Cahners
MICROPROCESSOR REPORT
www.MPRonline.com
THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

USB 2.0, Serial ATA Box Out 1394

New Initiatives at IDF Provide External, Internal PC Connectivity

By Kevin Krewell {4/3/00-03}

The Intel Developer Forum (IDF) is Intel's key event for industry initiatives, and it certainly was busy at this February's spring 2000 edition. In addition to several processor-related announcements, Intel began the call to arms for the Serial ATA program, kicked off the "beyond AGP 4x" initiative, and solidified the support behind USB 2.0. While the follow-on to AGP 4x will not be serial, peripheral interconnects certainly will be—and quickly.

Intel began its push toward serial channels with USB back in 1995 (see *MPR 4/17/95-01*, "Universal Serial Bus to Simplify PC I/O"). But Apple Computer began the trend toward serial connections with FireWire, which it designed to replace the SCSI parallel bus (see *MPR 3/7/94-04*, "FireWire Brings Fast Serial Bus to Desktop") years earlier. SCSI was based on 1980s-era technology, but it was easy to use and relatively scalable, and it worked for multiple device types, including disks, scanners, and printers. But SCSI had several disadvantages, including a rather unwieldy cable, a limited number of peripheral devices (eight), limited cable length (two meters), and no support for isochronous (timed interval) transfers. FireWire was a fresh attempt to address those needs. It offered a serial channel with scalable bandwidth and isochronous capability, requiring only a moderate amount of logic to implement. FireWire led to the development of the IEEE 1394-1995 open standard.

Coincidentally, USB's first specification was also released in 1995. USB was designed to be an inexpensive serial link for computer peripherals, but it was also touted for transporting MPEG-1 video and digital audio. But the 12Mb/s bandwidth offered by USB wasn't sufficient for more demanding applications, such as disk-drive connections and multichannel AV. There was another key difference: USB is *not* an IEEE (or ANSI) standard, which gives Intel and its partners greater control over the specification. USB also requires a root hub and a computer to control the bus, whereas 1394 is a peer-to-peer network, allowing devices to communicate without the need of a central computer. The 1394a spec allowed 100-, 200-, and 400Mb/s data rates, and USB 1.x supported only 1.5 and 12Mb/s, making 1394 more suitable to AV and data storage connectivity. But despite its advantages, there seemed to be an invisible barrier that was keeping 1394 from becoming widespread.

What should have been an appealing universal solution (1394) has languished because of a lack of support by influential hardware vendors. Specifically, without Intel's prodding, the computer industry was either incapable of—or uninterested in—driving 1394 into the mainstream. Sony eventually adopted 1394 to connect digital AV equipment to PCs, as did Compaq and Apple. But even these vendors are not using it for its original purpose of connecting disk drives within their boxes, and 1394 devices, other than AV equipment, have not caught on. Device bay, a key initiative for 1394, proved too expensive for OEMs and appears to have been scrapped.

Serial Provides Performance and Scalability

The trend to serial connectivity has been enabled by developments in high-speed networking. Modern process technologies now allow communication circuits to operate at gigahertz speeds without resorting to exotic IC processes.

While semiconductor process advances reduce die size, wide parallel buses use many I/O pads, which do not scale well with process shrinks. These are just two of the forces working to push system interfaces toward narrow, high-speed channels as replacements for wide parallel buses. The main obstacle has been the requirement for backward compatibility with existing parallel buses. This need has also held back PCI, IDE (ATA), SCSI, and others from moving to serial interfaces—that is, until now. The movement to eliminate legacy interfaces from the PC architecture, supported by both Intel and Microsoft, opens the door for greater flexibility in the development of new system interfaces.

USB 2.0 Gains Steam

The faster USB 2.0 specification had been in development by the USB 2.0 Promoter Group for roughly one year, with an initial speed target between 360Mb/s and 480Mb/s (see *MPR 3/29/99-04*, "Intel Forum Raises Standards Conflicts"). At the October 1999 USB 2.0 Developers Conference, the specification narrowed to 480Mb/s only (see *MPR 11/15/99-sp*, "The Future of Serial Buses"). (The USB 2.0 Promoter Group consists of Compaq, Hewlett-Packard, Intel, Lucent, Microsoft, NEC, and Philips.)

