PCIe™ 2.0 Logical Extensions

Debendra Das Sharma

Member, EWG

Agenda

- Overview of Logical Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Enhancements for Robustness
- Summary & Call to Action

Logical Extensions: Overview

- Changes limited to physical layer only
- Backwards compatibility with 1.1 spec
- Considerations
  - ✓RAS
  - ✓Power efficiency
  - ✓Robustness in design (HVM considerations)
  - ✓Ease of design and validation

Logical Extensions: Overview

Extensions	Explanation	Benefits
Speed Negotiation	Capability to upgrade or downgrade link speed	RAS (improved link uptime), dynamic link speed optimization, power savings (25%+)
Compliance Speed	Programmable as well as inband mechanism to select compliance pattern speed	Flexibility to perform compliance testing at multiple speeds with low cost
Electrical Idle Entry and Exit	Protocol changes to facilitate circuit design	Enhanced robustness, yield, power savings, ease of design (TTM)
Compliance Entry/Exit	Device Configuration despite link failures	High Availability, enhanced robustness, no compliance hang

Agenda

- Overview of Logical Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Enhancements for Robustness
- Summary & Call to Action

Speed Negotiation: Philosophy

- Ability to change link speed multiple times without taking the link down
  - ✓ Reliability:
    - Link goes Gen 1 if fails to operate at higher rate after being operational at higher rate for a while
    - Revert back to prior speed if a speed change did not succeed
  - ✓ Power improvements: can change to a lower data rate when bandwidth requirements are low to conserve power and upgrade later when one needs the higher bandwidth
- Ability to notify software on bandwidth change

Speed Negotiation

- Initially link trains to L0 in Gen 1 speed
- Supported speeds advertised in all TS ordered sets
  - ✓Supported speeds by the other component noted in Config.Complete and Recovery.RcvrCfg
- Speed changed through handshake in Recovery state
  - ✓New substate: Recovery.Speed
  - ✓Speed changed in Recovery.Speed
- Polling.Speed state obsolete

Speed Negotiation through Recovery

- LTSSM enters Recovery from L0 if a speed change is desired
  - ✓ Can be initiated by hardware or software
  - ✓ Both sides exchange speed information through TS ordered sets in Recovery.RcvrLock and Recovery.RcvrCfg
    - Includes supported speeds as well as intent to change speed
  - ✓ Speed changed in Recovery.Speed

