# System Co-Design and Co-Analysis Approach to Implementing the XDR™ Memory System of the Cell Broadband Engine™ Processor

## Realizing 3.2 Gbps Data Rate per Memory Lane in Low Cost, High Volume Production


Wai-Yeung Yip
Platform Solutions
Rambus Inc
Los Altos, CA 94022
Tel : 650-947-5000
Fax : 650-947-5001
e-mail : wyip@rambus.com

Scott Best
Architecture
Rambus Inc
Los Altos, CA 94022
Tel : 650-947-5000
Fax : 650-947-5001
e-mail : sbest@rambus.com

Wendemagegnehu Beyene
Signal Integrity
Rambus Inc
Los Altos, CA 94022
Tel : 650-947-5000
Fax : 650-947-5001
e-mail : wbeyene@rambus.com

Ralf Schmitt
Signal Integrity
Rambus Inc
Los Altos, CA 94022
Tel : 650-947-5000
Fax : 650-947-5001
e-mail : rschmitt@rambus.com



**Abstract – This paper describes the design and analysis of the 3.2 Gbps XDR™ memory system of the Cell Broadband Engine™ (Cell BE) processor developed by Sony Corporation, Sony Computer Entertainment, Toshiba and IBM. A System Co-Design and Co-Analysis Approach was applied where different components of the system are designed and analyzed simultaneously to allow trade-offs to be made to optimize system electrical characteristics at low overall system cost. The XDR memory interface circuit implemented in the Cell BE processor, the power delivery system design and analysis, and the interface statistical signal integrity analysis will be described to illustrate this design and analysis approach.**


## I. Introduction

As an ultra-high performance, general purpose microprocessor, the Cell BE processor requires a memory system which is high performance and yet manufacturable at low cost and in high volume. Therefore, the XDR memory system for the Cell BE processor must be optimized for cost-performance, and not for performance alone. It is not acceptable to design one component of the system in advance of the others, thereby restricting the design space of the components designed downstream. It is also not acceptable to design any component in isolation, thereby creating components which cannot be assembled to meet the stringent timing requirements of an ultra-high speed interface. Consequently, we adopted an approach where we co-designed and co-analyzed simultaneously the different components of the Cell BE processor memory system – the memory physical layer (PHY) in the form of a XDR I/O Cell (XIO) in the processor, the DRAM chip, the packages, and the printed circuit board (PCB). This allowed us to make the best trade-offs between cost and performance within the production constraints of each component.

Moreover, the System Co-Design and Co-Analysis Approach encompasses early system level power integrity and signal integrity analysis to uncover potential electrical issues, which allows timely resolution of such issues with the lowest cost solution. Oftentimes multiple solutions exist for resolving system level electrical issues, each requiring improvement to a different combination of the system components. Early discovery of a system level electrical issue in the design flow allows selection of the solution with the lowest system cost.

## II. Features of the XDR Memory Interface Circuit

The Cell BE processor houses an 8-byte wide XDR memory interface circuit split over 2 channels, wherein each bi-directional I/O nominally operates at 3.2 Gbps, yielding an interface bandwidth of 25.6 GB/s with ECC. The RX (receive) and TX (transmit) clock instants for each data I/O are set independently, determined from a digitally controlled, on-chip calibration process. By realizing all of the precision timing circuits in the memory controller, as described below, substantial cost savings can be realized in the XDR DRAM and PCB design. For example, no phase mixers are required in the XDR DRAM timing circuits. Furthermore, there are no data-to-data PCB trace length matching requirements. This results in less DRAM and PCB design complexity, less DRAM circuit area, and less PCB layout area. See Fig. 1 for an example layout of one 4-byte XDR memory channel of the Cell BE processor which was realized within an area of 35 x 43 $mm^2$ in a 4-layer PCB using high volume manufacturing rules.

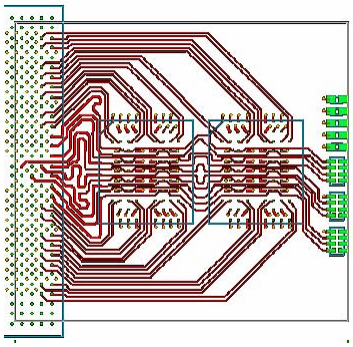

Fig. 1. An example layout of one 4-byte XDR memory channel of the Cell BE processor in a 4-layer PCB.

