

Overcoming the High Frequency Design Challenges of USB 3.0

By Joseph Spisak
Pericom

New digital interface standards require special attention to high speed/ high frequency performance in order to obtain the required throughput and maintain high reliability

The beauty—and primary market appeal—of USB is its simplicity. USB has an installed base of more than 6 billion devices, making it the most successful interface among PC and peripheral

devices. Part of the reason USB is nearly ubiquitous among these applications arises from the trust held by users that they can plug in devices using whatever cable is at hand. USB's relative low cost is also responsible for its unprecedented success.

Understandably, consumers have a fairly good idea of what to expect from USB 3.0 as it becomes widely available: low cost, simplicity, and reliable operation. Likewise, many manufacturers are confident that they can migrate their designs from USB 2.0 (High-Speed) to USB 3.0 (SuperSpeed) as easily they did from USB 1.0 (Full-Speed) to USB 2.0. However, SuperSpeed USB is no simple upgrade. The typical roadmap for most interfaces is to double their data rate, but the new USB 3.0 specification implements a full 10x increase in bandwidth. Moving to a 5 Gbps signaling rate introduces a wide range of signaling issues many developers have never had to manage before. With an actual signal frequency of 2.5 GHz—the 5 Gbps data rate is achieved by employing double data rate techniques—the overall system is far more sensitive to attenuation and jitter. In addition, USB 3.0 is no longer a simple bus-based architecture using host directed transactions. The new SuperSpeed spec utilizes a much more complex bi-directional and port-directed packe-

tized environment. Further complicating design is the fact that the spec requires backwards compatibility with the USB 2.0 connector, which was designed for optimal operation at much lower data rates.

There is already a strong market need and demand for USB 3.0. Today, High-Speed USB running at 480 Mbps is acceptable for many applications, but the increasing storage capacity of portable devices coupled with the rising availability of high definition video and desire for faster transfers makes SuperSpeed USB a powerful differentiating feature for manufacturers. In order to meet the reliability expectations of the market, developers will have to approach USB 3.0-based designs with the same care and attention to detail necessary for the architecture of any high-frequency system. Doing so, while controlling costs and maintaining ease-of-use, requires careful management of attenuation and jitter while driving high-speed signals through varying lengths of cables, across multiple connectors, and over extended board traces.

Consumers Don't Care About Signal Margin

Signal losses at the higher frequency of USB 3.0 will be more pronounced than many designers familiar with USB 2.0 expect. USB poses a fairly challenging environment, as can be illustrated with a typical real-world application connecting a laptop to a peripheral (see Figure 1). The distance from the notebook's USB controller to the connector is commonly on the order of 10 inches. Next, as per the USB 3.0 spec, signals will need to be able to travel over up to 3 meters of cable before crossing over another connector and more board traces

USB 3.0 DESIGN

to the peripheral controller. And, until USB 3.0 is integrated on notebook chipsets, an external host controller will also be required, potentially introducing additional signal losses.

Compliance testing measures signals after they have been equalized at the receiver, and the combination of trace, connector, and cabling losses quickly erodes signal quality. Keeping the signal eye sufficiently open will require careful budgeting of signal margin since the total distance over which high-frequency signals must be able to run begins to push the limits imposed by the USB 3.0 specification.

Much of the signal loss occurs in the cable. Signal integrity can be maintained by using high-quality, thick gauge cables. The difficulty here is that such cabling is significantly more expensive than the low-quality unshielded twisted pair (UTP) cabling typically associated with consumer applications. OEMs do have the option of requiring that consumers use high-quality cables, but this approach poses several challenges. First, bundling high-quality