Relevant LTSSM States

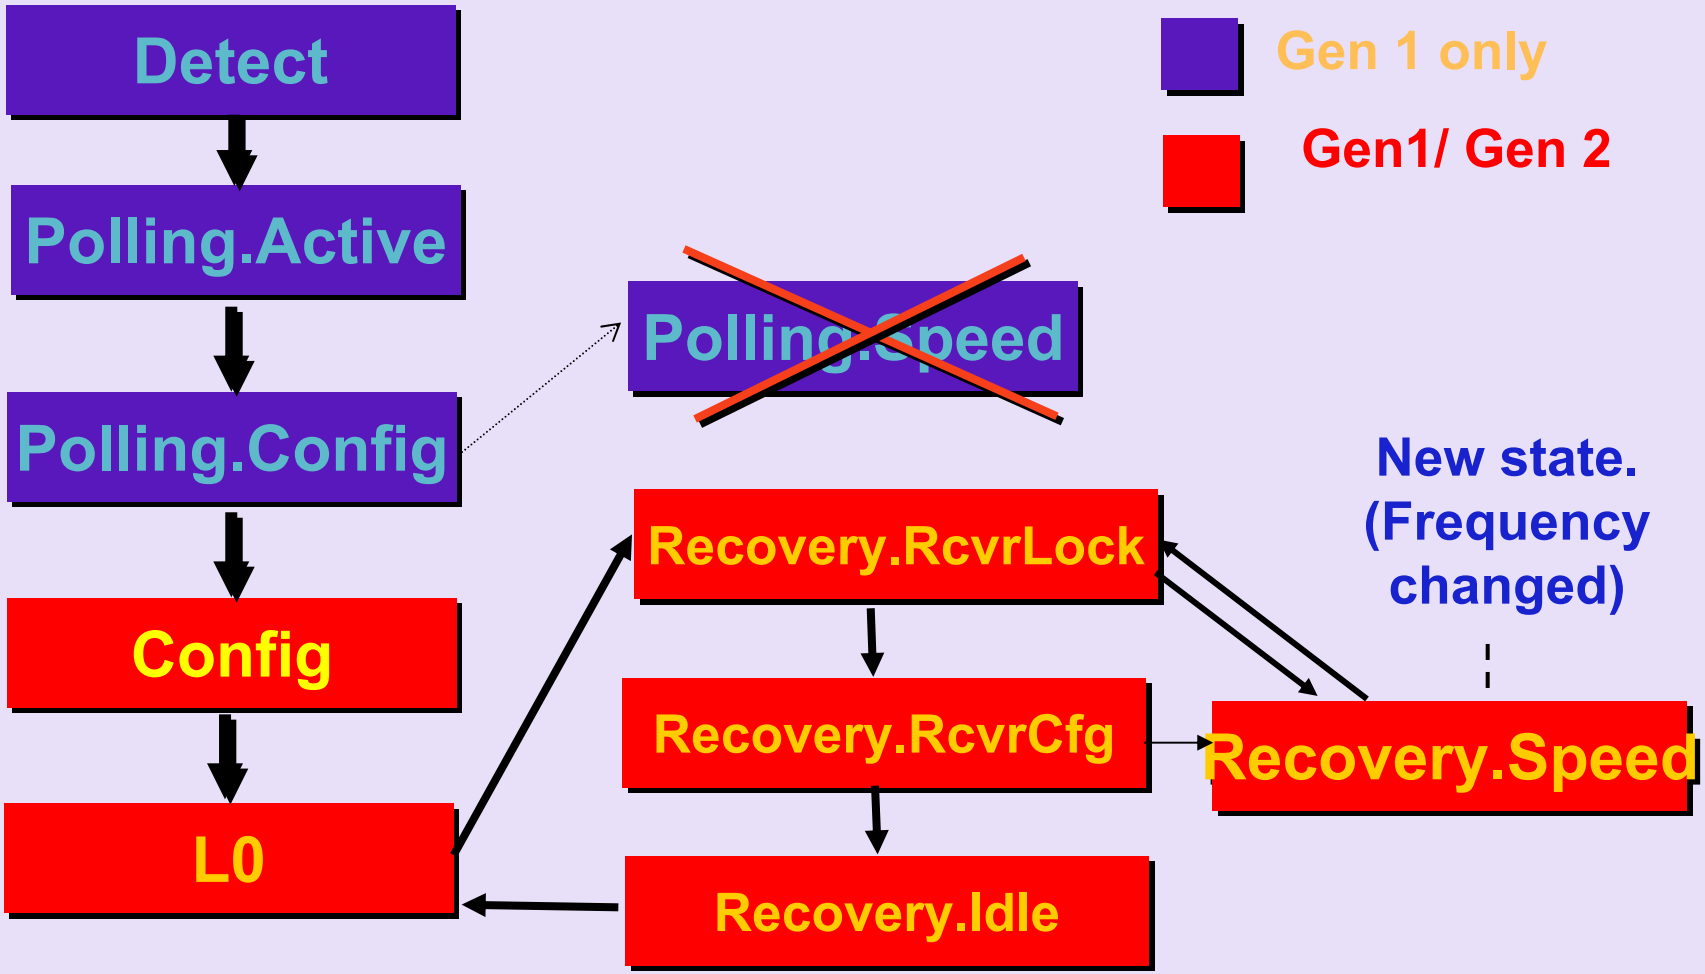

Training Sequence (TS1/TS2 changes)

Symbol Number	Allowed Values	Encoded Values	Description
4	2, 6, 130, 134	D2.0, D6.0, D2.4, D6.4	Data Rate Identifier Bit 0 – Reserved, set to 0 Bit 1 = 1, generation 1 (2.5 Gb/s) data rate supported Bit 2 = 1, generation 2 (5 Gb/s) data rate supported. Bit 3:6 – Reserved for future gen speeds past gen 2, set to 0 for devices that only support Gen 1 and/or Gen 2 speeds. Bit 7 (speed_change) = 1, Requesting to change the speed of operation. The remaining bits (0-6) specify the highest speed with which we can reliably operate. This bit can be set to 1 only during Recovery.RcvrLock state. All lanes under control of an LTSSM must transmit the same value in this symbol.

