

Xbox360™ Front Side Bus – A 21.6 GB/s End-to-End Interface Design



David Siljenberg, Steve Baumgartner, Tim Buchholtz, Mark Maxson,
Trevor Timpane
IBM Systems & Technology Group
Jeff Johnson
Cadence Design Systems, Inc.

January 23-26, 2007

Outline

- FSB & PHY Overview
 - CPU PHY
 - CPU Chip Description
 - GPU PHY Issues
 - GPU PHY Transmit Interface
 - GPU PHY High Speed Clock Interface
 - GPU PHY Development Logistics
- Physical Channel
 - Modeling and Simulation Approach
 - Impact of Reflections
- Bus Performance Measurements
- Summary & Conclusions

Front Side Bus (FSB) Overview

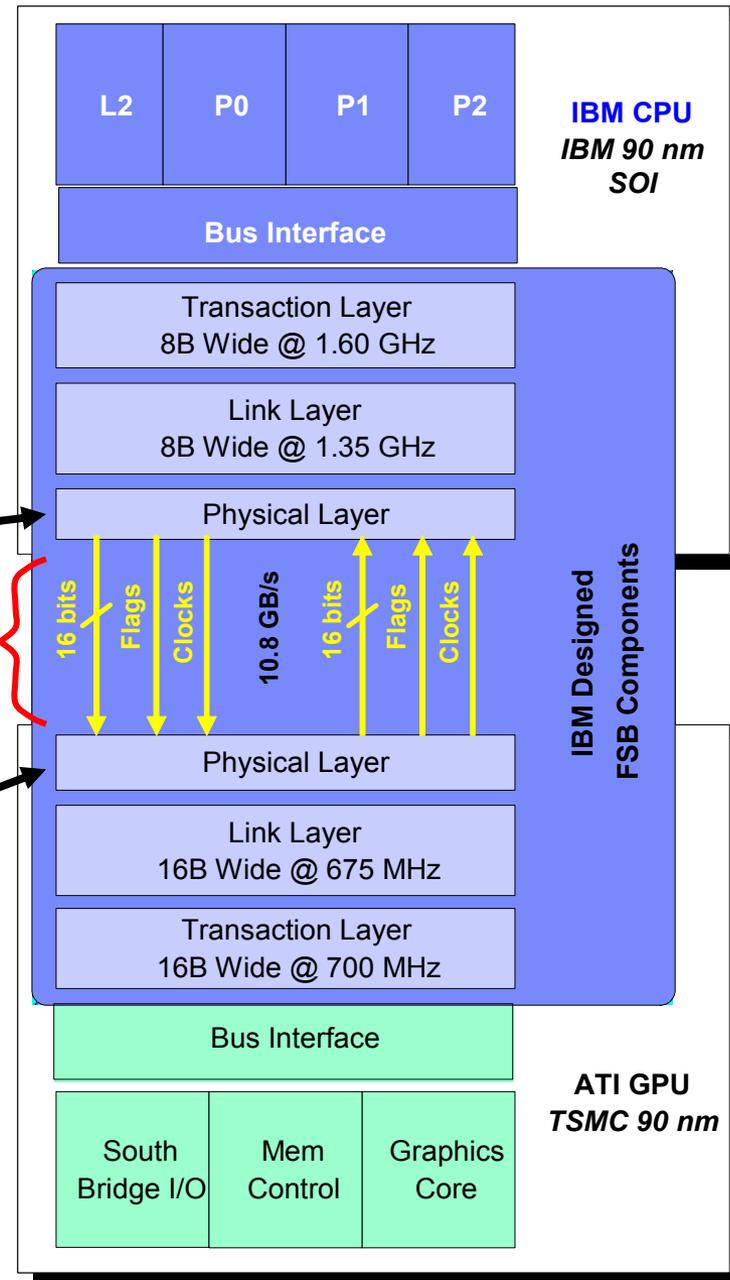
- FSB is the only interface to the Central Processing Unit (CPU)
- It connects to the Graphics Processing Unit (GPU)
- IBM owned the FSB end to end

