



USB 3.0: A SIMPLE IDEA FULL OF CHALLENGES

COMBINING 5 GBPS WITH THE CONVENIENCE OF USB SOUNDS LIKE A SURE WIN, BUT MANY ISSUES ARE HIDING BEHIND THE PREMISE.

BY RON WILSON • EXECUTIVE EDITOR

Super-speed USB (Universal Serial Bus) 3.0 sounds like a great idea. Just start with widely used, fast, and bulletproof USB 2.0 and graft in the PHY (physical-layer) interface from another common and reliable standard, PCIe (peripheral-component-interconnect express) Generation 2. Put two differential pairs into the USB connector to carry the high-speed serial signals from the Generation 2 PHY, and you have a rugged, flexible, inexpensive interface that can operate at 5 Mbps over consumer-priced cables and connectors with interfaces cheap enough to drop into a flash drive.

The idea promises to unleash new ways of using PCs with mobile devices and with storage. With application-level throughput approaching 400 Mbytes/sec and the ability to simply plug anything from a flash drive to 3m of USB cable into a host, users could link PCs and netbooks, quickly dump the contents of huge flash drives, or easily transfer HD (high-definition) video between devices. They could even create their own external storage networks (Figure 1). This promise of speed and flexibility, however, carries the seeds of a difficult challenge for chip, board, and system designers.

“The concept of simply using a PCIe Gen 2 PHY in a USB controller is appealing,” observes Mike Pennell, vice president of engineering at fabless-semiconductor supplier SMSC (Standard Microsystems Corp). “But you have to remember that, although the USB 3.0 PHY is based on the PCIe Gen 2 PHY, the media over which the two controllers must send data are very different. USB lives in a much more challenging world than [does] PCI.”

Navraj Nandra, director of product marketing at Synopsys, explains further: “The similarity between PCIe Gen 2 and USB 3.0 stops at the speed,” he says. “They both run at 5 Gbps. PCIe has to successfully work over only 20 inches of carefully designed PCB [printed-circuit board].”

Pennell says that the connection between a USB 3.0 host controller and a device controller could be far more complex than just using FR (fire-retardant)-4 material. The connection would include several inches of PCB, a connector pair, a short pigtail running to the host’s enclosure, another connector pair at the back of the enclosure, the 3m cable, another connector pair on the back of the target device, another pigtail, and another PCB. All together, those components make up a highly attenuating channel full of opportunities for strong reflections, and it is highly variable.

People don’t often discuss the variability issue in USB, but it becomes hugely important at speeds of 5 Gbps, according to Nandra. “Even at USB 2.0, some certified USB cables are better than others,” he warns. “Even wonderful, \$50 cables can degrade over time. We have already at the relatively low speeds of USB 2.0 seen poor cable performance impact the performance of the PHY. And, at 3.0, the problem will be much worse.”

Cable experts agree that variability is an issue. Peter Smyth, chief executive officer of active-cable vendor RedMere Technology, points out that you can make a USB 3.0 cable good enough to meet the specifications. Doing so also requires tight manufacturing

AT A GLANCE

▣ USB (Universal Serial Bus) 3.0 brings 5-Gbps speed to USB by blending USB with a PHY (physical-layer) interface it derives from PCIe (peripheral-component-interconnect express) Generation 2.

▣ Combining USB with PCIe Generation 2 presents significant design challenges.

▣ System designers may have difficulty distinguishing a great PHY from a poor one.

controls, which cost a lot of money, however. And, even with the best controls and within a production run, cables can vary significantly. Then there is the problem of those pigtails, which no one seems to notice much. “That pigtail that runs from the PCB to the back of the

box is typically just awful,” says Smyth. System designers may find that they have used up most of their eye opening just getting to the back of their own boxes. And, because cables that consistently meet the 3.0 specification will be expensive to manufacture, the door will be open for counterfeit cables. So chip and system designers both must assume the worst—mediocre board design, poor pigtails, cheap connectors, counterfeit cable, and even a wire-bond package for the PHY—when planning their approaches to the new standard.

Unfortunately, successfully pumping 5 Gbps through a messy channel isn’t the only problem. The channel is also highly variable. One minute, a port could have a thumb drive plugged into it, and, the next minute, someone could plug 10 feet of cable into a cage full of disk drives. So the PHY must be flexible. Just to make matters more interesting, USB 3.0 will start out expensive, but, by 2011, it will probably be standard in netbook computers and handheld consumer products, such as cameras, media players, and flash drives. The interface must have a migration path to becoming low cost, meaning that it will require an advanced process node. And—because a flash drive, to name one technology, draws its power from the USB cable itself—the interface must be low enough in power to allow cable-powered operation. The USB 3.0 standard will allow a device to draw as much as 900 mA during operation, but it must draw no more than 150 mA before configuration. That limitation itself demands a well-studied power-management strategy at the chip level, which the system designer must understand in detail. All of these issues represent significant departures from the requirements that spawned PCIe 2.0.

