M-PCIe Overview

Mahesh Wagh
Co-chair PCIe over M-PHY WG
Principal Engineer, Intel

M-PHY

M-PHY – a mobile-focused physical layer specification from the MIPI Alliance.

About MIPI Alliance

MIPI Alliance is a global, collaborative organization comprised of companies that span the mobile ecosystem and are committed to defining and promoting interface specifications for mobile devices. MIPI Specifications establish standards for hardware and software interfaces which drive new technology and enable faster deployment of new features and services. For more information, visit www.mipi.org.

M-PCIe (PCIe Over M-PHY)

- M-PCIe – an ECN that maps PCIe over M-PHY v2.0
  - ✓ Similar approach to the USB-IF SSIC Specification, which maps USB 3.0 over M-PHY
  - ✓ Extends PCIe to the mobile space and others where aggressive power management is required
  - ✓ Preserves the higher PCIe layers intact
  - ✓ Does not cover form-factor specific technology
  - ✓ Does not develop or enhance the M-PHY specification
  - ✓ Does not obsolete the existing PCIe PHY layer
- Implementers of this ECN must be members of both PCI-SIG and MIPI Alliance.
  - ✓ Licensing rights & track specification evolutions.

Key Features of M-PCIe

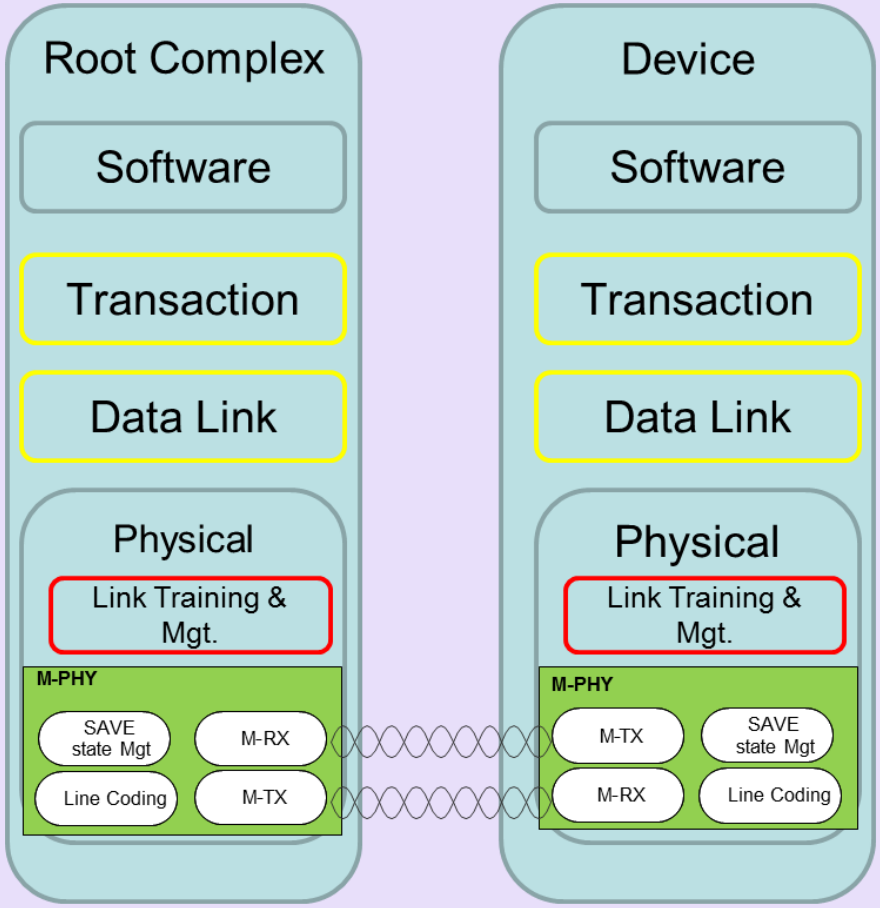

Source: Intel Corporation

- Maintains compatibility with PCIe programming models
- Supports multi-lane configurations as defined in PCIe
- Supports PCIe protocol
- Preserve PCI/PCIe stack with minor modifications to Transaction/Link Layer
- Revised Link Training & Management
- Supports asymmetric link width configurations
- Supports dynamic bandwidth scalability
- Optimized for RFI/EMI
- Enables short channel circuit optimizations
- Supports all M-PHY high-speed gears
- Supports M-PHY TYPE I MODULE only
- Supports M-PHY LS (Low Speed) gear for M-PHY parameter initialization
- Supports 8b/10b for data encoding
- Supports shared and independent reference clocks
- Reference Clock – 19.2 or 26 MHz
- Sideband signals not modified
  - ✓ PERST#, CLKREQ#, WAKE#

M-PCIe LINK

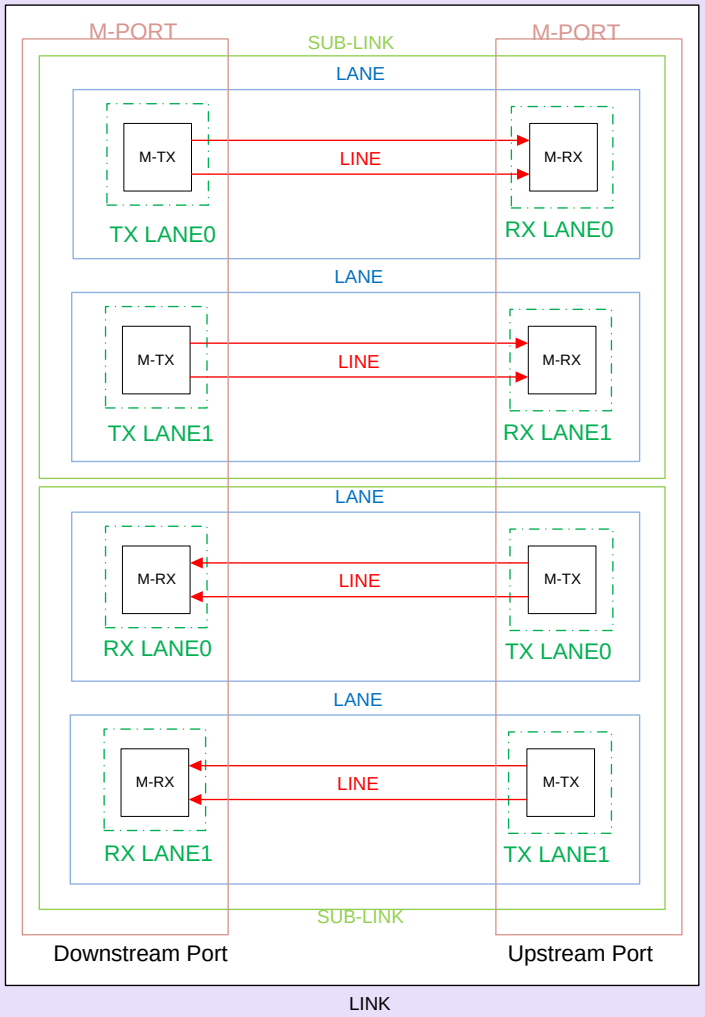

- LINK
  - ✓ A LINK shall have 2 SUB-LINKs.
  - ✓ SUB-LINKs shall operate at the same data rate.
  - ✓ SUB-LINKs are permitted to support different widths.
  - ✓ SUB-LINKs must be of the same RATE Series.
- SUB-LINK
  - ✓ A SUB-LINK shall have at least one LANE per SUB-LINK.
  - ✓ The number of LANEs per SUB-LINK is the same as supported by PCIe
  - ✓ All the LANEs within a SUB-LINK shall operate at the same data rate.
  - ✓ The modules of the SUB-LINK must be of the same RATE Series.
- Not Supported
  - ✓ Type-II modules are not supported.
  - ✓ LS-MODE for run-time operation is not supported.
  - ✓ Optical Media Converters (OMC) are not supported.