The PHY contains the circuits that interface to the physical channel

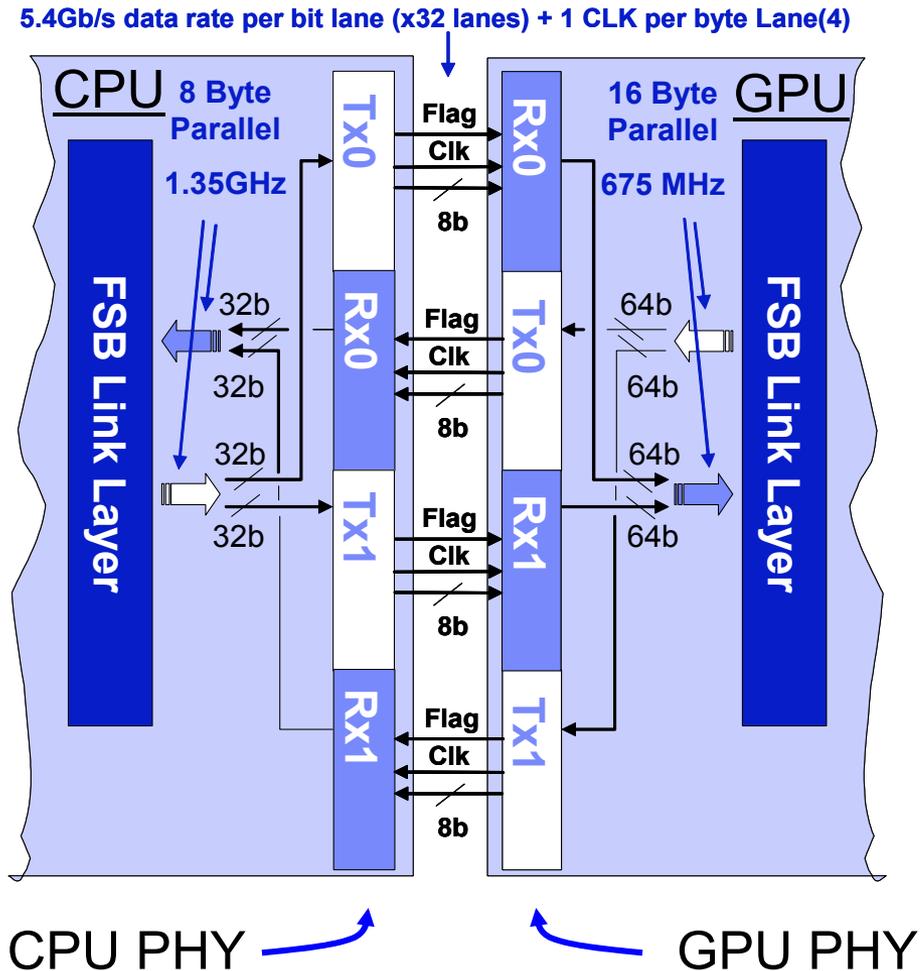
Physical channel

Chip package

Circuit board wires



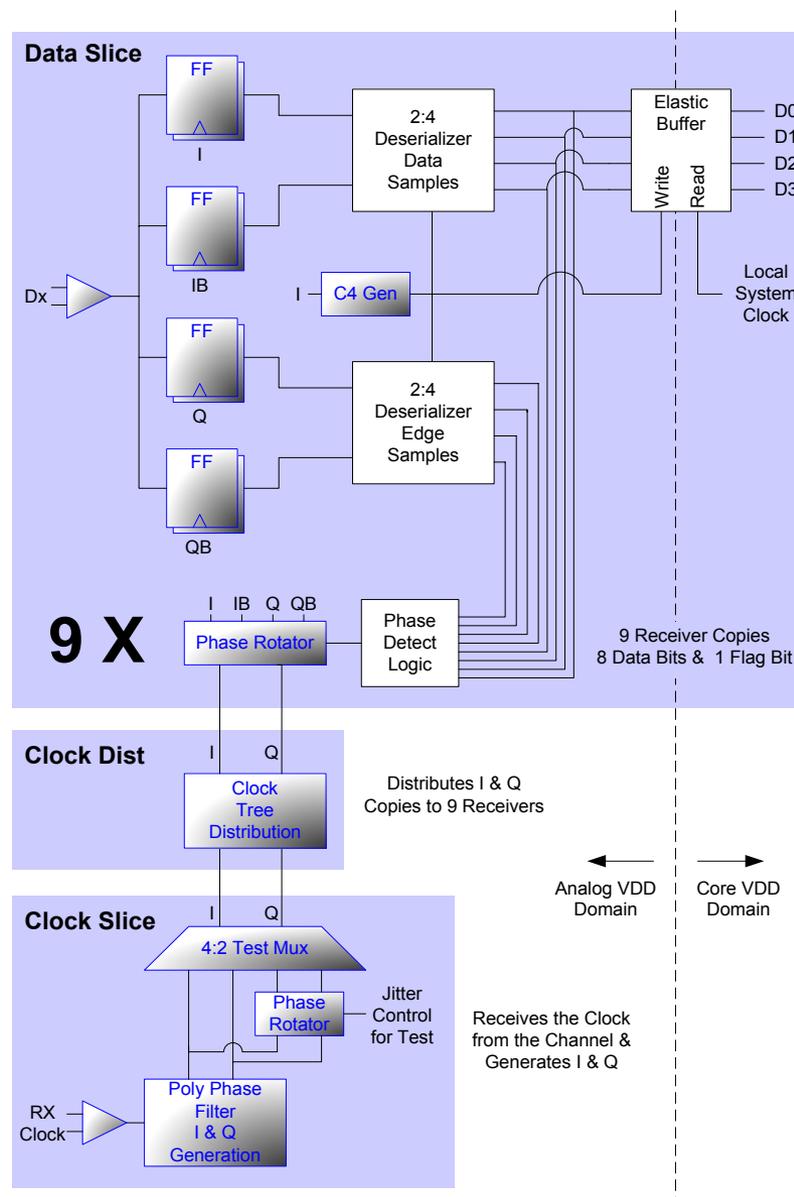
FSB Overview



- Point to point differential link
- 2 bytes in each direction
- Each bit lane runs at 5.4 Gb/s
- Source synchronous – clock is forwarded with the data for simple clock recovery
- The data is not coded for low latency
- Flag signal is used to delineate a packet

