

PCI

A stylized graphic element consisting of a blue swoosh that curves from the bottom left, loops upwards and to the right, and then curves back down to the right, passing between the 'PCI' and 'SIG' text.

SIG[®]



PCIe 2.0 Logical PHY Architecture

Debendra Das Sharma

Member, EWG



Agenda

- Overview of Logical Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Link Width Upconfigure
- Testability Enhancements
- Summary & Call to Action

Logical Extensions: Overview

- Changes limited to physical layer only
- Backwards compatible with PCIe 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)
Link width Up-configure	Dynamic Link width change	Power savings
Testability Enhancements	Receiver Compliance Transmitter Margining Loopback Enhancements	Cost effective way to test the transmitter and receiver Assumption based loopback for test equipment.

Agenda

- Overview of Architectural Extensions
- **LTSSM Speed Negotiation**
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Link Upconfigure Capability
- Testability Enhancements
- Summary & Call to Action

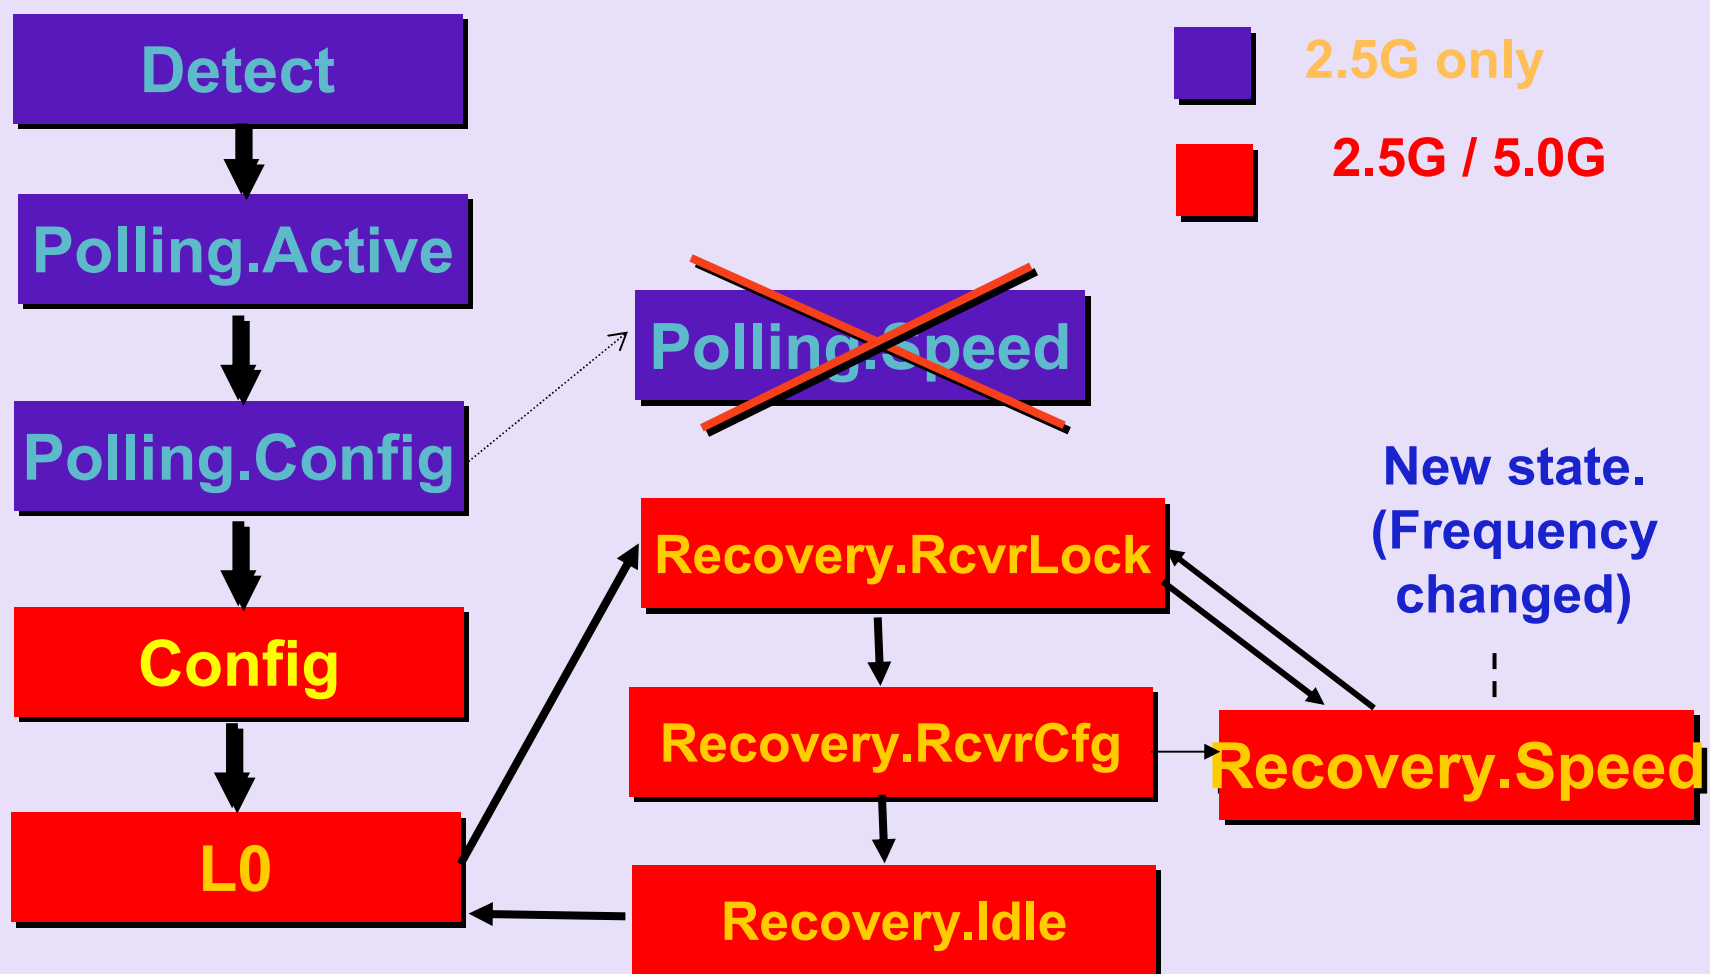
Speed Negotiation: Philosophy

- Dynamic link speed change
 - Link stays up
 - ✓ Reliability
 - 2.5 GT/s data rate as fall back speed
 - Revert back data rate if speed change unsuccessful
 - ✓ Power improvements
 - Lower frequency when demand drops
- Notify software on bandwidth change
 - ✓ Autonomous
 - ✓ Reliability / Software induced

