USB 3.0: Delivering SuperSpeed with 25% Lower Power

Power Management at the USB 3.0 Link Layer

The desire to extend battery life in the fast growing mobile computing market has placed a new spotlight on power management within portable systems. Developers of laptops, netbooks, smart phones, and tablets now scrutinize every amp of power usage at the system level in their drive for better power efficiency. The introduction of USB 3.0 brings new opportunities to boost battery life for both host and endpoint functions thanks to comprehensive power management features that operate autonomously at the hardware level.

Designed to overcome the drawbacks of the Advanced Power Management (APM) model, the Advanced Configuration and Power Interface, or ACPI, was introduced in1997. The specification brings some level of power awareness to the BIOS, system hardware and software. ACPI relies on tables in the BIOS to define the power modes for individual peripherals. The OS then uses these definitions to decide when to switch a device, or the entire system, from one power state to another. USB 2.0 has supported this software-based approach relying on suspend-resume commands to place the USB bus in a power reduced state. However, these ACPI based implementations have been plagued by stability and latency issues.

Implementing an effective power management policy for interfaces like USB presents additional challenges. USB is one of the only peripheral busses that allow different types of devices with varying usage frequencies to attach simultaneously. Many of these USB devices experience extended periods of idle. In addition, developers must contend with the growing popularity of devices that draw power or recharge batteries over USB.

The USB 2.0 power management model was enhanced with the introduction of Link Power Management (LPM) in the EHCI specification 1.1. The new LPM transaction is similar to the existing USB 2.0 suspend/resume capability however; it defines a mechanism for faster transition of a root port from an enabled state (L0) to a new sleep state (L1). Implementing LPM requires changes at both the chip and software layers which have slowed market adoption. The table below outlines the LPM entry and exit timing windows:

	Entry	Exit
L1 Sleep	host-initiated via LPM extended transaction Entry: ~10us	Device or host-initiated via resume signaling; Remote-wake can be (optionally) enabled/disabled via software Exit latency: ~70 us to 1ms (host-specific).
L2 Suspend	Implicitly entered after 3ms of link inactivity	Device- or host-initiated via resume signaling; (OS-dependent)

USB 2.0 Link Power Management (LPM) States

USB 3.0: Designed for Power Efficiency:

Recognizing that continued adoption of USB will require improved power efficiency, the USB-IF has made power management a cornerstone to its next generation interface: SuperSpeed USB. For backwards compatibility, USB 3.0 devices are required to support both 2.0 and 3.0 link speeds. USB 3.0 devices will maintain separate controllers and physical layers for High/Full speed and SuperSpeed links. To ensure power savings gained while operating in USB 3.0 mode are not lost when 3.0 hosts are connected to legacy 2.0 devices, all USB 3.0 ports (host & device) are now required to support the LPM feature above when operating at High/Full speed. Correct power management operation in both legacy USB 2.0 mode as well as SuperSpeed mode will be verified during USB 3.0 logo certification.

SuperSpeed USB uses dual simplex differential signaling operating at 5 GHz frequency to provide a 10x performance increase over High-Speed USB. The higher power required to drive the 5 GHz signaling in SuperSpeed mode is more than offset by the improved efficiency of 3.0 data transfers. The USB-IF estimates the system power necessary to complete a 20MB SuperSpeed data transfer will be 25% lower when compared to high-speed mode. This is possible because several architectural issues that hampered USB 2.0 power efficiency have been enhanced in the USB 3.0 specification below.

- Elimination of device polling by allowing devices to asynchronously signal when they need service from the host
- The ability for device ports to initiate low power states
- The ability for device ports to remove power from all or portions of their circuitry (function level suspend)
- The ability to use data streaming for bulk transfers
- More efficient token/data/handshake sequence
- The addition of packet routing eliminates the need to broadcast packets to all endpoints downstream from hubs

In addition to these changes, USB 3.0 improves efficiency by implementing power management at the link layer to provide greater speed and precision in managing power consumption.