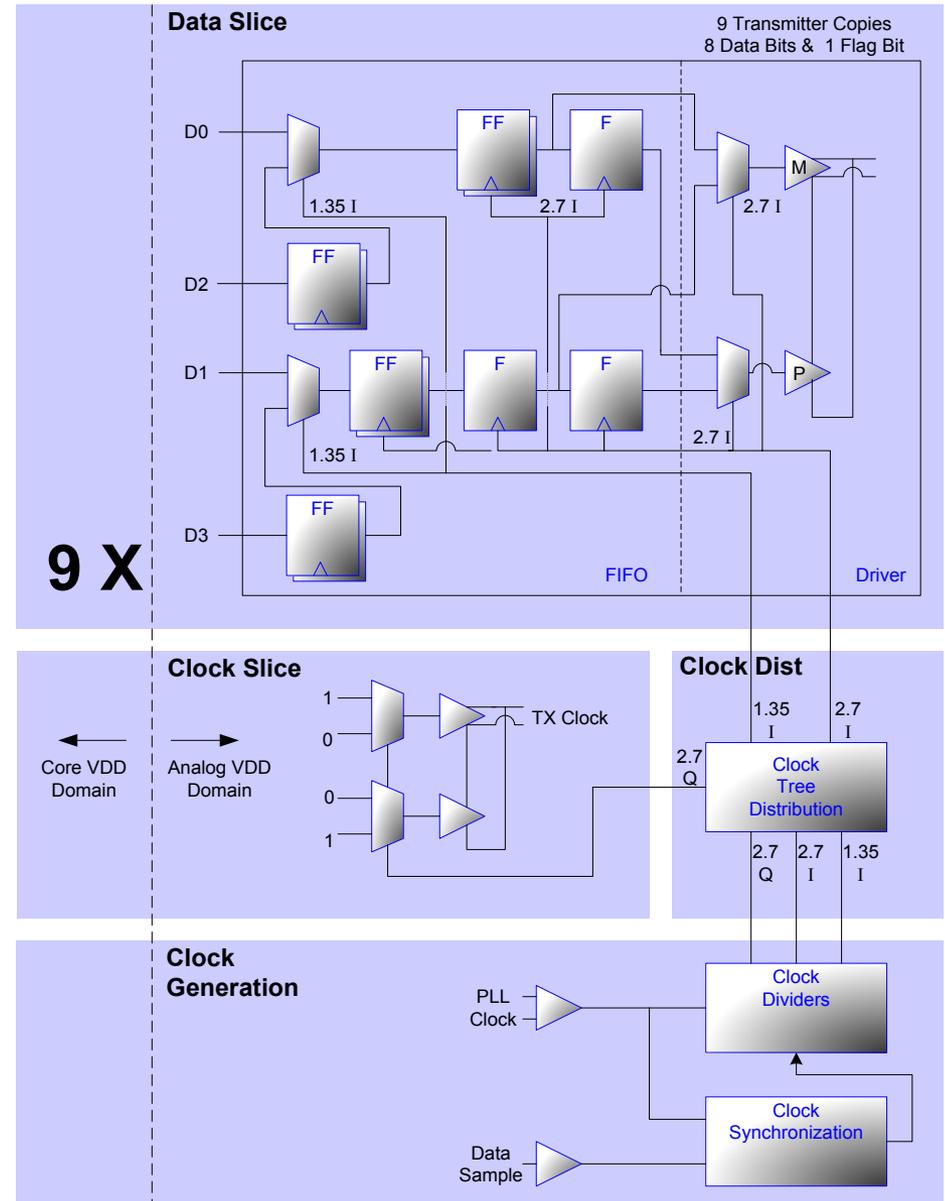
PHY Receiver

- Half rate (2.7 GHz) architecture
- Receives the serial data and forwarded clock
 - Clock receiver forwards clock copies to 9 data slices.
- Data slice phase rotators adjust the forwarded clock phase to correctly latch the data
- Deserializes the data to 4 bit parallel data in the CPU and 8 bit parallel data in the GPU

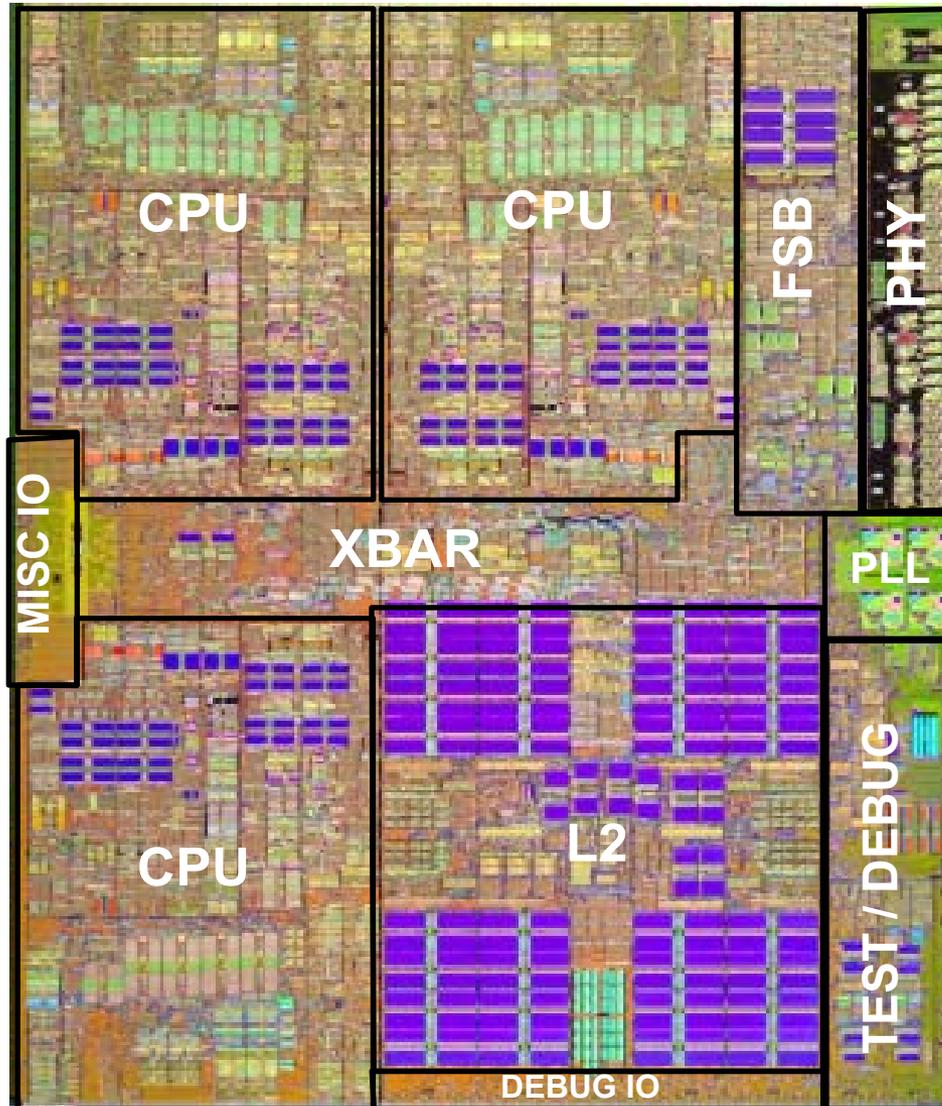


PHY Transmitter

- Low latency serializer architecture
- Serializes the 4 bit parallel data for the CPU and 8 bit parallel data for the GPU
- Clock synchronizer picks the phase of the half rate clock to latch the 4 bit word correctly.
 - Eliminates the need for a FIFO and reduces latency
 - Requires careful design of the PHY clock and ASIC clock as a system.