Speed Negotiation

- Initially link trains to L0 in 2.5G data rate
- Speeds advertised in TS ordered sets
 - ✓ Supported speeds by the other component noted in Config.Complete and Recovery.RcvrCfg
- Speed change handshake in Recovery
 - ✓ New substate: **Recovery.Speed**
 - ✓ Speed changed in Recovery.Speed
- **Polling.Speed** state obsolete

Relevant LTSSM States





Training Sequence (TS1/TS2 changes)

Symbol Number	Description
4	<p>Data Rate Identifier</p> <p>Bit 0 – Reserved, set to 0</p> <p>Bit 1 = 1, 2.5 GT/s data rate supported</p> <p><u>Bit 2 = 1, 5 GT/s data rate supported.</u></p> <p>Bit 3:<u>5</u> – Reserved <u>for future data rates past 5GT/s</u>, set to 0 <u>for devices that only support 2.5 GT/s and/or 5 GT/s data rates.</u></p> <p><u>Bit 6: 3 Purposes (direction/state dependent):</u></p> <ol style="list-style-type: none"><u>1. De-emphasis</u><u>2. B/W notification (upstream)</u><u>3. Upconfigure Capability</u> <p><u>Bit 7: Speed change Request</u></p> <p><u>All lanes must transmit the same value</u></p>