Significant progress has been made on USB 2.0 since then; at the spring 2000 IDF, Intel demonstrated a hardware implementation with prototype controllers and announced that the specification would be released for final review in April. The goal expressed at IDF is to have USB 2.0-compliant parts shipping by fall Comdex. Although initial silicon will not include chip-set integration, this is still an aggressive goal, considering that the prototype hardware at IDF was still in FPGA devices. The rush to get USB 2.0 launched will probably result in the availability of few peripherals and third-party add-in cards.

One big question mark is Microsoft support. At IDF, Microsoft presented its requirements to get USB 2.0 support into a Windows Millennium Edition patch. Microsoft requested three months of end-to-end testing of USB 2.0 driver stack with sufficient production-quality hardware. Production-quality hardware testing requirements included a minimum of two host controllers, three hubs, and six devices of various types (scanners, disks, printers, etc.), all with production-quality silicon.

Another challenge facing USB 2.0 is that of implementing the high-speed transceivers in a standard logic process. Data rates as high as 480Mb/s require custom circuits, which USB 2.0 member companies are developing for standard-cell libraries. The USB 2.0 group will draft a standard (intrachip) transceiver macrocell interface, which will be available in 8- and 16-bit-wide versions. The design of USB peripheral devices is based on the concept that new peripherals will support the high-speed standard (HS) and the existing (USB 1.1) 12Mb/s full-speed standard (FS) but will not be required to support the 1.5Mbp/s low-speed (LS) standard. Devices that currently require only low-speed support will be slow, cost-sensitive peripherals, such as keyboards and mice, that will never need to upgrade to 480Mb/s. A USB 2.0 hub will be required to support all three speed standards.

To get peripheral devices to market in a timely fashion, the USB committee is proposing three options: a discrete transceiver design, a microcontroller with integrated USB 2.0, and a transceiver macrocell for foundry libraries. The fastest time to market will be with the discrete transceiver, which is expected to be available in Q3. It is also the costliest option. The microcontroller design should follow shortly, sometime in 2H00. Transceiver libraries should be available before the end of 2000.

Tools are also required for USB testing and validation. To support USB 2.0 peripheral development, a peripheral developers kit (PDK) and bus analyzers will be available in Q2. Microsoft will provide the software stack for the PDK. The PDK also offers a PCI-based hub controller and transaction-generation software to create test traffic. The bus analyzer will be needed for both signaling analysis and packet-level transactions.

The fastest way for PC OEMs to add USB 2.0 will be with an add-in PCI card. The challenge for motherboard designers is that USB 2.0 will eventually be added to the I/O hub controller (IHC), which is almost always positioned on the ATX motherboard catty-corner from the back-panel connectors. Running 480Mb/s signals around a four-layer motherboard requires very careful layout—much more than was required with USB 1.1. This means that USB 2.0–enabled IHCs will not simply drop into existing motherboard designs. The routing of such signals requires careful attention to line impedance, the number of vias used, ground-plane crossings, and trace lengths. These are issues motherboard vendors have dealt with before on EV6 (Athlon's front-side bus) and Rambus designs, so the experience to accomplish this goal should be available.

All these factors must be put into a compliance plan and testing procedure, which Intel promised by mid-2000. Compliance testing will be performed at the motherboard connector to ease vendor testing, but much work remains to hit the target launch dates. If you base the likelihood of meeting those goals on past history, you might suspect that Christmas 2001 would be a better bet. If USB 2.0 does miss the Q4 launch date, it will not have a significant short-term impact. OEMs will still provide IEEE 1394a for digital-video camcorder connectivity; SCSI will still be available for devices such as scanners; and USB 1.1 will still be available for keyboards and mice. All these other standards are available today—although not as the unified, easy-to-use solution USB 2.0 promises to be.