CPU Chip Photo

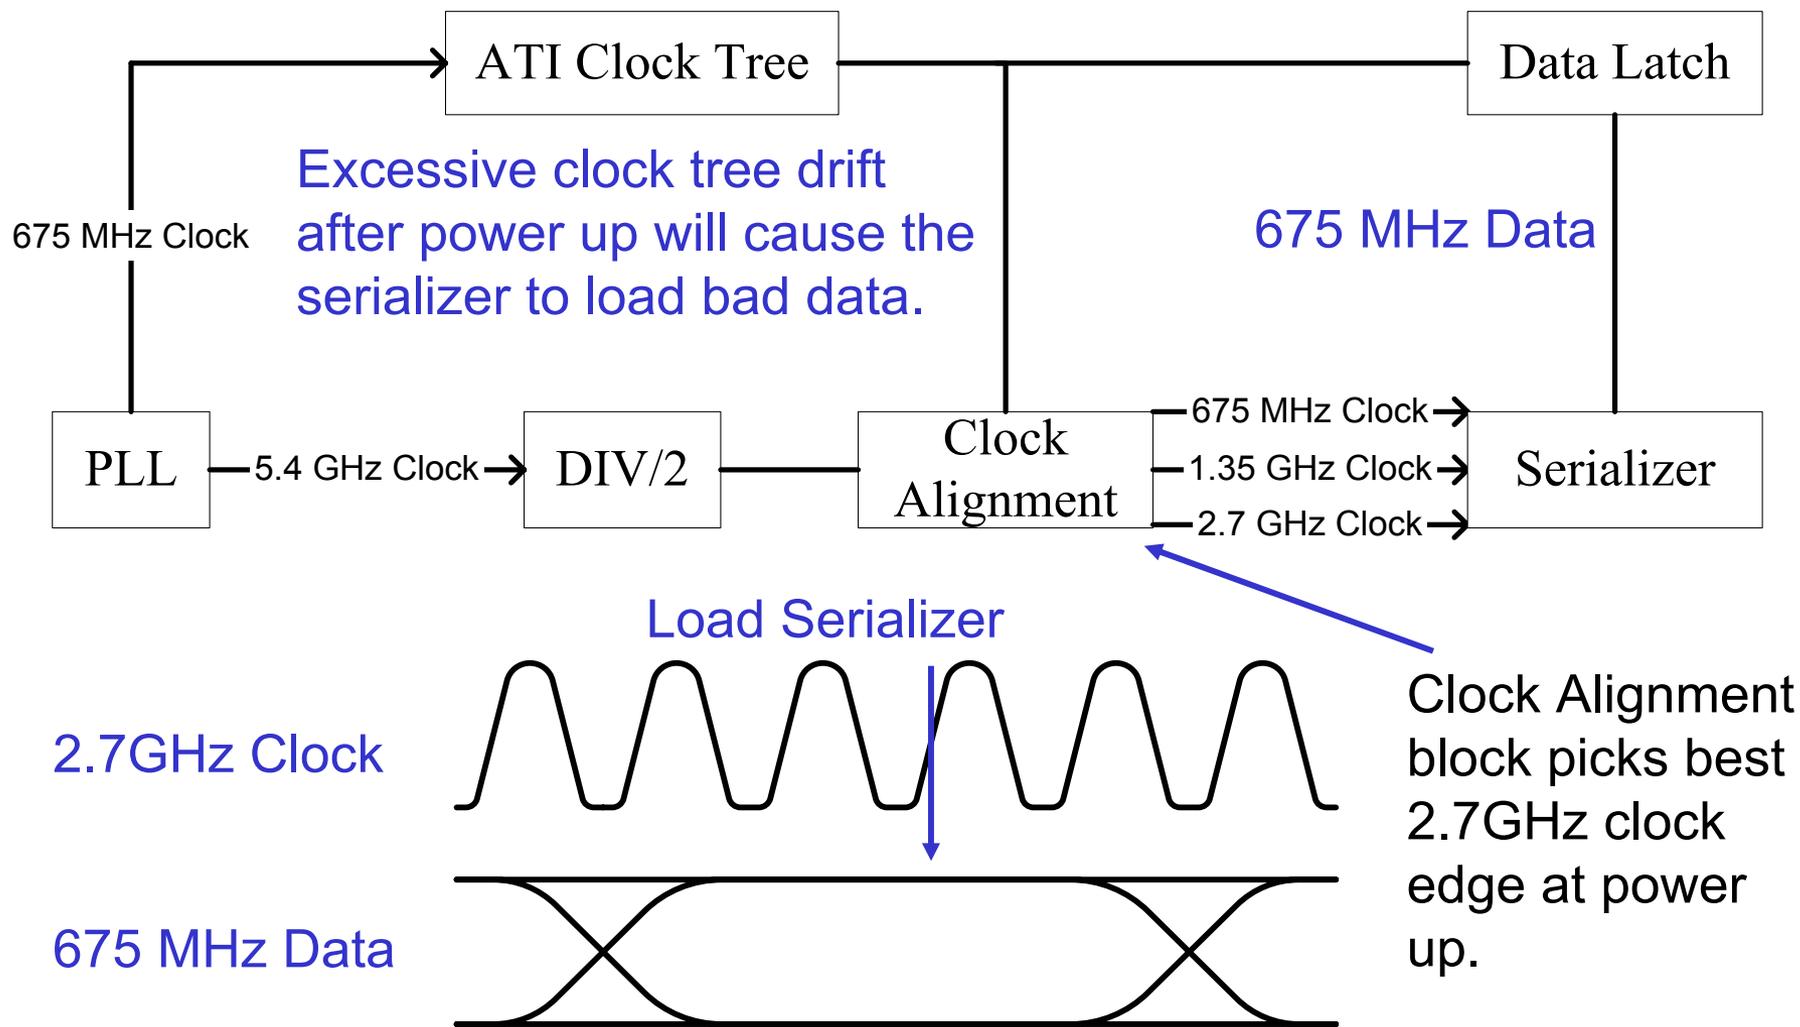


- 3-Way Symmetric Multi-Processor
 - IBM PowerPC™ Architecture
 - 3.2GHz CPU clock
- FSB – 21.6 GByte/sec aggregate bandwidth
- 90 nm CMOS SOI (silicon on insulator)

GPU PHY

- GPU – TSMC 90 nm bulk CMOS
- The GPU PHY uses the same architecture as the CPU except that the parallel data is 8 bits wide vs. 4 bits wide.
- The GPU FSB PHY was an integration challenge.
 - The GPU was designed by a graphics chip company.
 - The PHY designed by Cadence Design Systems.
 - The FSB PLL was designed by a PLL design company.
- GPU PHY Transmit Interface
 - As with CPU PHY a clock alignment circuit is used instead of a FIFO for lower latency.
 - Clock sync is a challenge because of 2 clock paths from 3rd parties.
 - Classic timing methodologies didn't apply.
 - Worst setup corner wasn't the slow corner.
 - Worst hold corner wasn't the fast corner.

Clock Alignment Block Diagram



GPU PHY Design Considerations

- GPU PHY High Speed Clock Interface
 - 5.4 GHz clock is provided from the PLL to the PHY for low jitter.
 - To assure correct operation the PHY developers received a PLL model and the PLL developers received a PHY model.
 - This included relevant circuits and wiring.
 - Identical circuits and references were used in the PHY & PLL 5.4 GHz paths.
- Other GPU PHY Development Logistics
 - Establishing communications and legal permissions across 4 companies was a challenge.
 - Example: Considerable communication was needed between the 4 companies to assure proper start up occurs for the GPU/CPU system.

