
PCI
SIG®

Transition to PCIe™ – Learning from PCIe 1.0 Challenges and Looking Forward to PCIe 2.0

Betty Luk

Senior Engineer

ATI Technologies

Agenda

- PCI Express® Transition In Review
- Electrical Challenges
- Logical Challenges
- Power Management
- Debugging PCI Express Systems
- Looking Forward – What's Next?

Transition In Review

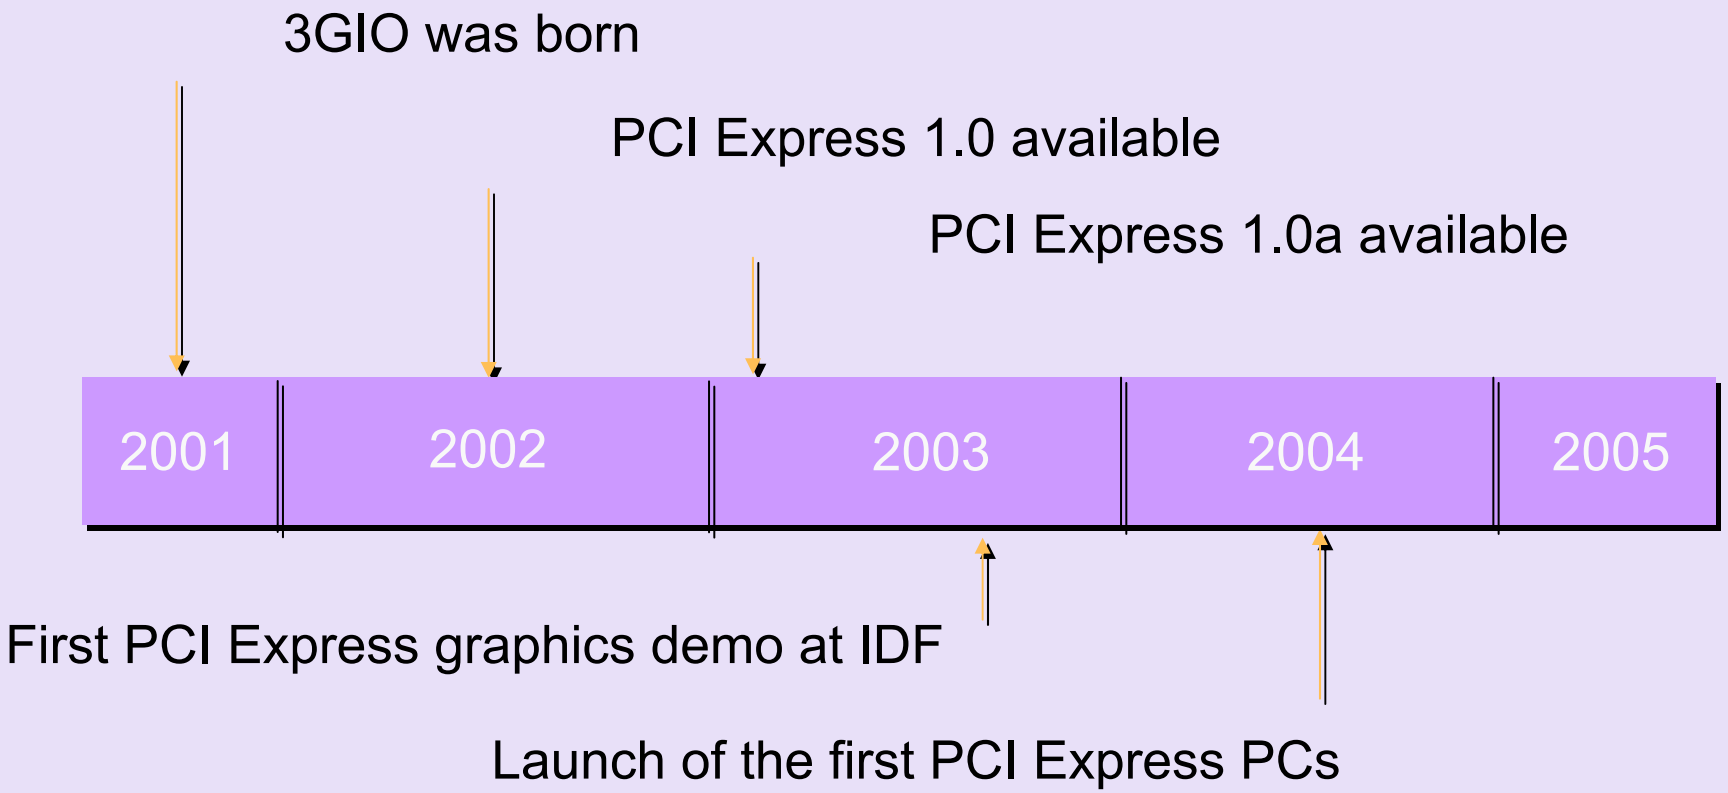

Where Are We Today…

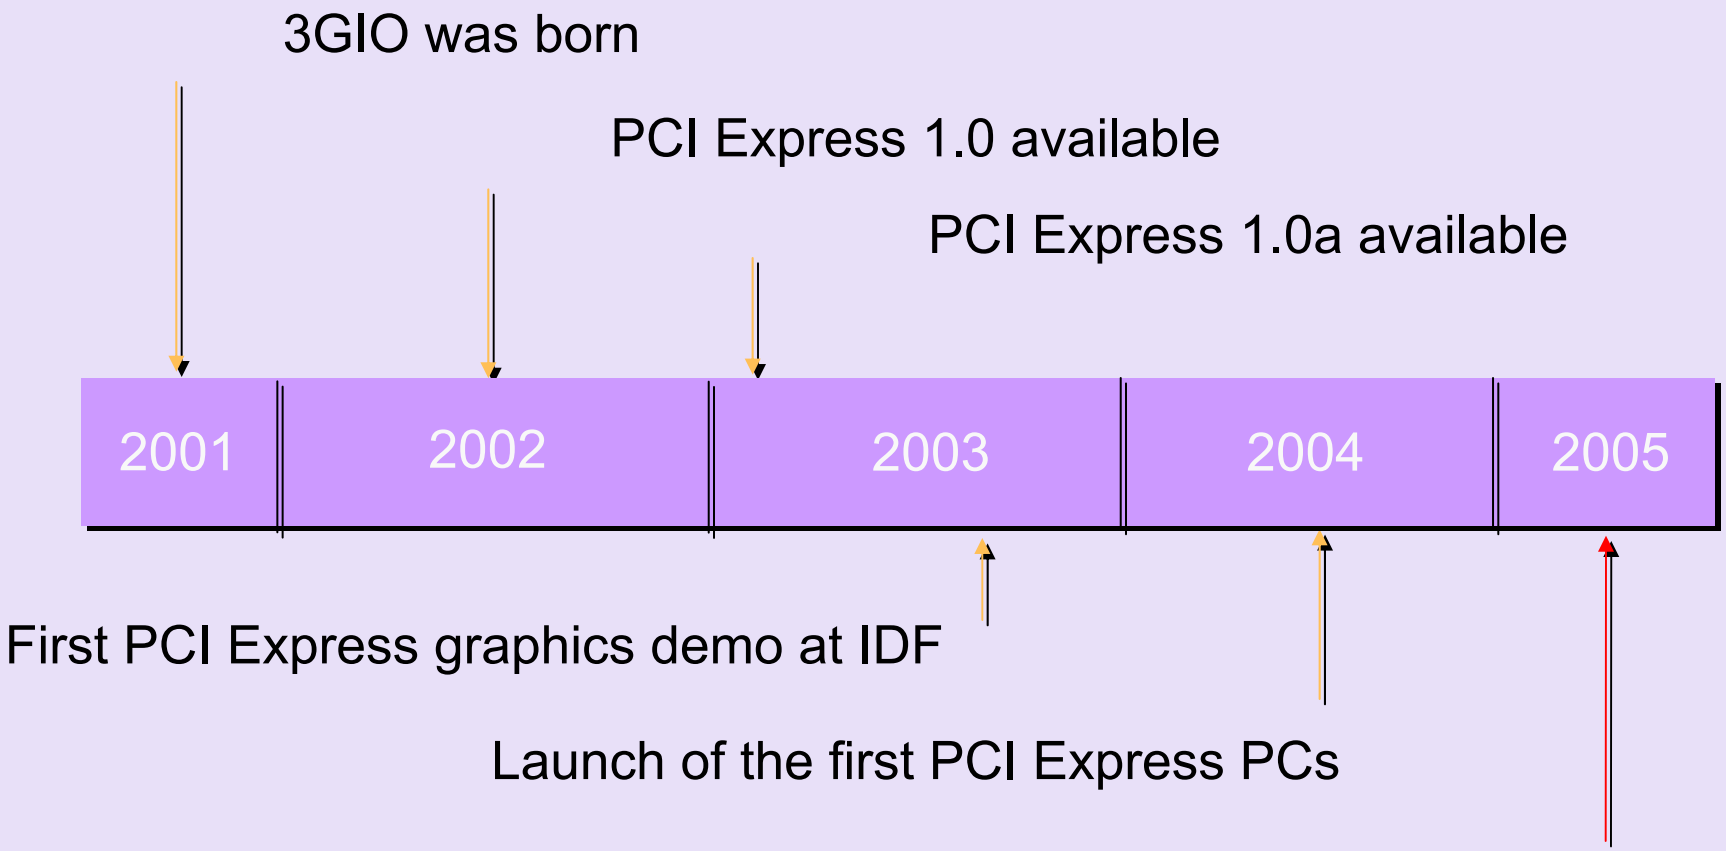

Over 13.5 million PCIe chips shipped by ATI!