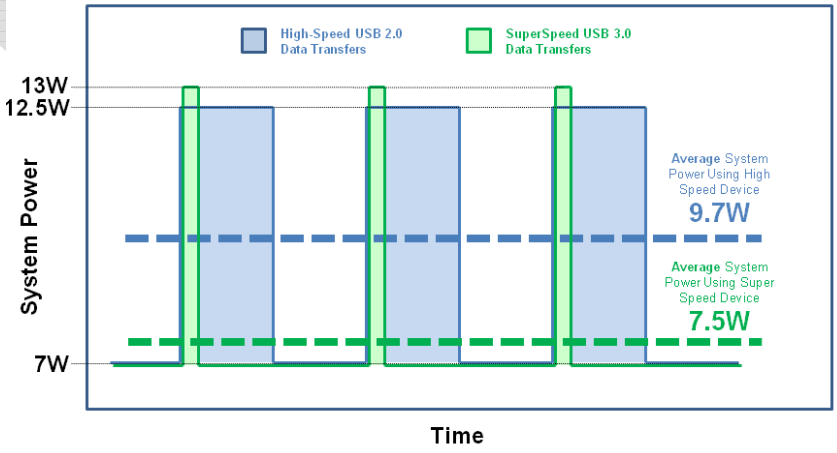

25% less system power is used during a SuperSpeed 20Mbyte data transfer compared to high-speed

The table below outlines the four power states in USB 3.0. Each state incrementally lowers power usage while increasing the allowed exit latency. This provides a more adaptive power management model that uses timers and link state awareness to reduce power usage. While the specifics of how devices will lower their power draw are left to the vendor, the table below outlines the link states defined by the USB 3.0 specification:

Link State	Description	Key Characteristics	Exit Latency
U0	**Link Active**		**NA**
U1	**Link Idle, Fast Exit**	**RX & TX Circuit Quiesced**	**µs range**
U2	**Link idle, Slow Exit**	**Clock Generation Circuit also Quiesced**	**Low ms range**
U3	**Suspend**	**Portions of device power removed**	**Higher ms range**

Logical Save Link States defined in USB 3.0

Most early 3.0 devices rely on inactivity timers to initiate entry into the U1 state. In the U1 state, these devices will typically reduce power to their SuperSpeed PHY. These devices will progressively lower power to other parts of the interface as the inactively period increases. In some cases, host ports will immediately request transition to the most aggressive power suspend state (U3) during idle periods. This more rigid approach to lowering power draw is generally initiated by higher layers and is based on expected usage patterns for specific device classes. USB 3.0 also preserves function suspend features from USB 2.0 allowing individual functions to be placed into a lower power state. The remainder of this article explores the SuperSpeed power management model and the power state transitions required by the USB 3.0 specification.

Configuring USB Devices for Power Management:

There are four steps involved in configuring a USB 3.0 device for power management.

1. **DEVICES MUST REPORT THEIR LEVEL OF SUPPORT FOR POWER MANAGEMENT WITHIN THEIR ENDPOINT DESCRIPTORS**

   While it is required for all devices to support power management to gain SuperSpeed certification, USB developers may elect to configure devices with this functionality disabled for specific applications.

2. **HOST MUST SEND SET_FEATURE TO U1/U2_ENABLE DURING CONFIGURATION**

   Alternatively, some peripheral devices that are used intermittently may aggressively direct their own links to the lower power state. Higher layers require a mechanism to enable (or

disable) the upstream port's ability to request low power entry. When asserted, U1/U2_ENABLE allows the upstream port to initiate entry to U1/U2.

3. Host must send Link Management Packet (LMP) to define the U1/U2 Inactivity Timeout

The U1/U2 inactivity timers allow the host to define the time interval between the U0 > U1 and the U1 > U2 power state transitions. These timers provide the flexibility to delay power state transitions for specific applications, such as Blu-Ray disk writers, that could suffer usability problems if response latency is introduced. The U1 and U2 inactivity timeout can be as long as 127 µs and 65 ms respectively. Sending an LMP with the U1 inactivity timeout value between the range 0x01-0xFE also serves to implicitly enable the host port to initiate U1/U2 transitions.

4. Host will inform the device of the U1/U2 System Exit Latency using SET_SEL

Reporting System Exit Latency (SEL) allows the host to more intelligently manage power state transitions for periodic endpoints, such as isochronous devices. SEL represents the total latency to transition the entire path of links between the device and host from U1/U2 back to U0. It provides a mechanism for higher layers to reduce or even disable U1/U2 entry if system exit latency exceeds the minimum service intervals reported by the device.