Fig. 2 shows a block diagram of the per-byte clock generation circuitry. Each byte has a supply-regulated, self-biased PLL with loop bandwidth tracking [1]. The small-swing PLL output vectors couple with 22 digitally

1-4244-0630-7/07/$20.00 ©2007 IEEE.

programmable (8-bit) phase mixers: 18 are used to generate the byte's 9 receive and 9 transmit clocks running at 1.6 GHz, also known as the fast clock rate, 1 is used for PLL feedback, 1 is used as the memory controller interface clock source, and the remaining 2 are used as DLL clock sources (1 for the command/control/address bus, and 1 for the memory controller/processor clock interface). Each phase mixer provides $0^o$ to $360^o$ of delay, and each consists of a supply-regulated phase-multiplexer, a phase-interpolator, and a small-swing to full-swing converter with duty-cycle correction. The PLL reference clock (CTM) running at 400 MHz, or 1/4 of the fast clock rate, is locked to the divided-down output of the feedback path. This phase-mixer associated with the feedback path has its 8-bit digital phase-selection code programmed to "00h" to establish the baseline for the other mixers: a phase code of "00h" yields a clock output which is also phase-aligned to CTM.

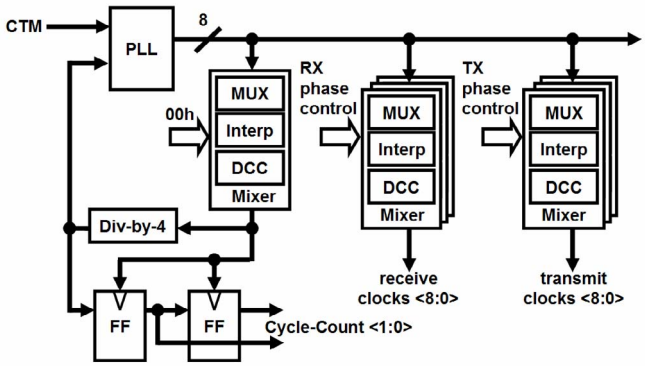


Fig. 2. Per-byte clock generation circuitry in the XDR I/O (XIO) cell.

Fig. 3 shows a DQ (XDR data) byte that comprises 9 bi-directional differential I/O's which receive and transmit data at 3.2 Gbps using double-data rate signaling. The interface between this circuit and the memory controller operates at the reference clock rate, yielding an 8-to-1 serialization ratio between the controller interface and the memory channel. During data reception, the transmitter is placed in a high impedance state by disabling the differential pair drivers. These drivers can be disabled or enabled within 1 reference clock cycle so as to minimize Read-Write turnaround delay.

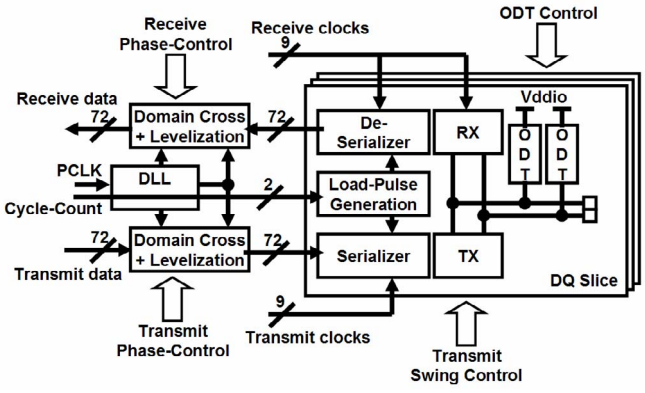


Fig. 3. DQ (XDR data) byte block diagram.

Our System Co-design and Co-Analysis Approach allowed us to uncover a potential VDDIO switching noise problem when drivers turn on and off, early in the design stage. To reduce the amount of switching noise, the current-source bias for the drivers is segmented so that it can be incrementally disabled or enabled to minimize noise impact on VDDIO. This simple improvement to the bias circuit resulted in significant cost savings in combating current switching noise. Alternate solutions may require a low inductance package design and an elaborate on-board bypassing solution resulting in high system cost.

A 50-ohm on-die termination (ODT) is included on the DQ pins for signal integrity. Desired transmitter swings are maintained across PVT by a closed-loop comparison of the swing of an internal reference transmitter to an externally provided, per-byte reference voltage. A similar closed-loop bias circuit is used to maintain desired ODT values across PVT variations.

A pair of 13-bit registers are used to control the receive and transmit timing instants of any one DQ bit slice's receiver or transmitter. Bits <7:0> of a phase-control register control the $0^o$ to $360^o$ phase position of the associated clock within a fast clock cycle (256 steps of $1.4^o$ increments). Bits <9:8> control the fast clock position within a reference clock cycle. Bits <12:10> control the choice of 1 of 6 reference clock cycle levelization values.

To determine the fast clock position within a reference clock cycle, a 2-bit cycle-reference (generated in the per-byte clock generation circuitry) is provided to every DQ slice within the byte. This Cycle-Count signal repeats a "00-01-11-10" grey-coded count pattern every reference clock cycle, where the "01-11" transition is aligned to the rising edge of the external reference clock signal (CTM). The memory controller clock interface DLL also uses this Cycle-Count signal as a reference for its clock-divider. This common use of the Cycle-Count signal by all of the DQ slices as well as the per-byte DLL simplifies clock domain crossings between the bit slices and the memory controller clock domain. By comparing bits <9:8> of one clock signal's phase-control value with another clock's <9:8> value, it can be determined how to best pass data reliably between the two clock domains.