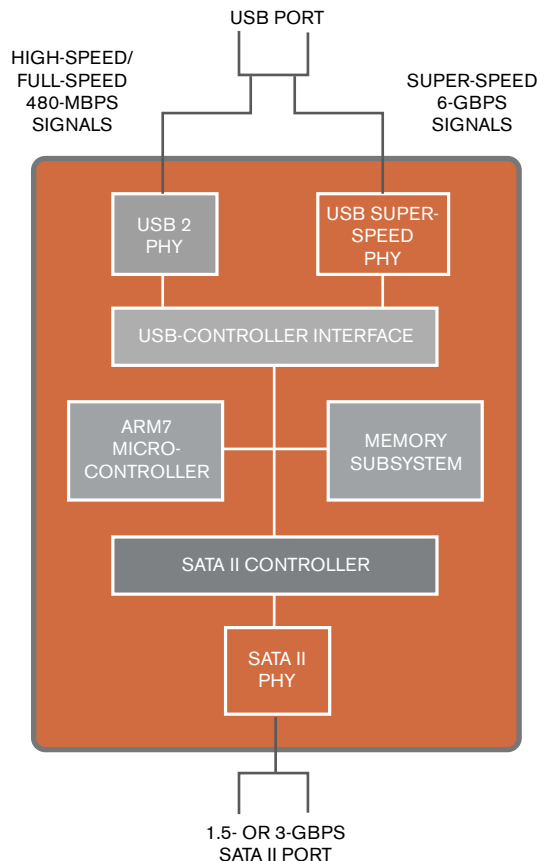


Figure 1 As USB 3.0 interfaces appear in mobile PCs, adapters such as this TI USB-to-SATA bridge will also show up.

CREATING A PHY

Much of the responsibility for getting a data stream through the travails of a USB 3.0 channel will fall on the transmitter pre-emphasis circuit and, especially, the receiver-equalization circuit. Interestingly,



the 3.0 specification appears to leave the design of these two blocks up to the chip-design team. "One of the big differences between PCIe Gen 2 and USB 3.0 is that the PCI document specifies an eye-diagram template at the receiver input," says Synopsys' Nandra. "You have to get the signal there at a specified quality. But, in USB 3.0, the eyes can be so closed by the time the signal gets through the cable that there is no opening to put a template into. So, instead, the 3.0 document

specifies an eye-diagram template at the output of the equalizer, not at the input to the receiver."

One of the greatest differences between PCIe Gen 2 PHYs and USB 3.0 PHYs will be in the receiver-equalization circuit. Many designers expect the quality of this block to be a major differentiator in the market, for PHY-IP (intellectual-property), chips, and the systems that use them. The equalizer must be both powerful in its action and adap-

tive. Otherwise, a PHY would be unable to handle the range of channel conditions that USB can throw at it. As an adaptive equalizer, this circuit will require a training sequence to lock onto. Yet, the equalizer must be low in power and compact to meet the needs of consumer-product applications. Those challenges are formidable.

Also, although PCIe Gen 2 is a fully synchronous interface, Nandra points out, USB 3.0 requires that the PHY use a spread-spectrum clock in a way that makes the transmitter and receiver essentially asynchronous. The receiver CDR (clock- and data-recovery) circuit must recover the transmitting clock without having access to the spread-spectrum signal that modulated it. "That requires a very elastic CDR," observes SMSC's Pennell.

Such issues will make previous expe-

WHY DO WE NEED A 5-GBPS USB?

Why would most USB (Universal Serial Bus) users today even want so much speed in a USB link?

After all, USB is supposed to provide flexibility and convenience for using mice, printers, flash drives, and the like. Consider flash drives, however. Capacity for the largest thumb drives has soared into the gigabytes, so transferring a significant portion of the contents of one of these drives through USB 2.0 can be a painful task. Meanwhile, marketers point to the supposed emergence of the home-media PC, with its enormous data files—think entire HD (high-definition) movies—and network-attached storage. Having a 5-Gbps version of USB would allow most home-media users to simply plug their external drives into their USB ports and not worry about the mysteries of SATA (serial-advanced-technology-attachment) or storage-area networks.

Also, forecasters expect a huge growth over the next few years in netbook computers and other mobile computing, Web-accessing, and media-gulping devices. Today, about the only option for creating a high-speed link between such a mobile device and your home PC is to set up a home network, complete with the hardware hassles and administration problems. A single USB cable would be a much simpler approach. And companies such as PLX Technology are working on protocols for using

SuperSpeed USB as a virtual network between computers.

Jimmy Chou, director of marketing for storage and USB products at PLX Technology, forecasts that consumer PCs will be the early adopters of USB 3.0, with discrete PHY (physical-layer) chips going into notebooks early next year. A lot of notebook vendors have such designs under way now, he suggests. Meanwhile, other sources suggest that makers of handheld devices that trade in large data files, such as videocameras and movie players, may be on a similar time line.