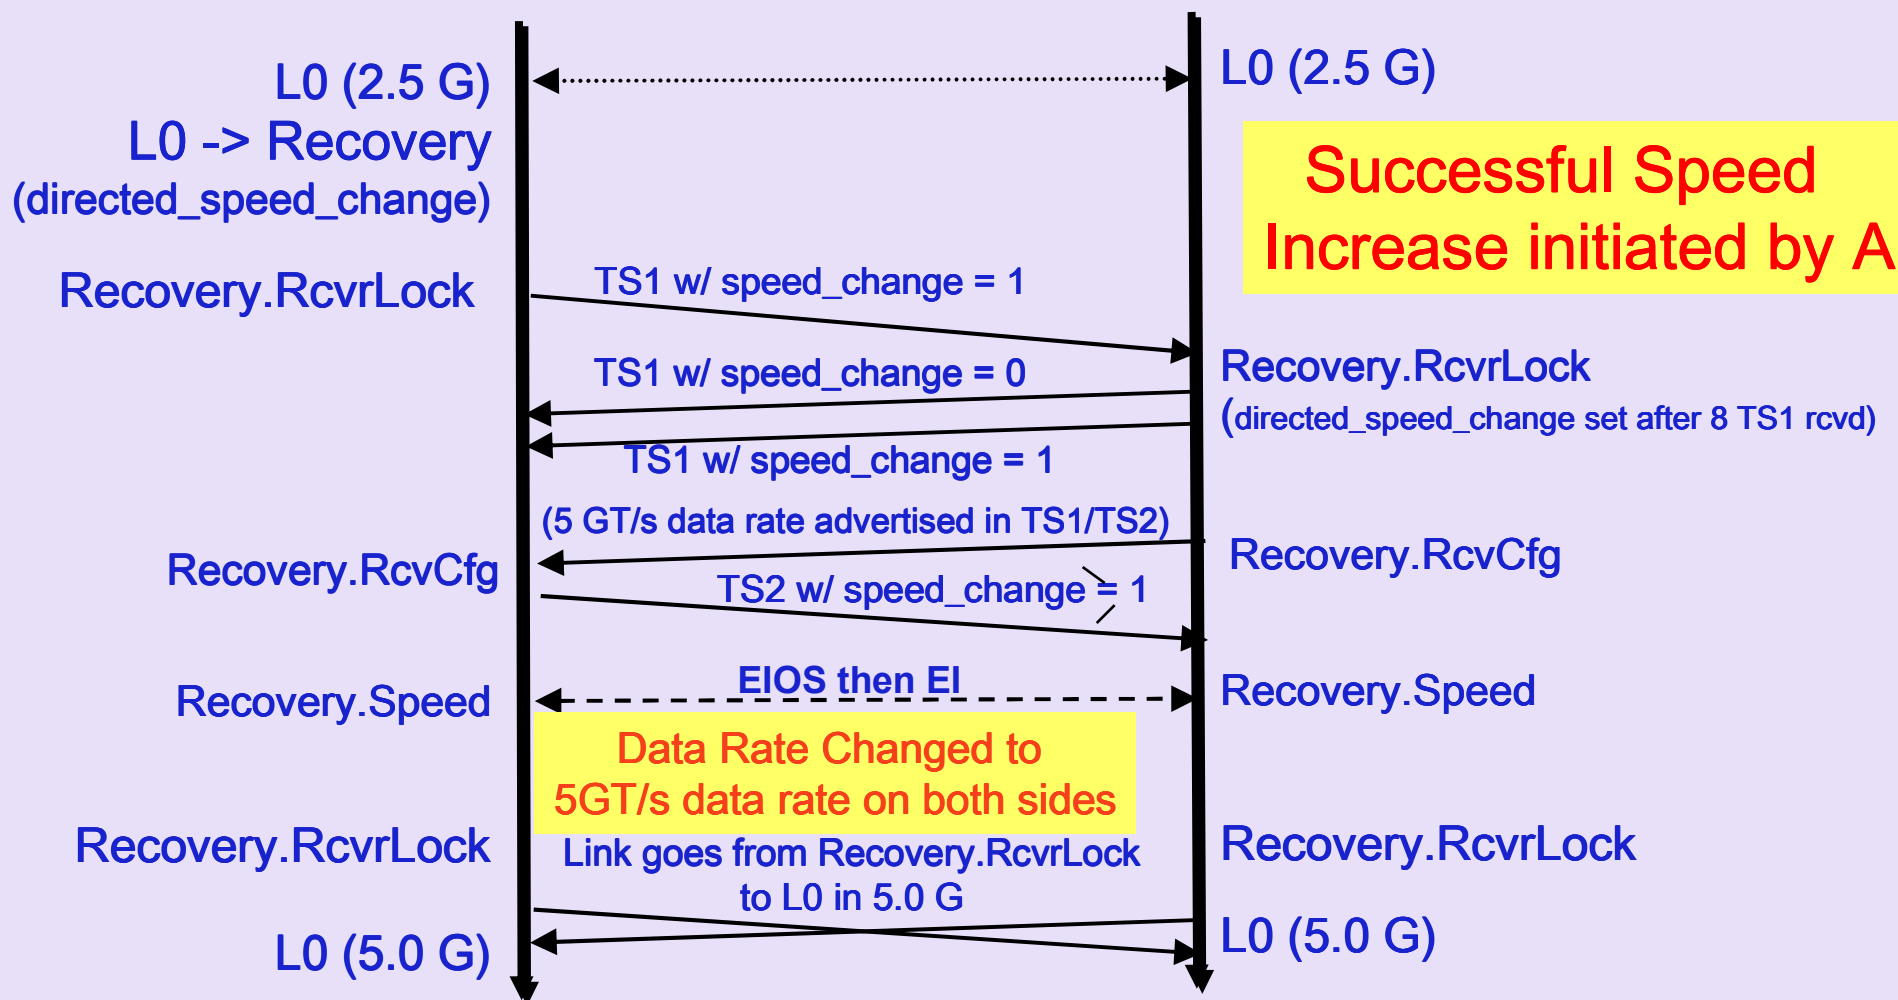
LTSSM Speed Change: Example 1



LTSSM Speed Negotiation

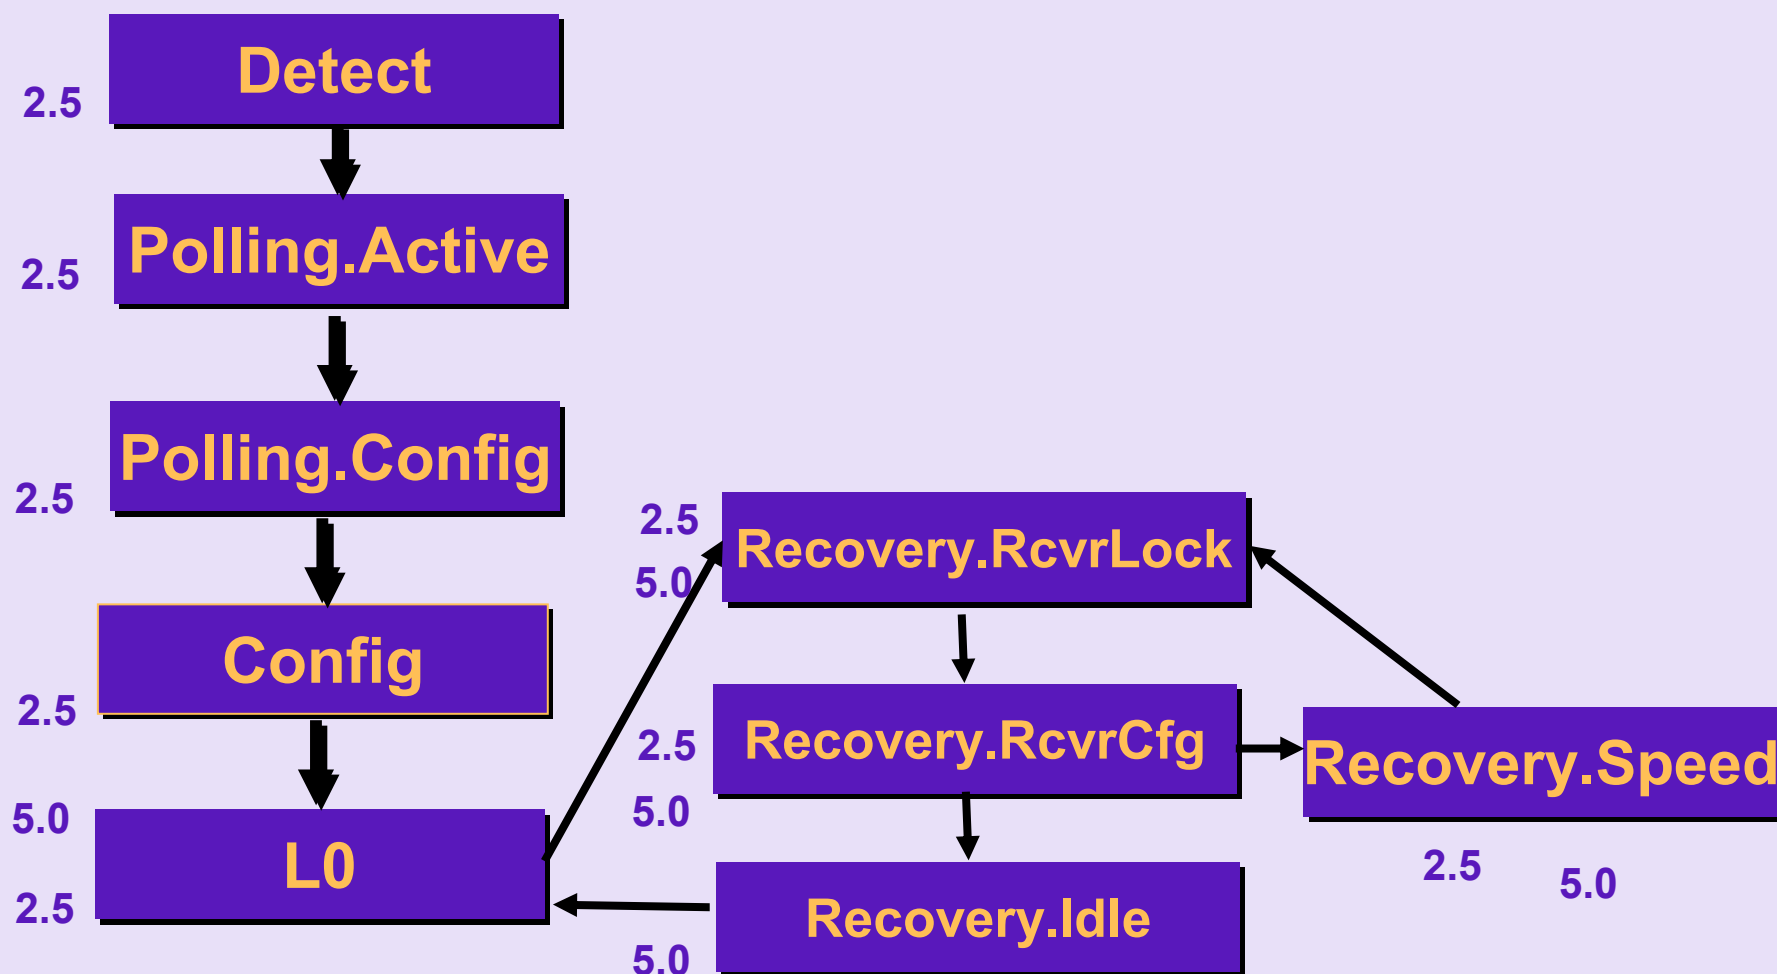
LTSSM in Device A

LTSSM in Device B



Speed Change : Example 1

LTSSM Speed Negotiation





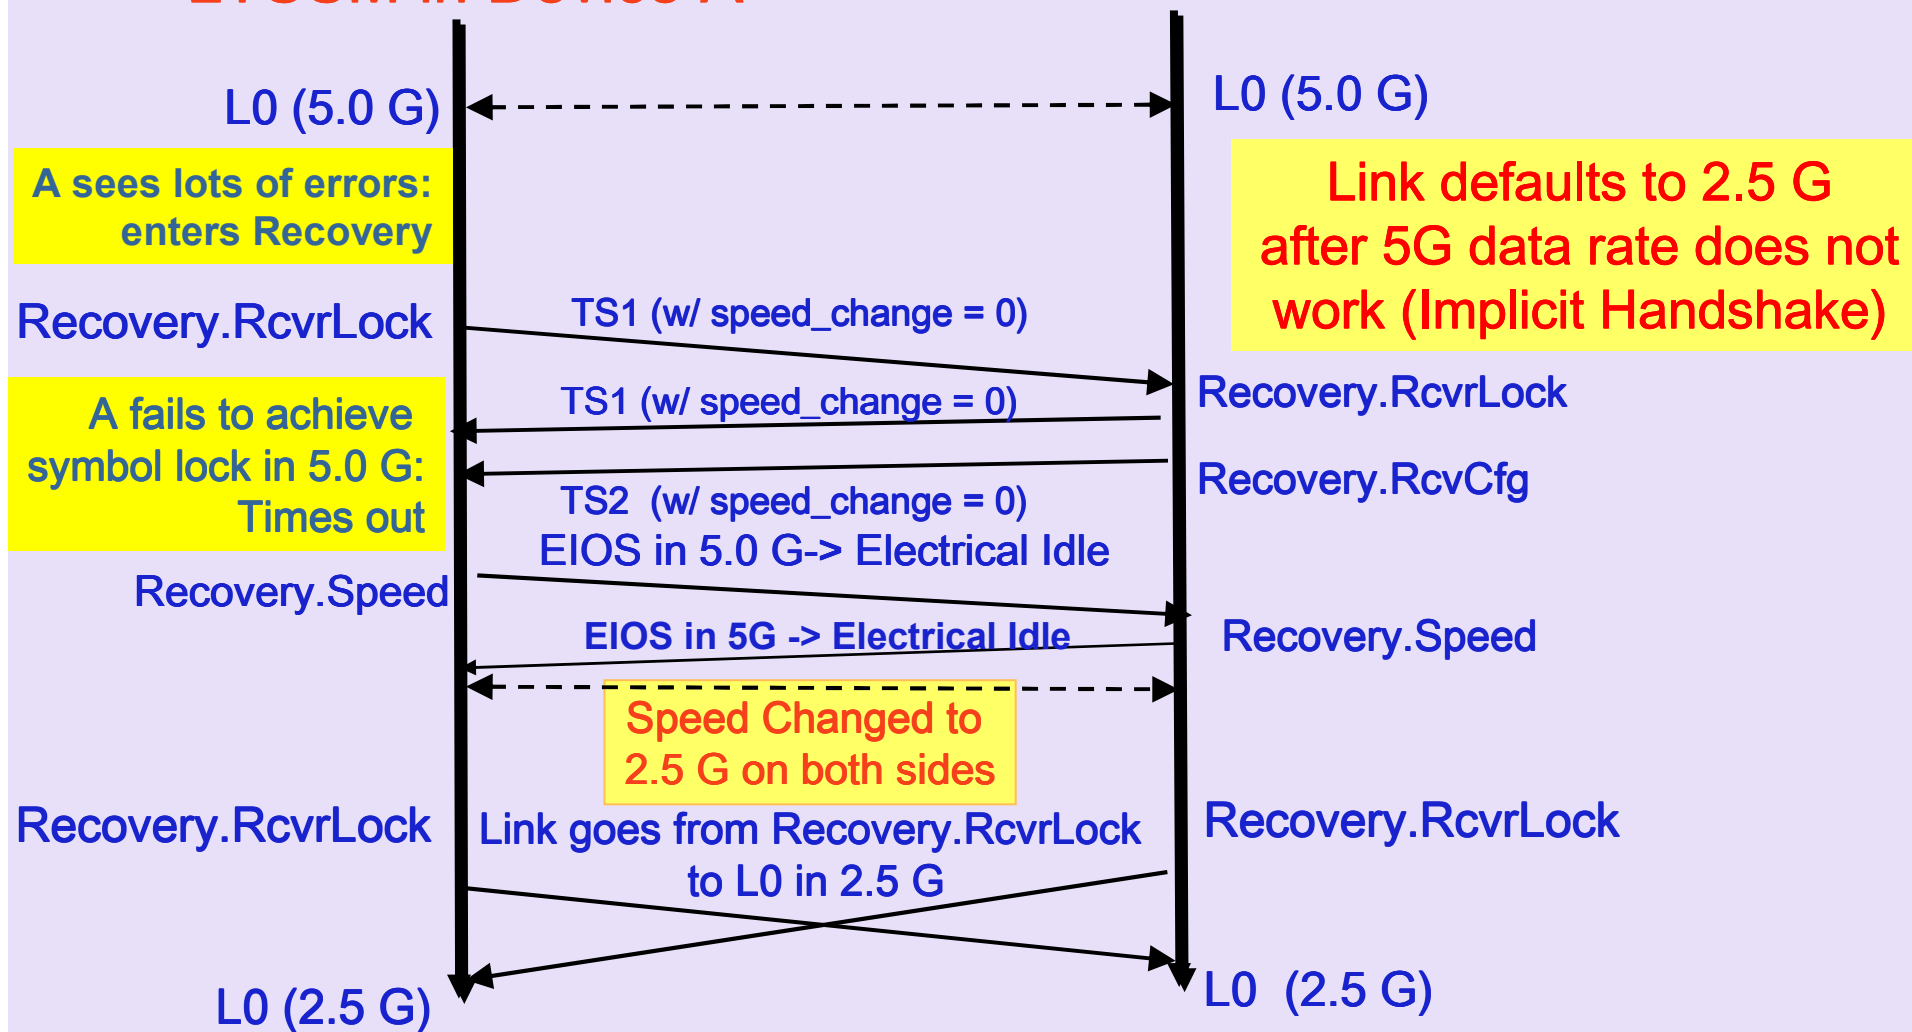
LTSSM Speed Change: Example 2



LTSSM Speed Negotiation

LTSSM in Device A

LTSSM in Device B



Selectable De-emphasis

- Pattern sensitivity at 5GT/s
 - ✓ 6.0 dB enables longer loss dominated channels
 - ✓ 3.5 dB enables shorter reflection/crosstalk dominated channels
- Upstream component sets the de-emphasis level for the link during speed negotiation in Recovery.RcvrCfg (bit 6, symbol 4 of TS2)
- Controlled by a HWInit CSR bit in Link Control 2 register in upstream component.
 - ✓ Implementation specific overrides