M-PCIe M-PHY Speed/Rate

- M-PCIe required Speed capabilities

M-PCIe LINK Speed Supported	Required HS-GEAR Support			
	PWM-G1	HS-G1	HS-G2	HS-G3
HS- G1	Required	Required	n/a	n/a
HS-G2	Required	Optional	Required	n/a
HS-G3	Required	Optional	Optional	Required

Source: Intel Corporation

- PWM-G1 is required for initialization and test
- Support for both Rate series is required

LOGICAL SUB BLOCK

Symbol Encoding & Framing

- M-PHY uses 8b/10b encoding and defines 8 special control symbols.
- M-PCIe uses 8b/10b encoding scheme.
- M-PCIe retains the special symbols as defined in PCIe.
  - ✓ Framing
  - ✓ Link Management

- M-PCIe mapping of special symbols to M-PHY special control symbols shown below

Control Symbols	M-PCIe Encoding	MIPI M-PHY Mapping
K28.5	COM	MARKER0
K28.3	IDL	MARKER1
K28.6	SDP	MARKER2
K23.7	SKP	MARKER3
K27.7	STP	MARKER4
K29.7	END	MARKER5
K30.7	EDB	MARKER6
K28.1	PAD	FILLER
K28.0	Reserved	Reserved
K28.2	Reserved	Reserved
K28.7	Reserved	Reserved
Others	Reserved	Reserved

Logical PHY Sub-Block

- Framing & Application of symbols to LANEs is same as PCIe with the following exceptions:
  - ✓ For multi-lane implementations, Logical Idle symbol is used for padding instead of PAD.
  - ✓ PAD (MIPI M-PHY FILLER) symbols are not permitted to be inserted within any link layer packet, any transaction layer packet or any ordered set.
- Data Scrambling is same as defined for PCIe
  - ✓ Data symbols transmitted during SYNC are not scrambled
  - ✓ M-PHY permitted to use internal SYNC data symbols as defined in M-PHY specification.
- Lane polarity inversion not supported
  - ✓ M-PHY does not detect polarity inversion
  - ✓ Transmitter polarity inversion is out of scope of ECN
- EIEOS, FTS not used for M-PCIe
  - ✓ M-PHY PREPARE and SYNC protocols are used.

SKEW & Clock Compensation

- Lane-to-Lane De-skew
  - ✓ MARKER0 or the COM symbol (Ordered Set, SKP OS) used to determine Lane-to-Lane De-skew
  - ✓ Lane-Lane De-Skew Timing parameters

Parameter	Description	Value
$T_{L2L\_SKEW\_HS\_TX}$	Permitted skew between any two LANEs measured at the transmitter's pins.	1600 ps
$T_{L2L\_SKEW\_HS\_RX}$	Permitted skew between any two LANEs measured at the receiver's pins.	4000 ps

- Clock Tolerance Compensation
  - ✓ Same scheme as PCIe (Skip ordered Set)
  - ✓ Transmission frequency is 250 Symbol Times for ± 2KPMM else between 1180 and 1538 symbol times.

TS1 & TS2 Ordered Set

Symbol Number	Description
0	COM (K28.5) for Symbol alignment.
1	LINK Bandwidth and GEAR support + Bandwidth Change Request and Bandwidth Change Confirm Bits
2	RATE Series Support
3	TX-LANE Number within LINK for Transmitter and RX-LANE Number within LINK for Receiver.
4	TX SUB-LINK Width.
5	RX SUB-LINK Width.
6	Training Control.
7 - 15	D10.2 (4Ah) as TS1 Identifier. D5.2 (45h) as TS2 Identifier.

LINK INITIALIZATION & DISCOVERY

M-PCIe LTSSM

- Detect
  - ✓ Initial state after power-on reset or fundamental reset de-assertion
  - ✓ Only Lane 0 operational, other lanes (multi-lane) are in HIBERN8
  - ✓ PWM-G1 only
- Configuration
  - ✓ Entered with Lane 0 in PWM-G1
  - ✓ Discover & Configure M-PHY
  - ✓ Exit configuration in HS Gear on all configured LANEs.
- L0
  - ✓ State to exchange data & control packets
  - ✓ State used to enter low power states
  - ✓ Tx permitted to enter STALL during idle periods
- L1
  - ✓ All M-PHY Lanes in HIBERN8
- L2
  - ✓ Same as PCIe
- Recovery
  - ✓ Retrain for recovering from transient errors
  - ✓ Change Link Bandwidth
- Disabled
  - ✓ Disable the configured Link
- Hot-Reset
  - ✓ Propagate Hot-reset
- LoopBack
  - ✓ Functional Loopback

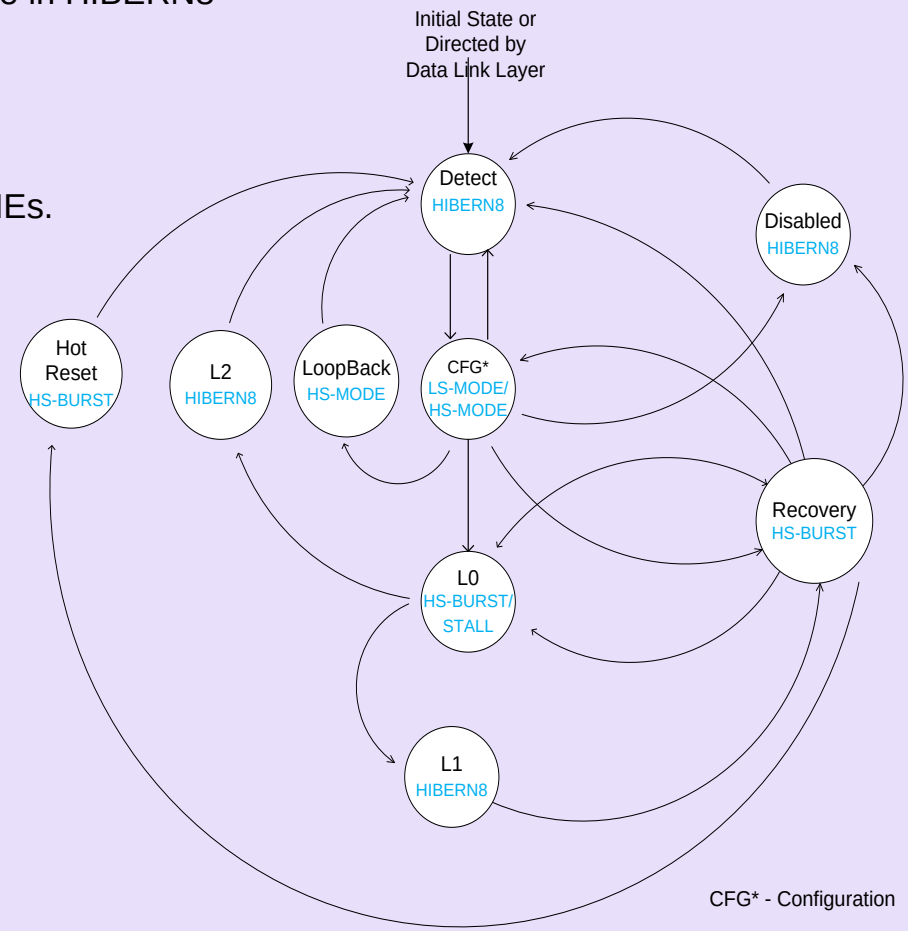

Configuration

- The M-PHY MODULE capabilities are discovered and MODULE attributes are configured in this state
  - ✓ Support for software assisted discovery & configuration is supported (Configuration.Software).
- The LINK operates in LS-MODE of operation and only TX-LANE(0) and RX-LANE(0) must be used. All other LANEs stay in HIBERN8 state
- The Upstream component must initiate the Initial LINK Discovery, Configuration and Effectuate process
- This process is needed for both Local Registers and Remote Registers
- The mechanism for Local Register Configuration (LRC) is not in the scope of the ECN (implementation specific).
- Remote Register Configuration (RRC) must use Remote Register Access Protocol (RRAP) as defined in the ECN

