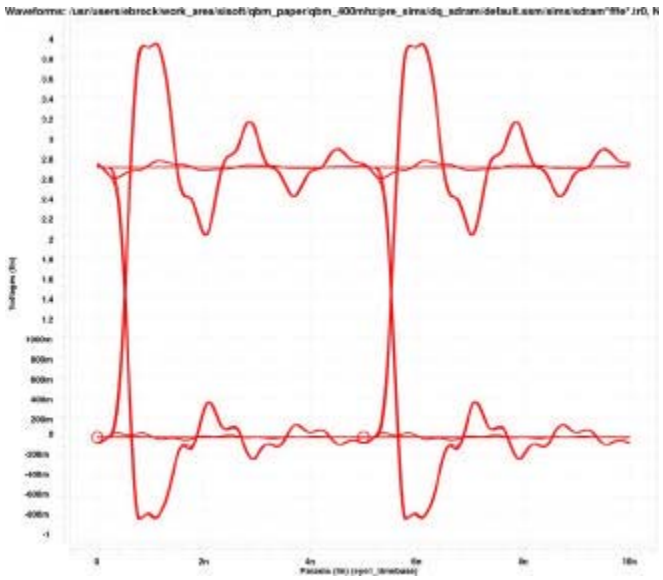

Dimm Write Transfers

Figure 9 and Figure 10 shows the eye diagrams for the DQ_MUX2SDRAM path during memory WRITE's for the SS and FF process. Each also demonstrates ample eye height and very little jitter. While there is signal overshoot and ringback, signaling margins and device limits are not exceeded.

Figure 9 DQ_MUX2SDRAM, WRITE, SS Process Corner

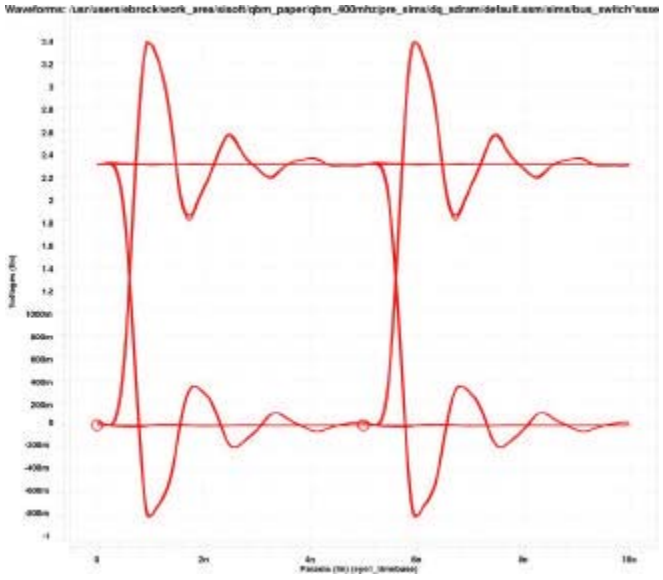

Figure 10 DQ_MUX2SDRAM, WRITE, FF Process Corner

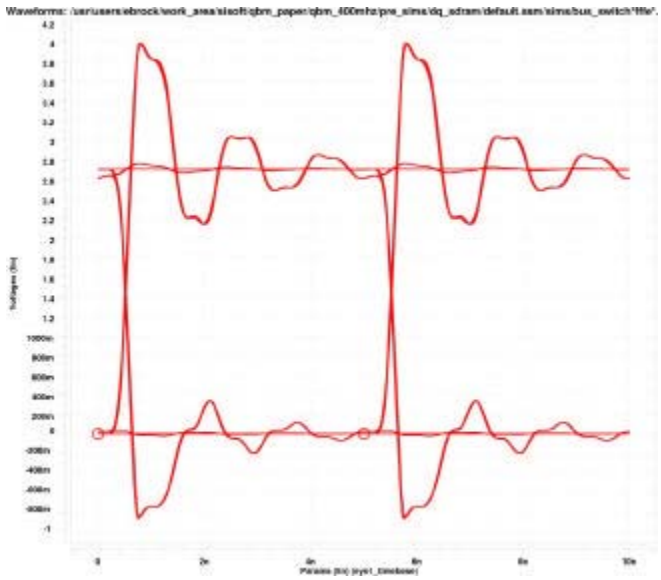

System Read Transfers

The DQ_CORE2MUX and DQS_CORE2MUX transfers are multi-drop networks running with a 400MHz data rate. These connect the DIMMs to the Corelogic chip. Careful balancing of the networks was required.

Figure 11 and Figure 12 show the resulting eye information for data READ's between the switch and corelogic. The eye opening provides sufficient width and height while overshoot and ringback are well controlled.