Config Changes

- Added encoding for 5G speed (Link Capabilities)
- New control field: Target Link Speed
 - ✓ Sets target & upper limit on link operational speed
 - ✓ Set only in upstream component
 - ✓ Speed change occurs by setting link retrain bit
 - ✓ Also sets speed for software initiated Compliance
- Hardware Autonomous Speed Control
 - ✓ Controls ability of Hardware to reduce speed dynamically
 - ✓ Does not affect speed reduced for reliability reasons
- Link Bandwidth change notification mechanism
 - ✓ Link B/W Management Status
 - ✓ Link Autonomous B/W Status
- Software initiated Compliance
 - ✓ Enter Compliance and Enter Modified Compliance

Agenda

- Overview of Architectural Extensions
- LTSSM Speed Negotiation
- **Compliance Speed Determination**
- Electrical Idle Entry and Exit
- Link Upconfigure Capability
- Testability Enhancements
- Summary & Call to Action

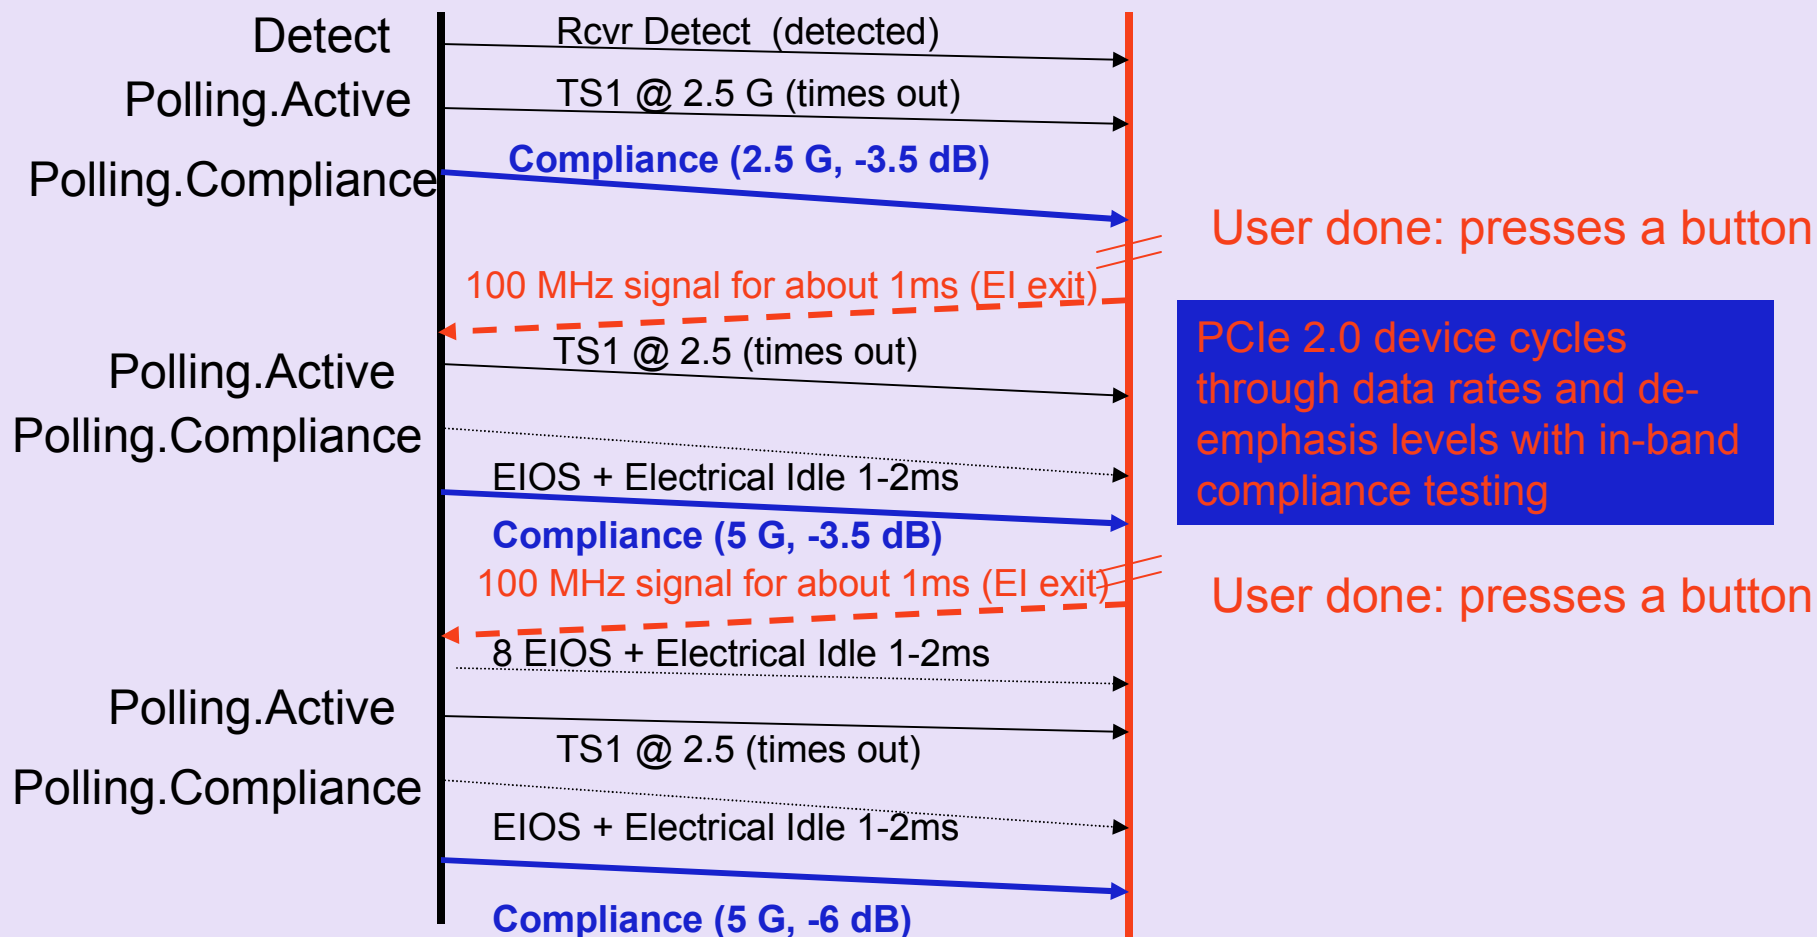
Speed of Polling.Compliance

- Two ways to set the speed
 - ✓ Inband Method: Compliance Load Board (Modified)
 - ✓ CSR method (new method)
 - Link Control2 Register Fields
 - Target Link Speed, Enter Compliance, Enter Modified Compliance, compliance de-emphasis, and Transmit Margin
 - Procedure:
 - Link trains to L0
 - CSR bits written on both sides to enter compliance
 - SBRE upstream.
 - LTSSM:Hot Reset -> Detect -> Polling.Active -> Polling.Compliance
- Speed, de-emphasis, transmit margin determined prior to entering Polling.Compliance

Polling.Compliance: Inband

PCIe 2.0 Device

Compliance Load Board



Agenda

- Overview of Architectural Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- **Electrical Idle Entry and Exit**
- Link Upconfigure Capability
- Testability Enhancements
- Summary & Call to Action

Electrical Idle Exit Challenges

- Exit from Electrical Idle a challenge in 5.0 G speed
 - ✓ Receiver sensitivity in 5G speed: 120 mV
 - ✓ Idle detection threshold in 5G speed: 175 mV
- Expect exit EI to be detected by circuits in 2.5 G speed
- Need a low-frequency recurring pattern for 5.0 G:
 - ✓ Electrical Idle Exit Sequence (EIES) defined
 - Effectively 5 1's followed by 5 0's 13 times => low frequency
 - ✓ Requirements relaxed to detect exit EI in *non-2.5 G* speeds
 - ✓ One EIES sent after every 32 TS1/TS2'es
 - ✓ Sent in Recovery and Config.Linkwidth.Start
 - ✓ Mix of low frequency (EIES) for EI exit detection and high frequency (TS1) for symbol lock.
- L0s Exit in 5GT/s speed
 - ✓ Send 4 to 8 K28.7 symbols to help with EI exit detection



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 Entry

- EI entry detection difficult. Required in:
 - ✓ L0, Loopback.Active, Polling.Compliance, and Recovery
- Have the logical layer do the *inference* to alleviate circuit design
- New *infer* EI as an alternative to detecting EI
 - ✓ L0: no SKP Ordered Sets in 128us.
 - ✓ Recovery.RcvrCfg and Recovery.Speed on successful negotiation:
 - no TS ordered sets in any configured lane in 1280 UI interval
 - ✓ Recovery.Speed with unsuccessful negotiation:
 - absence of exit from EI
 - 16000 UI in 5GT/s
 - 2000 UI in 2.5 GT/s
 - ✓ Polling.Compliance: No COM in 128 us window
 - ✓ Loopback Active
 - 2.5 GT/s: no exit EI in 128 usec
 - 5 GT/s: rely on EIOS only