A big milestone will arrive when Intel, in about 18 months, releases a south-bridge chip with integrated 3.0. With that announcement, momentum should grow until 3.0 has become universal by mid-2011, according to forecasts. Predictably, the fly in the ointment may be software. Chou says that the host-side driver for 3.0 is about an order of magnitude more complex than the device-side driver. Device vendors will also have to supply some code to the host to execute with the resident host driver. Microsoft does not plan to release any host driver for 3.0 before the Windows 7 service-pack 1, according to industry talk. So, it may turn out that, for all the added hardware complexity that USB 3.0 is imposing on chip vendors, the real delay may be in just getting the drivers right.

THE FIRST STEP IS TO RUN ALL THE CIRCUITS AS SLOWLY AS POSSIBLE TO CONSERVE ENERGY.

rience with PCIe Gen 2 a real asset for chip designers. "If a design team has really understood PCIe Gen 2, then they can adapt about 90% of their work to USB 3.0," says Scott Kim, manager of business development at Texas Instruments. Yes, the equalizer will need beefing up, but much of the circuitry will remain the same. For example, getting 5-Gbps performance from the PHY requires a careful trade-off between deep pipelining to meet throughput requirements and limited latency to meet bus timing.

Kim also emphasizes the importance of experience with power management. "The first step is to run all the circuits as slowly as possible to conserve energy," Kim says. "For instance, you have to keep the receiver listening for the LFPS (low-frequency-periodic-signaling) traffic so that you know when to wake up. But that [requirement] doesn't mean you have to run the whole receiver at full speed. We've figured out a way to filter

LFPS while running at very low power.” Kim also cites TI’s extensive low-power-design method, which now routinely includes clock gating and power gating. Such techniques allow designers to run the transmitter-pre-emphasis circuit at varying power levels, depending on the needs of the channel. Similarly, Synopsys 3.0 IP will reduce receiver-equalizer power to just the level the circuit needs for the required equalization.

All of these challenges represent a substantial investment, in both design and testing, according to Jimmy Chou, director of marketing for storage and USB products at PLX Technology. PLX has the benefit of an established presence in both USB 2.0 and PCIe Gen 2, but Chou says that the engineering investment was still substantial. In addition to the usual costs of chip design, he says, the company has spent about a year and a half in the lab testing and validating its PHYs, especially the equalizer algorithms.

LEGACY TROUBLES

One of the serious implementation challenges with USB 3.0 turns out to be the requirement that the PHY operate simultaneously in both legacy USB 2.0 mode and 3.0 mode. Most design teams have looked at the problem and concluded that the best approach is to simply drop a commodity USB 2.0 PHY-IP block into the design next to the newly designed 3.0 hardware. TI, for example, takes this approach, and most IP vendors will probably also do so. The cell exists in their libraries, so why not just use it? According to Synopsys’ Nandra, however, there are good reasons not to. At the layout level, you can combine the digital portions of the two PHYs and save some real estate. With cleverness, you can also reduce the pin count of the block and the frontage on the perimeter of the die. More important, according to Nandra, is the fact that a lot of signals must pass between the 2.0 and 3.0 PHY devices during operation. It may be unwise to expose these internal signals to the chip designers who don’t know the details of USB 3.0 operation.

Nandra also worries about crosstalk when both 2.0 and 3.0 modes are in simultaneous operation. He points out that, if you are using a continuous-time

+ Go to www.edn.com/090611df and click on Feedback Loop to post a comment on this article.

+ For more technical articles, go to www.edn.com/features.

linear equalizer, turning up the gain on the equalizer won’t help because doing so amplifies not just the signal, but also the crosstalk. So you may have to implement active crosstalk-suppression circuitry for simultaneous operation.

With complex channels varying in characteristics, power and cost constraints, legacy compatibility issues, and some lingering questions about the robustness of the standard, USB 3.0 is not a walk in the park. Given the range of applications (see sidebar “Why do we need a 5-Gbps USB?”) and the momentum behind the standard, however, IP and chip vendors will support the effort. They will present system designers with a variety of choices with different price and power points.

It may prove exceedingly complex for system vendors to determine exactly what kind of performance they are getting for their dollar and their milliwatt, however. Extensive testing with board layouts, different lengths and qualities of cabling, and varying connector quality may be the only way to tell the really top-quality PHY, which can deliver consistently high data rates and low bit-error rates, from the bargain-basement PHY, which will work well only under ideal conditions. **EDN**

FOR MORE INFORMATION

Intel Corp
www.intel.com

Microsoft
www.microsoft.com

PLX Technology
www.plxtech.com

RedMere Technology
www.redmere.com

Standard Micro-
systems Corp
www.smsc.com

Synopsys
www.synopsys.com

Texas Instruments
www.ti.com

You can reach Executive Editor **Ron Wilson** at 1-510-744-1263 and ronald.wilson@reedbusiness.com.