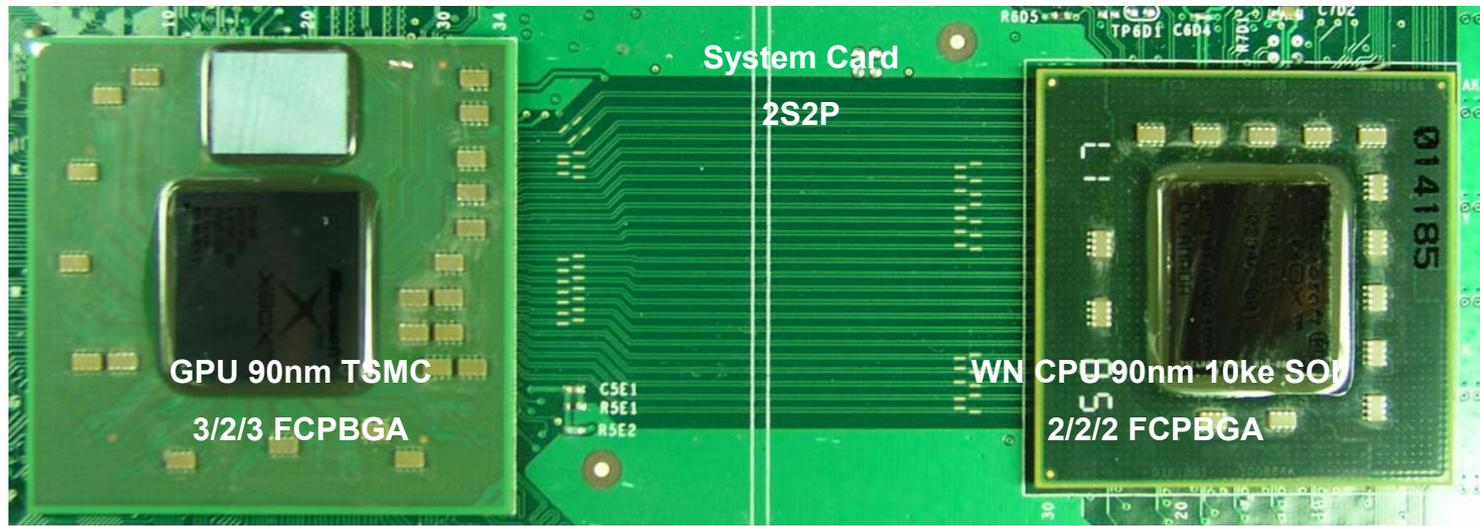
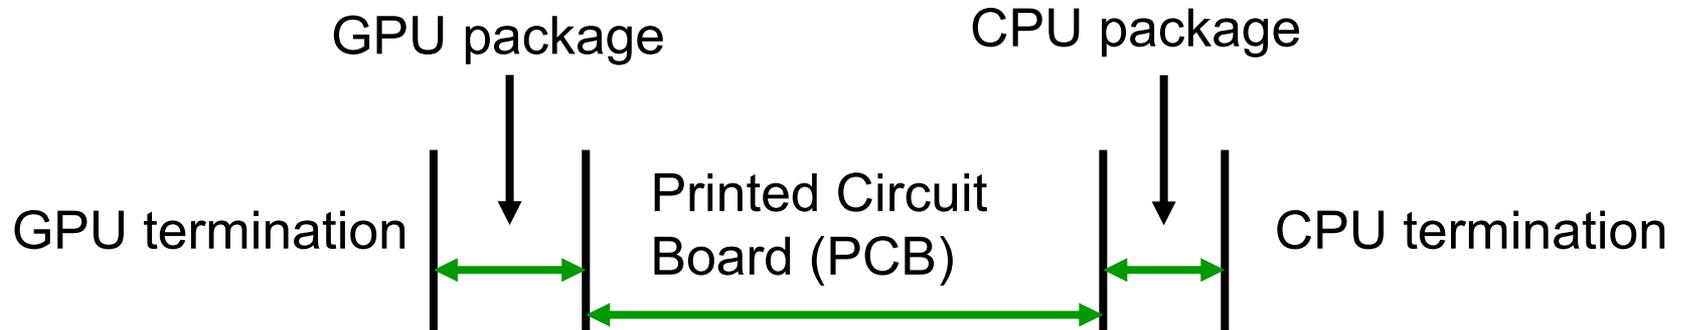
Physical Channel

- High volume low cost packaging from multiple vendors
- 2 signal wiring layers on the board and GPU package
- 1 signal wiring layer on the CPU package
- Low cost means that the tolerances can't be tight
- Table specifies worst case values

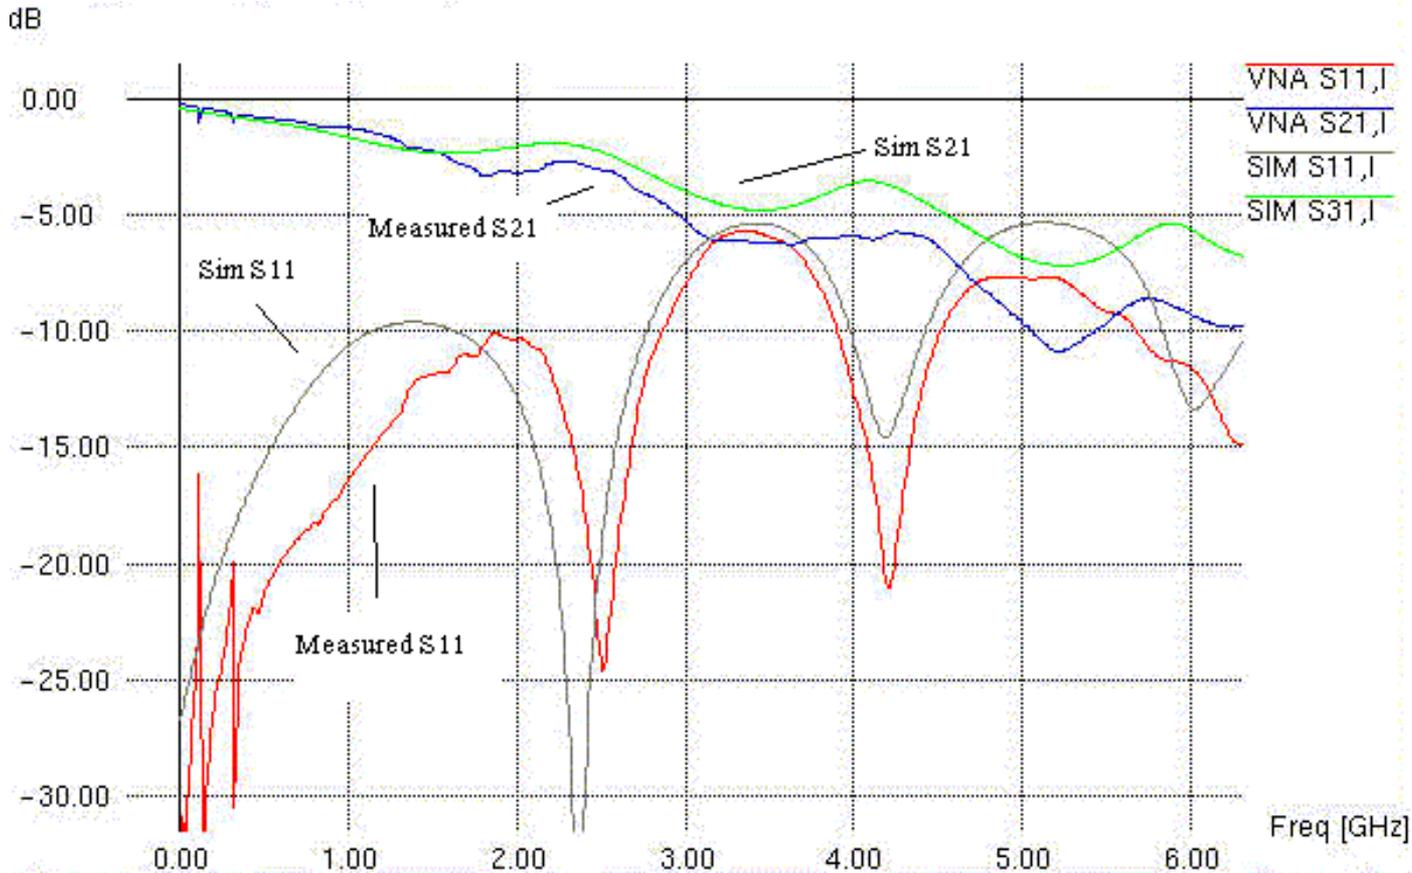
Electrical Characteristic	Chip Carrier Wiring	PCB Wiring
Differential Impedance	100 +/- 18 Ohms	100 +/- 15 Ohms
Differential Attenuation	-1.7dB @ fBaud/2	-1.5dB @ fBaud/2
Return Loss	-10dB @ 0.75*fBaud	NA
Pair-to-pair Isolation	-30dB @ 0.75*fBaud-	-30dB @ 0.75*fBaud
Maximum Length	18mm	50.8mm