Figure 11 DQ_CORE2MUX, Switch Read, SS Process Corner

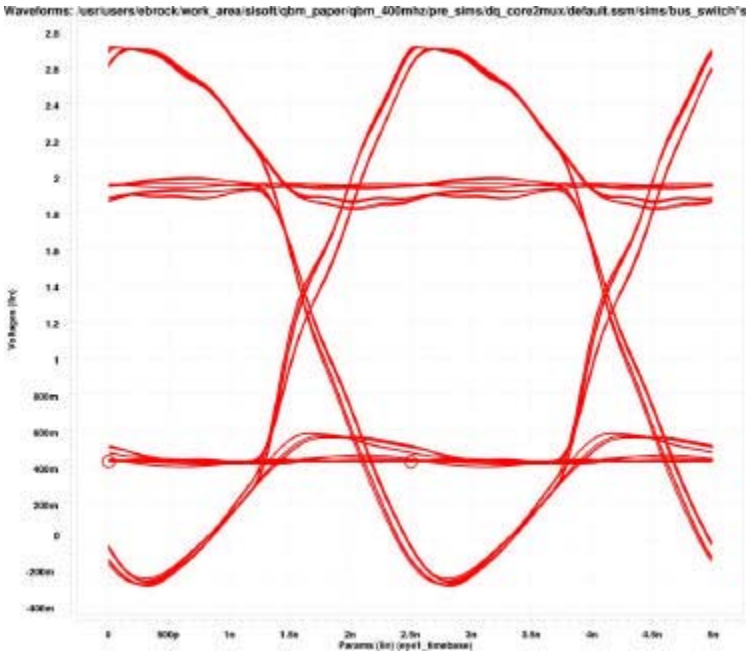

Figure 12 DQ_CORE2MUX, Read, FF Process Corner

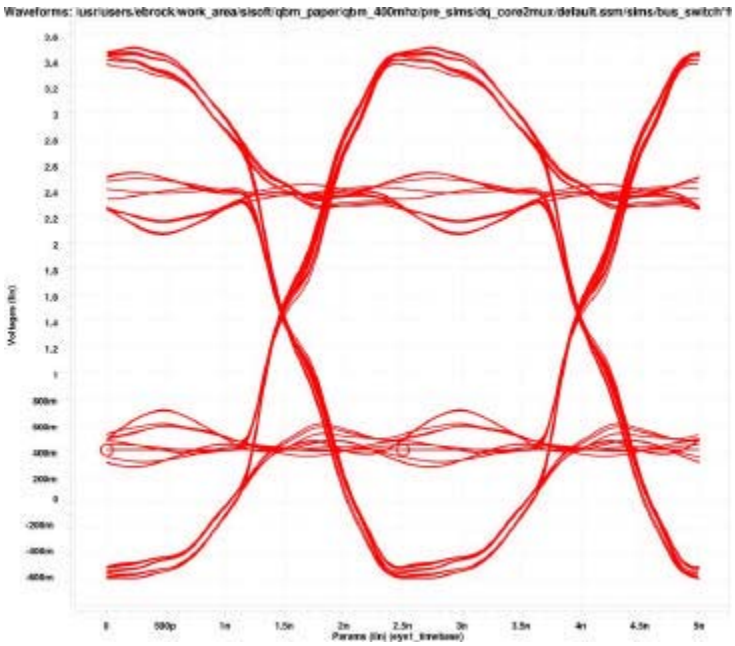

System Write Transfers

Figure 13 and Figure 14 depict the eye data for WRITE's between the corelogic and the switch. The eye's provide sufficient design margin to the system requirements.