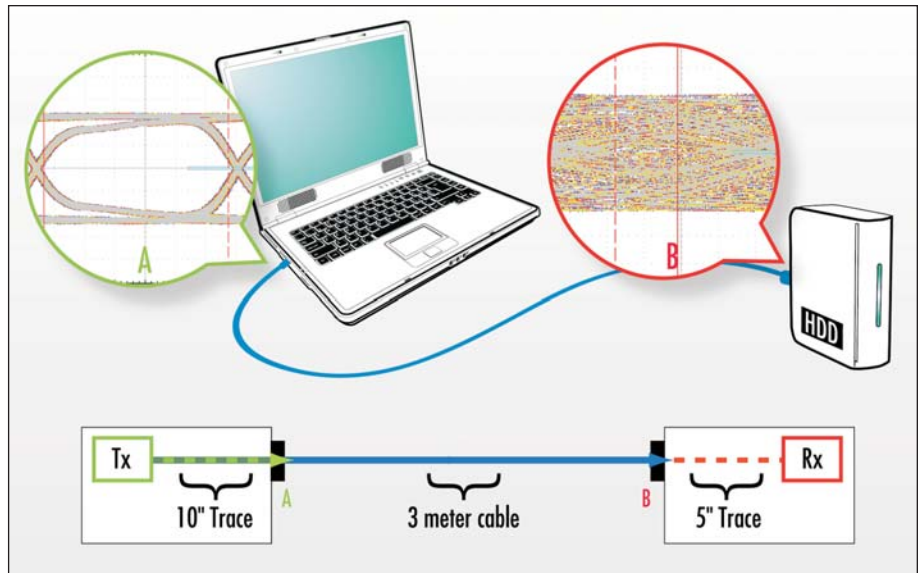


Figure 1 · USB 3.0 signals must not only run reliably over varying lengths of cable, but also across multiple connectors and traces between the transmitter and receiver.

cables with a product can substantially increase the bill of materials for a product, making its price less attractive and less competitive when compared to products that don't bundle cables. However, even if a high-quality cable is included with a product, consumers are used to mixing

their USB cables together and may inadvertently use a lower quality cable provided with another device.

Alternatively, OEMs can employ product labeling to signify that a specific type of cable be used with a product. Consumers, however, have never had to worry about cable quality

USB Encroaches on PCI Express

Engineers who have worked with PCI Express 2.0 will find the migration to USB 3.0 to be more straightforward than it might otherwise be. While there are differences between the two standards, USB 3.0 and PCI Express 2.0 have equivalent data rates and use similar transmit and receive blocks including scrambling, 8b/10b encoding and serialization/de-serialization of data, while also sharing many characteristics at the physical layer.

One consequence of the similarity between USB 3.0 and PCI Express 2.0 is the encroachment of USB 3.0 into a number of PCI Express 2.0 applications. In a docking station, for example, using USB to connect to the notebook rather than PCI Express 2.0 eliminates the need to bridge USB traffic. In effect, the docking station effectively becomes a USB hub that fans out the signal, thus simplifying the overall architecture while lowering system cost.

The extensive penetration of USB into consumer

electronics devices also gives platforms using USB 3.0 a tremendous advantage in the ability to immediately connect to the installed base of USB 2.0 devices already deployed. Certainly, PCI Express 2.0 will hold its own in applications where it is already the explicit off-chip interface, including internal bus interfaces from processors and in server, storage, and various embedded applications. However, this incumbent advantage only holds so long as PCI Express 2.0 remains the primary integrated interface.

As USB 3.0 is integrated onto chip-sets and processors, it will compete with other interfaces as well. With its high bandwidth and low cost—given its tremendous market volumes—USB 3.0 can be expected to compete with display interfaces such as HDMI, DisplayPort, PCI Express, and DVI. Effectively, USB 3.0 can be expected to markedly impact the market share of other protocols in any application where there is already at least one USB interface.

when using Full-Speed and High-Speed USB. Part of the appeal of USB is its ease-of-use and its perceived guaranteed interoperability. To ensure a quality end-user experience, OEMs of SuperSpeed devices will want to avoid forcing users to have to think about matching hosts, cables, and endpoints.

The crucial concern for OEMs is that if a low-quality cable is used and performance is degraded, it is the device that consumers will assume is at fault, not the cable. Such a perception could result in a higher rate of USB 3.0-based product returns, substantially eroding profit margin. It could also affect the overall adoption rate of USB 3.0 as an emerging technology, if the average user perceives that getting USB 3.0 to work reliably is difficult, market uptake will be adversely affected as well.

Accommodating Cabling Losses Through Signal Conditioning

Given that OEMs cannot control the type or length of cable that consumers will use with their devices, engineers must assume worst case losses. Cable losses, however, are fairly deterministic, and reductions in signal quality through noise and attenuation can be restored through signal conditioning technology.

Signal conditioning devices—often referred to as redrivers since the signal is restored and “redriven” across the channel—provide the clean passage of signals across longer distances and additional connectors through the use of emphasis and equalization techniques that reduce signal attenuation and jitter. Redrivers are used across a wide variety of high-speed interfaces and protocols to effectively open a closed eye so that data can be recovered at the receiver with an overall effect similar to breaking the signal path into several legs and restoring signal quality at each leg (see Figure 2).