What Were The Challenges?

- PCIe is a layered serial interface
- Need to start building from the bottom up
- Electricals are the foundation of the interconnect

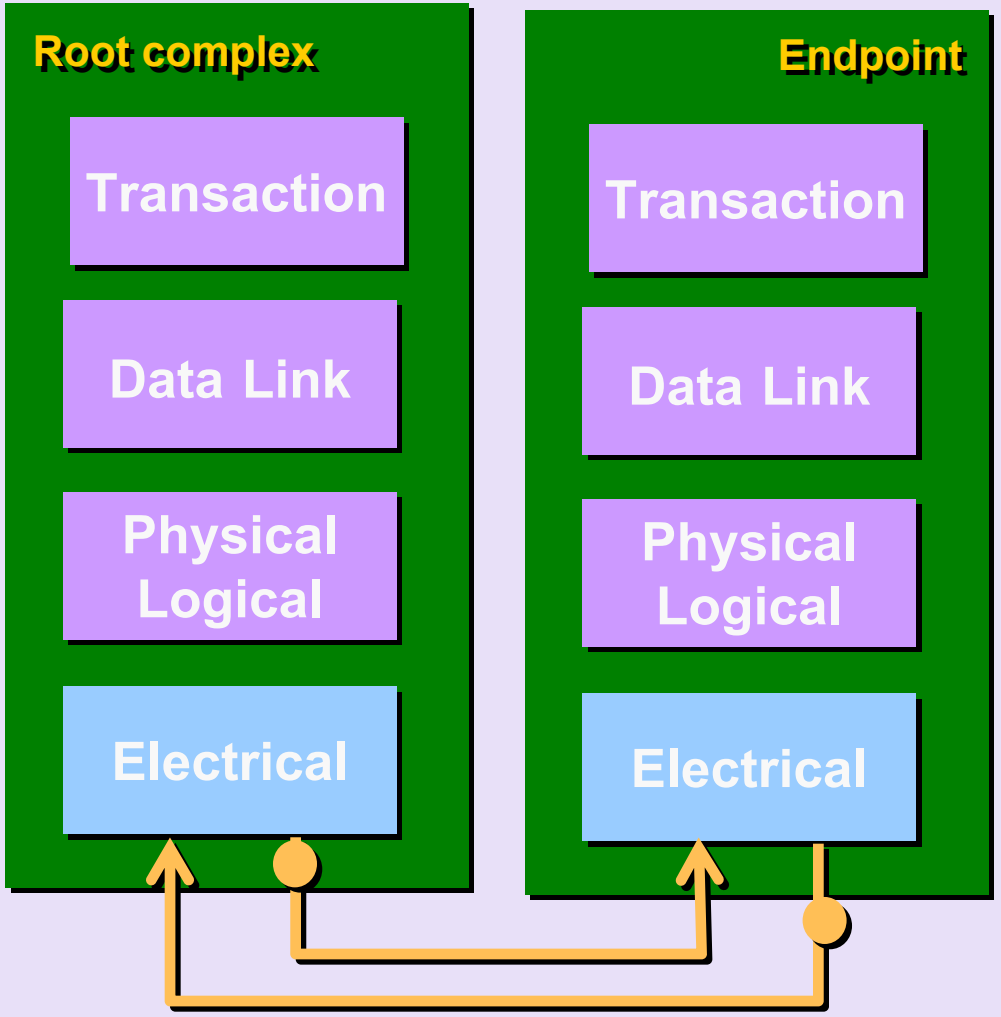

Basic Electrical Parameters

- Basic electrical parameter defined by the PCI Express specification is the eye mask
  - ✓ Defines amplitude margin
  - ✓ Defines jitter margin
- Transmitter eye diagram testing is done as part of the compliance tests