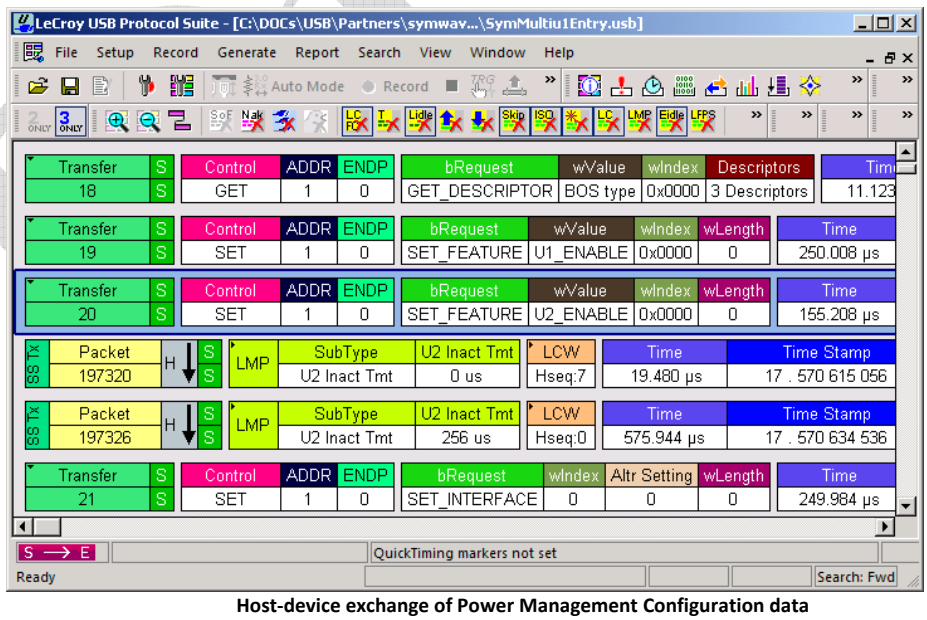

Host-device exchange of Power Management Configuration data

Transitioning from U0 >U1

Either link partner can initiate a transition from U0 >U1 based on the expiration of the PORT_U1_TIMEOUT timer. Alternatively, some devices may attempt to save power by proactively initiating U1 mode more aggressively by setting their U1_Enable feature selector and reporting their U1 Inactivity Timeout equal to 0.

Initial entry into a low power state is always negotiated between ports using the LGO_U*n* followed by LAU (accept) or LXU (reject). The port sending the LAU should wait until it receives a single LPMA (accept response) which serves as a final handshake before transitioning to any of the low power states. To maximize power savings, ports are required to respond to power management commands within the PM_LC_TIMER timeout. If the port initiating the state change does not receive an LAU or LXU before the PM_LC_TIMER expires (3us), it's considered a link error and should initiate recovery.

Alternatively, if after sending the LAU, the device does not receive the LPMA or any other valid traffic (TS1, LFPS or Link Command, etc…) before the PM_ENTRY_TIMER expires (6us), it should proceed to the low power state anyway. In this event, it is assumed the LPMA was corrupted and the port issuing the LGO_U1 has already entered U1.

Transitioning from U1 > U2

The transition from U1 to U2 is generally triggered by a second timer called the U2_Inactivity_Timer which, when enabled, will silently move the link to the lower power U2 state. This U2 inactivity time out value is reported by the endpoint's configuration descriptor. It's the host responsibility to enable this timer using the U2 Inactivity Timeout LMP. When a link enters U1, this starts the U2 inactivity timer and provides a mechanism for the port to autonomously move to the U2 state.

For some devices, it may not be practical for individual endpoints to enter U1 (ie: composite devices that may have a shared PLL). Some devices may bypass the U1 mode altogether and instead transition the link from U0 directly to U2 using the LGO_U2 link command thus allowing a larger portion of the SuperSpeed interface to be suspended. A device can be configured to support U2 exclusively with SET_FEATURE: U1_DISABLE.