Physical Channel

5 reflection interfaces

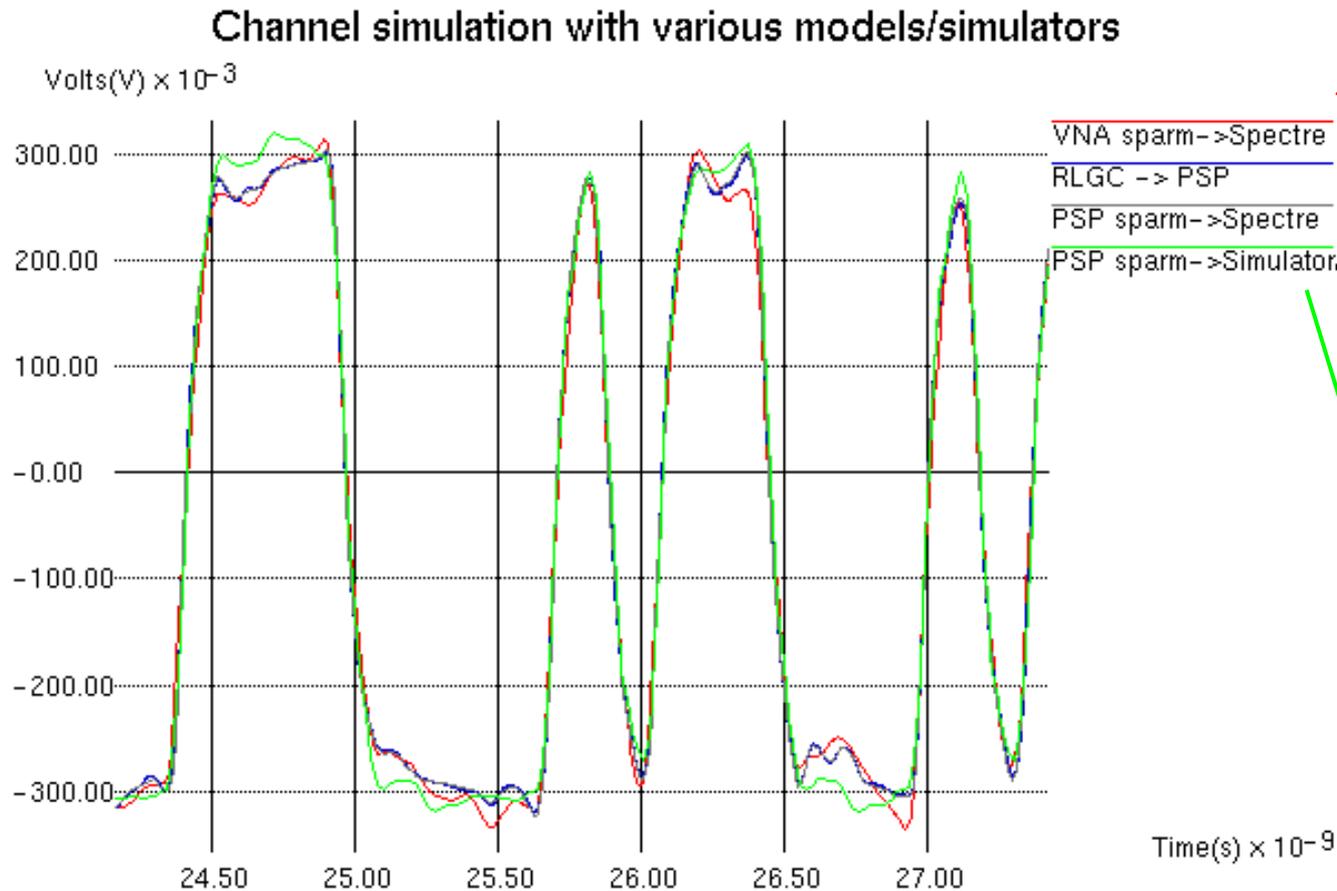


Model to Hardware Correlation



- Simulation and measurement of the entire channel, chip solder ball to chip solder ball.