Figure 13 DQ_CORE2MUX, Writes, SS Process Corner

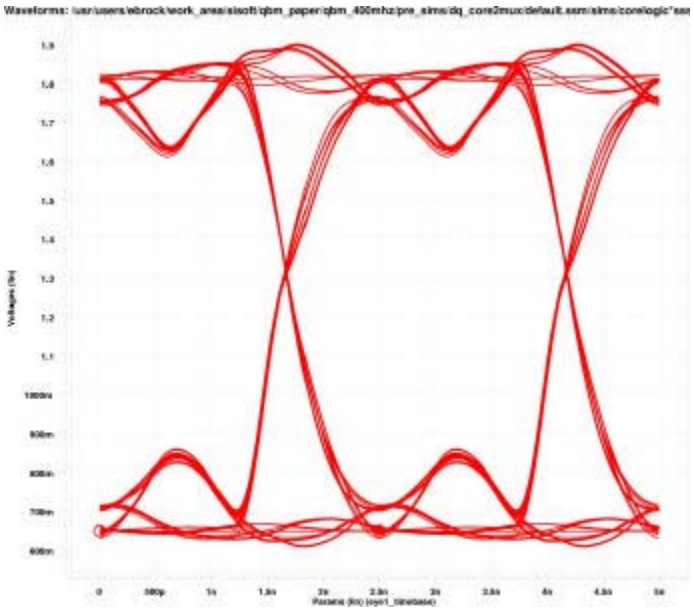

Figure 14 DQ_CORE2MUX, Write, FF Process Corner

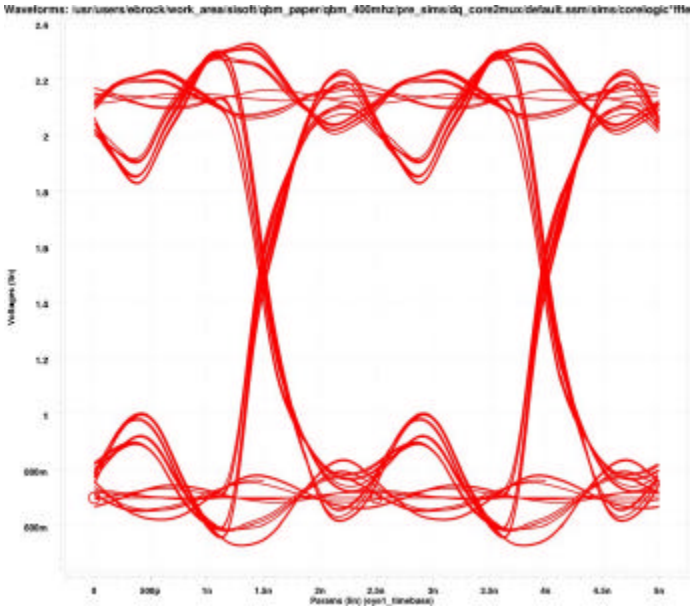

Timing

System simulations were performed at 400MHz data-rates. The SDRAM specifications were from existing 200MHz data-rate devices and theCorelogic specification was defined by an ASIC vendor. As is shown in Table 1, positive timing margins were obtained for all data paths with significant design margin. This could be a huge win for system designers. The QBM memory, using readily available devices, can provide a 2x increase is system bandwith TODAY!!!

Table 1 QDR System Timing Results (@400MHz)

Microsoft Excel - timing_report7.xls

File Edit View Insert Format Tools Data Window Help

H2

	A	B	C	D	E	F	G
1	Setup Margin (ns)	Hold Margin (ns)	Min Delay (ns)	Max Delay (ns)	Transfer Net	Driver	Receiver
2	0.617	0.508	1.231	1.477	dq_core2mux	bus_switch1	corelogic
3	0.617	0.508	1.231	1.477	dq_core2mux	bus_switch2	corelogic
4	0.617	0.508	1.231	1.477	dq_core2mux	bus_switch3	corelogic
5	0.455	0.285	0.726	0.878	dq_core2mux	corelogic	bus_switch1
6	0.455	0.285	0.726	0.878	dq_core2mux	corelogic	bus_switch2
7	0.455	0.285	0.726	0.878	dq_core2mux	corelogic	bus_switch3
8	0.477	2.81	0.312	0.392	dq_mux2sdram	bus_switch	sdram
9	0.129	0.215	0.015	0.171	dq_mux2sdram	sdram	bus_switch

Timing Report Log / Timing / By Transfers / By Variation / Synchronous Details / Source Synchronous Details

Ready

Opportunities

This design was re-run at 533MHz data rate. No changes were made to the network and all component specifications were unchanged. The timing results from the 533MHz analysis are in Table 2.

Table 2 QDR System Timing Results (@533MHz)

Microsoft Excel - timing_report3.xls

File Edit View Insert Format Tools Data Window Help

A10 = No AC specs

	A	B	C	D	E	F	G
1	Setup Margin (ns)	Hold Margin (ns)	Min Delay (ns)	Max Delay (ns)	Transfer Net	Driver	Receiver
2	0.393	-0.031	1.21	1.665	dq_core2mux	bus_switch1	corelogic
3	0.393	-0.031	1.21	1.665	dq_core2mux	bus_switch2	corelogic
4	0.393	-0.031	1.21	1.665	dq_core2mux	bus_switch3	corelogic
5	0.035	-0.073	0.682	0.956	dq_core2mux	corelogic	bus_switch1
6	0.035	-0.073	0.682	0.956	dq_core2mux	corelogic	bus_switch2
7	0.035	-0.073	0.682	0.956	dq_core2mux	corelogic	bus_switch3
8	0.166	1.902	0.312	0.392	dq_mux2sdram	bus_switch	sdram
9	0.129	-0.384	0.041	0.171	dq_mux2sdram	sdram	bus_switch

Timing Report Log / Edge Details / Delay Details / Timing / By Transfers / By Variation / Synchronous Details

Ready

All timing failures can be eliminated of reduced with simple re-tuning of the clocks. Of the remaining failures, discussions with SDRAM vendors have shown that critical device specifications (Output skew window, Setup, Hold) could be improved. Reducing the setup and hold requirements from 500ps to 450ps and improving the DQS to DQ skew from +/-750ps to +/-650ps would result in positive 533MHz operation. 533MHz SDRAM devices will not be available until late 2003, thus QBM systems would be able to take advantage of the higher performance memory 9 to 12 months earlier.

With specifications similar to the 400MHz DDR parts, calculations show that 666MHz and faster operation may be achievable. While more simulation work is needed, this trend is encouraging for high end memory system designs.

Conclusions

QBM technology holds the promise of obtaining memory systems with twice the performance of conventional designs. Utilizing existing design styles, it has been shown that 400MHz operation, with margin, is possible with existing 200MHz devices. Additional work has shown that adjustments in the system implementation will support 533MHz operation with 266MHz devices. With proper design specifications, higher speed QBM systems are possible. The benefits of this approach are clear: Higher performance with lower cost devices.