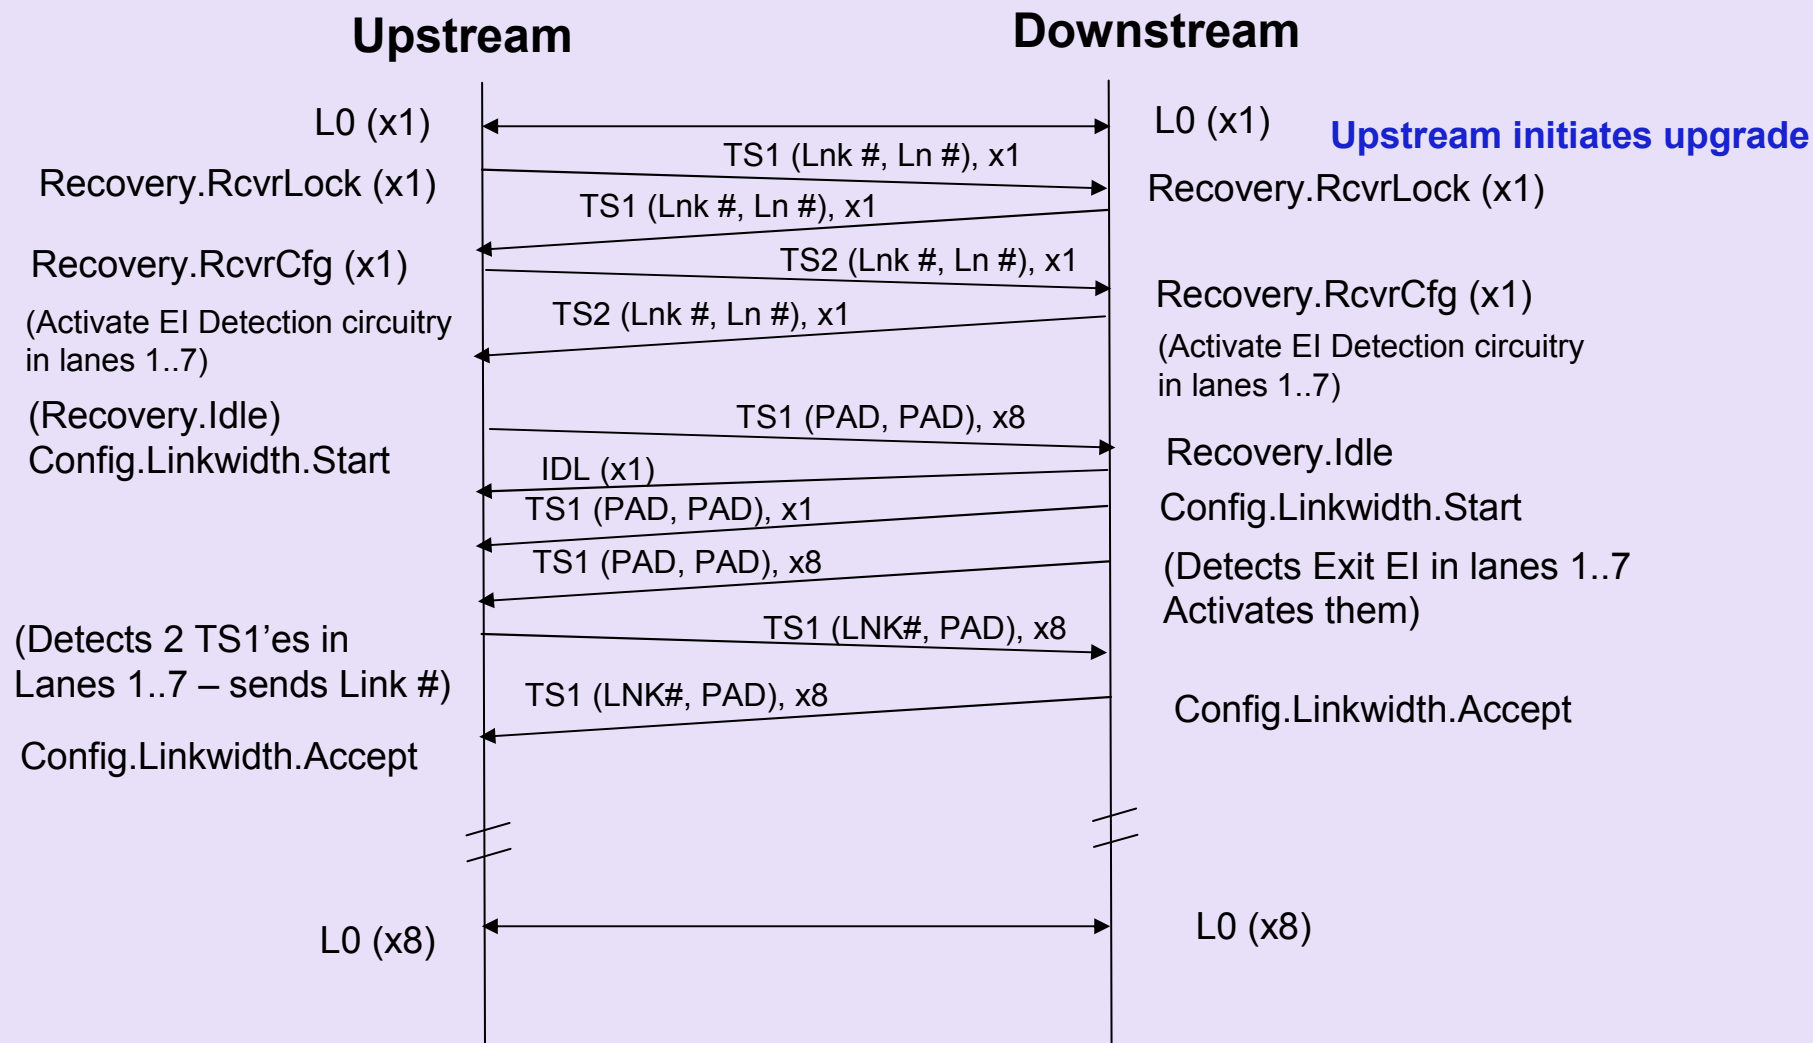
Agenda

- Overview of Architectural Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- **Link Upconfigure Capability**
- Testability Enhancements
- Summary & Call to Action

Link Width Upconfigure

- New capability to upconfigure link width
- Power savings during low demand
- Up to the width initially negotiated
 - ✓ E.g.: x8 -> x4 -> x1 -> x2 -> x8
- Capability advertised in the TS2 ordered sets in Config.Complete state (Symbol 4, Bit 6 of TS2)
- Link width upconfigure and planned downconfigure:
 - ✓ L0 -> Recovery.RcvrLock -> Recovery.RcvrCfg -> Recovery.Idle -> **Config** -> L0
- Link width is changed in Config State
 - ✓ Initiator of upgrade delays sending link number till the other side wakes up on those lanes and sends TS ordered sets

Link Upgrade Example: x1->x8



Agenda

- Overview of Architectural Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Link Upconfigure Capability
- **Testability Enhancements**
- Summary & Call to Action

Testability Enhancements

- Receiver Compliance:
 - ✓ A new mode in Polling.Compliance
 - ✓ Receiver sends out modified compliance patterns; includes receiver error count
 - ✓ Intended to margin receiver
 - ✓ Two mechanisms: inband and CSR
- Transmitter Margining:
 - ✓ Multiple voltage levels for Tx; CSR setting
- Loopback:
 - ✓ Speed change (deemphasis) in Config -> LB
 - assumption-based training
 - ✓ Ability to loopback even without achieving lock

Receiver Compliance: Inband

PCIe 2.0 Device

Compliance Test Fixture

Detect
Polling.Active

Rcvr Detect (detected)

TS1 @ 2.5 G

TS1 @ 2.5 G (rcvr compliance = 1,
2.5 GT/s and 5 GT/s rates advertised,
Selectable de-emphasis value set)

EIOS + Electrical Idle 1-2ms

Modified Compliance Patterns (5GT/s)

Electrical Idle – Speed changed to 5GT/s

Modified Compliance Patterns (5GT/s)

Modified Compliance Patterns (5GT/s)

Test Patterns (5GT/s)

Modified Compliance Patterns (5GT/s)

Example of Receiver
Testing at 5GT/s using
Assumption based
Training

TS1'es sent for about 36 ms

(Modified compliance sent ~ 2ms)

(Send test patterns.
Margin Transmitted patterns as needed)

(Measure error count advertised)