LTSSM Speed Change: Example 1

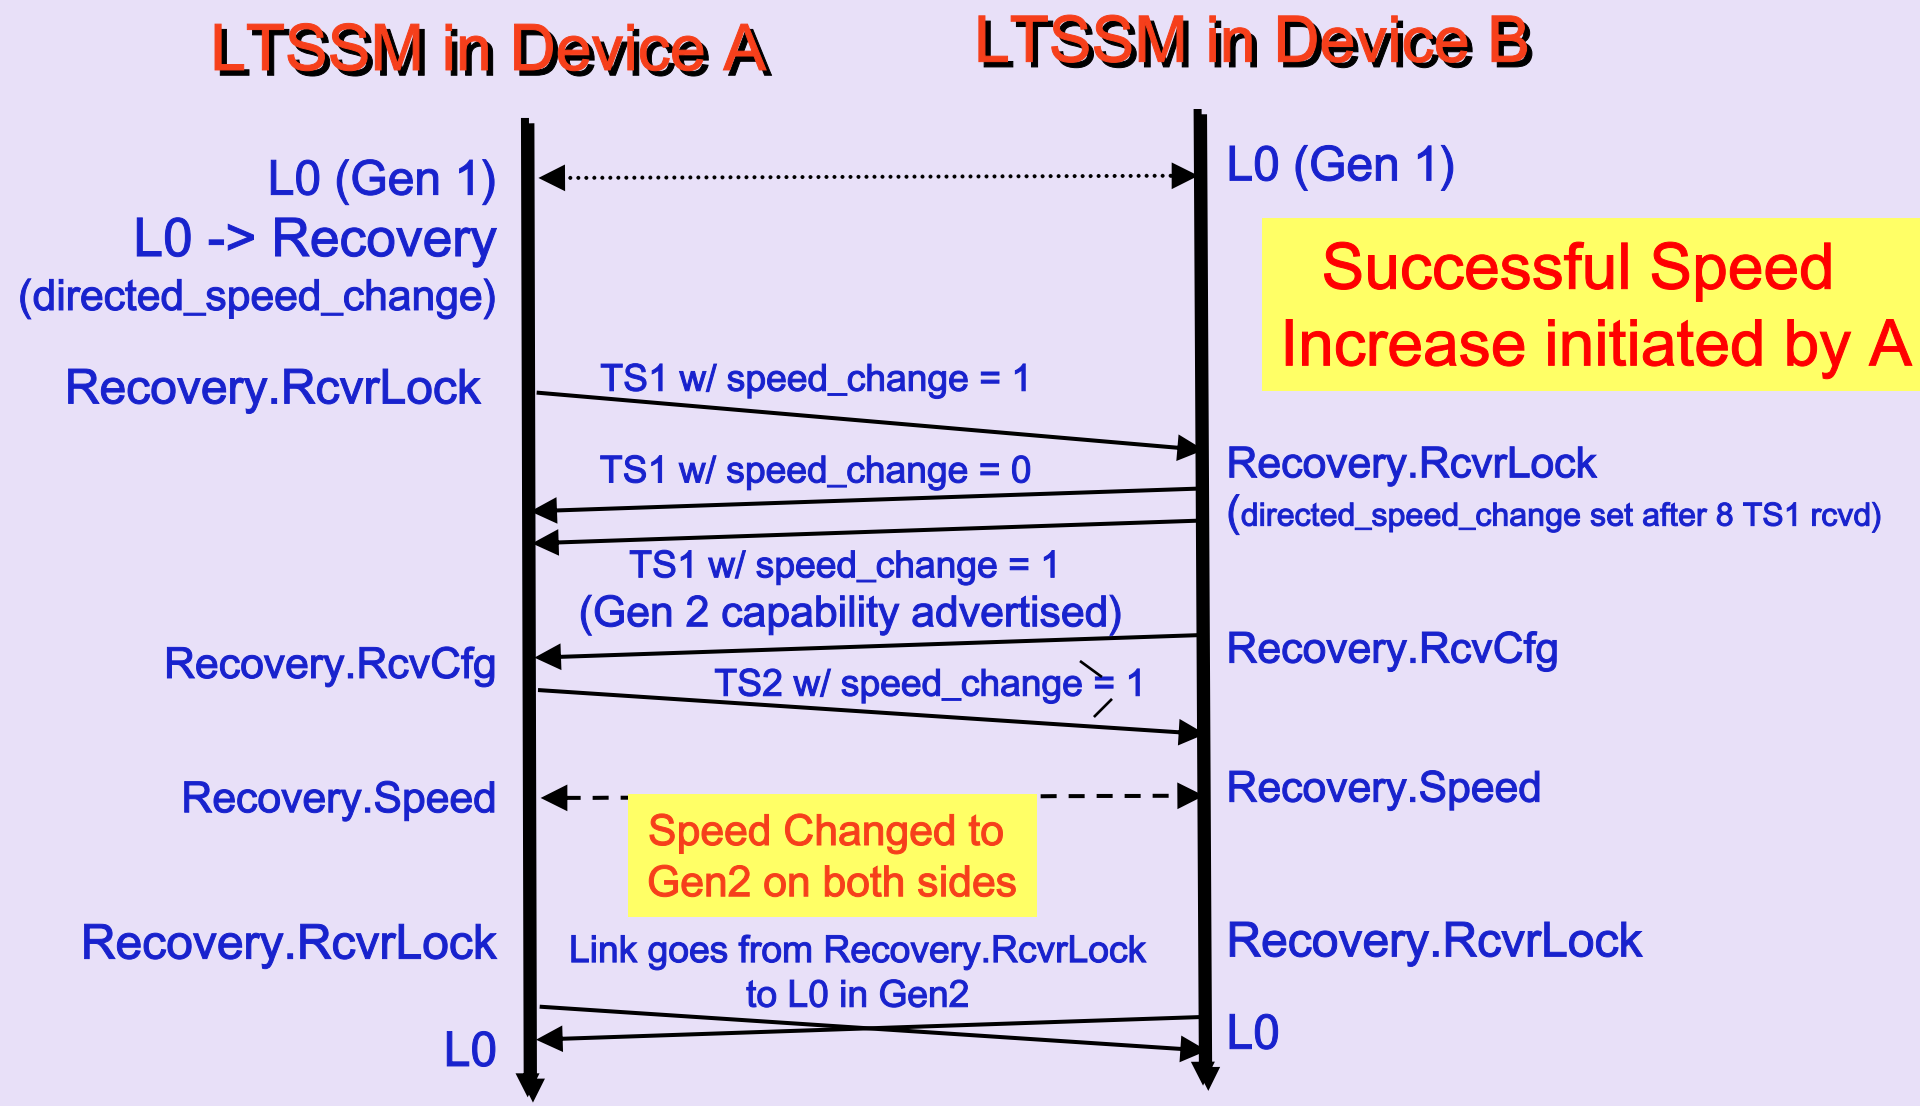

Speed Change Gen 1 -> Gen 2: Example 1

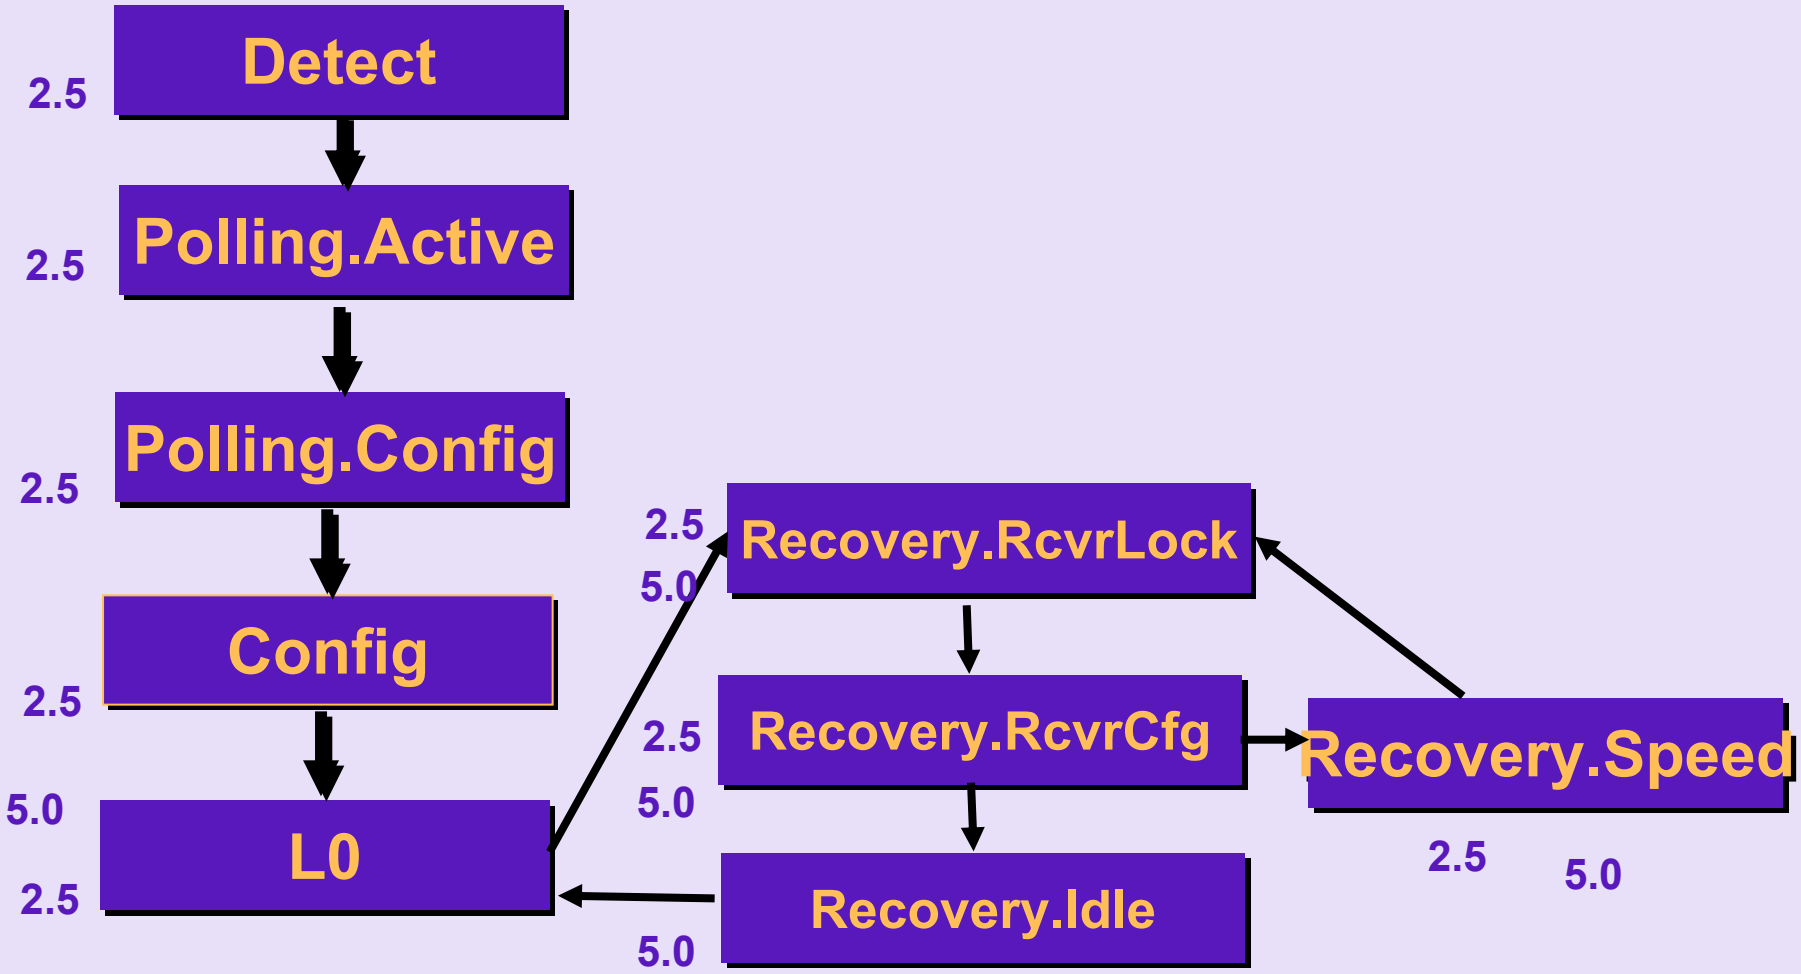

LTSSM Speed Change: Example 2

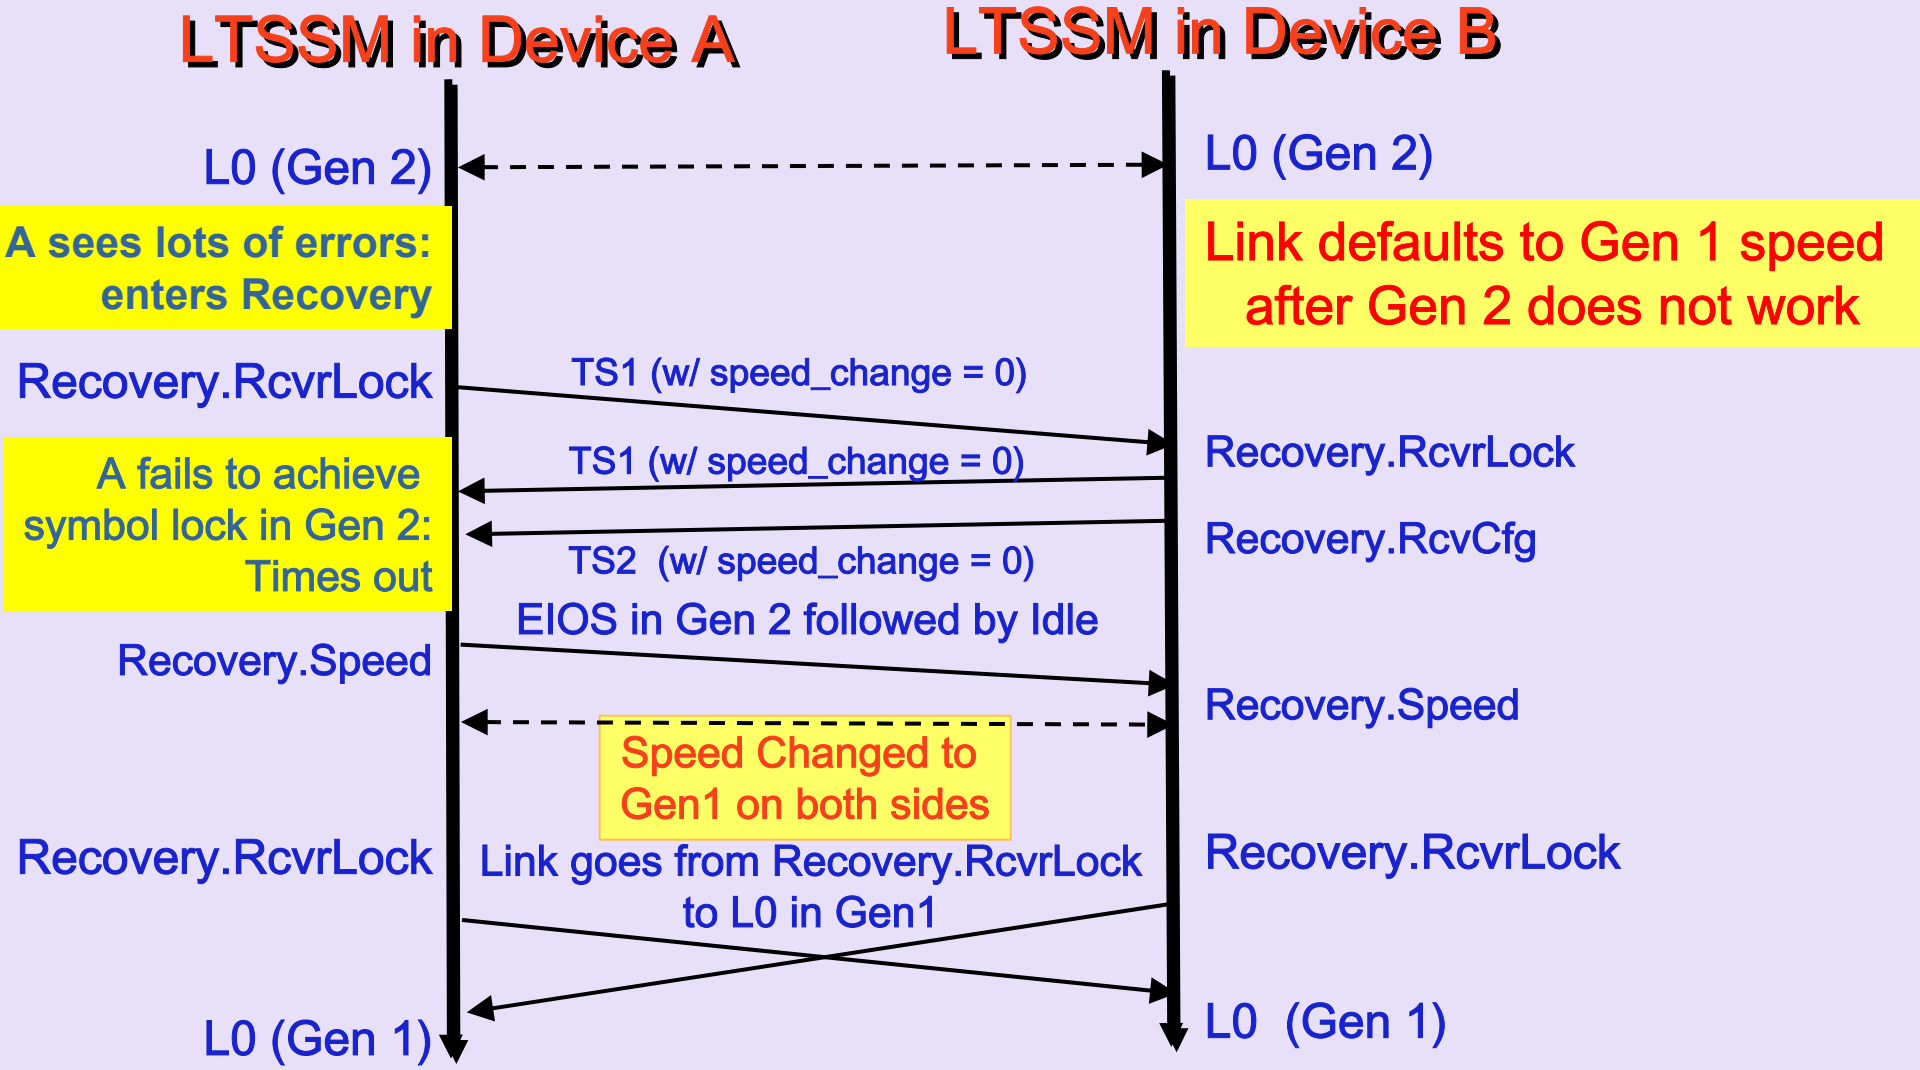

Config Changes for Gen 2

- In Link Capabilities:
  Maximum Link Speed → Supported Link Speeds
  - ✓ Reports component capabilities – encoding added for gen2 speed
- New control field: Target Link Speed
  - ✓ Sets target & upper limit on link operational speed
  - ✓ Also sets speed for software initiated Compliance