Serial ATA Presents a Smooth Transition

While USB 2.0 addresses external connectivity, it is not targeted at internal mass storage. To address this issue, another initiative was begun—Serial ATA. Once again, 1394 has been bypassed as a solution. The founders of the serial ATA

effort include Intel, of course, along with major disk- drive manufacturers IBM, Maxtor, Quantum, and Seagate. System OEM Dell and adapter manufacturer APT Technology round out the supporter list.

The goal for Serial ATA is to provide a new, scalable serial interface that maintains software compatibility with the existing parallel ATA interface. This initiative should not be very controversial, because its goals are clear: to provide very high performance with the lowest possible overhead and cost. The performance goals were set on the basis of projections of media rate (data transfer rate off the drive head), with the goal of minimizing the amount of buffering required (see Figure 1). Late this year or early next, that rate should drive the need for ATA-100. Intel projects that in 2002 those data rates should exceed ATA-100 speeds. Rather than push the parallel interface another notch to ATA-133, the technology will be available to accomplish the equivalent of ATA-150, using a 1.5GHz serial data link. Intel showed a roadmap at IDF that takes Serial ATA from 150MB/s in 2001 to 300MB/s in 2004 and 450–600MB/s in 2007.

The technique used to convert the parallel ATA interface into a high-speed serial interface is reminiscent of one used almost a decade ago in a product from AMD called the TAXI chip set. Parallel (byte) data was latched into the TAXI transmitter chip, which would serialize the data using an internal PLL running at 10 times the parallel data rate. The one version of the TAXI sent data using an 8b/10b encoding scheme used in Fibrechannel, the same encoding scheme to be used in 1394b and now proposed for Serial ATA. While this encoding scheme introduces a 25% overhead, it allows sufficient bit transitions to synchronize a PLL in the receiving chip and recover the data with low bit-error rates. The 8b/10b-encoding scheme maps the 256 possible data states into 1,024 possible symbols. Symbols with sufficient transitions are used to represent the data values and maintain frequency lock at the receiver side. Additional symbols provided in-band signaling: valid data start, error codes, PLL start-up protocol, and so on. The TAXI receiver then deserialized and presented the data at its outputs. At the time, the TAXI chip set was defined, 125MHz (100Mb/s) was state of the art; Serial ATA will operate more than an order of magnitude faster, inexpensively. Because of the simplicity of the interface, it will be possible to interface the existing parallel bus to the new serial standard and to support mixed serial and parallel solutions in systems (see Figure 2). It should also be possible for a device to support both parallel and serial interfaces.

The Serial ATA group is so serious about making the transition transparent to software—including the operating system—that it included wording to that effect on its mission statement. The interface design has so far focused on connectivity within the chassis to mass-storage devices, and it has ignored external and peer-to-peer connections. The one-meter cable length restriction is about double today's parallel cable length and should more than suffice for tomorrow's smaller PC and appliance form factors. The interface will also be designed with mobile computers in mind. The interface has low voltage and current requirements, and it achieves high noise margins by using unidirectional, differential signaling. Intel has estimated that the energy required to send one byte of data on parallel ATA wires (UDMA) is 1.85×10^{-9} joules, but only 8.5×10^{-11} joules are required for serial ATA. The four-wire cables (two unidirectional, differential serial lines) that replace the existing