RRAP

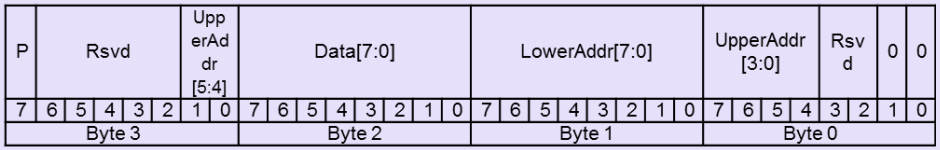

RRAP Write Command Format

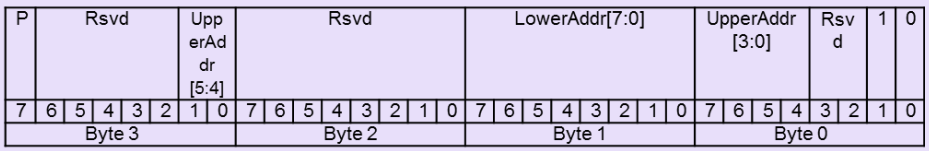

RRAP Read Command Format

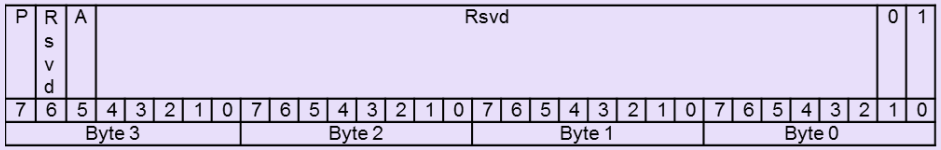

RRAP Write Response Format

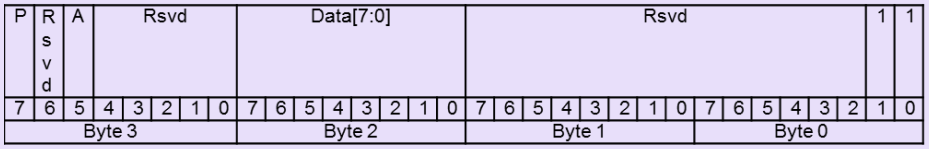

RRAP Read Response Format

- While in the LS-MODE (PWM-G1), communication is achieved using the RRAP command & response packets
- An RRAP Master must be capable of issuing Command packets while an RRAP Target must issue response packets
- The Downstream port acts as the master (when not in TEST mode) and the Upstream port acts as the target
- Only one outstanding command at a time, with response time-limit specified for Target, and response time-out specified for Master
- Protocol supports Parity Check, Completer Abort, and Write Broadcast
- Supports a mechanism to shorten configuration times (Link Configuration Retain)
- Two Address Maps are defined
  - ✓ Protocol Specific Address Map – Upper Address – 00 thru 0x24h
  - ✓ Test Address Map with Upper Address – 0x2Eh

