Intel Forum Raises Standards Conflicts

IDF Showcases Intel's I/O Roadmap and New Platform Components

by Peter N. Glaskowsky

Intel has really learned how to put on a show. IDF, Intel's semiannual developer forum, has settled into posh digs at the Palm Springs Convention Center, growing to eight parallel technical tracks and seven simultaneous lab sessions. About 2,000 people attended the three-day February show, where they were offered seven keynote speeches by Intel executives (which may be viewed online at *http://intel.broadcast.com/intel/idf99*) and 110 technical presentations.

By comparison, April's Windows Hardware Engineering Conference (WinHEC) will provide just one track each morning and two each afternoon of its three-day run. Clearly, Intel has more to say than Microsoft on the topic of PC-hardware development.

Some say that Intel is taking too much control of the PC platform, however, pointing to the disastrous consequences of IBM's attempt to dominate the PC industry in the 1980s. That attempt, in which IBM touted its MicroChannel bus and OS/2 operating system as the next wave of computing technology, led to a major power shift in favor of independent system developers such as Compaq.

Intel's vision for the PC of the future has a new version of USB replacing SCSI and all other external peripheral interfaces on mainstream PCs. The company also sees its next-generation I/O (NGIO) interconnect replacing PCI on backplanes and Fibre Channel on server mass-storage and networking subsystems.

Intel is also looking to inject a little iMac-style excitement into the plastics that house the PC. Pat Gelsinger, Vice President of Intel's Desktop Products Group, closed his keynote with a "fashion show" of unconventional-looking PCs carried down the runway by Intel's Bunny People.

Gigahertz, Geyserville Demos Attract Attention

Other keynote presentations also displayed Intel's growing talents of showmanship. During his keynote speech, Intel's Albert Yu hosted the first public demonstration of a Pentium III operating at more than 1 GHz. Intel says the demo used a 0.25-micron processor and "special cooling techniques" to achieve the record speed.

By carefully choosing the fastest part off the line and applying extreme cooling and overvoltage, vendors can often obtain speeds 50% to 100% above normal. The IDF demo lasted only a few seconds at 1.002 GHz and the chip ran no applications at that speed, indicating possible instabilities. Although good for public relations, this demo has little bearing on Intel's ability to deliver such clock speeds in production, which we still expect to occur no sooner than 2H00.

In another presentation, Robert Jecmen, VP of Intel's Architecture Business Group, offered the first public look at Intel's Geyserville technology for mobile computers. As we reported a year ago (see MPR 3/30/98, p. 4), the Geyserville technology permits laptop CPUs to run faster (and hotter) when plugged in. To save power when operating from a battery, a Geyserville-equipped system will run its processor at the voltage and clock speeds typical of mobile systems. When the system is connected to an external power source, the voltage to the CPU is increased, allowing higher clock rates.

In the IDF demo, Intel showed a laptop operating at 400 MHz in "battery optimized" mode, switching automatically to 500 MHz when AC power was connected. Jecmen said that by the end of the year, Intel Mobile Pentium III processors fabricated in 0.18-micron technology will be capable of 600-MHz operation in this "maximum performance" mode.

The Geyserville technology will allow laptops to more closely approach the performance of desktop systems. CPU power consumption in mobile systems can't match that of desktops, however, due to thermal limits. Intel had originally planned to support the use of an external fan in notebook docking stations to help remove the heat generated in maximum-performance mode, but that plan has been dropped.

I/O Roadmap Favors Home-Grown Technology

Intel's plans for PC I/O highlight how the company is relying more on its own developers and less on independent standards bodies to meet future PC technology demands. Intel plans to develop faster versions of today's Ultra-ATA and Universal Serial Bus specifications, expecting these two interfaces to replace all other forms of internal and external peripheral interfaces on future mainstream PCs.

Though Intel was touting the IEEE 1394 standard (see MPR 3/7/94, p. 18) for external mass storage as recently as last fall's IDF and at one time planned to integrate 1394 into its chip sets, the company now believes 1394 will never be widely used on PCs. Intel recognizes that 1394 (aka FireWire) has become a de facto standard in the consumer-electronics industry and acknowledges that some multimedia-oriented PCs will need 1394. It now appears, however, that Intel will never offer chip sets with 1394 support. OEMs that want to offer 1394 support will have to use separate 1394 peripheral-controller chips.

Intel disapproves of the royalty structure for 1394 technology, which was developed by Apple. Apple was reportedly asking about $1 per node for 1394 devices. Discussing 1394 at IDF, Gelsinger said "broad deployment requires reasonable licensing." Intel doesn't seem to have a problem with the

high license fees for some of its own preferred technologies, however. Intel downplays the significance of the royalties charged by Rambus for Direct RDRAM, for example, which come to several dollars per system. Intel also points out that 1394 is inherently expensive, since electrical-isolation issues mandate the use of a separate chip for 1394's physical-layer interface. The same argument applies here, however. Direct RDRAM's complex interface also adds cost to DRDRAM chips and controllers. Intel's objections to the cost of 1394 seem more political than substantive in nature.

