SuperSpeed USB Developers Conference

San Jose, California

November 17-18, 2008

SuperSpeed USB Power Management

Jim Walsh
USB 3.0 PM WG Chair / Mobile Platform Architect
Intel Corporation

Agenda

- SuperSpeed Power Management Overview
- Link States
- Device PM
- Host Support for Bus PM
- Hub PM
- Latency Tolerance Messaging
- Summary

Power Management Overview

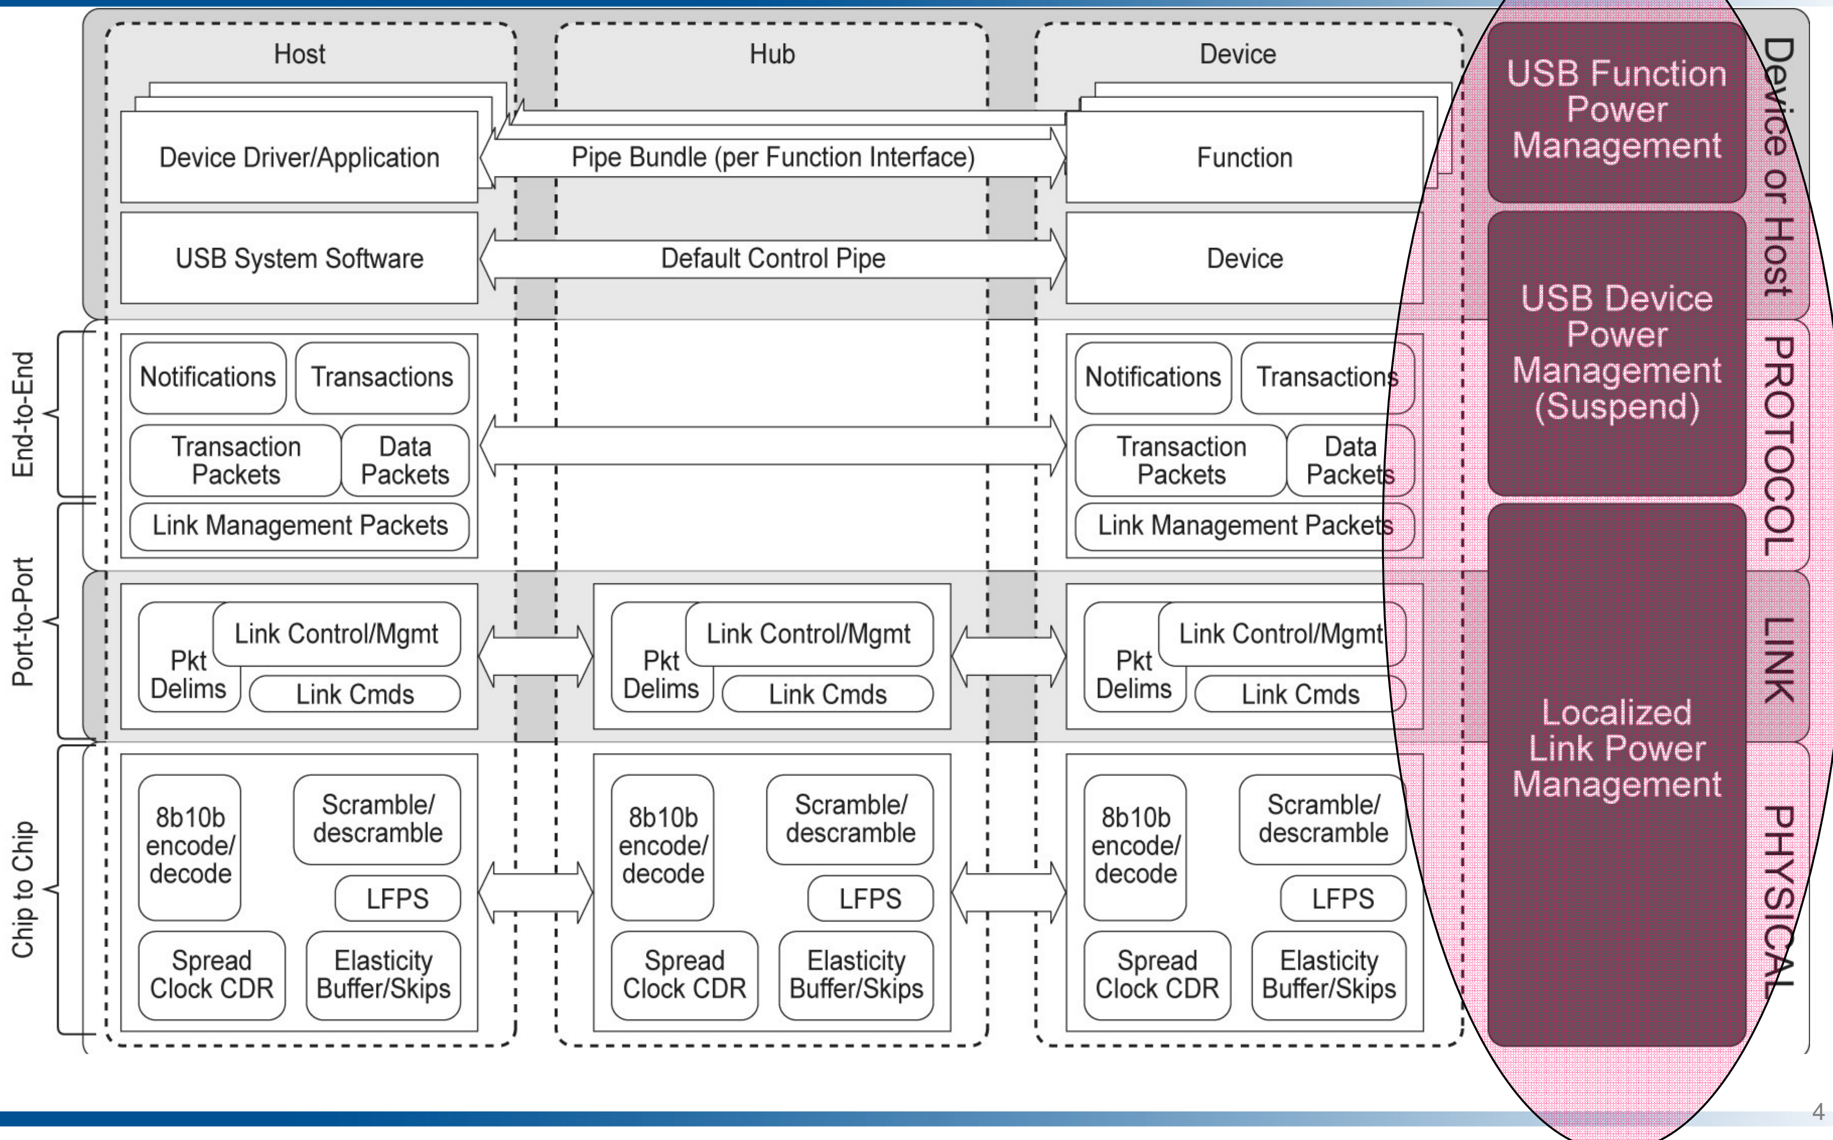

Power Management Overview

- PM is embedded at all levels
  - PHY layer, e.g. remote wakeup signaling
  - Link layer, e.g. low power link state entry & exit
  - Protocol layer, e.g. endpoint busy / ready notifications
  - Devices, e.g. function suspend
  - Hubs, e.g. “bubble up” link PM
  - Hosts, e.g. ping / ping response messaging
- Power efficiency at *system* level
  - Async endpoint busy / ready notifications – no polling
  - Packets routed, not broadcast