Initialization – Reset to L0

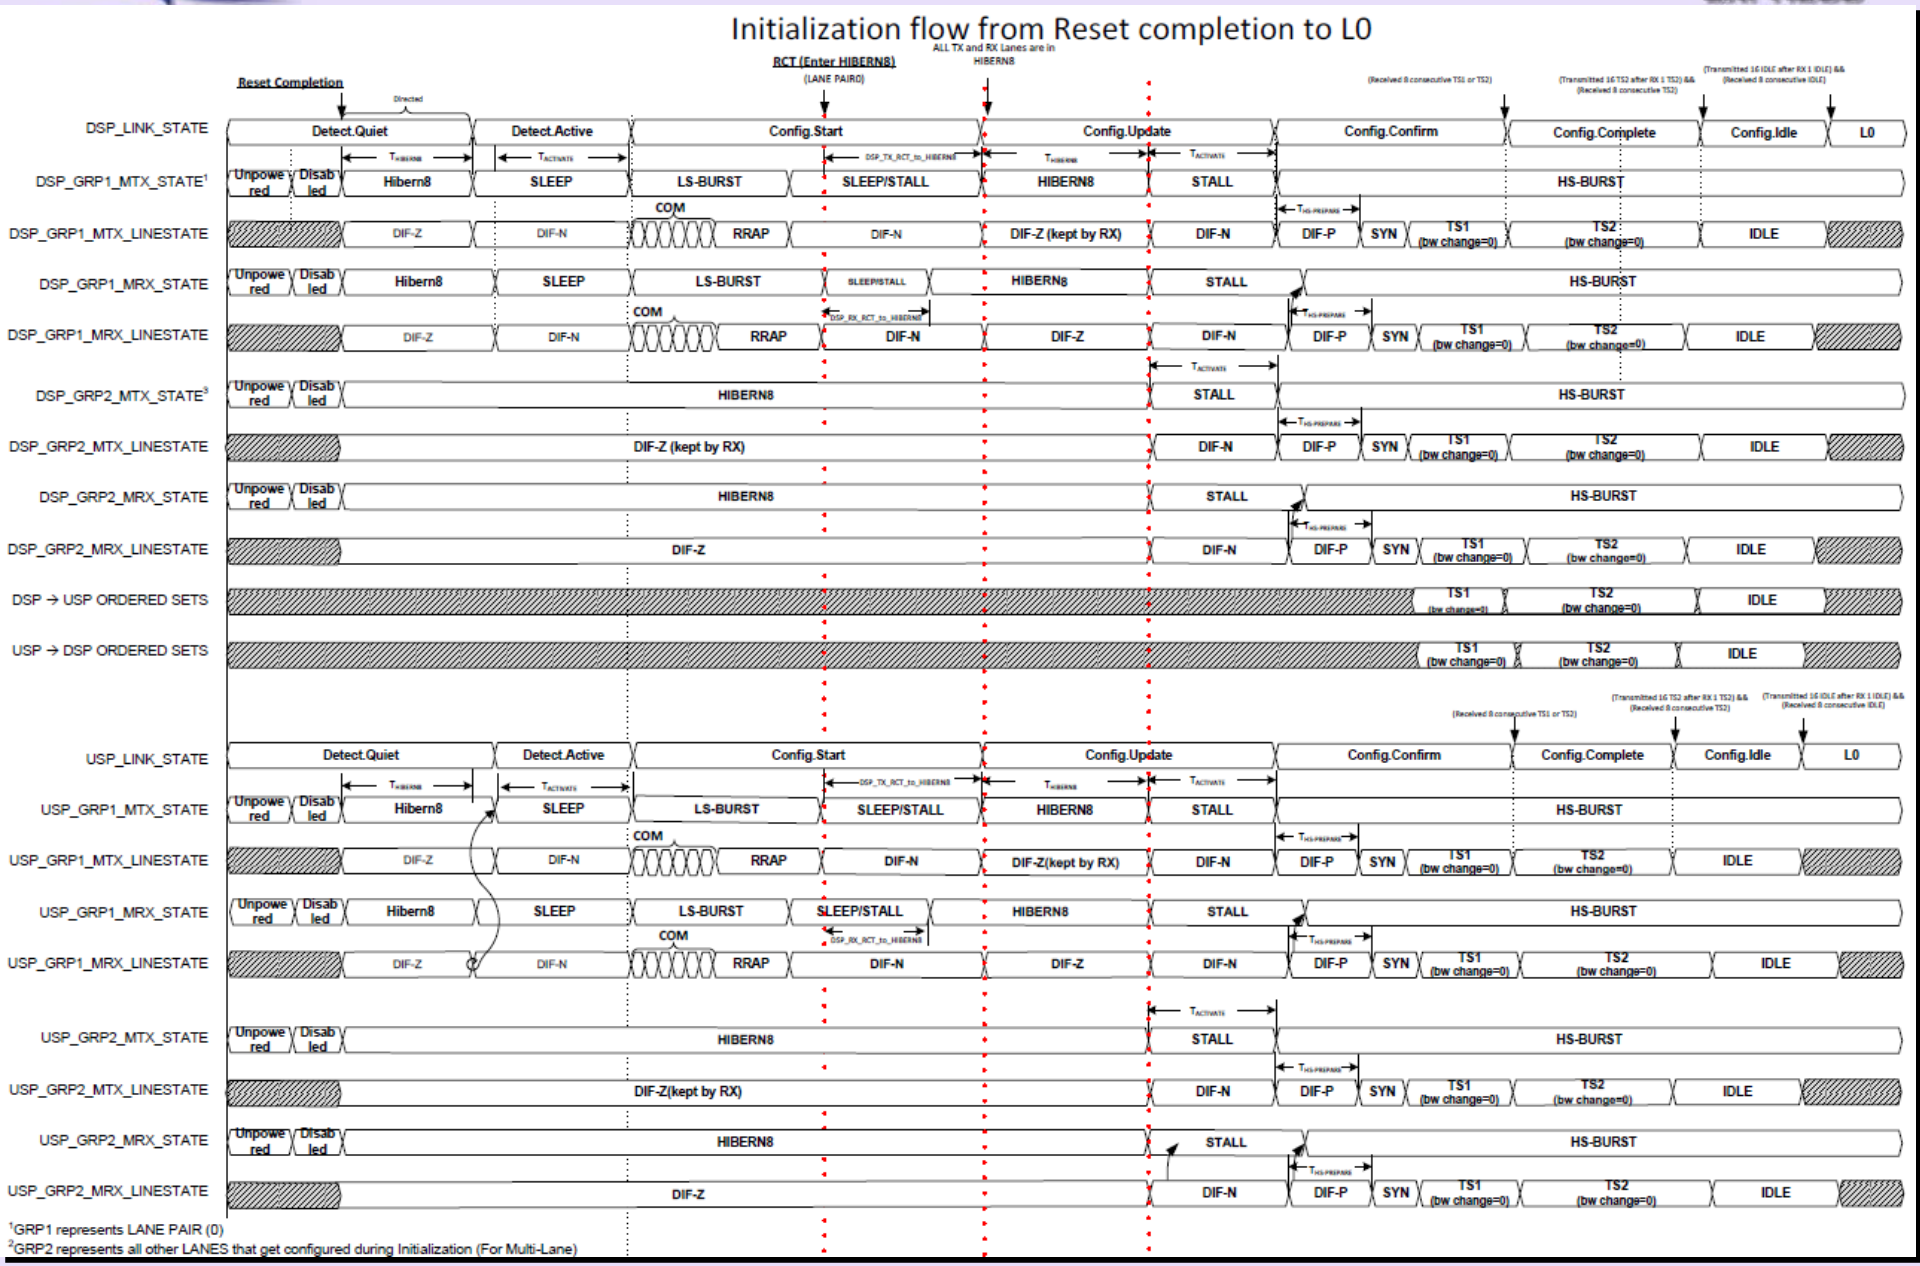

Transmitter in STALL

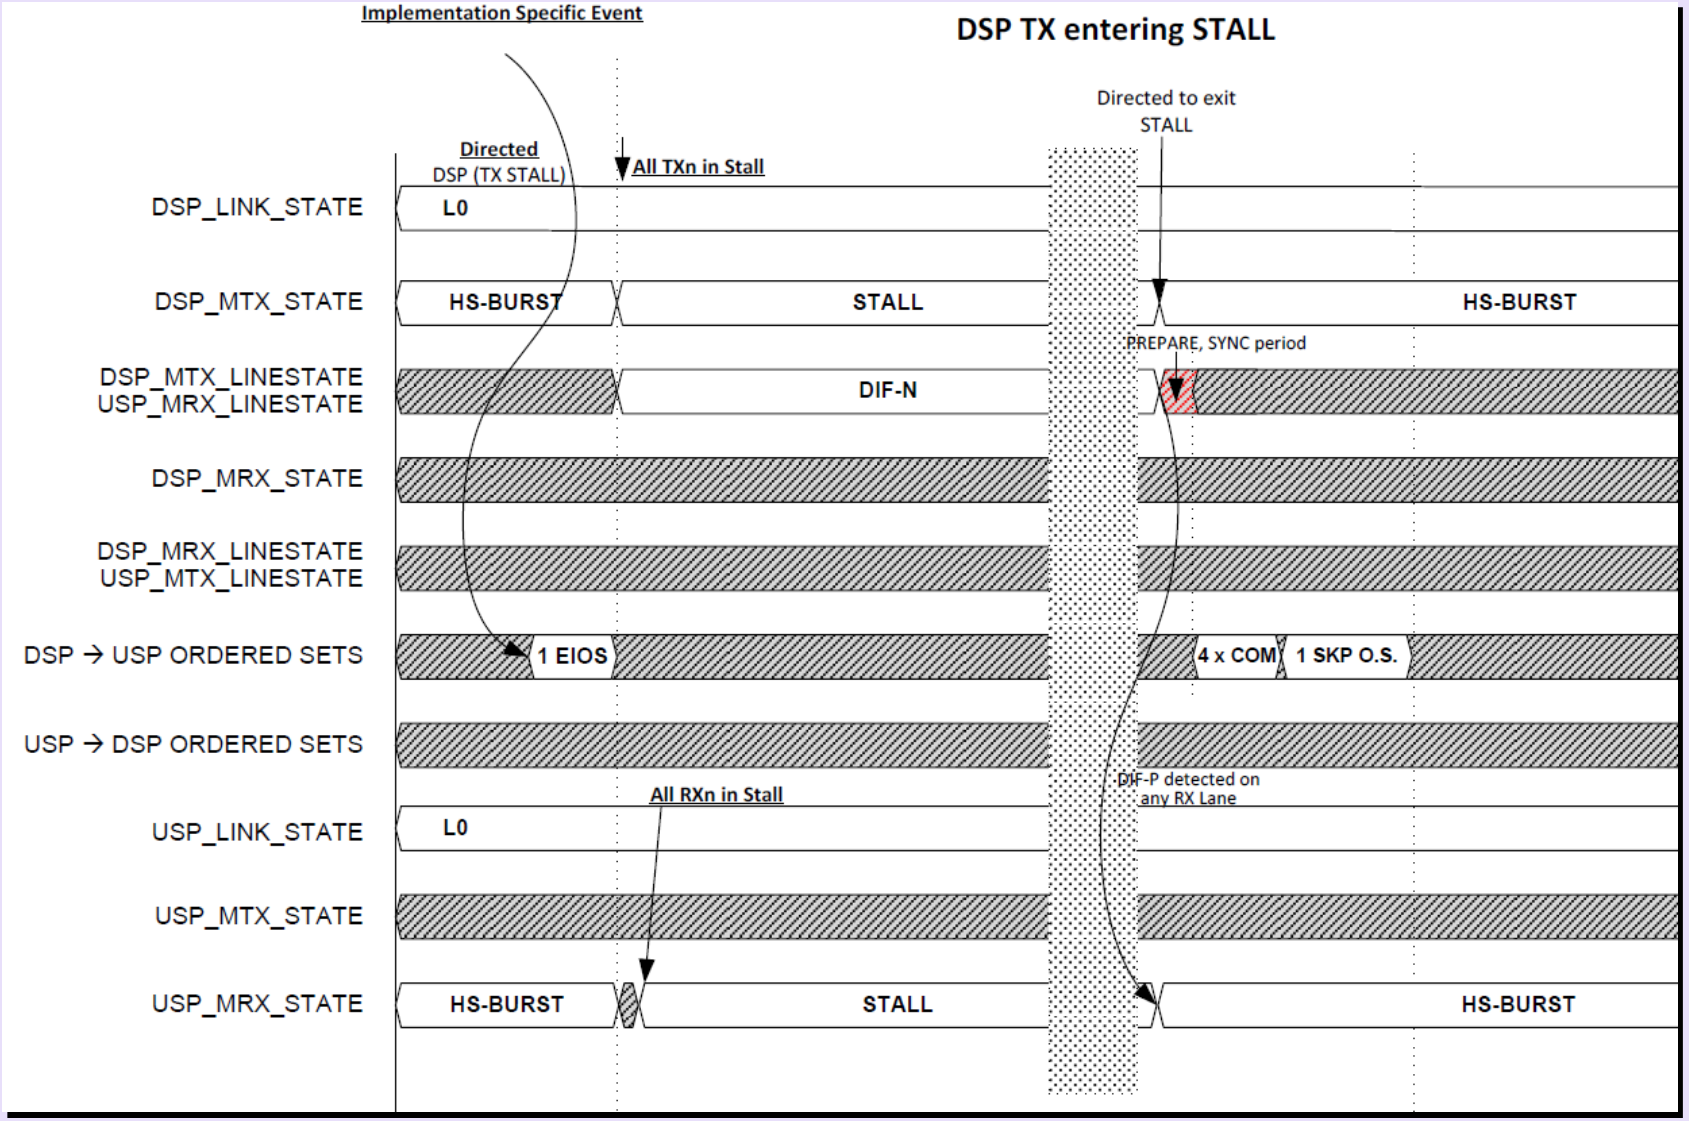

HIBERN8 Entry/Exit

- Special conditions apply when transitioning configured LANEs to HIBERN8
- M-PHY requires the receiver lanes to be ready to enter Hibern8 before the remote transmitter can enter Hibern8
  - ✓ Requires protocol to guarantee safe transitions.
- M-PCIe LANEs enter HIBERN8
  - ✓ LANE(0) after completion of configuration
  - ✓ LTSSM timeout conditions
  - ✓ L1, L2, Hot-Reset and Disable.