Model Correlation



• S-params from VNA measurements

• RLGC model simulated with PowerSPICE™

• S-params simulated w/ Spectre™

• S-params simulated with another commercial SPICE simulator

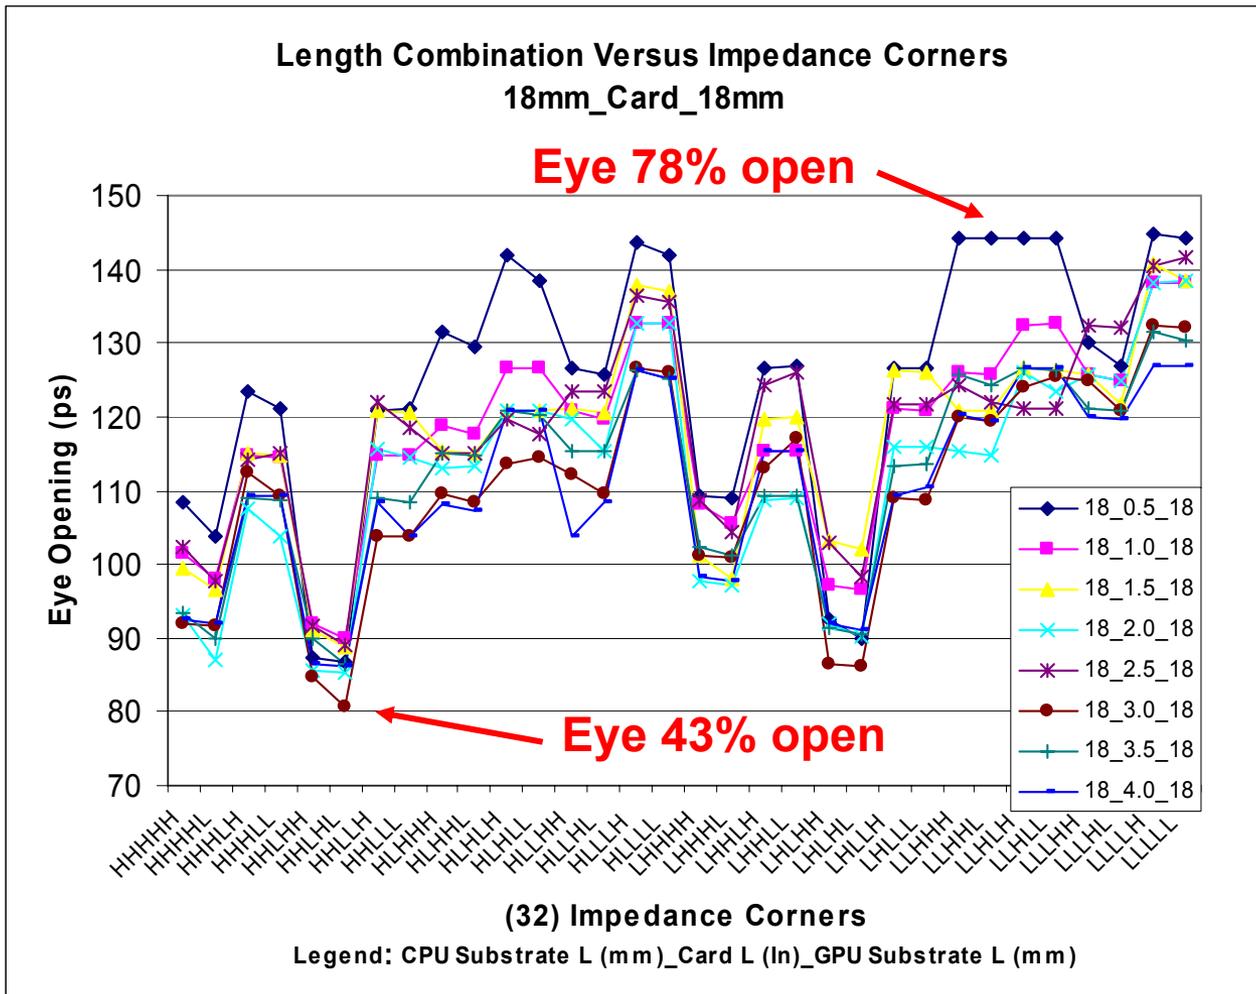
Impedance Discontinuities

- Short channel and low attenuation means reflections keep bouncing back & forth
- Some reflections with some data patterns interfere destructively
- Spice doesn't simulate sufficient data pattern lengths

	Channel Impedance		
	Min	Nom.	Max
Driver return loss			-5.0 dB
Driver package	85.4Ω	93.5Ω	109.7Ω
PCB	85.0Ω	100.0Ω	115.0Ω
Receiver package	85.4Ω	93.5Ω	109.7Ω
Receive return loss			-5.5 dB

- Return losses are high due to high customer ESD requirements and SOI ESD capabilities.
- Some paths were tuned to reduce the reflection impact.

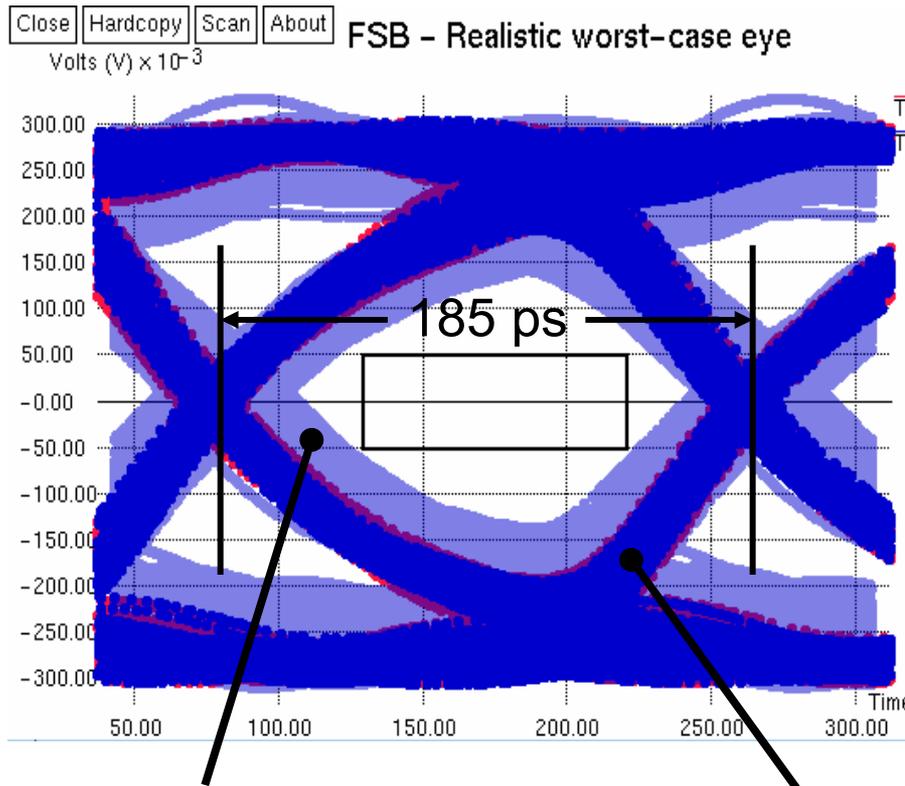
Eye Opening Variation due to Reflections



- IBM communications based link simulator tool used to simulate millions of bits to find worst case data patterns.
- Spice simulates 100's of bits.
- 2^5 Impedance corners
- Trace lengths from 0.5 to 4 inches