- New control bit: Hardware Autonomous Speed Control
  - ✓ Controls ability of Hardware to reduce speed

Config Changes for Gen 2

- New control bit: Enter Compliance
  - ✓ Software initiated Compliance mode entry
- In Link Status:
  Link Speed → Current Link Speed
  - ✓ Reports the current operational link speed when link is up
- New status bit:
  Link Speed Change Attempt Failed
  - ✓ Set to indicate that hardware was forced to reduce target speed during negotiation

Link Speed Controls

- Hardware automatically trains to the greatest link speed possible
  - ✓ Link trains to 2.5GHz in L0 and then attempts to train higher
  - ✓ Hardware and/or software can place an upper bound on the speed
  - ✓ Hardware and/or software can change the speed for power management
- New mechanism for PCIe-aware software to be notified when bandwidth (speed or width) changes
  - ✓ Hardware notification to software when marginal link retrains to a lower bandwidth

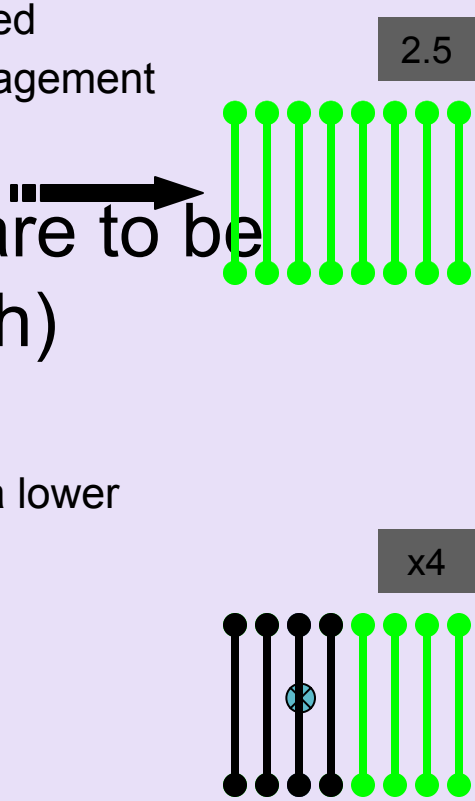

Agenda

- Overview of Logical Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Enhancements for Robustness
- Summary & Call to Action