- Entry into HIBERN8
  - ✓ All transmit LANEs must be transitioned to enter STALL state.
  - ✓ All Configured RX-LANEs must trigger RCT to enter HIBERN8 state no later than $T_{RX\_RCT\_to\_HIBERN8}$
  - ✓ All Configured TX-LANEs must trigger RCT to enter HIBERN8 state no earlier than $T_{TX\_RCT\_to\_HIBERN8}$
- M-PCIe HIBERN8 requirements
  - ✓ Minimum Hibern8 residency – 1 to 100µs
  - ✓ Min. Tactivate period – not to exceed 100µs

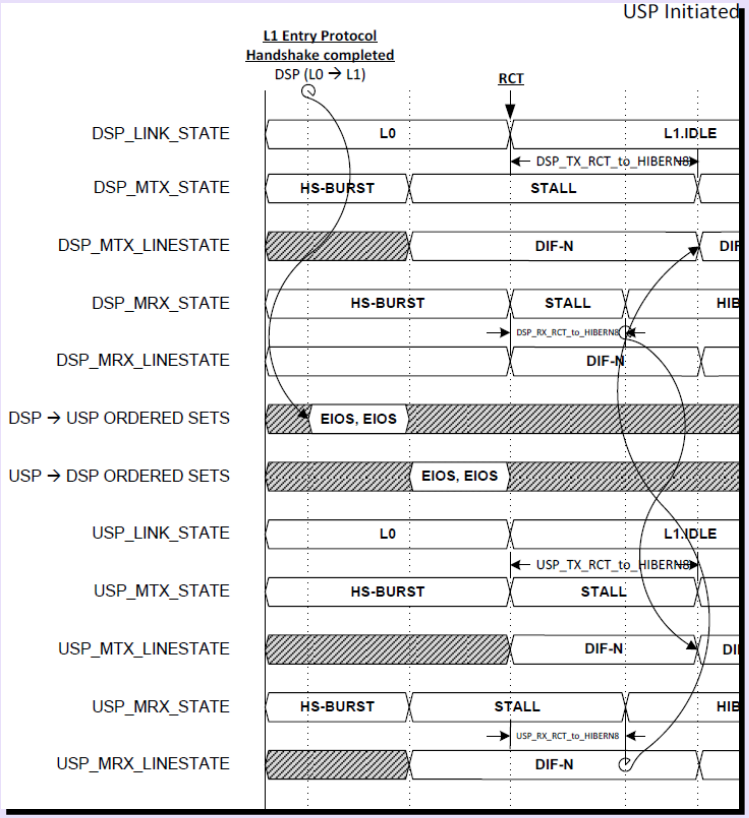

Parameter	Value
$T_{RX\_RCT\_to\_HIBERN8}$	800 ns
$T_{TX\_RCT\_to\_HIBERN8}$	1 us