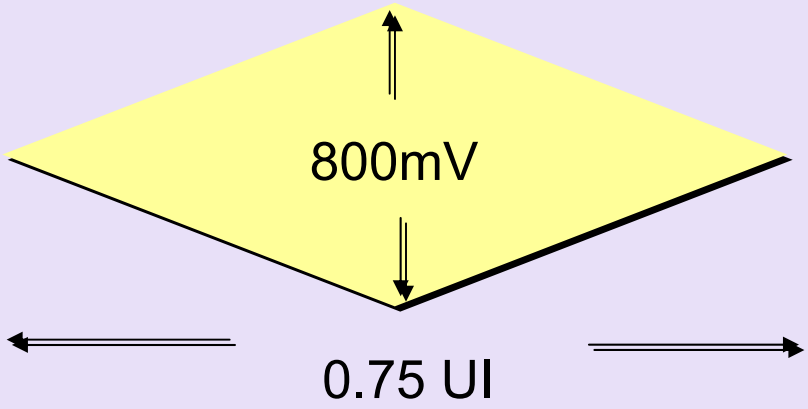

Transmitter eye mask

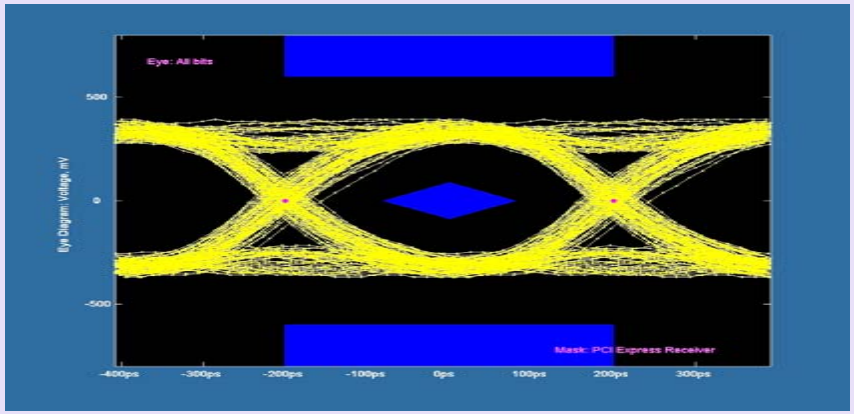

Sample eye diagram

What About The Receiver?

- No standard tests are defined for the receiver
  - ✓ Need robust transmitter and receiver for a functional PCIe link!
- Receiver testing is more important
- Ensure receiver design is robust by:
  - ✓ Loopback tests
  - ✓ Injecting minimum eye opening
  - ✓ Injecting jitter
- Need built-in pattern checkers
  - ✓ E.g. toggle pattern, compliance pattern, PRBS

Don’t forget to validate the receiver!

Is That Enough?

- Testing the PHY in a CBB environment shows a clean eye mask with good jitter and amplitude margin
- This is an “ideal” environment for the PHY
  - ✓ Rest of the chip is quiet, not a lot of simultaneous switching noise
  - ✓ Power delivery is stable
  - ✓ Thermals are not a concern
  - ✓ Updated specs for transmitter are with an ideal refclk

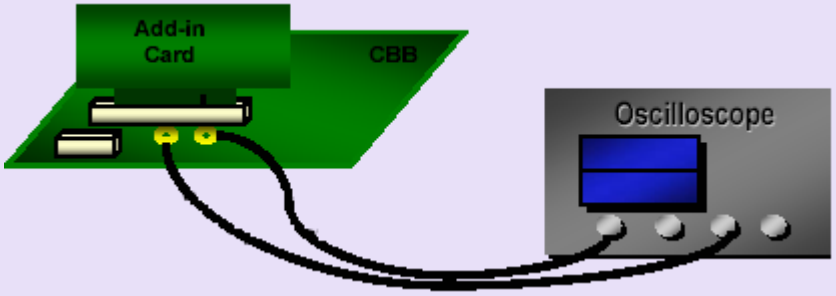

- What happens in a real system?