Prior to normal system operation, all 72 DQ I/O's receive and transmit timing instants must be phase calibrated. The calibration sequence begins by writing a test pattern into the DRAM via a low speed serial bus. A series of READ instructions are then issued via the command bus to the DRAM to transmit the pattern repeatedly to the controller at full speed. The controller sweeps the phase-control setting for each pin's receive clock, storing the phase-setting values corresponding to the "Fail-to-Pass" (FP) transition as well as the "Pass-to-Fail" (PF) transition. The receive clock is placed as the numerical average of the 13-bit FP and PF phase-control values. Once the read link is calibrated, the write link is calibrated utilizing a full speed write and read back from the DRAM in a manner similar to the receive calibration process.

To minimize the time for the calibration process, known-good FP and PF values can be loaded from a system BIOS, and programmed into the calibration control circuitry.

This allows the full 13-bit sweep to be required only once per system; subsequent power-on calibration can then be restricted to a much smaller range.

The XIO memory interface just described is realized on a first generation Cell BE processor fabricated in a 90nm, 8 layer metal SOI process. The interface occupies a total die area of 11 $mm^2$ with typical power of 5 W at 25.6 GB/s data. A block diagram of the entire memory interface of the first generation Cell BE processor is shown in Fig. 4.

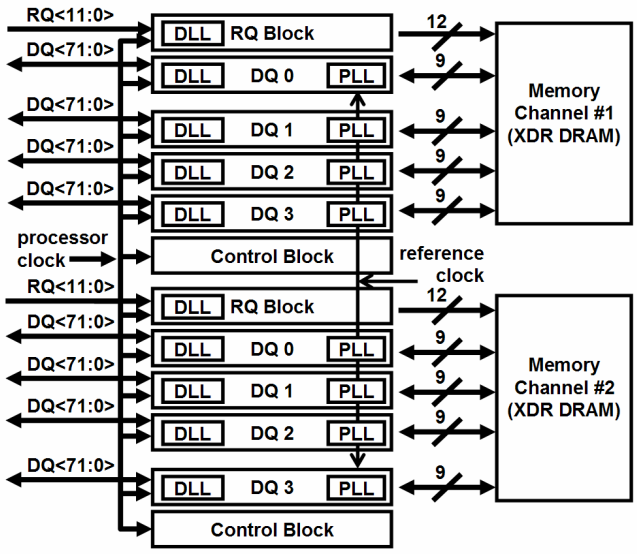


Fig. 4. XDR memory interface block diagram for the Cell BE processor.

## III. An Optimally Accurate Power Distribution System Design and Analysis Approach

For multi-gigabit I/O systems, power integrity is a design parameter that has to be addressed from the very start of the design process and at the system level. In these high-speed designs the impact of supply noise on timing and voltage margin can have a significant impact on the total system margin and has to be included into the channel budget.

A thorough, wide-frequency-range (DC-to-Daylight) power integrity analysis requires that every component of the power distribution system be included in the model. Moreover, the analysis needs to start at the design planning stage and be iterated to guide the design of individual components to generate the best and lowest-cost overall design with acceptable power supply quality. However, such analysis requirements increase multi-fold the cost of the analysis. Therefore, we employed an optimally accurate analysis methodology to render the power integrity analysis of the Cell BE processor XDR memory system computationally practical.

The power distribution of an electronic system contains numerous components on different system hierarchy levels, from voltage regulator to on-chip distribution. The contributions of the different components vary in frequency, and usually only a subset of components dominates supply noise in a particular frequency range. Therefore, power integrity analysis can be divided into separate analysis steps, each of these steps addressing a noise component at a different frequency range.

In each of these analysis steps only the components that contribute significantly to this particular supply noise component are modeled accurately, while the other components are either omitted or replaced with much simpler reduced-order models. This keeps the complexity of the supply model reasonable in each of the analysis steps. This 'divide and conquer' strategy allows the analysis of entire interface systems in a single analysis. At the same time, it maintains the interaction between the different components of the supply system. In this way, trade-offs between different components can be explored to find the best and most cost-effective system solution for the power distribution design.

### *A. Analysis of On-Chip IR Drop*

Due to the large resistance of on-chip wires there is a noticeable resistive voltage drop (IR drop) on the on-chip power distribution. The spatial distribution of IR drop on the chip is usually not affected by package and PCB design, since basically the same voltage is provided to all chip pads (flip-chip bumps or wirebond pads) due to the much lower resistance of package and PCB traces and planes. In typical design flows the IR drop is checked late in the design cycle, when most of the layout is available (post-layout). At this late stage of the design, only minor improvements are possible to fix local IR drop problems.