  - Low power link states entered automatically when idle

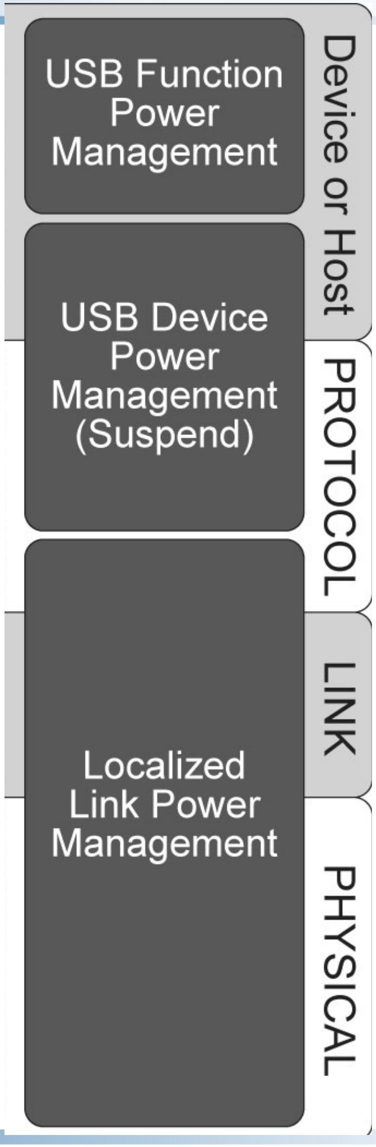

Agenda

- SuperSpeed Power Management Overview
- Link States
- Device PM
- Host Support for Bus PM
- Hub PM
- Latency Tolerance Messaging
- Summary

Link PM States

Link State	Description	Key Characteristics	Exit Latency
U0	Link active		N/A
U1	Link idle, fast exit	Rx and Tx circuitry quiesced	μs range
U2	Link idle, slow exit	Clock generation circuitry (e.g. PLL) may also be quiesced	typically low ms range, can be μs range
U3	Suspend	Portions of device power may be removed, e.g. much of PHY	typically ms range, can be high μs range

- Mechanisms for U1 & U2 entry:
  - Downstream port inactivity timeout, programmed by s/w
  - Device h/w initiated based on implementation specific knowledge, Packets Pending flag
- Mechanism for U3 entry: SetPortFeature(PORT_LINK_STATE U3)