Real Life Electricals

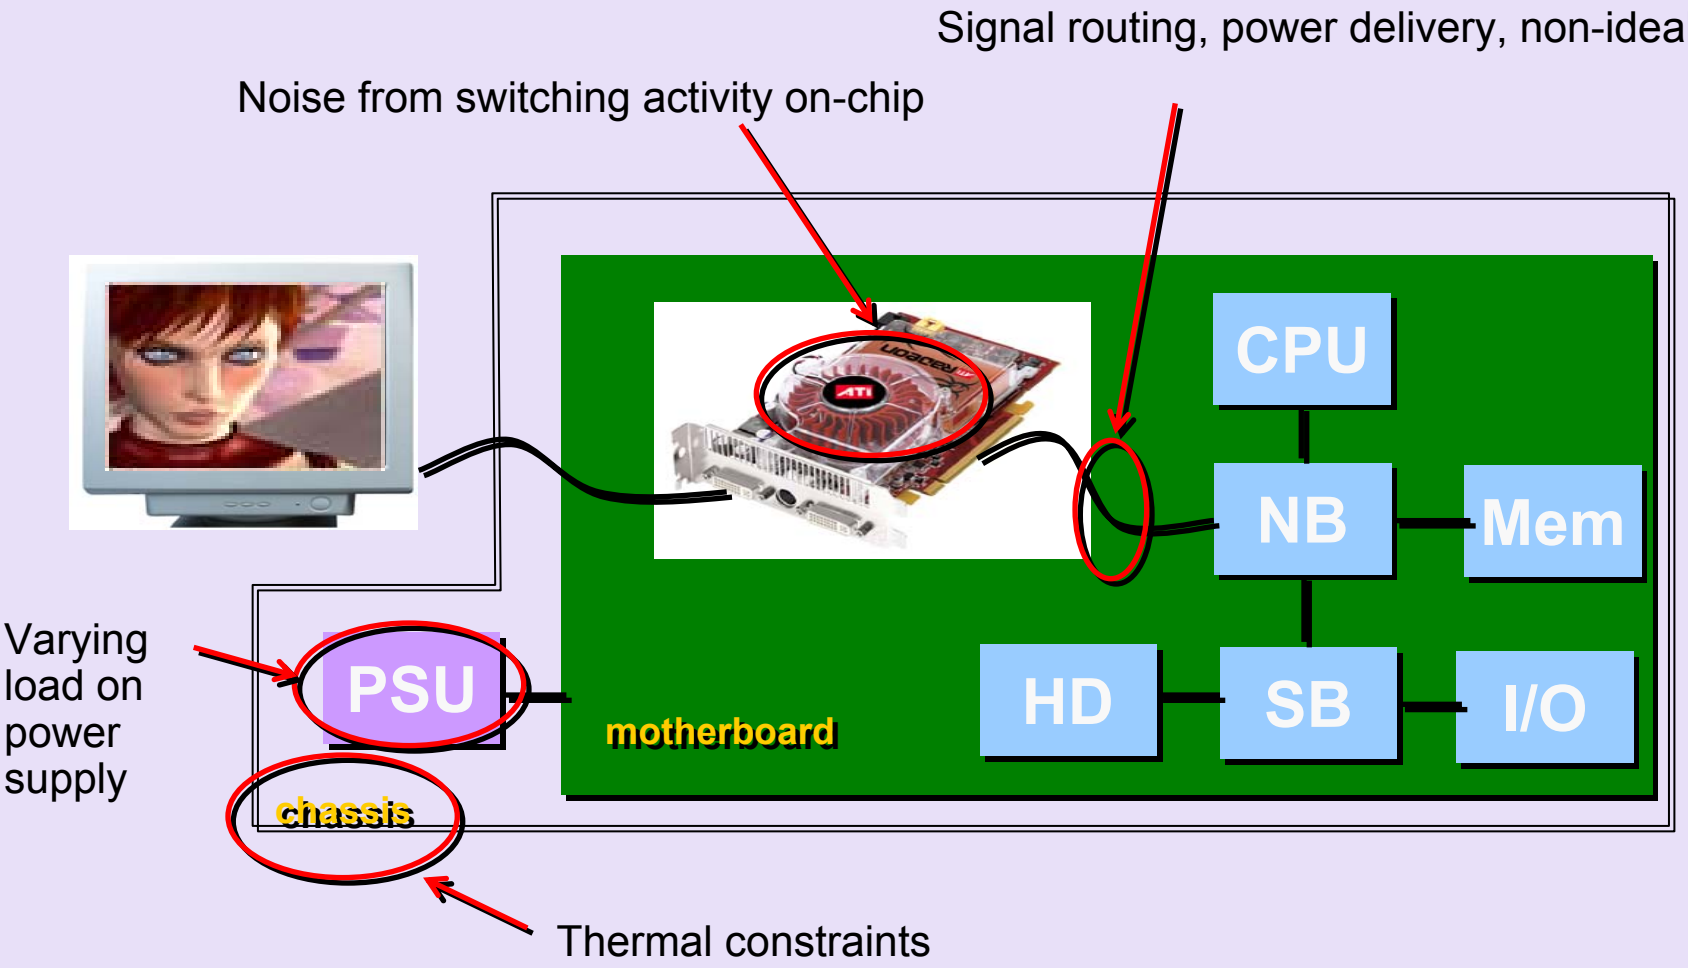

Mobile Systems

- On top of all the system concerns discussed
  - ✓ More sharing of power rails to save real estate
  - ✓ Power noise from switching regulators
  - ✓ More routing constraints
  - ✓ Thermal constraints are even tighter
- Standard electrical compliance procedure not available