Gen 2: Speed of Polling.Compliance

- Two ways to set the speed
  - ✓ Inband Method:
    - Gen 2 Speed: Tie n and p. Pull up to a DC voltage. Voltage at receiver to be 500-700 mV
    - Gen 1 Speed: Default. Also can tie n and p and pull down to GND
    - Need a comparator for determining speed
    - Multilane Link: Speed determined by the highest reported speed among the lanes of the link that has not received an exit from electrical idle
  - ✓ CSR method (new method)
    - Write to newly defined CSRs to get to compliance at desired speed
    - Link comes up to L0; CSR written on both sides to enter compliance with same speed; SBRE done in upstream; causes link to go to Detect and eventually to Compliance

Gen 2: Speed of Polling.Compliance

- Speed determined prior to entering Polling.Compliance
- Polling.Active is always Gen 1
  - ✓Transition to/from Polling.Compliance operating in Gen 2 speed:
    - EIOS needs to be transmitted followed by
    - Electrical idle between 1 ms to 2 ms before changing speed.

Speed Determination Circuit

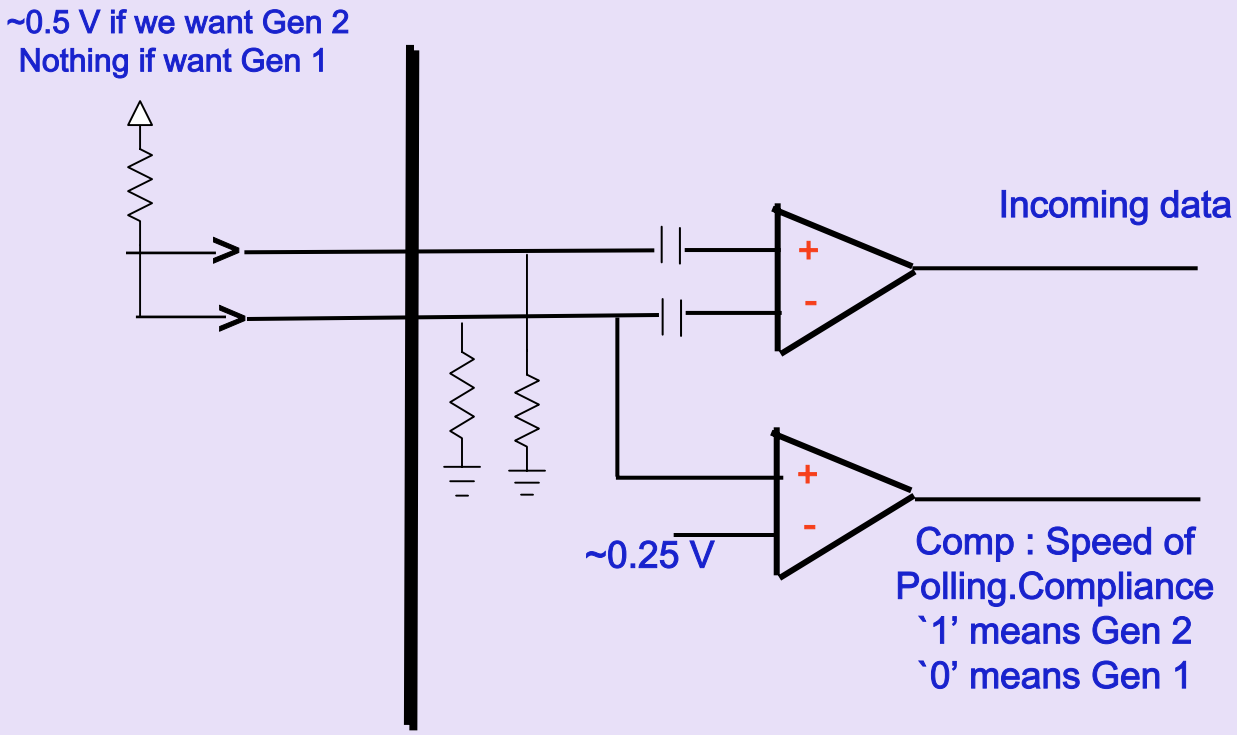

Agenda

- Overview of Logical Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Enhancements for Robustness
- Summary & Call to Action

Gen 2: Electrical Idle Exit Challenges

- Exit from Electrical Idle a challenge in Gen 2
  - ✓ Receiver sensitivity in Gen 2: 120 mV
  - ✓ Idle detection threshold in Gen 2: 175 mV
  - ✓ COM with a run length of 5 does ensure that to cross the idle detection threshold but the frequency of COM is not high enough
  - ✓ Expect exit EI to be detected by circuits in Gen 1 speeds
- Need a low-frequency repeating pattern
  - ✓ Electrical Idle Exit Sequence (EIES) defined
  - ✓ COM followed by 14 K28.7 followed by TS1 identifier
  - ✓ Effectively 5 1's followed by 5 0's 13 times
  - ✓ Requirements relaxed to detect exit EI only on these symbols at *non-Gen 1* speeds
  - ✓ One EIES sent after every 64 TS1/TS2'es.
  - ✓ EIES sent before the first TS1 in a TS1/TS2 sequence.