Link State Transitions

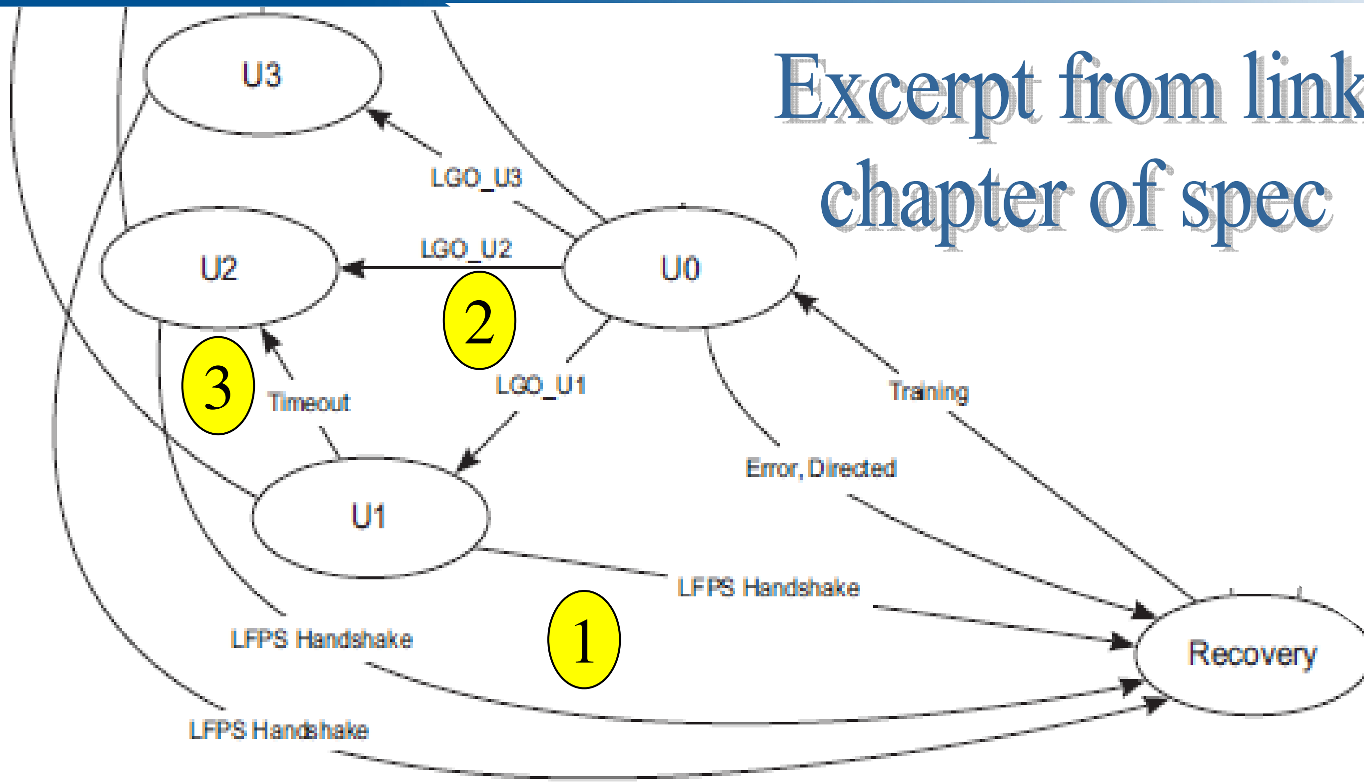

Excerpt from link chapter of spec

1 U1/U2/U3 exit: Low Freq. Periodic Signaling (LFPS) handshake → Recovery → U0

2 U1/U2/U3 entry: LGO_U* followed by LAU (accept) or LXU (reject) followed by LPMA

3 Direct U1 → U2: silent drop by both link partners independently

- S/W programs timeout in downstream port, which sends timeout to link partner via LMP packet

LFPS Handshake Process for U1/U2/U3 Exit

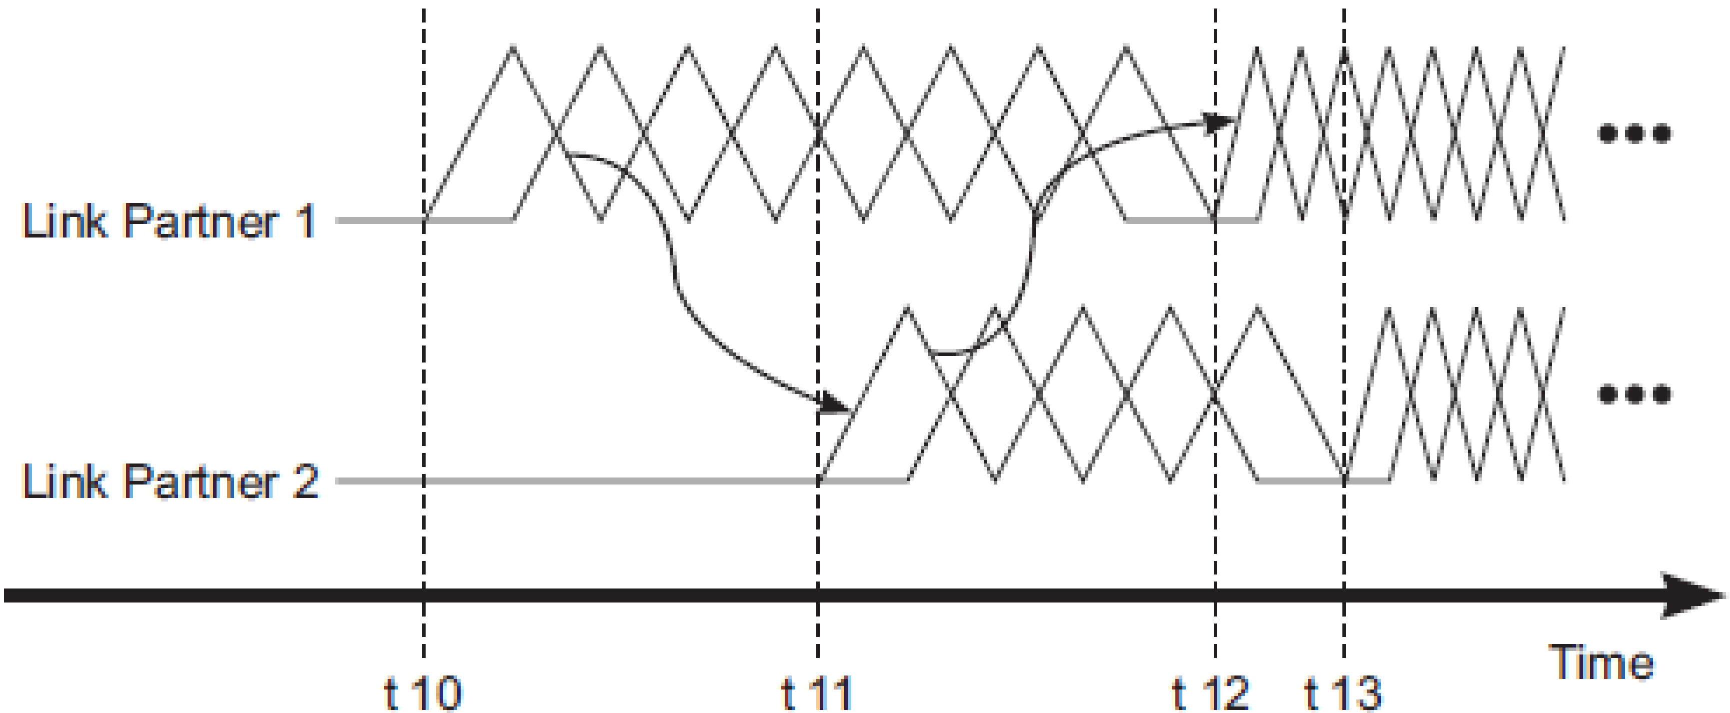

- Handshake process logically same for U1, U2, and U3
  - Timing just required to be faster for U1 vs. U2 vs. U3

Agenda

- SuperSpeed Power Management Overview
- Link States
- Device PM
- Host Support for Bus PM
- Hub PM
- Latency Tolerance Messaging
- Summary

Function Suspend and Device Suspend S/W Interface

- Function suspend
  - Individual functions* placed into *function* suspend independently
  - Controlled by FUNCTION_SUSPEND feature selector
- Device suspend
  - Device-wide state coupled to U3
    - Entered / exited intrinsically as a result of U3 entry / exit
      - SetPortFeature(PORT_LINK_STATE U3)
  - Device suspend entered regardless of function suspend state
- Selective suspend also supported
  - System software may initiate device suspend when all of a device's functions are in function suspend