Identifying Link Quality Issues

- Measurement with an oscilloscope
  - ✓ Good for catching gross electrical issues and eye violations
  - ✓ Good for catching any spectral noise component in the signal
  - ✓ Has limitations
    - Difficult to find probe points close to the receiver
    - Difficult to observe intermittent violations
- Need hardware to monitor link quality
  - ✓ Link quality is only as good as viewed from the receiver!
  - ✓ Monitor NAK rate, link recoveries, and link retrains
  - ✓ Currently no standard defines this
  - ✓ Vendors should provide software utility for interoperabiity test

Lessons Learned

- Robust electrical design with…
  - ✓ Margin! Margin! Margin!
- Conduct interoperability testing as early as possible to flush out issues
  - ✓ Test with non-ideal, cost-reduced systems
  - ✓ Test in all corners of operating range (voltage and temperature)
  - ✓ Monitor link quality during testing
- Must have hooks within your design to give visibility into the link quality

Poor link quality will have a negative impact on overall system performance!

Other Key Components

- Receiver detection
  - ✓ First step in establishing a working link
  - ✓ All lane widths supported by a multi-lane device must be tested
- Electrical idle detector
  - ✓ Not part of any compliance test, but is a crucial component!
  - ✓ Used in initial link training, power state transitions
  - ✓ Gate “noise” that may be present on the link
- Electrical behavior must also be reliable after multiple resets & power cycles
  - ✓ No sensitivities to power-on and reset sequence

Receiver detection and electrical idle detector must be reliable across all operating conditions!

Beyond the Electricals – Link Training

- Ensure all arcs are implemented and tested
  - ✓ Link training may not take the most direct path
  - ✓ State machines being out-of-sync can result in timeouts and a long link train
  - ✓ All states must have a graceful escape condition
- In-band presence detect
  - ✓ Presence detect bit is cleared in Detect state
  - ✓ Software must poll this bit in case there is a timeout back to Detect

Logical Layers

- Common assumption – the PHY is perfect and all symbols received are correct
  - ✓ WRONG!
  - ✓ Non-ideal link quality in a real-life system
- Another assumption – only single bit errors occur
  - ✓ WRONG!
  - ✓ E.g. morphed symbol, corrupted symbol, bursts of errors
- Logic needs to be able to handle received errors
  - ✓ Corrupted symbols should not be propagated up the receive data path

Nak, Replay, Link Recovery

- Nak, replay and link recovery must be handled properly
- Logic must be able to handle even a high Nak rate
- Challenge is – how do we test this?
  - ✓ Simulation
    - Random test cases to insert errors to the receiver
    - Targeted test cases to insert errors at a specific point in the sequence
    - Specific end cases must be covered
    - Make use of code coverage and assertion based tools
  - ✓ Silicon validation
    - Artificially induce errors in the receiver
    - Bus exerciser

Productizing PCIe

- Motivation for PCI Express is to provide a high bandwidth interface
- This increased bandwidth comes at a cost
  - ✓ More POWER!
- PCI Express dissipates much more power than PCI / AGP designs
- System thermal budgeting becoming increasingly difficult
- Aggressively use ASPM L0s / L1 power savings features