EIES Sequence

Symbol Number	Encoded Values	Description
0	K28.5	COMMA code group for Symbol alignment
1-14	K28.7	K Symbol with low frequency components for helping achieve exit from electrical idle
15	D10.2	TS1 Identifier

Electrical Idle Exit Usage

- Exit from EI needed in
  - ✓Detect.Active (TS1)
  - ✓Polling.Active (TS1/ TS2)
  - ✓Polling.Compliance (TS1)
  - ✓ Rx_L0s.Idle (FTS/ TS1 on errors)
  - ✓ L1.Idle (TS1)
  - ✓ L2.TransmitWake (TS1)
  - ✓ Disabled (TS1)
  - ✓Recovery.Speed (Gen 2) for inferring EI (TS1/ TS2)

Gen 2: Changes to ensure EI Exit

- In all states except Rx.L0s, the LTSSM gets the initial EIES to detect an exit EI
- EIES repeated after every 64 TS1/TS2 for robustness
  - ✓ Also helps in inferring electrical idle condition in Recovery for the failure cases
  - ✓ Lack of exit EI in a 16000 UI interval can be inferred as EI
- TS1 ordered set has enough high frequency components to achieve bit lock

Gen 2: Changes to ensure EI Exit

- FTS changes for non-Gen1 speeds
  - ✓Prepend the FTS sequence with 4 K 28.7
  - ✓4 sets of K28.7 helps detect exit from EI
  - ✓FTS used to achieve symbol lock
  - ✓If circuits did not exit EI with the prepended sequence (e.g., missed K28.7)
    - FTS has enough low frequency components to cause exit EI
    - Designs can optimize (N_FTS) at non-Gen1 frequency if they can exit EI with 4 K28.7

Gen 2: Electrical Idle Entry Challenges

- EI entry detection difficult. Required in:
  - ✓ L0 (for surprise detach)
  - ✓ Loopback.Active (as slave to know when master terminates)
  - ✓ Recovery.Speed (in Gen 2)
- Gen 2 Spec allows to *infer* EI as an alternative to detecting EI in all speeds
  - ✓ L0: EI may be inferred if receiver did not receive COM in 128us. SOS has COM
  - ✓ Recovery.RcvrLock and Recovery.RcvrCfg: absence of COM in 1280 UI interval
  - ✓ Recovery.Speed: absence of exit from EI in 16000 UI interval
  - ✓ Loopback.Active: Under discussion

Agenda

- Overview of Logical Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Enhancements for Robustness
- Summary & Call to Action

Enhancements for Robustness

- Changes to Polling.Compliance for HA
  - ✓ Entry condition relaxed from any lane that detected a receiver but did not get an exit from Electrical Idle to multiple lanes
  - ✓ Must go to Polling.Compliance if all lanes that detected a receiver did not get an exit from electrical idle
- Electrical Idle Ordered Set extension for non-Gen 1 speeds:
  - ✓ Two consecutive sets of COM, IDL, IDL, IDL
- Electrical Idle detection (optional):
  - ✓ Received signals switching at a frequency greater than 125MHz

Agenda

- Overview of Logical Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Enhancements for Robustness
- Summary & Call to Action

Summary & Call to Action

- Track Gen 2 development work to be ready for future speed increase
  - ✓ 1.1 Components not directly affected
- Take advantage of the flexibility offered by LTSSM speed change to optimize for power and HA.
- Take advantage of the flexibility offered by compliance : inband as well as CSR mechanisms
- Take advantage of a simple electrical idle detection circuitry with HVM and low power advantages by adopting the protocol changes associated with electrical idle .
- Make robust designs by adopting HA related enhancements.

Thank you for attending the PCIe Technology Seminar 2005.