L0 to L1 Timing

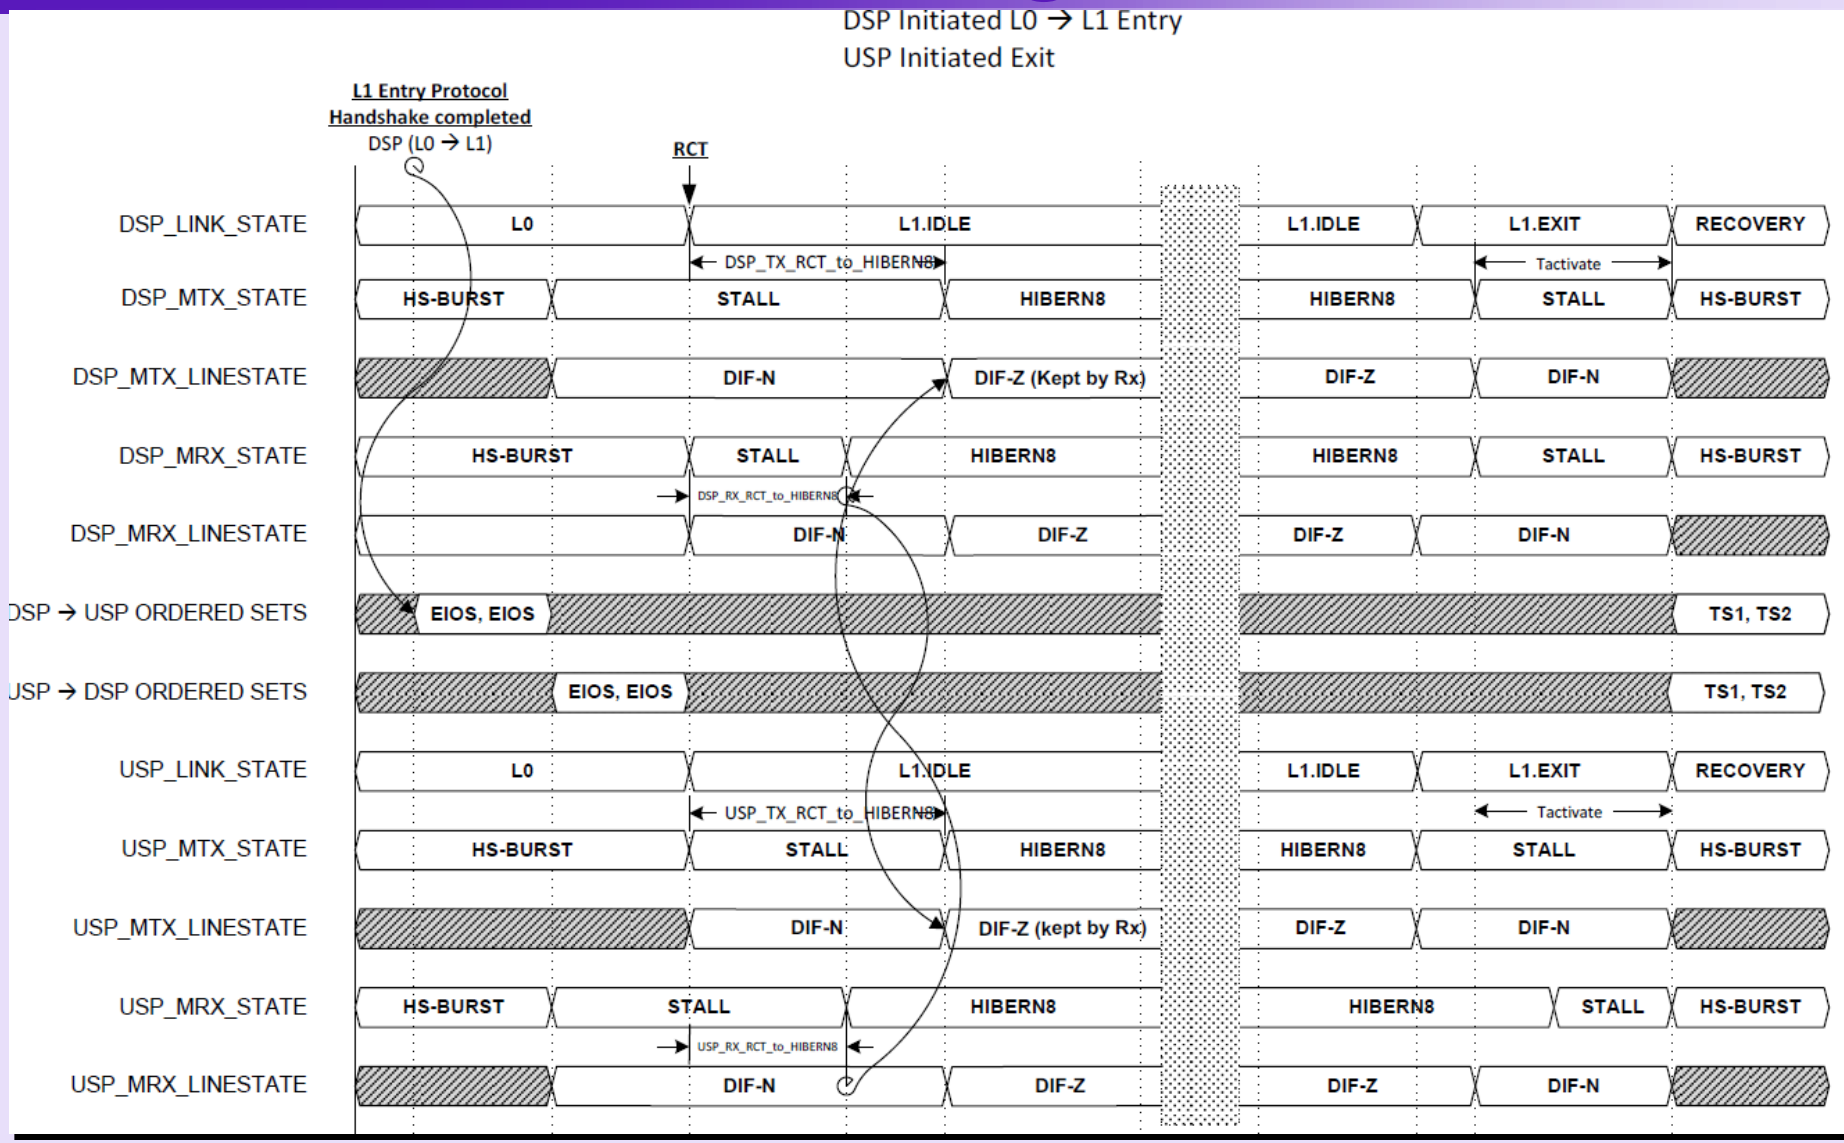

LINK BANDWIDTH MANAGEMENT

Link Bandwidth Management

- Dynamic LINK Re-Configuration is an optional normative functionality that supports changing the LINK Bandwidth without going through Link Down.
- The LINK Bandwidth parameters that can be re-configured during run-time are:
  - ✓ RATE Series – A or B
    - Symmetric between sublinks
  - ✓ High-Speed GEARs
    - Symmetric between sublinks
  - ✓ SUB-LINK width
    - Allows asymmetric widths between sublinks
- Supports hardware autonomous as well as software driven bandwidth re-configuration.
  - ✓ Bandwidth Re-Configuration can be initiated by both upstream and downstream device.
  - ✓ Link retraining can also be triggered by writing 1b to the Retrain Link bit.

Link Speed, Rate & Width Management

LINK Speed

- If the Hardware Autonomous Speed Disable bit of both components is clear, the component is permitted to autonomously adjust the LINK High-Speed GEARs.
- Either component is permitted to lower or increase the LINK speed by removing or adding the corresponding LINK High-Speed GEAR from the list of supported High-Speed GEARs.
- The criteria for removing or adding High-Speed GEARs is implementation specific.
- The Target Link Speed field in the Link Control 2 register sets the upper bound for the LINK speed.

LINK RATE

- If the Dynamic RATE Series Reconfiguration of both components is clear is enabled, the component is permitted to autonomously adjust the LINK RATE Series.
- Either component is permitted to change the LINK RATE.
  - ✓ When requesting to change the RATE Series, the requesting component must set the corresponding RATE Series support bit for the new RATE Series and clear the RATE Series support bit for the current RATE Series.
  - ✓ The responding component must set the corresponding RATE Series support bits for the RATE Series that it supports, which may be more than one RATE Series.
- The criteria for changing RATE is implementation specific.

LINK Width

- If the Hardware Autonomous width Disable bit in the Link Control register of both components is clear, the component is permitted to lower or increase the SUB-LINK width.
- Either component is permitted to change the SUB-LINK width by removing or adding the corresponding TX-LANE width and RX-LANE width setting from the list of supported SUB-LINK width settings.
- The criteria for removing or adding LINK width settings is implementation specific.

Link Re-Configuration Rules

- If the new High Speed GEAR requested is different than the current High Speed GEAR
  - ✓ The component on both sides of the LINK has at least one commonly supported High Speed GEAR, the new High Speed GEAR will be the highest commonly supported High Speed GEAR.
  - ✓ The component on both sides of the LINK does not have at least one commonly supported High Speed GEAR, the new High Speed GEAR will remain at the current High Speed GEAR.
- If the new SUB-LINK width requested is different than the current SUB-LINK width
  - ✓ The component on both sides of the LINK has at least one commonly supported SUB-LINK width in the request, the new SUB-LINK width will be highest commonly supported SUB-LINK width.
  - ✓ The component on both sides of the LINK does not have at least one commonly supported SUB-LINK width in the request, the new SUB-LINK width will remain at the current SUB-LINK width.
- If the new RATE Series requested is different than the current RATE Series
  - ✓ The component on both sides of the LINK supports the new RATE Series, the new RATE Series must be the requested RATE Series.
  - ✓ The component on the other side of the LINK does not support the new RATE Series, the new RATE Series will remain at the current RATE Series.

Dynamic Link Re-Configuration Sequence

- Component initiates Dynamic LINK Re-Configuration
  - ✓ Send TS1 with the bandwidth change bit set.
  - ✓ Set the fields in TS1 with the requested RATE Series, HS GEARs, SUB-LINK width.
- Corresponding component upon successfully receiving TS1
  - ✓ Respond with Training Ordered set TS1 with the bandwidth change request bit set.
  - ✓ Set the TS1 fields based on the support RATE Series, HS GEARs and SUB-LINK width.