Power Management

- L1 is the most beneficial power saving feature

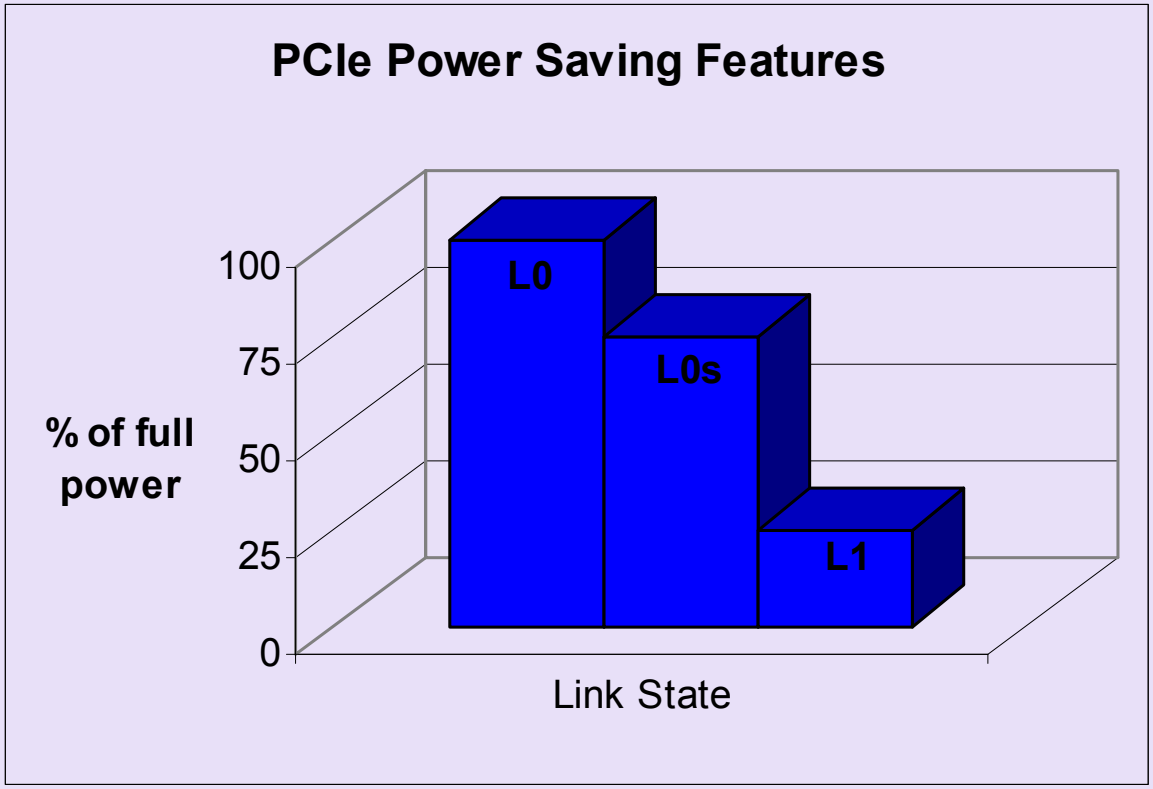

Get L0s and L1 working!

Challenges with L0s and L1

- Logic – PHY interface
  - ✓ Must be defined clearly
- Enabling / disabling transmitter and receiver
  - ✓ Current surges
  - ✓ Noise
- Exit latency
- PLL
  - ✓ How long does it take for PLL to start up?
  - ✓ Graceful start and stop of the clock
    - No glitches
    - Stable

L0s, L1 cont'd…

- Programming model
  - ✓ Enable L1 on the root complex first
- Design flexibility
  - ✓ Tune entry and exit strategies to get maximum power savings without sacrificing throughput

Do extensive testing with L0s and L1!

Power Savings Innovations

- Look for other ways to save power
- In a multi-lane device, not all lanes need to be used at all times
- ATI Dynamic Lane Count Switching™
  - ✓ Use less number of lanes when more power savings is required

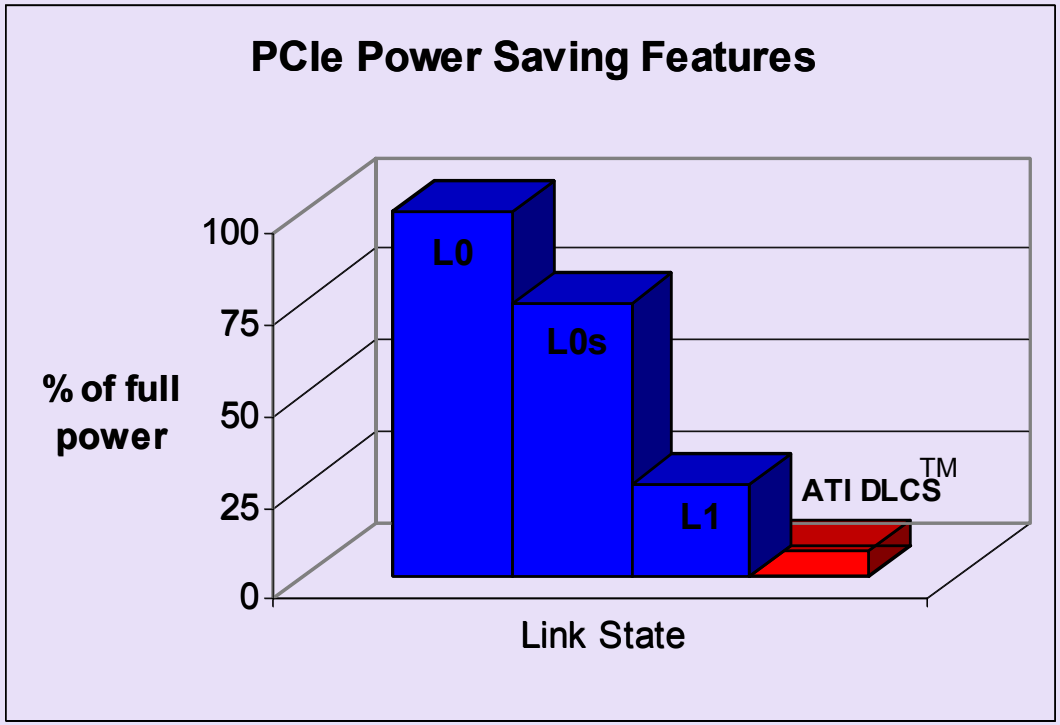

Debugging PCIe Systems

- PCIe analyzers
  - ✓ Interposer probing solution
    - Additional trace length
    - Additional connector
  - ✓ Ability to turn off spread spectrum clocking in the target system
  - ✓ Remember logic analyzer is just another receiver!
    - Need to capture reliably
    - Non-intrusive – cannot degrade the signal
    - Need to support L0s / L1

Electrical margin is key!