We started the analysis of the Cell BE processor XDR memory system on-chip IR drop early in the design cycle, during floor planning and pad assignment of the XIO. The IR drop is dependent on the pad placement, the power routing on the chip, and the current distribution over the chip area. Signal escape routing requirements in the package restrict the placement of supply pads, and several different power rails are competing for pad locations and routing resources. The simultaneous progression of chip design and package design permitted iterations of pad placement and package signal escape routing to arrive at an optimized signal and supply pad placement early in the design cycle. This optimized pad placement together with careful routing of the various power rails and floor planning of the circuit blocks on the chip guaranteed a high-quality power supply.

### *B. Analysis of High-Frequency Switching Noise*

The second supply noise component beside IR drop affecting the design of the on-chip power grid is high-frequency switching noise. Switching circuits in the I/O and core region can cause current peaks on the supply rails. Since the package typically is primarily inductive, the current provided by the package cannot change fast enough to provide the charge for this high-frequency switching. Instead, on-chip decoupling capacitors have to provide the switching current to prevent high-frequency noise on the power rails.

The inductive nature of the package impedance creates an effective low-pass filter, separating PCB and chip at high frequencies. This filter significantly attenuates high-frequency switching noise leaking from the chip into the PCB, where it could excite resonances in the PCB supply planes. It also prevents high-frequency noise from any other source or device

leaking from the PCB into the chip. This filter makes it possible to analyze high-frequency supply noise as an effect limited to the on-chip supply distribution and circuits. Since package and PCB have little impact on the high-frequency noise, their models can be simplified for this analysis to an effective inductance connected to each pad.

### C. Medium Frequency AC Noise

Medium frequency noise is often the dominating supply noise component in I/O systems. One contribution to medium-frequency noise in I/O systems is Simultaneous Switching Noise (SSN) on the supply rail of output drivers.

In order to predict medium frequency supply noise on the supply rails of the circuits, the frequency dependent impedance $Z_{PDN}$ of the supply network (Fig. 5), seen by the circuits on the chip, has to be analyzed. Additionally, the current changes causing supply noise have to be modeled.

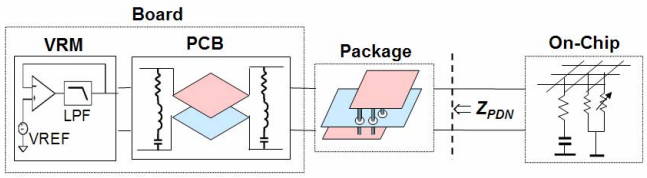


Fig. 5. Power distribution impedance $Z_{PDN}$.

#### 1) Modeling the Power Distribution Network

For medium frequency noise analysis the on-chip power distribution model can be much simplified. At high frequencies, the position of a circuit block inside the power grid and its proximity to on-chip decoupling capacitors determines the local supply noise for this circuit. At low and medium frequency, however, the proximity of a circuit block to on-chip decoupling cells has little impact as the cycle time of the current excitation becomes much larger than the RC time constant of the on-chip power. As a result, the on-chip power grid and the decoupling capacitors in this grid can be combined to a lumped equivalent RC network. This equivalent on-chip RC network can be modeled with a small number of passive elements, reducing the complexity of this model drastically.

At medium frequencies, the package and PCB are dominant components of the power delivery system and have to be modeled accurately. Traces and vias as well as circuit elements in the package and PCB can be easily modeled using traditional equivalent models. Power planes require complex models to reflect wave propagation inside plane pairs and two-dimensional current distribution effects at the ports of the planes.

Fig. 6 shows the typical impedance profile of $Z_{PDN}$ over frequency, seen by the circuits inside the package. It also shows the frequency range where the different supply system components influence the supply noise seen at the circuit. At low frequencies the impedance is determined by the output impedance of the Voltage Regulator Module and by low-frequency decoupling capacitors on the PCB. Different types of board capacitors with different frequency responses create smaller ripples in the impedance profile in this frequency range.

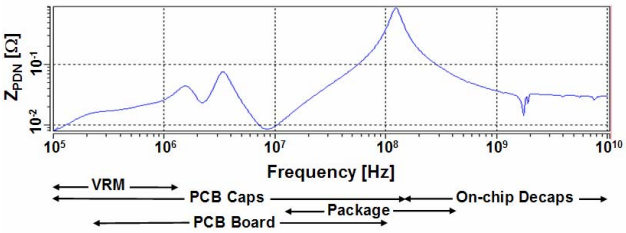


Fig. 6. Typical power distribution impedance profile.