A similar conflict is looming between the OpenBoot firmware standard (IEEE 1275, aka Open Firmware) and Intel's Extensible Firmware Interface (EFI). Intel presented its EFI plans at IDF as an enabling technology for IA-64 compatibility, allowing boot devices—typically PCI add-in cards—to provide IA-32 or IA-64 native code in firmware.

Sun, Apple, and others developed OpenBoot in 1994 to provide exactly the same sort of processor-independent boot capability. Unlike EFI, which requires two or more sets of machine-code drivers, OpenBoot uses hardware-independent Forth-language drivers to allow any OpenBoot-compatible system to boot from the same firmware. As the operating system boots, it loads machine-specific drivers to replace the slower interpreted Forth code. Intel declined to participate in the OpenBoot development effort and now seems intent on reinventing this wheel.

Intel's decision to limit support for 1394 puts the DeviceBay initiative (see MPR 5/6/96, p. 12) into limbo. DeviceBay requires USB and 1394 interfaces, but if OEMs must add an expensive 1394 controller to support DeviceBay, it is unlikely to ever see wide use. Intel did not address the DeviceBay issue at IDF, but we believe the company may be planning a defeatured DeviceBay with only the USB interface. Such a solution would be much less attractive than the original 1394-equipped configuration, especially in high-end desktop and laptop systems, where reconfigurable mass storage and compatibility with consumer-electronics devices are needed most.

USB 2.0 Seeks To Replace SCSI

Version 2.0 of the USB specification, currently under development at Intel and expected to ship in Intel core logic in 2H00, will boost the specification's current 12-Mbit/s data rate to speeds 10 to 20 times faster. USB 2.0 will preserve complete backward compatibility with existing USB products; the same cables, connectors, and software interfaces will be used in the new specification.

So far, so good—but Intel believes that USB 2.0 can and should displace all existing external peripheral interfaces on mainstream PCs, such as the popular and widely supported SCSI. SCSI, however, is already much faster than USB 2.0, and it is supported by all major disk-drive manufacturers. If Intel pushes its belief hard enough to discourage SCSI support on performance-oriented PCs, users will find themselves limited to slower devices. Such an outcome is unlikely, of course; it's more likely that OEMs and end users will force Intel to change its plan.

Intel has no plans for USB 2.0 to solve perhaps the most serious shortcoming of the current USB specification: its lack of support for peer-to-peer operation. USB today cannot be used to connect two PCs or to dock a USB-equipped notebook computer or PDA to a desktop system. As presently envisioned by Intel, USB 2.0 will not provide this feature—nor will any other Intel-supported mainstream interface technology. Although Ethernet and the wireless Bluetooth standard (see MPR 6/1/98, p. 22) support peer-to-peer operation, neither is likely to be as universally supported as USB.

For internal storage devices, the current ATA-33 specification, which permits a peak throughput of 33 MBytes/s, will be replaced by the ATA-66 specification this year. From 2000 to 2005, Intel plans to introduce faster versions of the ATA interface with fewer signal pins. The new versions will enable longer cables to improve cabling flexibility. ATA serves an important role in today's PCs, allowing low-cost host controllers and low-cost disk drives with acceptable performance. We hope Intel will preserve this essential simplicity and not add unnecessary sophistication in pursuit of its recent fixation on serial communications interfaces.

For now, Intel views ATA as suitable only for storage devices internal to the PC enclosure, while USB is reserved for external devices—but we believe this may change. Even the ATA-66 specification is about twice as fast as USB 2.0. If Intel produces an ATA derivative with greater cabling flexibility and greater throughput, it may choose to use it for external mass-storage devices as well.

For example, some of the stylish small-case PCs displayed at the IDF fashion show were originally conceived as network computers. These systems have room for, at most, one internal storage device, typically a hard-disk drive. An external variant of ATA may be needed to support external DVD-ROMs or other drives while maintaining performance parity with conventional desktop and minitower PCs. This illustrates that even within Intel's own roadmap, there are potential conflicts among USB 2.0, ATA, and independent standards for external peripheral interfaces.

NGIO Slated To Replace PCI

The need for a successor to the successful PCI bus is becoming acute, with new disk and network interfaces needing more bandwidth than today's PCI buses can provide. The PCI slots in current PCs have about 133 MBytes/s of peak bandwidth but can typically sustain only about 80–100 MBytes/s. That's the speed of Wide Ultra2 SCSI or Gigabit Ethernet, but Ultra3 SCSI and Fibre Channel are even faster.

The PCI standard includes 64-bit and 66-MHz variants of the bus, but these have problems of their own. Designers do not consider the 64-bit versions cost-effective as they require about 100 signal pins and complex motherboard layouts. The 66-MHz PCI alternative has electrical characteristics that limit it to just two slots; systems that need more

Intel Previews New 440MX Mobile Chip Set

At IDF, Intel released preliminary details of a forthcoming chip set for mobile systems. Unlike Intel's recent mobile core-logic products, the Mobile TX and BX, the 440MX will include features not found on current desktop chip sets.