For more information please go to www.pcisig.com

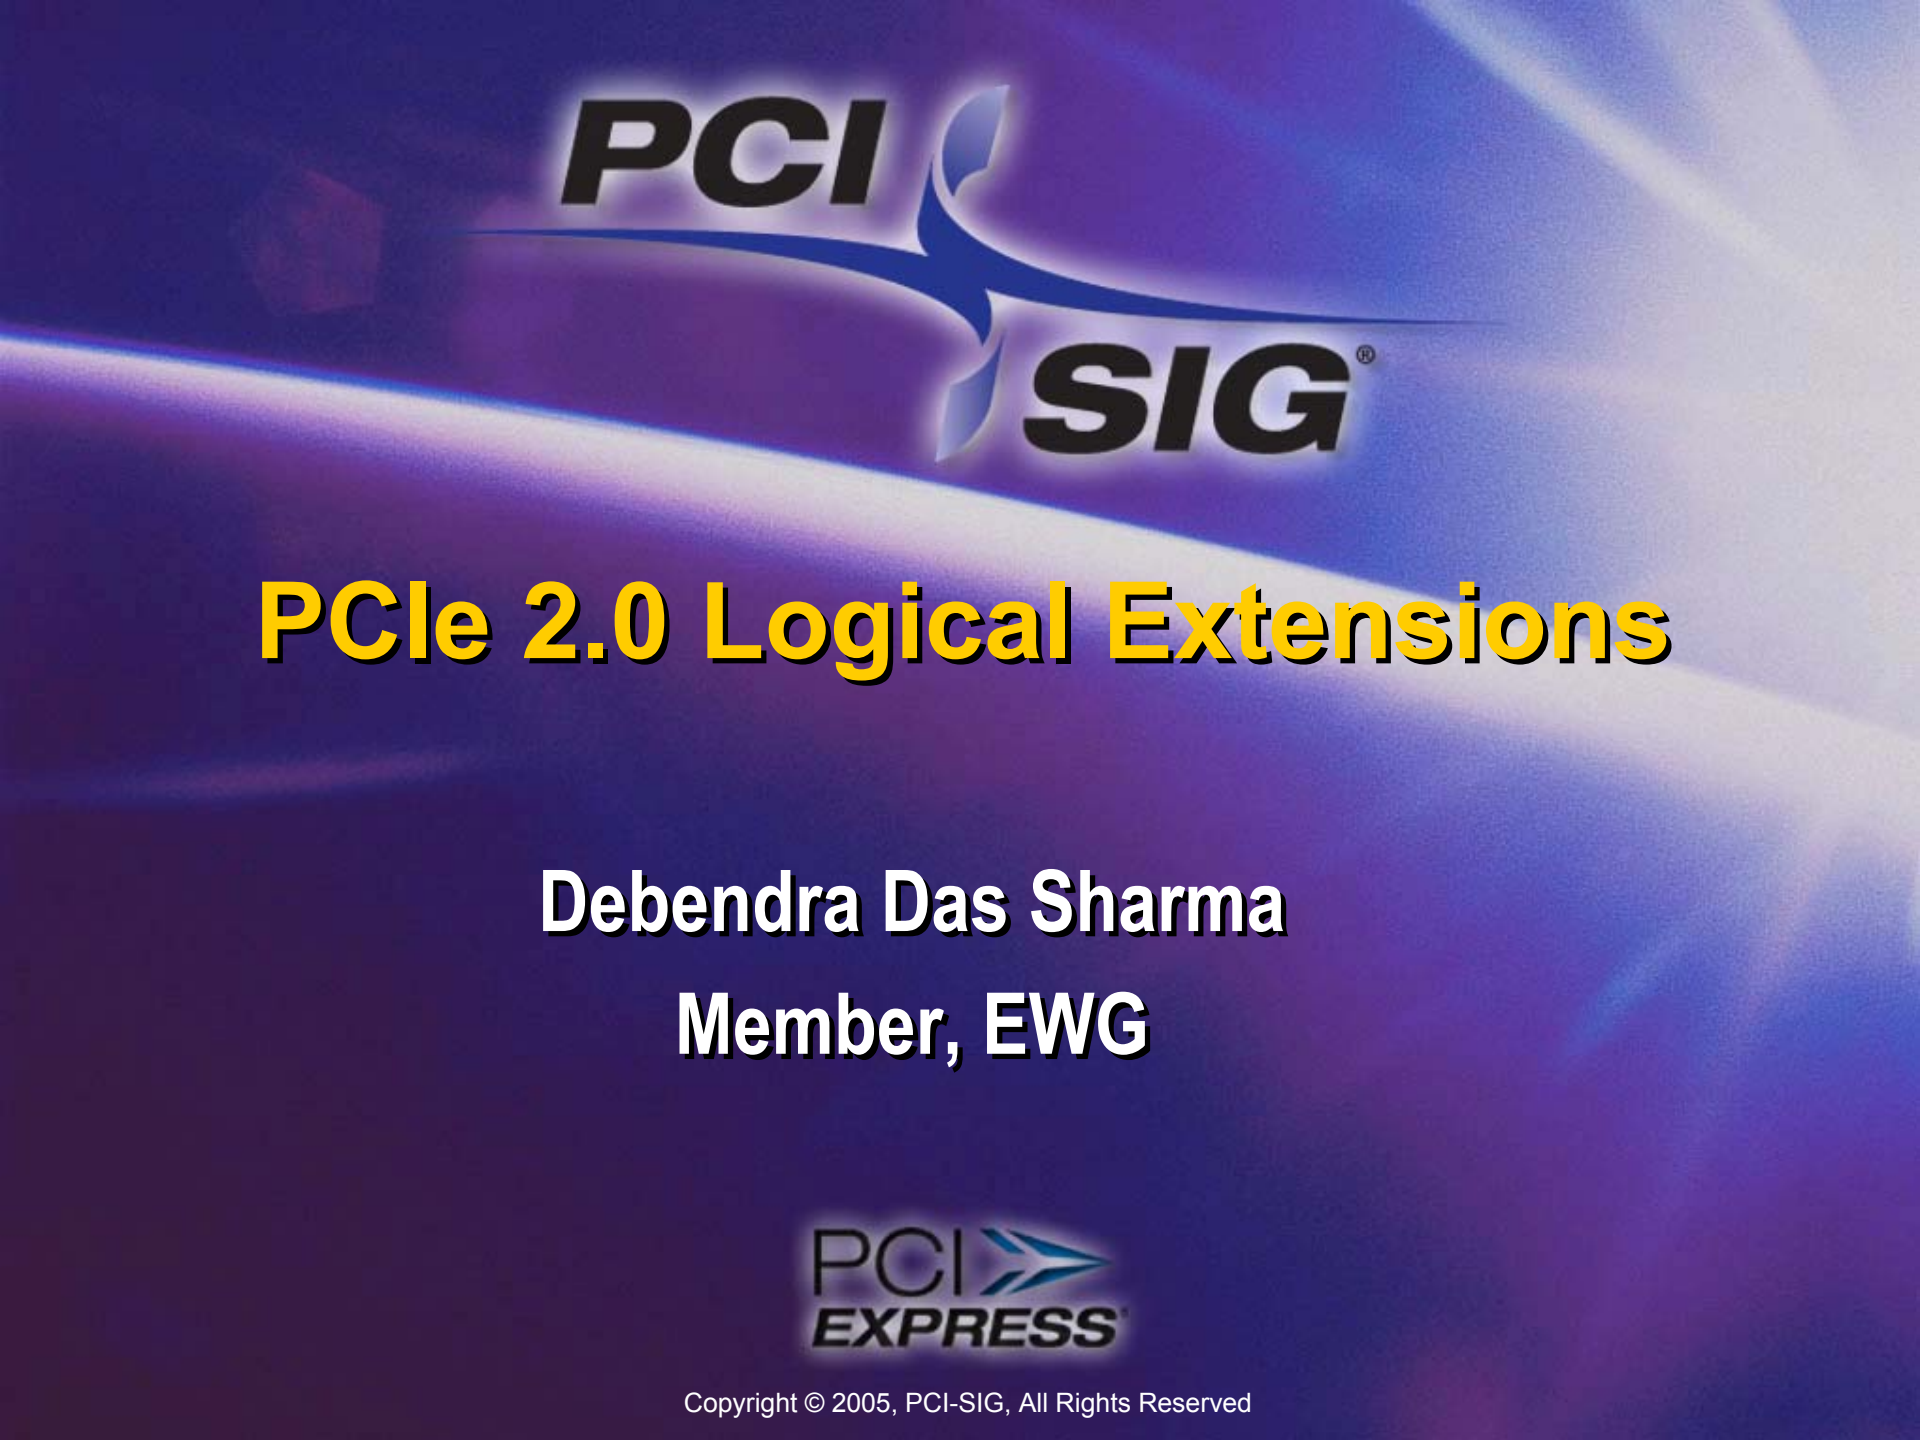

PCIe 2.0 Logical Extensions

Debendra Das Sharma

Member, EWG


PCI
SIG®