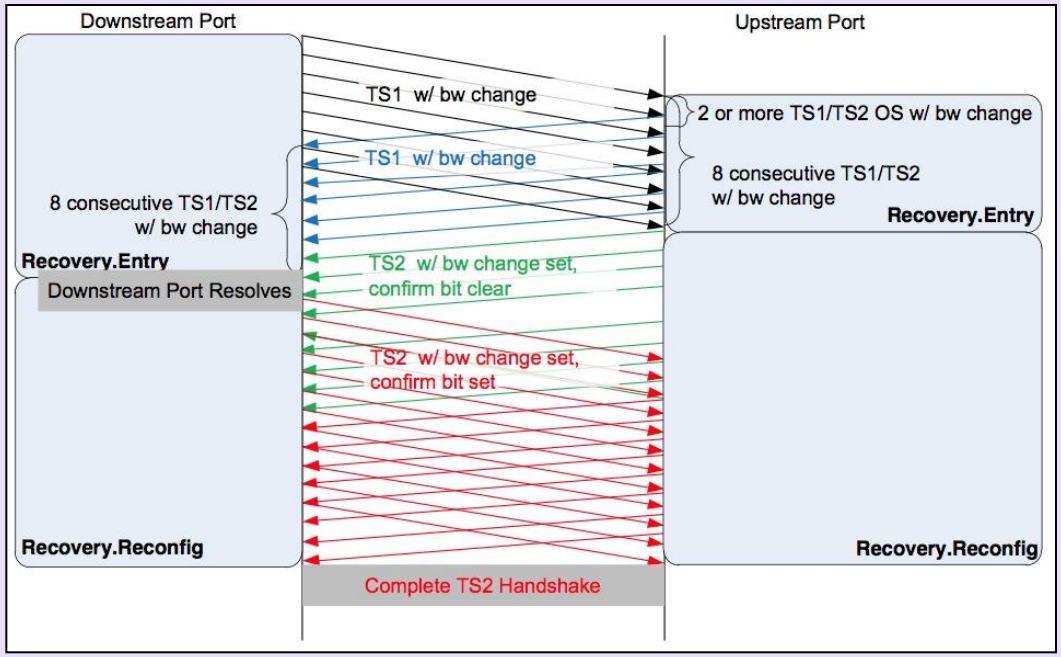

- The Upstream component is responsible for determining the new RATE Series, HS GEARs and SUB-LINK width as per the Link Re-configuration Rules.
- The result of Dynamic LINK Re-Configuration shall be communicated by TS2 ordered set transmitted by the Upstream component.
  - ✓ LINK Bandwidth confirm bit in TS2 when Set indicates that the LINK shall operate at bandwidth different than current bandwidth
  - ✓ LINK Bandwidth confirm bit in TS2 when Clear indicates that the LINK shall operate at current bandwidth.

Upstream Port Initiated BW change

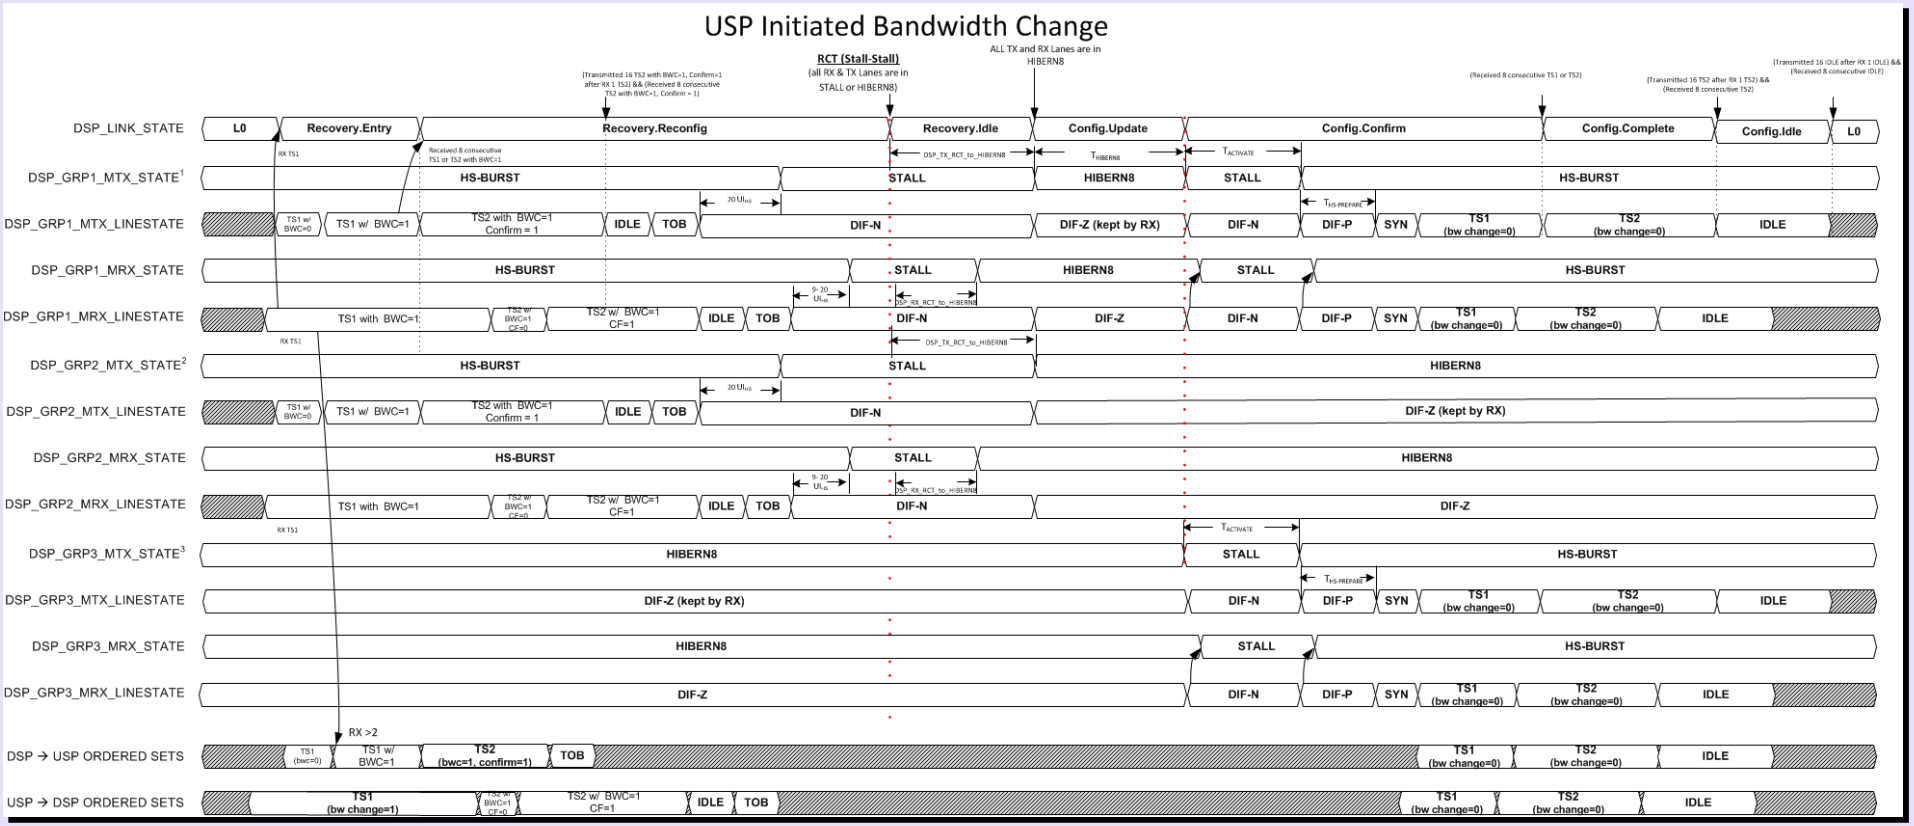

CONFIGURATION REGISTERS

Configuration Registers

- **PCI Express Capability and Extended Capability Registers:** Standard PCI Express Capability Registers are updated for M-PCIe functionality.
  - ✓ Compatibility with existing software.
  - ✓ Ex. Link Speeds vector updated to map M-PHY speeds.
- **New Extended M-PCIe Capability**: Capability, control and status fields to make effective use of M-PCIe features.
- **M-PCIe Capability & Configuration Attributes**: M-PCIe Capability & Configuration Attributes defined in RRAP Protocol specific Address Map. This space is access locally through the LRC and remotely using RRAP.