At medium frequencies, the impedance associated with the package dominates the PDN impedance, which is primarily inductive. This inductance together with the on-chip decoupling capacitors generates a 'Package-Chip-Resonance' peak in the impedance profile. Exciting a system at the frequency of this peak typically leads to the highest supply noise amplitudes in the system. As shown in Fig. 6, several components of the power supply system affect the impedance at this range, and the resonance frequency as well as the amplitude of the peak can be influenced in different ways. One possible way to reduce the amplitude of the resonance peak is increasing the amount of on-chip decoupling capacitors. As shown in Fig. 7, however, similar reduction can be achieved, for example, with a change of the PCB board from a 4-layer technology to a 6-layer technology. Thus, our power analysis methodology allows us to explore solutions on different hierarchy levels of the design, identifying the best and most cost-efficient solution for the overall system design.

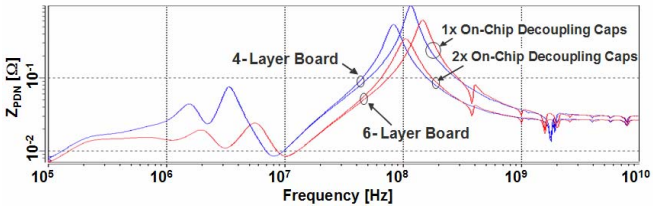


Fig. 7. Trade-off analysis between power supply components on different design hierarchy levels.

At high frequencies, on-chip decoupling capacitors dominate the PDN impedance, reducing the impedance with increasing frequency. As the frequency increases, the impedance approaches a steady value determined by the resistance of the on-chip power distribution and the series resistance of the decoupling capacitor cells itself.

#### 2) Finding the Worst-Case Excitation Pattern

Self-induced supply noise is generated when the current dissipation of the circuit changes, causing a change in voltage drop over the impedance $Z_{PDN}$. The generated noise is dependent on the spectrum of the current excitation and the frequency profile of $Z_{PDN}$. In order to find worst-case supply noise, a current pattern is constructed with large components at frequencies with high $Z_{PDN}$. Combining this worst-case current profile with the model of the power distribution system makes it possible to simulate the worst-case self-induced noise on this supply rail.

#### 3) Simulating Margin Loss due to Supply Noise

The important parameter that has to be derived from the supply noise analysis of an I/O system is the margin loss due to noise. Predicting this parameter requires a combination of the supply noise generated in the system with the circuit sensitivity to this noise.

In high-speed I/O systems the channel itself has an impact on the current profile of the output drivers as well as the power distribution impedance $Z_{PDN}$. Therefore, for SSN analysis the power supply model for the output driver rail is combined with a channel model. This provides an accurate current excitation for the supply and ground rail, and therefore an accurate supply noise prediction on both rails. Using a driver model that reflects the impact of supply noise on the conductivity of the driver allows predicting the timing loss on the channel due to supply noise for an excitation causing maximum supply noise in the system. With this simulation, the impact of supply noise on the channel timing is verified.

## IV. Statistical Signal Integrity Analysis to Design for Low Cost Manufacturing Tolerances

The design of interconnect systems of high-speed parallel interfaces, such as XDR systems, using conventional PCB and package technologies can be challenging. The key design objective is to ensure robust system operation under the permitted worst-case component variations. This is achieved with balanced specifications of all system components and architectural improvements.

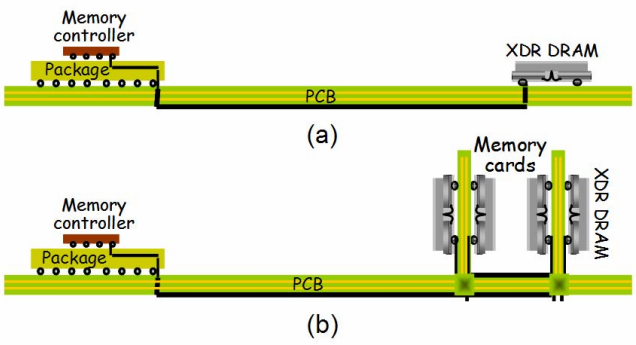


Fig. 8. Example XDR memory systems: (a) a short channel with DRAM chips mounted on the PCB; (b) a long channel with memory modules plugged into the PCB.

An example Cell BE processor XDR memory system for a short channel and a long channel implementation with two modules are shown in Fig. 8 (a) and (b), respectively. Both XDR systems support point-to-point bi-directional differential topologies with on-chip termination to minimize reflections and controller-based FlexPhase™ technology to compensate for timing differences due to skews introduced by the interconnects and devices. The channel between the controller and memory consists of a large number of traces in packages and PCBs and other structures in connectors. To satisfy low-cost manufacturing constraints, conventional packaging and PCB technologies are used in the XDR system. The interconnect design goals include the use of conventional package and PCB design rules, use of 100 Ohm differential impedance PCB and package traces in microstrip line or stripline configurations, support for four-layer conventional PCBs, and for both wirebond and flip-chip BGA packages. PCB and package trace impedances can vary by ±15% and ±20% of their nominal values, respectively [2]-[4].