Given the dynamic nature of the typical PC-peripheral environment—

devices can be removed or attached at any time with varying cables lengths—the characteristics of the overall USB signal path can change over time. Over- and under-equalization can become a problem since the equalization ideal for a 3-meter cable will compromise signal integrity when applied to a thumb drive with

an effective cable length of zero meters. Adaptive signal conditioning is required to continuously monitor and tune the channel to ensure that the redriver is adjusted to optimally match the actual signal path. However, for stable applications where the signal path does not change (i.e., a storage server environ-

ment where cables are never removed) continuous retraining is unnecessary and may even negatively impact performance. In these environments, configurable equalization that is set to match the channel is more appropriate.

Redrivers are introduced between the transmitter and receiver so that

incoming signals can be restored, adjusted, and retransmitted, resulting in greater signal margin. A single redriver can add up to 11 dB of gain, thus significantly increasing the ability of a system to accommodate losses through SDP cables. This added signal margin will ensure that devices are able to reliably transmit data

over even the lowest quality cables.

Typically, redrivers are placed before the receiver to open the signal eye. In an implementation such as an internal bus where the engineer controls both the transmitter and receiver, receive-side signal conditioning is quite effective. With USB, however, engineers have to create products that connect to devices they did not design.

Given the reality that consumers expect USB 3.0 to be available at near-USB 2.0 prices, there is considerable pressure on manufacturers to lower costs in any way possible, most notably leaving signal conditioning as a concern for the receiving device. The result will be the availability of inexpensive hubs and peripherals that do not have signal conditioning, which have passed compliance testing using high quality cables, but perform marginally when connected with the lower quality cables that the average consumer is likely to have on hand.

The challenge for manufacturers designing reliable USB-based products is that their products will be expected to work reliably when connected to these marginal devices. Consider a hard drive with receive-side signal conditioning. While the drive will receive data accurately, lack of transmit-side conditioning will result in poor signal quality from the drive as well as frequent errors and lower throughput. Users will perceive that the problem is with the drive, not the low quality cables or cheap hub that seem to work well enough with other peripherals. The result will be a returned drive and damage to the manufacturer's brand.

In order to ensure reliability, signals need to be conditioned both as they are transmitted and received. By properly restoring signals, both coming and going, devices will be able to interoperate with any other device across any cable length, even with those devices that do not condition signals on the receive side.




Figure 2 · Redrivers with emphasis and equalization signal conditioning technology ensure the integrity of high-frequency USB 3.0 signals by opening closed signal eyes to recover data and meet strict compliance testing requirements. Increased signal margin also supports longer drive lengths over even low-quality cables.

Lowering System Cost Through Higher Signal Margin

The extra signal margin gained through signal conditioning not only ensures interoperability with other devices, it also translates directly to greater flexibility and lower cost. Specifically, signal conditioning allows engineers to run signals longer distances and/or have more flexibility in how signals are routed across a printed circuit board (PCB). Such margin can enable engineers to achieve sufficient signal and ground isolation with fewer board layers, thus reducing system complexity and manufacturing cost. Improved signal quality also increases reliability by lowering the overall bit error rate (BER). This leads to higher system performance and an increase in effective throughput by reducing the number of transmission errors and retransmissions required.

The combination of emphasis, equalization, and increased drive also supports the use of longer, thinner, and lower cost cables while maintaining reliability and throughput. Thinner cables—which tend to be lower in cost, have greater flexibility

and are typically more aesthetically pleasing compared to thicker gauge cables—are the likely choice of consumers.

Redrivers also provide protection against electrostatic discharge (ESD), such as when a user accidentally shocks a device. Consider that the cost of an interface is lowest when the controller is integrated onto the primary processor or chipset. When this controller is directly connected to the port connector, not only is the controller vulnerable to an ESD event, the entire chipset/processor is at risk, potentially rendering the entire system inoperable. Because a redriver sits between the port connector and controller, the redriver effectively isolates the chipset/processor. In the case of an ESD event, the redriver will protect the USB controller and chipset/processor with a worst-case outcome of losing the single USB port while leaving the rest of the system otherwise operational.

Clocking and Switching Integrity Issues

Clocks play an important role in

high-frequency systems, for as signal rate increases, jitter budgets drop. Given the expense of high-frequency clocks, it is more cost-effective for engineers to multiply the clock signal from a lower frequency clock to supply the appropriate high-frequency clock required for USB. However, when a low-frequency clock is multiplied into a high-frequency signal, the jitter will be multiplied by an equal factor as well. The challenge for engineers is to balance induced jitter with input frequency. To minimize cost, USB clock buffers and generators need to introduce very low jitter.