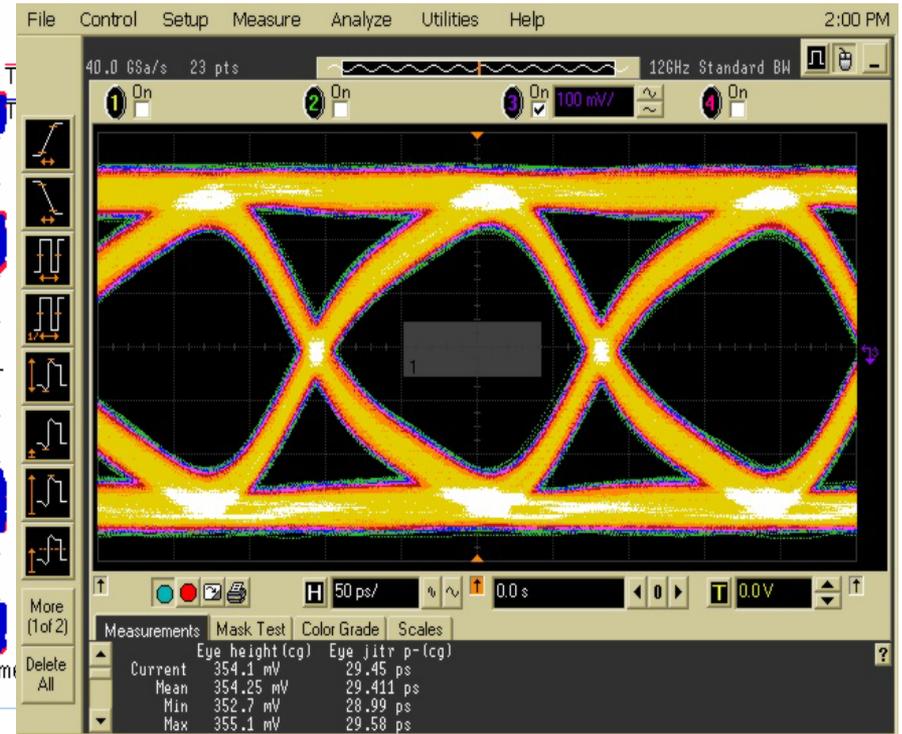
Simulated vs. Measured Eye Diagram

Simulated TX Eye



Non-ideal impedance combination

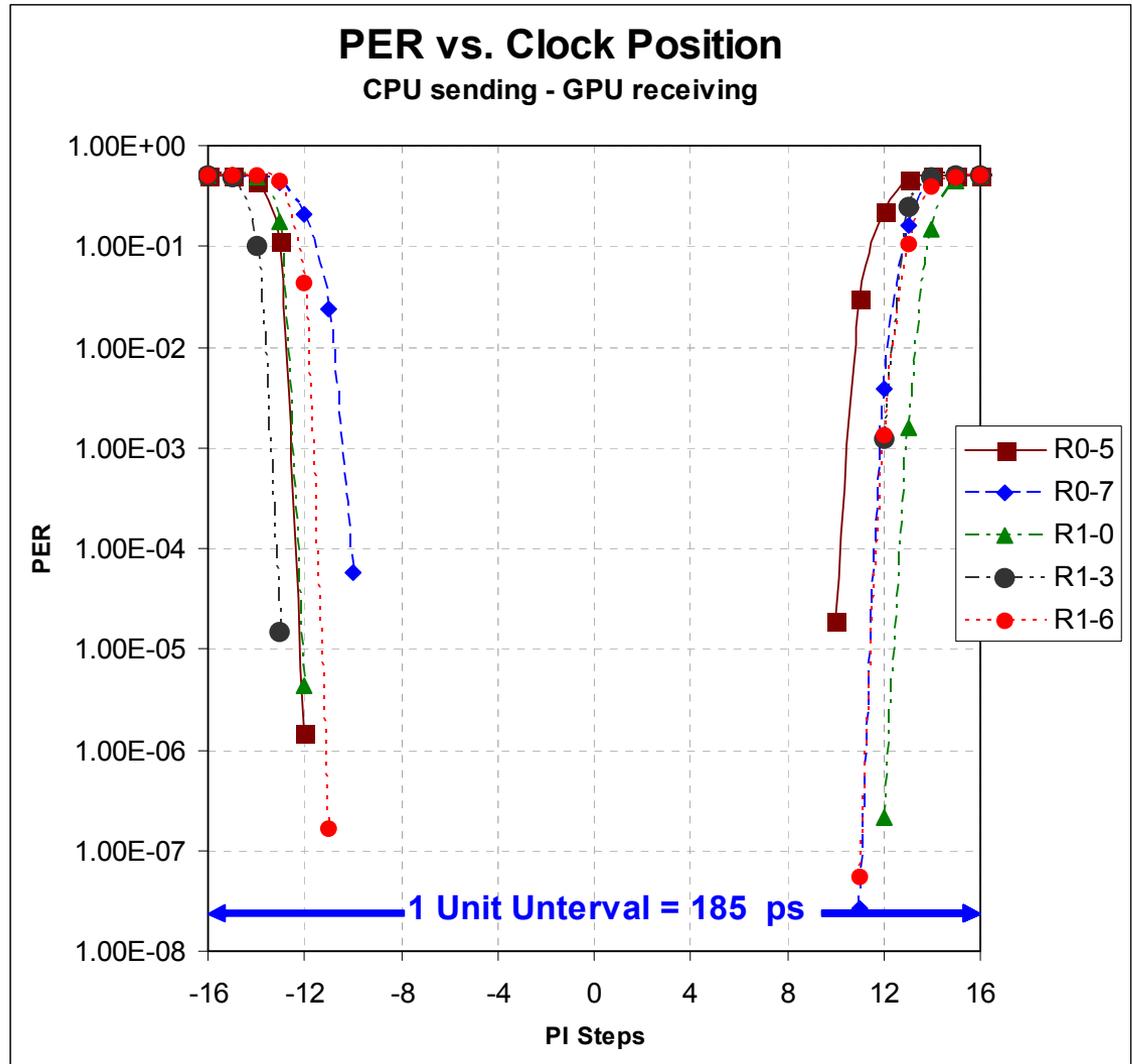
Measured TX Eye



Nominal impedances

Packet Error Rate vs. Clock Position

- Phase rotators are manually adjusted
- The packet error rate is plotted vs. clock position
- This is done on a product level system board without any special hardware
- Errors are checked with the existing error checking logic in the link layer



Summary & Conclusions

- The entire on chip clock paths were analyzed as a system to enable a low latency transmit architecture.
- In-system error rate curve measurement enabled characterization and verification of production hardware.
- Reflections are a large detractor for short links.
- It was important to have one FSB link owner.
 - This enabled a robust link design over large channel variation.
 - Over 10 million units have shipped as of year end 2006.