Even though PCBs, packages, connectors, and terminations are designed to certain specifications to obtain the target system performance, the components can randomly deviate from their nominal values. For example, the impedance of traces in low-cost conventional PCB and packaging technologies can show normal distributions with standard deviations of over 7% from their nominal values, as shown in Fig. 9 (a) and (b), respectively. In this figure, each case corresponds to a different supplier. The Normal curve is derived by fitting all data to a normal distribution. Other channel parameters such as termination, device input capacitance, edge rate, etc. also show similar random deviations as a result of manufacturing variations. Therefore, most parameters can be assumed to have normal distributions with $3\sigma$ equal to the desired tolerance of the nominal values. Consequently, the voltage and timing budgets of an interconnect system can be degraded by these deviations in channel parameters and operating conditions.

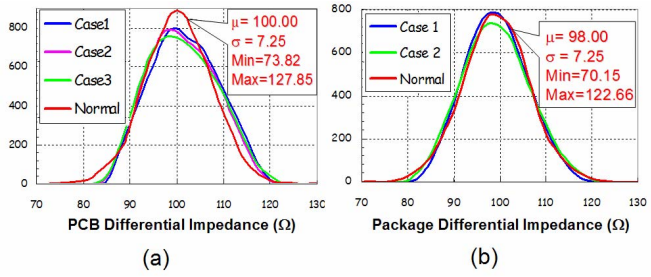


Fig. 9. Normal disbributions and distributions of differential impedances of (a) PCB and (b) PBGA package traces based on manufacturing tolerances.

In order to investigate the influence of interconnect parameter variations, it is crucial to build accurate broadband models. The channel model is first created based on physical design. Fig. 10 shows the channel model corresponding to one pair of XDR interconnect nets. Each net consists of many lossy transmission lines representing the subsections of traces in a PCB and packages. Lumped and distributed models of the solder balls and vias are generated using field solvers. Behavioral models are used to represent the drivers and the termination networks, while the input capacitance and substrate resistance of the receivers are determined from measurements of prototype chips [2], [4].

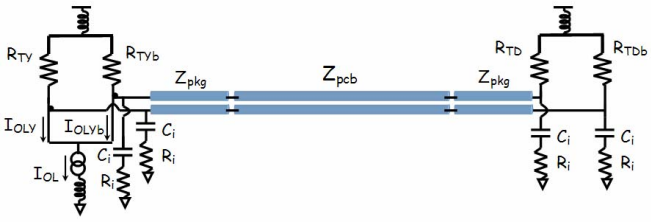


Fig. 10. The channel model of a XDR memory system.

To validate the channel models and to understand the channel physical design requirements, several test boards, packages, and test chips were designed. Excellent correlation was achieved between measurement and simulation as exemplified by the waveform comparison in Fig. 11.

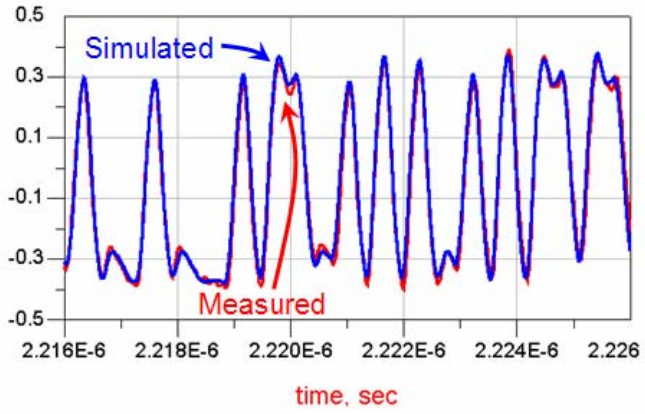


Fig. 11. Measured and simulated read data waveforms at XIO package pin.

The study of these prototypes also allows us to understand second-order effects, improve the models, maximize voltage and timing margins, and consequently develop system design guidelines.

In order to perform tolerance analysis, the key parameters are first identified. Then, the channel model is used to reliably predict channel behavior and sensitivity of the desired performance to parameter variations. In evaluating channel performance at multi-gigahertz frequencies, the first quantity of major interest is the insertion loss. In addition to the discontinuities caused by the large solder balls and package core vias that severely limit the bandwidth of the channel, the effects of parameter variations can easily be seen from the channel transfer function shown in Fig. 12. For the memory interconnect system in Fig. 8 (a), the parameter variations can affect the insertion loss by as much as 20%. Similarly, these variations can show horizontal uncertainties over 0.2 UI (unit interval) in timing jitter of received signals.