* Composite devices contain multiple functions

Resume (Wake from Suspend) S/W Interface

- Host initiated wakeup
  - SetPortFeature(PORT_LINK_STATE U0)
- Remote wakeup (device initiated wakeup)
  - Enabled / disabled per function, encoded within FUNCTION_SUSPEND feature selector
    - Device-wide remote wake feature selector deprecated for SS devices
  - Remote wakeup enable / disable can be read using GetStatus request
  - Function sends Function Wake device notification to host
    - Notification sent after link returns to U0, if not already in U0
    - Notification causes host to generate an interrupt to system s/w

Device Initiated U1 and U2

- Two mechanisms for initiating U1 & U2
  - Downstream port link inactivity timers
    - Controlled by PORT_U1_TIMEOUT, PORT_U2_TIMEOUT feature selectors
  - Peripheral device driven
    - Enabled/disabled by U1_Enable, U2_Enable feature selectors
- Peripheral devices save power by initiating U1 & U2 more aggressively than inactivity timers
  - Peripheral devices have more information available
    - Device implementation and state
    - Host controller schedule information for device endpoints
      - Packets Pending and End of Burst flags, for bulk and interrupt endpoints
      - Last Packet flag, for isochronous endpoints

Example – Device Initiated U1 (Isochronous)

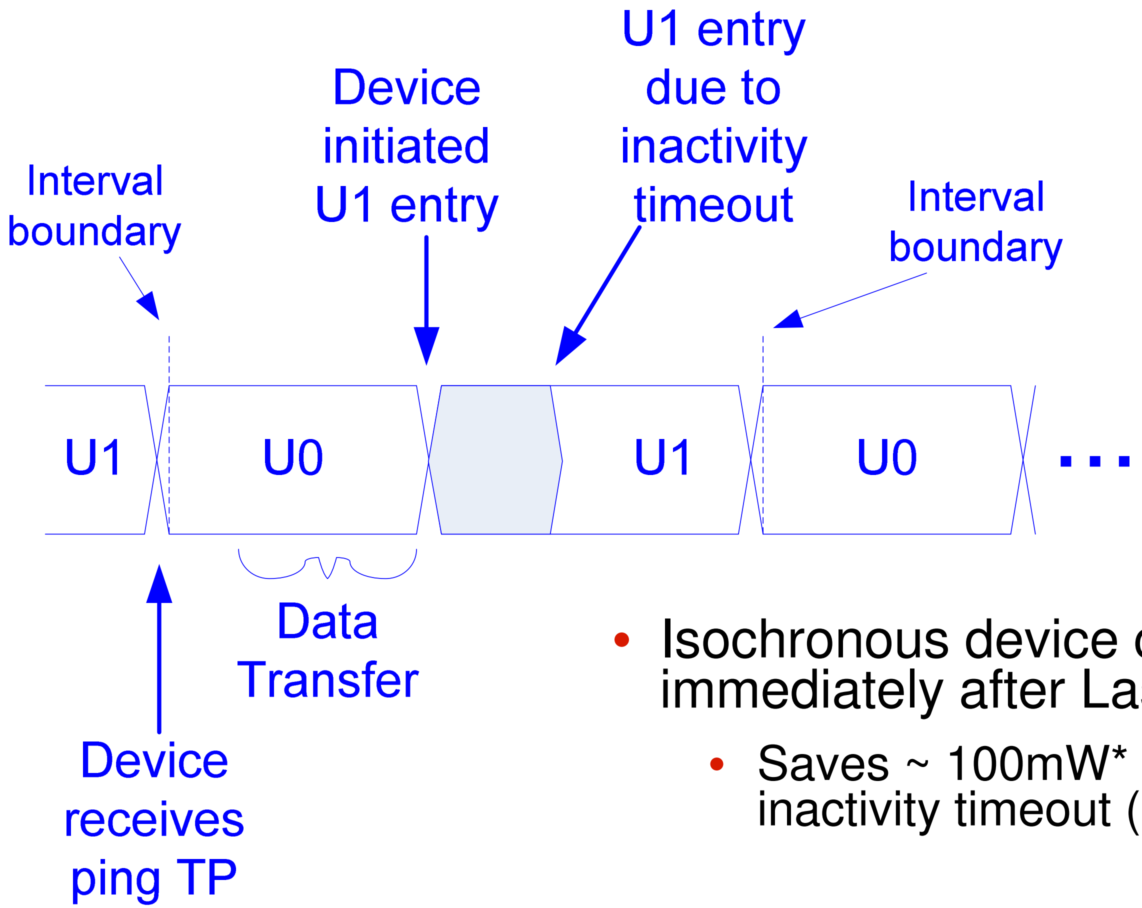

- Isochronous device can initiate U1 immediately after Last Packet flag is asserted
  - Saves ~ 100mW* compared to waiting for inactivity timeout (more if hubs are present)

* Based on PHY power difference of 250mW between U0 and U1, two PHYs per link, a 125us service interval, and a 25us U1 inactivity timeout: (2 PHYs) (250 mW per PHY) (25 / 125) = 100 mW.

Agenda

- SuperSpeed Power Management Overview
- Link States
- Device PM
- Host Support for Bus PM
- Hub PM
- Latency Tolerance Messaging
- Summary

Host Support for Bus PM: Isochronous Endpoints and Timestamp TPs

- Problem statement:
  - Isochronous transfers must get completed within service interval, but still want to use U1 / U2
    - Must comprehend U1 / U2 exit latencies
- Solution: Ping / ping response messaging
  - Host sends a ping to isochronous device ahead of an isoch transfer
    - Gets all links in path to device back to U0 prior to transfer
  - Device sends ping response to host (in response to ping)
    - Host then schedules isochronous transfer
    - Devices keeps link in U0 until transfer occurs
      - Due to hub link PM rules, this keeps all links in path to host in U0
  - Host can perform other transfers while waiting for ping response

- Timestamp packets sent at bus interval boundaries
  - Only sent on downstream ports in U0
  - U1 / U2 link inactivity timers ignore timestamp packets