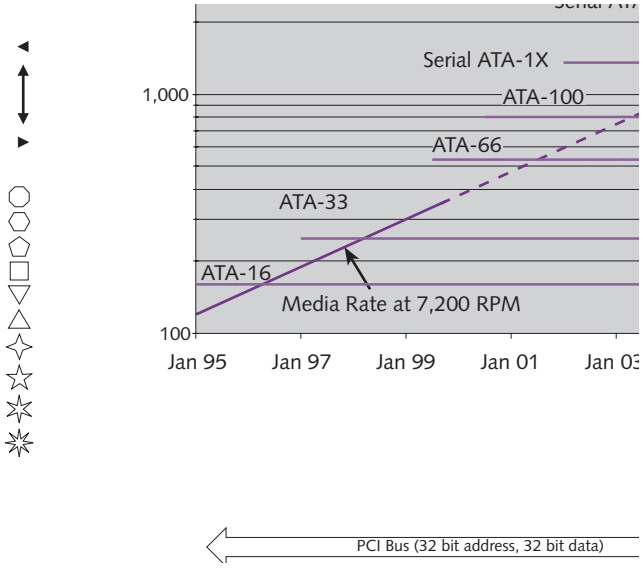

Figure 1. This roadmap was presented at IDF to show the growth in media rate (based on 7,200RPM spindle speeds) from January 1995 through projections for January 2005. Media rates are shown to be growing at an order of magnitude every 10 years—from 125Mb/s in 1995 to more than 1,000Mb/s by 2005.

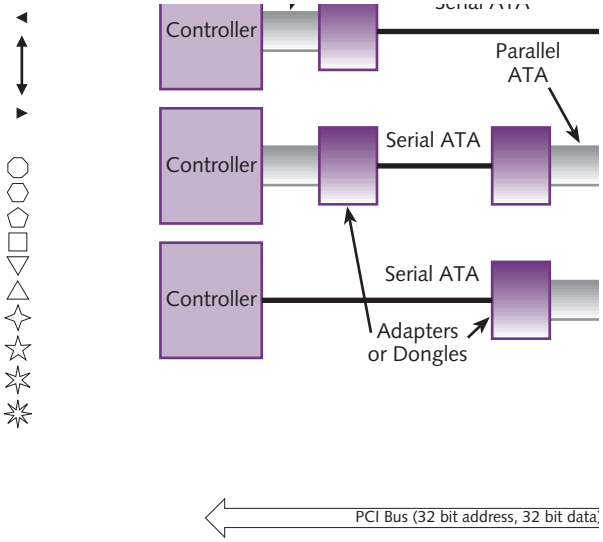

Figure 2. Serial ATA can connect to legacy parallel ATA devices by using an adapter or dongle. This will allow easy mixing and matching of the serial and parallel storage devices and controllers.

Availability

The USB 2.0 specification will be made available for final review in April, and products are targeted to be available in Q4 of 2000. A USB 2.0 Developers Conference is scheduled for May 15–17 in Anaheim, Calif. For more information, see the USB Web site at *www.usb.org*.

Serial ATA Working Group information can be found at *serialata.org*.

The 1394 Trade Association is alive and well at *www.1394ta.org*.

bulky parallel cables will allow cleaner, less cluttered chassis and notebook designs. The new cable connector will be designed to provide a more secure connection than the existing parallel cable. The specification is also expected to allow extensions for hot plugging, very large drives (>137GB), efficient command queuing, and first-party DMA.

Is Serial AGP Next?

Wide parallel buses will not be replaced on motherboard (planar) designs anytime soon. Processor front-side buses may actually get wider, as well as faster, in order to feed forthcoming monster processors such as Intel's McKinley IA-64 processor and IBM's Power4. Graphic buses also need to speed up, as graphic performance may grow an order of magnitude in the next few years—a fact that has led to the formation of a "beyond AGP 4x" initiative. The next generation of graphics buses could abandon entrenched standards and move toward very-high-speed narrow buses using signaling technology from JAZiO (see *MPR 2/21/00-02*, "JAZiO: Slow Edges can Run Fast"), Rambus, AMD's lighting data transport (LDT), or Motorola's RapidIO.

The march toward new serial buses seems unstoppable, but where does that leave 1394? For consumer AV equipment connectivity, 1394 will continue to be a must-have item, and it may still find a use in PCs connecting to external storage. But when USB 2.0 replaces USB 1.1 in 2001, Intel will likely position 1394 as a redundant and unnecessary PC interface. Can a USB 2.0-to-1394 adapter be far off? ◇