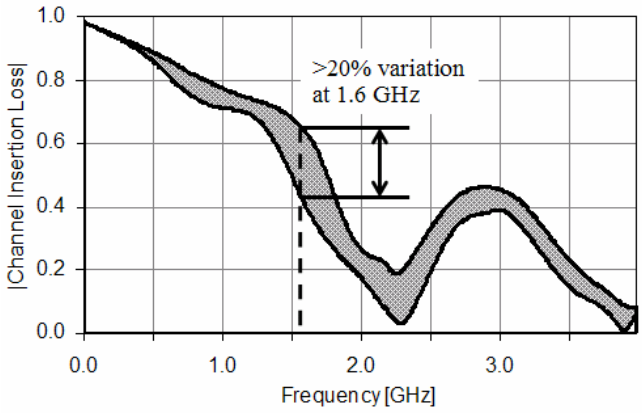


Fig. 12. Variations in magnitude of the insertion loss of the interconnect system in Fig. 8.

For systems with bit times and voltage swings of a few hundred picoseconds and millivolts, the traditional deterministic approaches to calculating the voltage and timing budgets of high-volume systems can be pessimistic and costly. It is no longer possible to define quality of high-speed systems as having "zero defects" or being "within specification." A more realistic approach to evaluate the effects of random parameter variations is to use a statistical approach. The approach can be used to combine tolerances in a realistic way so that performance and yield can be improved.

The histograms of the timing jitter and eye height of the memory interconnect system in Fig. 8 (a) from Monte Carlo simulations with a sample size of 5000 are shown in Fig. 13 (a) and (b), respectively. The timing jitter and eye height are normalized to the bit time and transmitted swing, respectively. The channel parameters are assumed independent and normal.

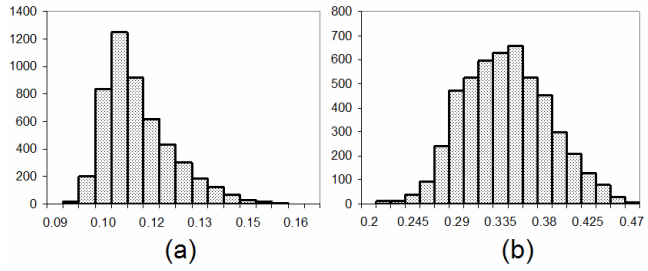


Fig. 13. The distributions of (a) timing jitter and (b) eye height of the received signal of a memory interconnect system.

A tolerance analysis of voltage and timing budgets of high-speed interconnect systems can be performed in several ways. However, the System Co-Design and Co-Analysis Approach requires that the variations of all components of the system be accounted for in the analysis, which can exponentially increase the computational cost of some of the analysis methods. Monte-Carlo and worst-case analyses are traditionally performed by running numerous simulations of the system using the channel models. Monte-Carlo analysis for systems with 15 parameters is computationally intensive and very time consuming. Worst case analysis based on the full factorial analysis requires tens of millions of simulations. Often only a few thousand cases are selected using heuristic approach. This worst-case approach does not explore the entire design space and may not unveil the worst performance of the system. However, techniques based on Taguchi method uses very few carefully selected simulations. The Taguchi method uses orthogonal arrays to systematically and efficiently explore the design space and build a regression model that relates the parameters variations on boards, connectors, packages, and chips to the performance [3]. Then, the model is used to generate the distributions of voltage margin and timing jitter.

Table 1 shows the descriptive statistics for the received eye heights using the three methods: Monte Carlo, worst case, and Taguchi-based techniques. The table demonstrates the strong similarities between the results of the three methods. The results from Taguchi-based model show the mean, and standard deviation similar to those of Monte Carlo analysis. However, the Taguchi-based regression model and worst-case analysis give dissimilar skewness, minimum, and maximum values. The Taguchi-based regression model shows a tight and optimistic distribution, while the worst case analysis shows large and pessimistic variations for eye height. Similar results

are also obtained for the timing jitter.

TABLE 1

Statistics of eye openings of the memory channel using Monte Carlo, worst case, and Taguchi-based methods.

| Statistics | Monte Carlo | Taguchi Based | Worst Case |
|---|---|---|---|
| Mean | 0.341 | 0.334 | 0.333 |
| Standard | 0.001 | 0.001 | 0.001 |
| Median | 0.340 | 0.335 | 0.331 |
| Mode | 0.341 | 0.334 | 0.332 |
| Std. Deviation | 0.038 | 0.036 | 0.052 |
| Sample | 0.001 | 0.001 | 0.003 |
| Kurtosis | -0.473 | -0.583 | -0.547 |
| Skewness | 0.187 | -0.043 | 0.260 |
| Range | 0.212 | 0.197 | 0.267 |
| Minimum | 0.244 | 0.237 | 0.216 |
| Maximum | 0.456 | 0.434 | 0.483 |
| Sum | 1705.347 | 1672.479 | 727.632 |
| Count | 5000.000 | 5000.000 | 2187.000 |

The three methods were also compared in terms of CPU time, accuracy, and generality when applied to high-speed systems. For the Monte Carlo analysis, 5000 simulations are performed, while for the worst-case, a total of 2187 ($3^7$) simulations are run. In practice, the CPU time for worst-case analysis is much smaller because often only a small subset of cases is run. The computation cost to generate the Taguchi-based regression model from the simulation results is negligible because it requires only 27 simulations. The computation time associated with generating the performance distribution using the regression model is also negligible. For systems with numerous channel parameters, both the Monte Carlo and worst-case analyses are CPU-intensive and impractical. Regression model based on Taguchi method is a very viable approach for tolerance analysis of high-speed interconnect systems.