Host Support for Bus PM: Interrupt Endpoints

- Problem statement:
  - Interrupt transfers must get completed within service interval, but still want to use U1 / U2
    - Must comprehend U1 / U2 exit latencies
- Solution:
  - Host sends transfer far enough ahead of time to compensate for worst case link exit latency

- Host stops interrupt endpoint activity upon receipt of an NRDY
  - Resumes upon receipt of an ERDY
    - Device (EP) will experience link exit latency in addition to (a maximum of) twice the subscribed service latency
  - No polling – links can enter U1 / U2 when there is no activity

Agenda

- SuperSpeed Power Management Overview
- Link States
- Device PM
- Host Support for Bus PM
- Hub PM
- Latency Tolerance Messaging
- Summary

Packet Deferring

- Problem statement (example)
  - Host initiates transfer with Device1, and Link1 is in U0 but Link2 is in U1 or U2
  - Host should not sit around waiting for Link2 to return to U0
    - Waiting would effectively waste a lot of bandwidth

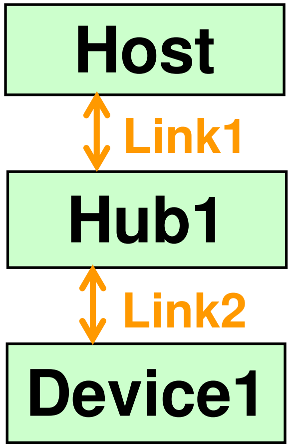

- Solution
  - Whenever a hub receives a packet targeted to a port in non-U0:
    - Hub sends deferred command back to host
    - Hub forwards original packet header to device, with deferred field asserted
  - Host treats deferred command similar to an NRDY from device
    - Can move on to initiate transfers with other devices
  - When ready, device sends ERDY to notify host it's ready for original transfer

Hub Power Management
U1 and U2

- Hubs "bubble up" low power link states
  - Upstream port enters lowest power state it can, but never lower than any of its downstream ports
  - Peripheral device link states "bubble up" through hierarchy
  - Devices can use this to limit link exit latency for entire path to host
- U1 and U2 exit
  - Exit to forward an upstream flowing packet
    - Hub upstream port exits in parallel with downstream port; 1us max delay
  - Exit to forward a downstream flowing packet
    - Hub receives packet first, then only targeted downstream port exits
- U1 and U2 entry
  - Software programmable downstream port inactivity timeouts for U1 and U2
  - Upstream ports can also initiate U1 and U2 (see earlier slide)

Hub Power Management
U3 and Hub Initiated Remote Wakeup

- U3 entry by SetPortFeature(PORT_LINK_STATE U3)
  - Also initiated automatically by downstream port if device is connected when hub / host is suspended
- Host initiated U3 exit by SetPortFeature(PORT_LINK_STATE U0)
- Device initiated U3 exit (remote wakeup) using hardware automated process (see next slide)
- Hubs can originate remote wakeup due to connect, disconnect, and over-current events
  - Enabled on a port by port basis (downstream ports only)
  - If remote wakeup is disabled, events are still detected
    - The last event is reported after hub is resumed
    - Eliminates need to wake computer just to service these events

Hub Power Management Wakeup Signaling Forwarding

- Hubs always forward remote wakeup signaling upstream
  - Upstream signaling terminates at Controlling hub
    - Controlling hub is either the root hub, or the hub nearest the device that is not in suspend
- Controlling hub reflects wakeup signaling on port that received wakeup signaling
  - This begins resume (U3 exit) of device under controlling hub
- Hubs exiting U3 initiate U3 exit on downstream port(s) that received wakeup signaling
  - This resumes all hubs in path to device, and the device itself

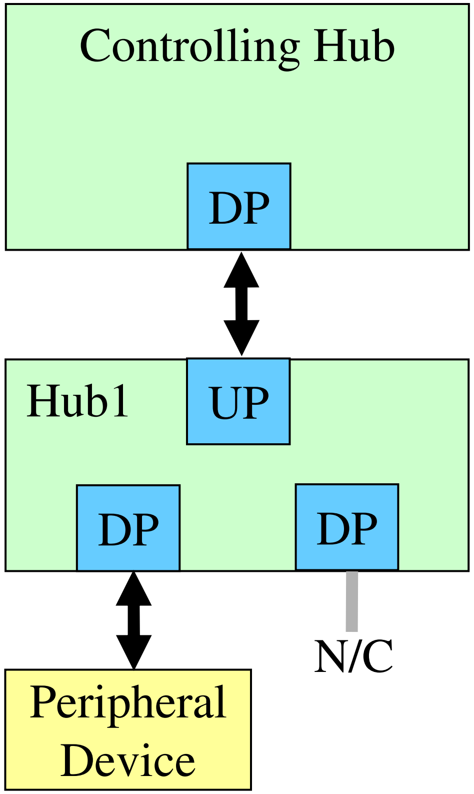

Animation: Wakeup Signaling Forwarding

Build

- Peripheral device and Hub1 in suspend
- Peripheral device transmits remote wakeup LFPS signaling
- Hub1 reflects remote wakeup LFPS signaling upstream
- Controlling hub reflects wakeup LFPS signaling downstream
  - Hub1 begins resume
- Hub1 reflects wakeup LFPS signaling on downstream port that received remote wakeup LFPS signaling
- LFPS handshake processes complete, and peripheral device and Hub1 complete wakeup process

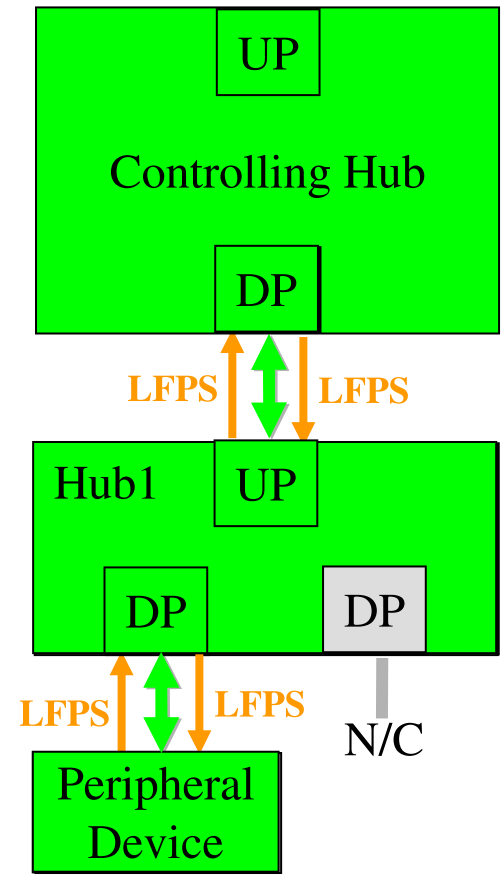

Agenda

- SuperSpeed Power Management Overview
- Link States
- Device PM
- Host Support for Bus PM
- Hub PM
- Latency Tolerance Messaging
- Summary