As mentioned previously, some devices may attempt to save more power by immediately transitioning to U1 or U2, using the U1/U2_Enable feature selector. For example, storage devices may immediately issue an LGO_U2 after each transfer if the packets pending bit is de-asserted in the previous transaction packet.

Transitioning from U0 > U3

The U3 state is a deep power saving state where interface power may be removed. It's the equivalent of Suspend state in USB 2.0 and it can only be initiated by a downstream facing port using the LGO_U3 followed by LAU (accept). Upstream facing ports are not allowed to reject the LGO_U3. While the goal is to conserve as much power as possible, while in U3, a port must still maintain its Warm Reset detect, U3

wake detect, (for host initiated wakeup) as well as wake transmission (for remote_wake capable devices).

Transitioning from U1/U2 >U0

Returning a link from U1 to U0 active state mandates the shortest recovery time in the range of 10us. This transition is normally initiated when a packet needs to be transmitted, such as an IN message from the host, or an ERDY message from the device. Ports in lower power states need a mechanism to signal its link partner to begin the link recovery process. Low Frequency Periodic Signaling (LFPS) is a 50 MHz side-band signal that provides a port with a low-power mechanism to send a "wake signal" to a link partner. Both sides must receive an LFPS "handshake" to avoid entering the Recovery link state before the far-end receiver is ready.

To deliver acceptable performance, SuperSpeed devices utilize a low-latency recovery sequence that provides a streamlined way to re-train links when exiting these low power conditions. SuperSpeed ports may also enter the Recovery state when errors are detected during data transfers. In both cases, only TS1 and TS2 ordered sets are exchanged with the goal of returning the link to U0 as quickly as possible.

Resolving conflicts between power management commands

There are numerous rules and conditions defined in the USB 3.0 specification to preserve the integrity of the link during power state changes. This includes obvious requirements such as disallowing devices from starting low power transitions unless they have transmitted and received all pending data packets, acknowledgements, flow control link commands, header and buffer credit advertisements. There are also rules to ensure links maintain coherency in the event an expected power management response is not received. For example, a port that sends U1 or U2 exit signal but does not receive an LFPS handshake from its link partner should transition to the SS.disabled state (assumes the sleeping device is removed from the system). Because power state changes can be initiated by both host and peripheral device ports, there are several rules designed to manage link state race conditions and potential conflicts between ports. For example, peripheral devices that have sent an LGO_U1 or LGO_U2 and also received an LGO_U3, should wait until they receive an LXU from the host and then send an LAU accept for the U3 request. In the case of a host port that has been directed (by a higher layer) to initiate a transition to U3 while a transition to U1 or U2 has been initiated but not yet completed, the host port should complete the in-process transition to U1 or U2, then immediately return to U0 and request entry to U3.

Test and verification of USB 3.0 power management

To ensure USB 3.0 devices properly implement these power management behaviors, they will be verified during the USB-IFs SuperSpeed certification program. Testing devices to ensure reliable operation in power managed environments raises a substantial verification challenge. Post-

silicon functional test teams may struggle to simply initiate power management transitions as the necessary commands occur at the lowest layers making them difficult to control using software. Entrance and exit from these low power states must occur within rigid pre-defined time limits. This task is greatly simplified by protocol layer test tools that have the following capabilities:

- **Low Level Link Tester –** To test many of the link states outlined above requires special test systems that can control and manipulate the logical link layer. Most functional test teams rely on traffic generators that emulate real device behaviors to perform this testing. Verification systems, such as LeCroy's Voyager USB tester platform provide fine grained control of link layer handshaking. These tools should be capable of creating intentional timing violations and invalid state transitions to test error recovery on the device-under-test.
- **Accurate capture of U1 recovery sequence** - The SuperSpeed transition from U1 to the active state (Ux_EXIT_TIMER) mandates both ports should enter U0 within 6ms or the link will enter SS.disabled. Unlike Power-on link training, recovery from U1 uses a fast link training sequence without the added TS.EQ symbols. This frequent re-training places considerable pressure on the analyzer front-end as it must seamlessly capture the LFPS handshaking followed by 5 GHz signal lock and the U1 recovery sequence.
- **Triggering on power link state changes** – Traffic at the logical link layer is invisible to the upper layers of USB 3.0 protocol making it impossible to see Link Commands using software based tools. This mandates using an inline protocol analyzer capable of accurately time stamping link layer traffic between devices. Triggering on link commands such as the LGO/LAU exchange and the LFPS wake signals are critical for efficient power management debug.
- **Triggering on power management timeouts** - Returning to U0 from the U1 low power state has proven to be a common problem area for early devices. This transition in particular can occur hundreds of times in only a few seconds. To minimize latency at the application layer, devices are required to enter and exit power save modes within very short timing windows. For example, during the low power exit sequence, both link partners must exchange an LFPS exit handshake within 2ms (tNoLFPSResponseTimeout) . If either side fails to send the required response, the opposite link will go to SS.disabled and reverts to USB 2.0 mode. Test systems, such as LeCroy's USB Voyager analyzer, provide independent event timers that can trigger when either a handshake or the required state change is late. This offers firmware engineers a mechanism to capture rare or intermittent timing violations during power management transitions.

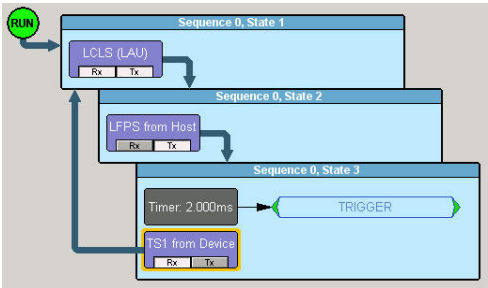

LeCroy's USB 3.0 analyzer provides sequential state detection with independent timers in each state for identifying timing violations

- **Monitoring v_{BUS} power draw** - V_{BUS} power supplied by the downstream facing port can represent a significant source of battery drain for mobile platforms. LeCroy's Voyager analyzers are available with a special multimeter option that can correlate actual V_{BUS} power usage with protocol layer state changes. The Voyager M3i Power Tracker™ option monitors V_{BUS} power draw and displays voltage graphically in a time line format. This power information is synchronized to the trace allowing users to correlate power usage at the electrical layers with commands occurring at the link layer.

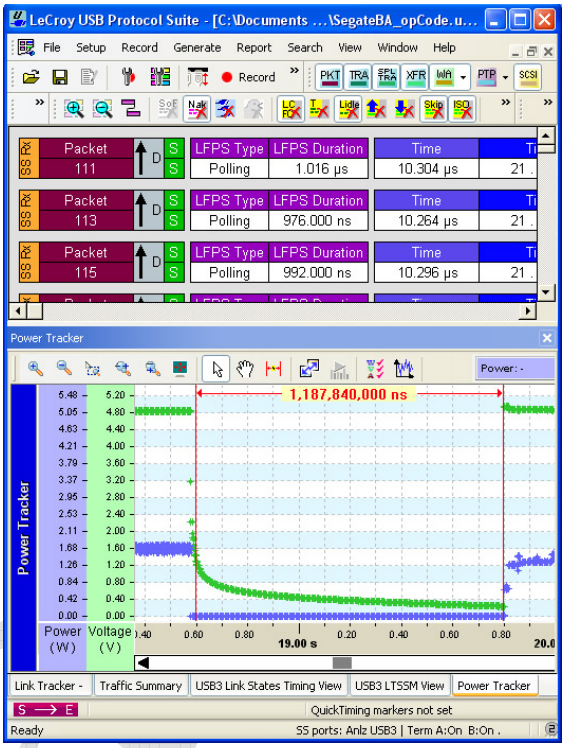

Power Tracker™ monitors V_{BUS} power draw and displays voltage graphically in a time line format

About the author:

Mike Micheletti is the senior product marketing manager at LeCroy with over 10 years of experience defining high speed serial data acquisition solutions for USB, SAS, SATA and Fibre Channel. Mr. Micheletti is regular contributor to the USB-IF Compliance Working Group.

Sources:

Ethier, Sheridan, "Application-Driven Power Management", 2004, QNX Software Systems Ltd.

"Universal Serial Bus 3.0 Specification", USB Implementers Forum Inc., 2008, http://www.usb.org

Walsh, James; SuperSpeed USB Power Management, 2008, http://www.usb.org