Taguchi method effectively finds the sensitivity of the system performance to channel parameters. If each parameter can assume a nominal, low, or high, {N, L, H}, value within its specifications, Fig. 14 is a graphical approach to show the effect of each parameter. The timing jitter is very sensitive to PCB and package variations and less sensitive to termination and current drive as shown in Fig. 14 (a). Low PCB impedance reduces timing jitter. On the other hand, the voltage margin is more sensitive to current drive and less sensitive to PCB and package impedance as well as termination as shown in Fig. 14 (b). High current drive increases the voltage margin. Therefore, Taguchi method can be used to determine the best setting for each factor that would optimize the performance of the system.

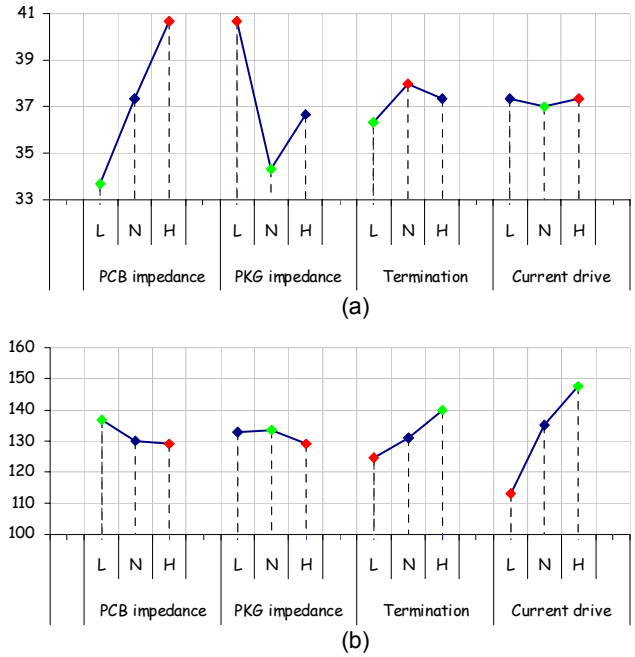


Fig. 14. Factor effects plots (performance vs. parameter variations): (a) timing jitter in ps and (b) voltage margin in mV.

## V. Summary and Conclusions

Various features are implemented in the XDR interface circuit in the Cell BE processor to allow for simpler and lower cost DRAM and PCB design. Our strategy is to distribute complexity in accordance with the component cost sensitivity. The ASIC process generally has lower cost sensitivity to circuit complexity than the DRAM process. Therefore, circuit innovations for precise timing adjustment were implemented in the Cell BE processor instead of the DRAM chip.

High speed interfaces such as the XDR interface face many power integrity challenges. A divide-and-conquer, optimally accurate analysis approach, where different model simplifications are applied in the analysis of power supply noise in different frequency ranges, provides us the means to tackle such challenges. It renders practical our System Co-Design and Co-Analysis approach where every component of the system is included in the analysis.

The design of low-cost, high-volume product must accommodate significant manufacturing deviations and tolerances. A Taguchi-based statistical analysis method allows us to efficiently quantify signal integrity errors as a result of manufacturing tolerances, which can then be properly accounted for in the interface timing and voltage budget.

The end result of our design and analysis approach is the successful implementation of an ultra-high data rate of 3.2 Gbps per memory lane at a low enough cost for practical production of the high volume expected of Cell BE processor based products.

## References

[1] S. Sidiropolous, D. Liu, J. Kim, G. Wei, M. Horowitz,

"Adaptive bandwidth DLLs and PLLs using regulated supply CMOS buffers," *VLSI Circuit Symposium, 2000*.

[2] W. Beyene, C. Yuan, N. Cheng, and H. Shi, "Design, modeling, and hardware correlation of a 3.2 Gbps/pair memory channel," *IEEE Transactions on Advanced Packaging Technologies*, Vol. 27, No. 1, pp. 34-44, February 2004.

[3] W. Beyene, N. Cheng, J. Feng, and C. Yuan, "Statistical and sensitivity analysis of voltage and timing budgets of multi-gigabit interconnect systems," *DesignCon 2004: High Performance System Design Conference*, Santa Clara, CA, February 2-5, 2004.

[4] W. Beyene, N. Cheng, J. Feng, H. Shi, and C. Yuan, "Performance analysis of multi-gigahertz parallel bus with transmit pre-emphasis equalization," *IEEE Transactions on Microwave Theory and Techniques*, vol. MTT-53, No. 11, pp. 3568-3577, November 2005.