Latency Tolerance Messaging (LTM) Overview

- LTM enables system to dynamically enter deeper power saving states reliably, with cooperation of devices
- Devices report how much latency they can tolerate from the system – only for best effort services such as bulk
  - LTM notification packet containing latency tolerance value
- System enters deeper sleep states – when all devices in system can tolerate it
- USB 3.0 implementation
  - All devices support a default (idle) Best Effort Latency Tolerance (BELT) of 1ms
  - Any change in BELT communicated via LTM notification packet
    - Devices that can always tolerate default BELT never need to send an LTM packet
  - Hubs just forward LTM packets upstream
  - Device capability field, LTM enable/disable feature selector

Agenda

- SuperSpeed Power Management Overview
- Link States
- Device PM
- Host Support for Bus PM
- Hub PM
- Latency Tolerance Messaging
- Summary

Summary

- Additional resources
  - Power Management appendix in the specification
    - PM overview – ties all the pieces together
    - Policies for devices to initiate U1 and U2
    - Latency Tolerance Messaging device implementation example
    - Power considerations for SuperSpeed vs. High Speed interface for new device designs (also backup slide)
- Q & A

Backups

Back-up

Device PM – SuperSpeed vs. High Speed Interface Selection Example

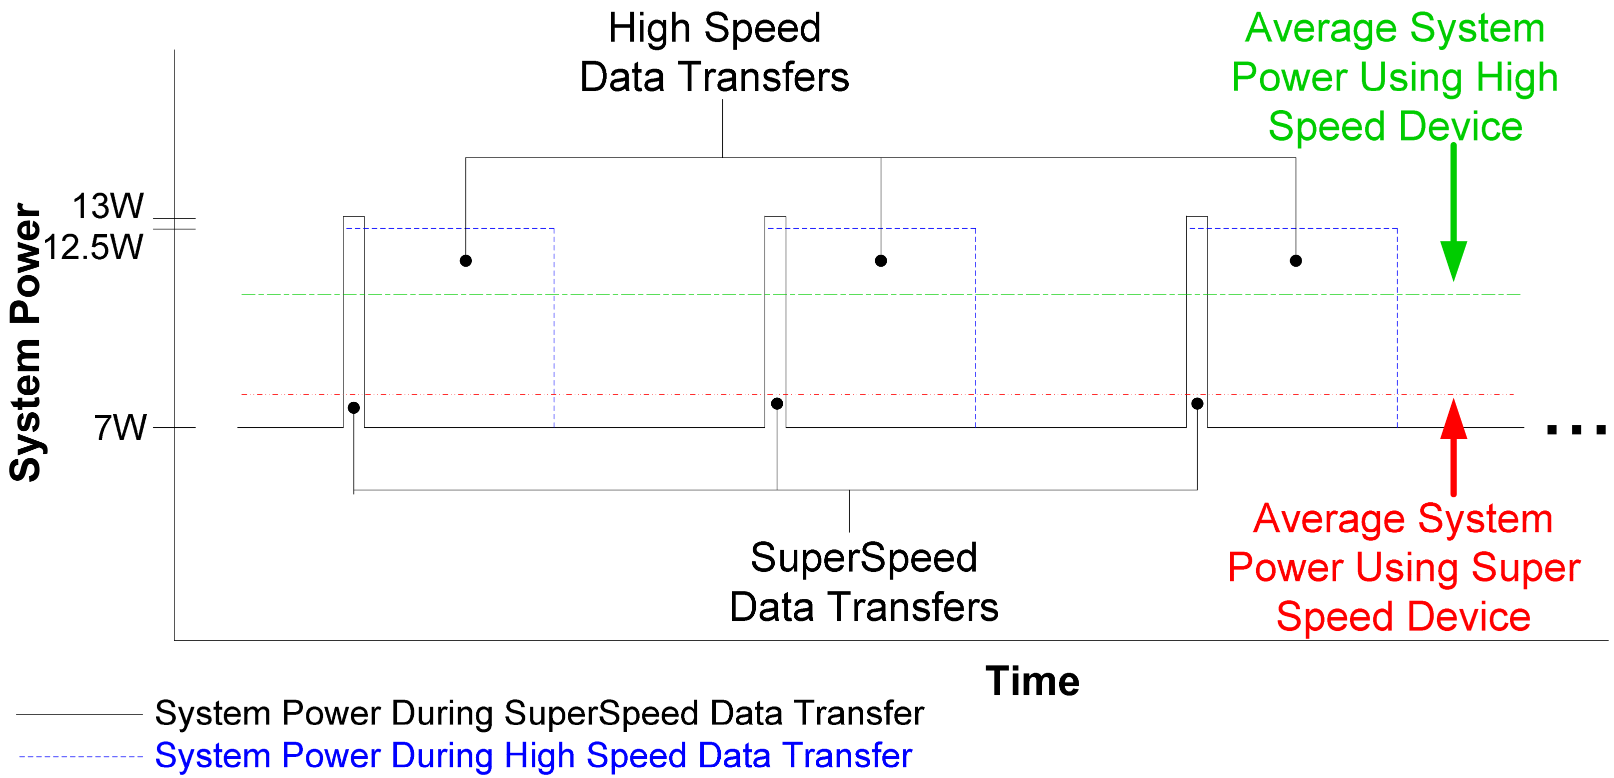

- System power can be far less on SuperSpeed vs. High Speed
  - System power is ~ 25% less during a 20MByte/s data transfer
    - High Speed ave. system power: (20MB/s)/(40MB/s) = 50%; (50%)(12.5W) + (50%)(7W) = 9.7W
    - SuperSpeed ave. system power: (20MB/s)/(333MB/s) = 6%; (6%)(13W) + (94%)(7W) = 7.5W