Designs without a PCIe Connector

- Solutions for connector-less designs
  - ✓ Mid-bus probes
  - ✓ Solder-on probes
  - ✓ Dropping down to a lower link width (x1)

Built-in Visibility

- Key to debugging any issue
  - ✓ VISIBILITY!
    - Have to be able to see what is wrong!
- Visibility inside the hardware itself
  - ✓ Build debug signals into the design
  - ✓ Isolate failing block
  - ✓ Enables tracing over long time periods
  - ✓ Enables lower cost debugging
  - ✓ Examples
    - Link state
    - Lane deskew

Industry Initiatives

- Compliance
  - ✓ Devices must interoperate
  - ✓ Increase compliance test coverage

***STOP: 0x000000D1 (0x00000000, 0xF73120AE, 0xC0000008, 0xC0000000)

A problem has been detected and Windows has been shut down to prevent
further damage to your computer

Industry Initiatives

- Compliance
  - ✓ Devices must interoperate
  - ✓ Increase compliance test coverage
- Avoid the dreaded BSOD

Hardware Initiatives

- Consider a system with poor link quality
- Link recovery always completes successfully
- Consequences:
  - ✓ System performance is very slow
  - ✓ Can cause software to timeout and system to hang!
- Improve HW ability to recover the link

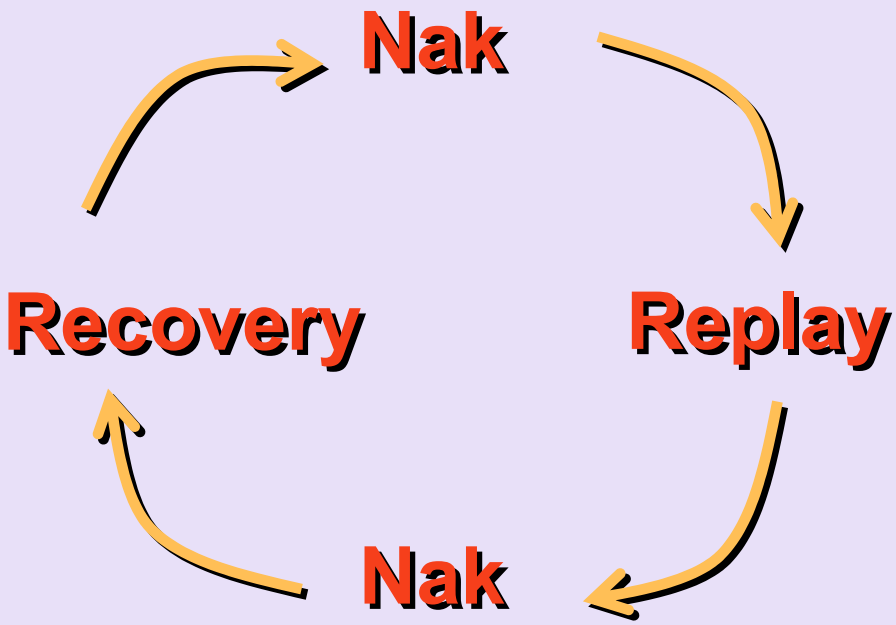

Software Initiatives

- Make use of error reporting features to recover from errors
- Need software support to manage detected errors
  - ✓ With current software, link down must not occur
  - ✓ Example: Completion timeout
    - Results in all F's being processed as data
    - Likely outcome – system hang or BSOD!
    - Better scenario – resend the request again

Preparing for Gen 2

- 5GB/s signaling rate
- Electricals are even more challenging!
- Gen 2 test equipment available
  - ✓ Requirements:
    - Must be non-intrusive
    - Must have margin to debug non-ideal systems
- Thoroughly test link speed negotiation
  - ✓ Must be backward compatible with Gen 1 speed

Call to Action

- Ensure good link quality
- Receiver validation is imperative
- Robust Nak, replay, and recovery mechanism
- Validate L0s and L1
- Increase compliance test coverage
- Implement fault tolerant PCI Express systems
- Prepare for Gen 2 – design and test

Thank you for attending the
PCI-SIG Developers Conference 2005.

For more information please go to
www.pcisig.com


PCI
SIG®