- M-PHY Configuration
  - ✓ M-TX configuration attributes must be set appropriately to match the remote M-RX Capability Attribute values.
- The following attributes are implementation specific

Attribute Name	AttributeID	Description	Reset Default
TX_HS_SlewRate	0x26	0 to 255	Implementation specific.
TX_LCC_Enable	0x2C	"0": No	TX_LCC_Enable must be disabled before entering PWM-Burst.
TX_DRIVER_POLARITY	0x2F	"0b": Normal "1b": Inverted	Implementation specific.

COMPLIANCE & TEST

Test Modes

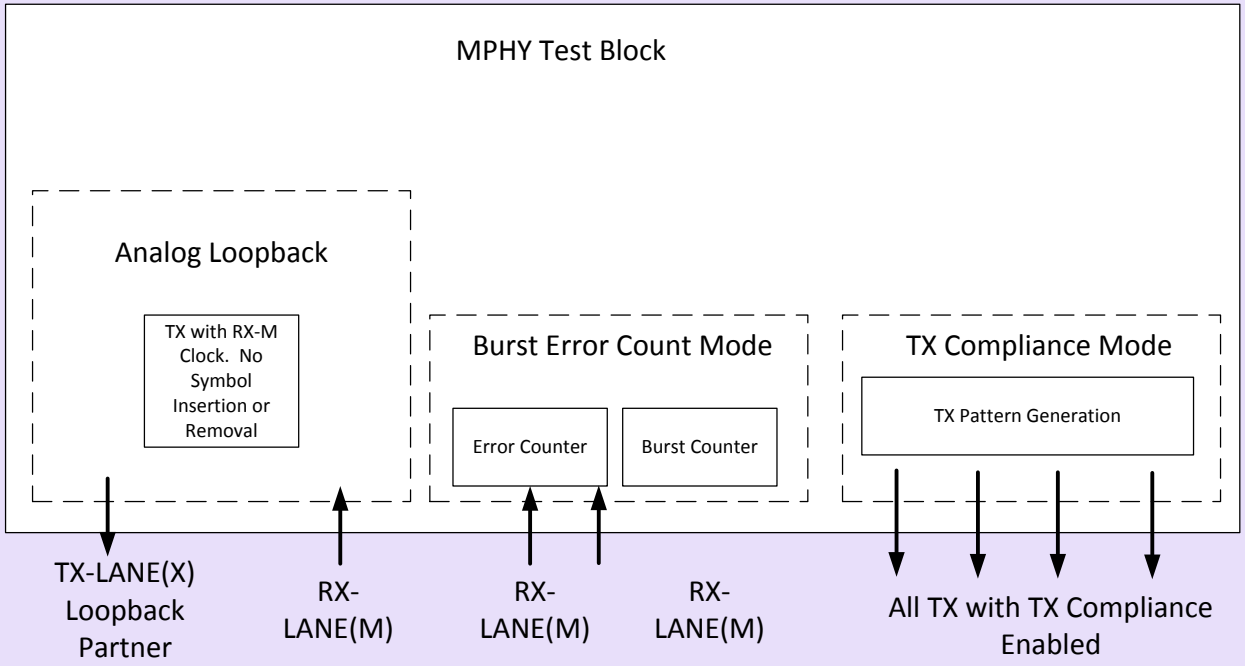

- Protocol Agnostic PHY Compliance & Test
- Standard entry through PWM-Gear 1 register protocol (RRAP)

The Motivation for Analog Loopback

- "Analog loopback" == TX loopback using RX clock
- **Required** to achieve protocol independence
  - ✓ SKPs and elastic buffer management differ by MPHY protocol
  - ✓ Already defined protocols have different elasticity buffer requirements
  - ✓ Standard PHY elasticity buffer block not defined
  - ✓ Protocol agnostic test block and procedures across M-PHY implementations

TX Compliance

Multi-Lane CRPAT

#	Lane 0 8B Data Hex	Lane 1 8B Data Hex	Lane 2 8B Data Hex	Lane 3 8B Data Hex	Lane 4 8B Data Hex	Lane 5 8B Data Hex	Lane 6 8B Data Hex	...
1	be	23	6b	b3	5e	35	be	...
2	d7	47	8f	14	fb	59	d7	...
3	23	6b	b3	5e	35	be	23	...
4	47	8f	14	fb	59	d7	47	...
5	6b	b3	5e	35	be	23	6b	...
6	8f	14	fb	59	d7	47	8f	...
7	b3	5e	35	be	23	6b	b3	...
8	14	fb	59	d7	47	8f	14	...
9	5e	35	be	23	6b	b3	5e	...
10	fb	59	d7	47	8f	14	fb	...
11	35	be	23	6b	b3	5e	35	...
12	59	d7	47	8f	14	fb	59	...

Continuous Mode CRPAT (Ex. LANE 0)

Burst Mode CRPAT (Ex. LANE 0)

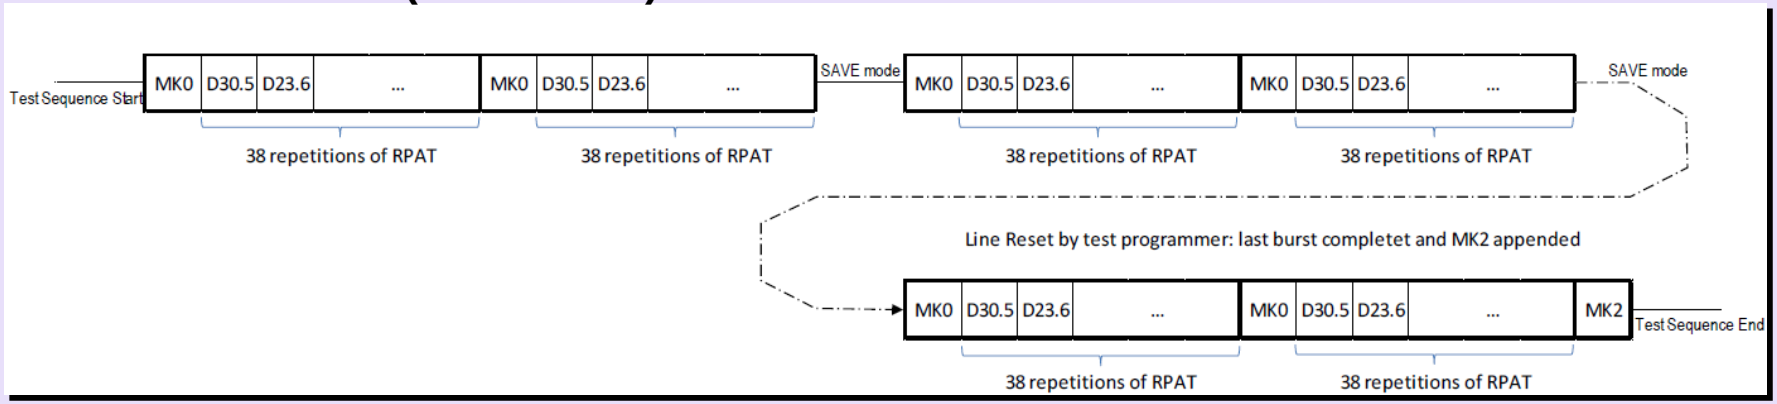

Thank you for attending the
PCI-SIG Developers Conference 2013

For more information please go to
www.pcisig.com