Removing Portions of Device Power During Suspend

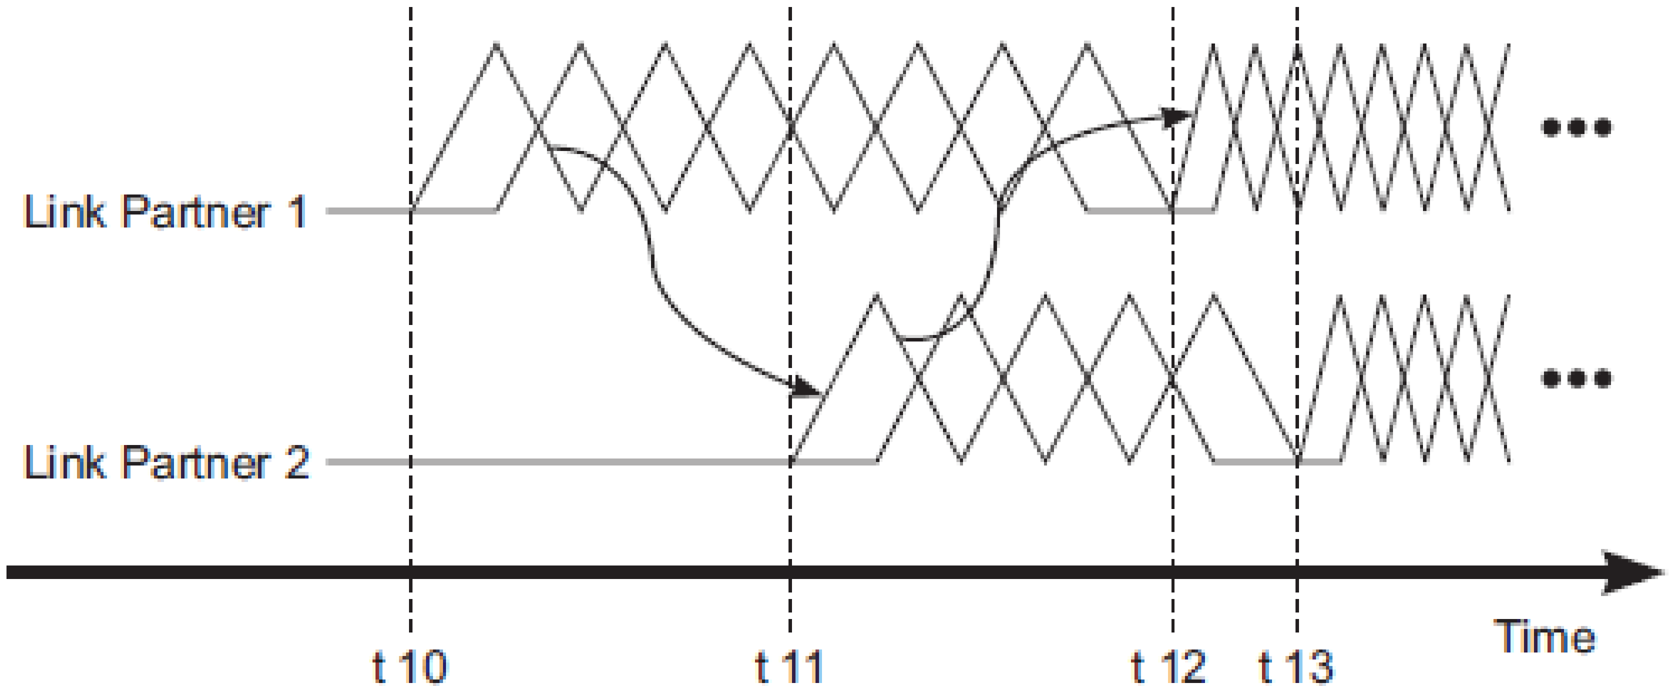

- U3 exit handshake process allows 10ms (t11 – t10) for device to respond, plus another 10ms to complete (t13 – t11)
  - Enough time for a device to turn on a power supply
  - Allows devices to remove portions of their power during suspend

BELT and System Exit Latency

- Device has a total intrinsic Latency Tolerance (LT)
  - Determined solely based on device implementation, e.g. amount of buffering
- Total LT must be allocated among system components
- Total LT is sum of t1, t2, t3, and t4:
  - t1: device-to-host path link exit latency
  - t2: ERDY propagation time from device to host
  - t3: host response time (from ERDY to transmission of request packet)
  - t4: packet propagation time from host to device
- BELT value is portion allocated to t3
  - Device subtracts U1SEL / U2SEL* from its pre-determined total LT
    - Reports result as LTM BELT value