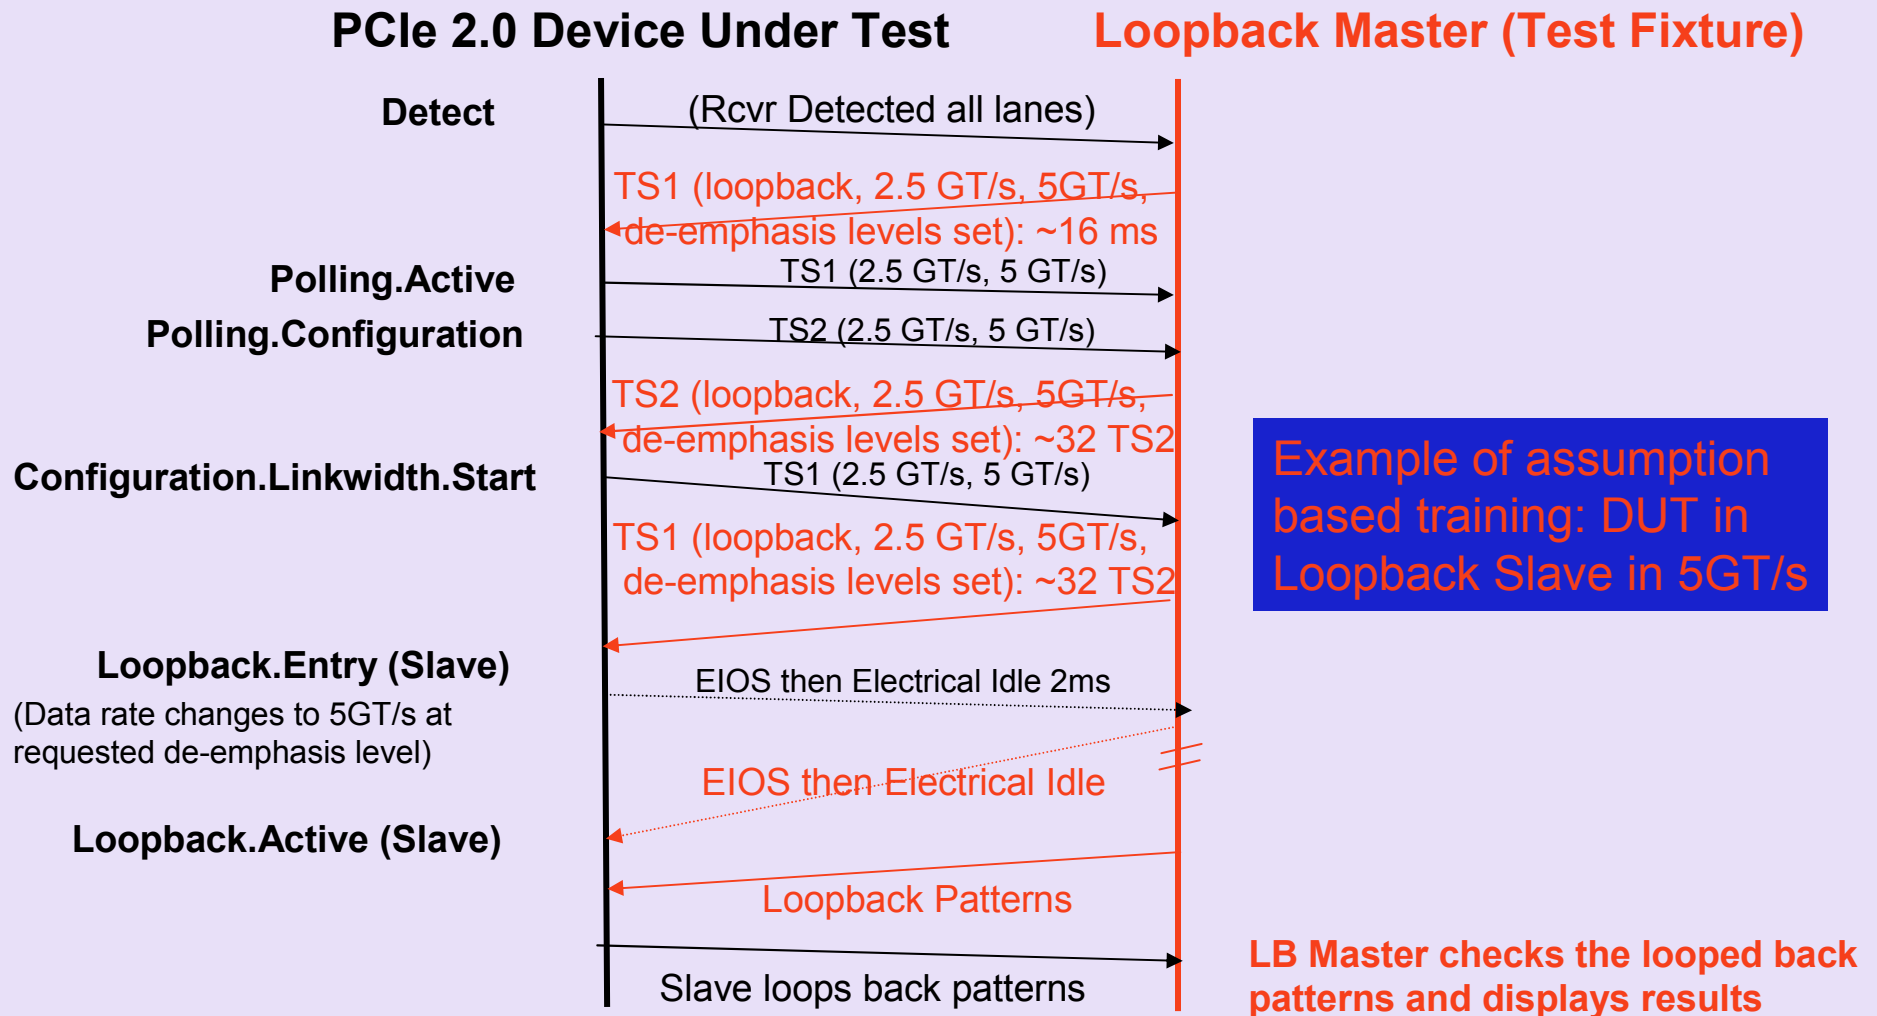
Polling.Compliance

(Data rate changed to
5GT/s at requested
de-emphasis levels)

(Expected to lock within 1ms
after the first modified
compliance pattern received)

Receiver Compliance enables margining the receiver in a cost-effective way

Loopback Testing at 5GT/s



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 >2.5 G 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 Architectural Extensions
- LTSSM Speed Negotiation
- Compliance Speed Determination
- Electrical Idle Entry and Exit
- Enhancements for Robustness
- Link Upconfigure Capability
- **Summary & Call to Action**

Summary & Call to Action

- Track PCIe 2.0 development work to be ready for PCIe 2.0 components
 - ✓ Rev 0.9 released for membership review
 - ✓ PCIe 1.1 Components not directly affected
- Take advantage of the flexibility offered by LTSSM speed change and link upconfigure capability 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
PCI-SIG PCIe Technology Seminar.

For more information please go to
www.pcisig.com



PCIe 2.0 Logical PHY Architecture

Debendra Das Sharma
Member, EWG



PCI

A stylized graphic element consisting of a blue swoosh that curves from the bottom left, loops upwards and to the right, and then curves back down to the right, passing between the 'PCI' and 'SIG' text.

SIG[®]