Similarly, USB switches need to provide smooth transitions. Switches are an elegant way to reduce system cost in applications such as docking stations by reducing the number of physical ports a system has while increasing the number of devices that can be connected at any one time. Maintaining low resistance and impedance are the key concerns here to keep insertion losses low and prevent glitches from being reflected back to the transmitter so that signal quality is preserved.

Intelligent Power Management

The USB 3.0 spec introduces several new low power operating modes—including idle, sleep, and suspend—to support lower power consumption and longer battery-life in hosts and endpoint devices. To maximize power savings, signaling circuitry needs to be turned off when not in use.

USB 3.0 uses a bi-directional, differential interface so the transmit and receive channels can be powered down individually. Transmit channel power management is fairly straightforward: sleep when there is no data currently being sent. USB 3.0's hot-pluggable capabilities make receive-side power management somewhat more involved since data can arrive at any time. In addition, the front-end signal detect mechanisms in redrivers needed to support and

maintain transparency are complex (i.e., redrivers are not USB endpoints and so cannot terminate signals but must pass them through).

First, the system determines whether a device is plugged in or not. Since USB receivers have a single-ended 50-ohm termination, if no termination is present, no endpoint is present either. In this case, the redriver can shut down and wake infrequently to poll the port to see if a device has been plugged in. Power efficiency and wake up latency is determined by how frequently the interface is polled.

When a device is plugged in, the redriver can check if the link is in electrical idle or transmitting data by measuring the level of the signal detector. When this level drops below a set threshold, there is no signal being sent and the redriver can drop into deep sleep or slumber, depending upon the application. Typically this threshold is set to 100 mV differential, although engineers may want to adjust the threshold to increase sensitivity for low signals sent over longer cables. In many applications the USB port connector may be relatively far from the actual controller; for example, in a digital TV, the circuitry may be centered in the TV while the port connector is located some distance away on the side of the screen.

Receive sensitivity also plays an important role in restoring a closed signal eye. While engineers know the drive capability of the transmitter in their device, they typically do not know the characteristics of the far-end transmitter. Lowering the electrical idle threshold is an excellent way to increase the sensitivity of the receiver. A lower threshold is also often desirable when USB signal traces may be relatively long, such as in a storage server.

The SuperSpeed 3.0 spec promises to bring significant change, both to how consumers use and access digital data as well as to how engineers design USB into the future (see the sidebar on page 28, "USB Encroaches on PCI Express"). The average consumer will expect USB to perform as it always has: reliably, easily, and inexpensively. Maintaining signal integrity and system performance across any quality cable and with marginal devices requires signal conditioning to compensate for

REMEMBER!

All past articles in High Frequency Electronics are available in the 'Archives' section of our web site:

www.highfrequencyelectronics.com

Articles are placed in the Archives one month after publication. To read current articles electronically, simply download the present month's Online Edition.

board, connector, and especially cable losses. While consumer electronics devices cannot afford to bear the price of high-frequency components and quality cables that maintain signal integrity, they also cannot afford to allow signal quality—and thus application reliability and performance—to suffer.

Using redrivers to restore signal quality at both the receiver and transmitter using emphasis and equalization signal conditioning techniques enables engineers to increase signal margin. Not only does this permit more flexibility in routing while allowing signals to run over longer distances, it reduces overall system cost while increasing reliability. In this way, engineers can ensure not only that their USB 3.0-based products will pass strict compliance testing but also reliably interoperate with all other USB devices over even the lowest quality cables.

Author Information

Joseph Spisak is a product marketing manager for Pericom's PCI Express and USB3.0 products. He holds a B.S. in Electrical Engineering from Michigan State University and an MBA and MS in Finance from the University of Denver. He can be contacted at jspisak@pericom.com.

Additional Information

The USB 3.0 specification was developed by the USB Implementers Forum, Inc., (www.usb.org) and was released in November 2008.

An information package for developers is available for interested designers:

“Universal Serial Bus Revision 3.0 Specification (.zip file format, size 3.80 MB) provides the technical details to understand USB 3.0 requirements and design USB 3.0 compatible products. Modifications to the USB 3.0 specification are made through Engineering Change Notices (ECNs). Enclosed in this zip file are the following documents:

- The USB 3.0 Specification released on November 12, 2008
 - USB 3.0 Adopters Agreement
 - New Addition of Link Command LDN Engineering Change Notice as of April 4, 2009
 - USB 3.0 Standard-B and Standard-B Crosstalk Engineering Change Notice as of April 4, 2009
 - Reset Propagation Engineering Change Notice as of May 8, 2009
 - Reset Propagation Figure
 - Q1 09 USB 3.0 Errata as of May 15, 2009”
- <http://www.usb.org/developers/docs/>