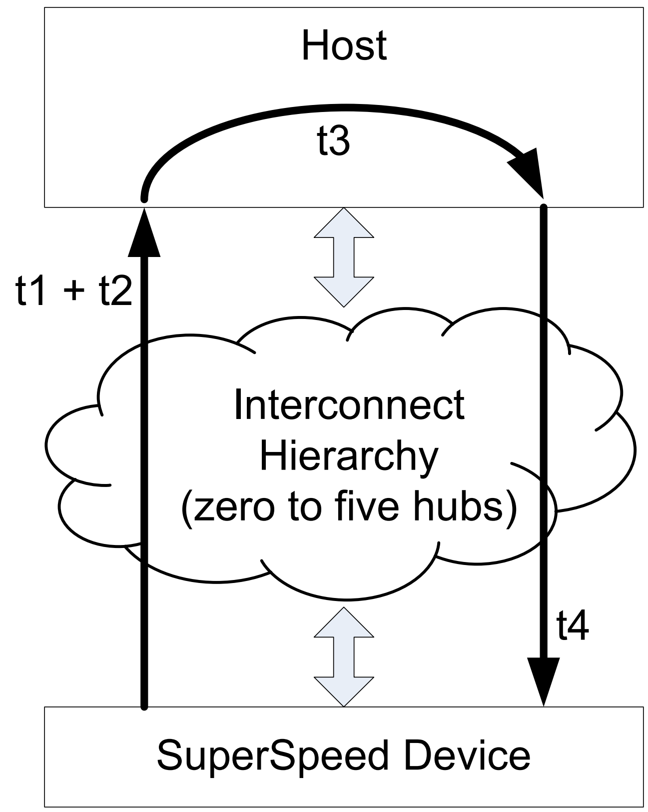

* U1SEL / U2SEL programmed using SET_SEL device request

PM Software Interface Summary

- Standard feature selectors for device recipients
  - U1_ENABLE, U2_ENABLE
  - LTM_ENABLE
- Standard feature selectors for interface recipients
  - FUNCTION_SUSPEND
- Standard device requests
  - SET_SEL
- Device GET_STATUS fields
  - U1_ENABLE, U2_ENABLE
- Interface GET_STATUS fields (first interface in function)
  - Function Remote Wake Capable
  - Function Remote Wake Enable
- SuperSpeed Device Capabilities Descriptor fields
  - U1DevExitLat, U2DevExitLat
  - LTM Capable
- Hub class feature selectors
  - PORT_U1_TIMEOUT
  - PORT_U2_TIMEOUT
  - PORT_REMOTE_WAKE_MASK
  - PORT_LINK_STATE
  - C_PORT_LINK_STATE
  - FORCE_LINKPM_ACCEPT
- Hub Port GET_STATUS fields
  - PORT_LINK_STATE
  - C_PORT_LINK_STATE
- SuperSpeed Hub Class Descriptor fields
  - HubHdrDecLat

PM Software Interface Detail (1)

- U1_ENABLE and U2_ENABLE feature selectors
  - Enable a device (upstream port) to *initiate* U1 and U2
  - Device can accept or reject U1 and U2 entry requests (from its link partner) regardless of U1_ENABLE and U2_ENABLE
- FUNCTION_SUSPEND feature selector
  - Two bit fields:
    - Bit 0 – low power (function) suspend state enable
    - Bit 1 – function remote wake enable
  - Device wide remote wake feature selector deprecated for SuperSpeed devices
- U1_ENABLE and U2_ENABLE device GET_STATUS fields
  - Return state of respective feature selectors
- Interface GET_STATUS fields:
  - Function Remote Wake Capable – indicates whether function is capable of remote wakeup
  - Function Remote Wake Enable – returns state of function remote wake enable bit in FUNCTION_SUSPEND feature selector

PM Software Interface Detail (2)

- U1DevExitLat and U2DevExitLat SuperSpeed Device Capabilities Descriptor fields
  - Return U1 and U2 device exit latencies respectively, when exit latency is limited by the device and not by its link partner
- LTM_Capable SuperSpeed Device Capabilities Descriptor field
  - Indicates whether device supports LTM
- LTM_ENABLE feature selector
  - Enables device to send Latency Tolerance Messaging packets
- SET_SEL standard device request – four values:
  - U1SEL – U1 System Exit Latency for use with LTM
  - U2SEL – U2 System Exit Latency for use with LTM
  - U1PEL – U1 device-to-host Path Exit Latency for use with link PM
    - Exit latency for entire path of links between device and host, when exit is initiated by the device
  - U2PEL – U2 device-to-host Path Exit Latency for use with link PM

PM Software Interface Detail (3)

- PORT_U1_TIMEOUT feature selector
  - Set U1 link inactivity timeout, and enable / disable U1
  - 0x00: U1 disabled
    - Port does not initiate U1, and rejects all U1 entry requests
  - 0xFF: U1 enabled, timer disabled
    - Port does not initiate U1, but accepts U1 entry requests (unless it has a packet to transmit)
  - 0x01 – 0xFE: U1 enabled, timer enabled
    - Port initiates U1 upon inactivity timer timeout, and accepts U1 entry requests (unless it has a packet to transmit)
    - Inactivity timeout is value of feature selector (MSB of wIndex), in units of μs (up to 127 μs)

PM Software Interface Detail (4)

- PORT_U2_TIMEOUT feature selector
  - Set U2 link inactivity timeout, and enable / disable U2
  - 0x00: U2 disabled
    - Port does not initiate U2 from U0, does not transition from U1 to U2, and rejects all U2 entry requests from U0
  - 0xFF: U2 enabled, timer disabled
    - Port does not initiate U2 from U0, does not transition from U1 to U2, but accepts U2 entry requests from U0 (unless it has a packet to transmit)
  - 0x01 – 0xFE: U2 enabled, timer enabled
    - If PORT_U1_TIMEOUT = 0x00
      - Port initiates U2 from U0 upon inactivity timer timeout, and accepts U2 entry requests from U0 (unless it has a packet to transmit)
    - If PORT_U1_TIMEOUT > 0x00
      - Port transitions from U1 to U2 upon inactivity timer timeout, and accepts U2 entry requests from U0 (unless it has a packet to transmit)
    - Inactivity timeout is value of feature selector (MSB of wIndex), in units of 256 µs (up to 65 ms)

PM Software Interface Detail (5)

- PORT_REMOTE_WAKE_MASK feature selector
  - Three bit fields:
    - Bit 0 – enable remote wake due to connect events
    - Bit 1 – enable remote wake due to disconnect events
    - Bit 2 – enable remote wake due to over-current events
- PORT_LINK_STATE feature selector
  - Used for U3 entry and exit
  - Can also be used for test purposes to initiate U1 and U2 entry
- C_PORT_LINK_STATE feature selector
  - Port status change event due to a change in PORT_LINK_STATE
  - A ClearPortFeature(C_PORT_LINK_STATE) clears this port status change event bit
  - Primary PM purpose is to indicate completion of a transition from U3 to U0 due to a SetPortFeature(PORT_LINK_STATE U0) request
- PORT_LINK_STATE GET_STATUS field
  - Returns link state of port
- C_PORT_LINK_STATE GET_STATUS field
  - Returns state of port status change bit

PM Software Interface Detail (6)

- HubHdrDecLat SuperSpeed Hub Class Descriptor field
  - Hub Packet Header Decode Latency
  - Worst case latency for hubs whose upstream link is in U0 to decode the header of a downstream flowing TP or DP packet and initiate a transition to U0 on the relevant downstream port
- FORCE_LINKPM_ACCEPT feature selector
  - Used for test purposes, to force a port to always accept U1 and U2 entry requests from a link partner