The primary goal of the 440MX is to reduce notebook cost and power consumption. To achieve this goal, the 440MX relies more heavily on soft (host-based) implementations of features typically provided as hardware peripherals on desktop systems. The 440MX features Intel's first integrated AC'97 audio controller, requiring only an inexpensive AC'97 codec chip to enable software audio synthesis and V.90 modem support.

The 440MX, due out by midyear, will be a single-chip device, Intel's first to combine both the north-bridge and the south-bridge functions in one component. This configuration is likely to migrate to the desktop, but it is most valuable in mininotebook computers, where space is at a high premium.

Intel did not announce pricing information for the 440MX, but we expect the price will be similar to the $52 Intel asked for the Mobile 440BX when that chip set was introduced.

would require multiple north-bridge chips, adding further unwanted cost to the platform. Though these two options can be combined to achieve a peak bandwidth of 533 MBytes/s, this configuration is even less efficient, offering a sustained bandwidth of only 300–350 MBytes/s.

The PCI special interest group (SIG) is working on an enhanced PCI standard known as PCI-X that will boost clock speeds to 133 MHz and peak bandwidth to more than 1 GByte/s (see MPR 10/5/98, p. 4). PCI-X will permit up to four slots per controller (albeit with a reduced clock rate). The new specification suffers from the same inefficiencies found in standard PCI, but with 600–800 MBytes/s of sustained throughput, PCI-X should be fast enough for the entry-level and midrange server markets through 2002.

Many in the industry, including Intel, are working on new narrow, high-speed point-to-point switched communications channels to replace peripheral-interface buses. Years of research into related technology for high-speed asynchronous transfer mode (ATM) and Ethernet networks, Fibre Channel, and other serial interfaces provide fertile ground for the development of new PC-platform technology.

These new channels offer many potential advantages. With fewer signal pins (as few as four, for full-duplex differential serial interfaces), skew is reduced or eliminated, more channels can be supported on a single ASIC, and interconnect costs are greatly reduced. Some of these cost savings can be traded for more expensive driver, receiver, and interconnect technology, allowing much greater speeds. Point-to-point connections are electrically cleaner than buses, simplifying signal termination and permitting even higher speeds. Buses, on the other hand, benefit from statistical multiplexing effects; all of the devices on a bus are unlikely to be busy at the same moment. With point-to-point links, however, bandwidth not needed by one device is not available to the others.

Intel's NGIO proposal (see MPR 12/7/98, p. 4) encompasses all the latest advances in I/O technology. NGIO is a serial interface that operates at 2.5 Gbits/s with 8B/10B coding for a peak bandwidth of 250 MBytes/s in each direction simultaneously. (NGIO was originally planned to debut at 1.25 Gbits/s but now will not appear until the 2.5-Gbit/s version is ready.) At this speed, each NGIO channel offers somewhat better throughput than 64-bit 66-MHz PCI, but it will be slower than PCI-X. To help alleviate this limitation, NGIO supports "bonding" of multiple channels, allowing power-of-two combinations of NGIO channels to provide higher sustained throughput.

NGIO Plan Has Problems

At IDF, Intel proclaimed that NGIO is "the successor to PCI," at least for servers, but there are some problems with this plan. PCI-X will certainly be ready before NGIO, and many OEMs and peripheral makers need PCI-X in the near term. Intel believes PCI-X will have a "brief, transitional role" in servers before being replaced by NGIO, but customers are unlikely to be interested in a new interface architecture that is slower than its predecessor.

NGIO is being developed initially as a means to connect the multiple chassis found in large servers—CPU cabinets, disk-drive cabinets, network switches, and so on. This is a function currently handled by Fibre Channel, which is conceptually similar to NGIO. Indeed, NGIO uses the Fibre Channel physical layer, though the two protocols are very different. A backplane version of NGIO will not be ready until some time after the chassis-interconnect version.

Even Intel admits that some of today's PCI-based products can't easily be ported to the NGIO architecture. Add-in cards that rely on memory-mapped random access, such as graphics cards, cannot be readily adapted to NGIO. Those that use periodic CPU polling to manage transactions and observe device status will have to find other ways to perform these functions. Vendors of these products will eventually be able to make the change, but this transition is likely to take many years.

Intel could have put as much effort into developing a faster, more flexible derivative of Fibre Channel through the Fibre Channel Association (FCA), but this probably would have taken longer than developing a new standard itself—and Intel does not participate in the FCA. Intel has had political clashes with some FCA members and chose to go its own way with NGIO.

NGIO is on a collision course with the Future I/O initiative, originally proposed by Compaq, HP, and IBM as the successor to PCI-X, which they also developed. Future I/O (*www.futureio.org*) will be four times faster than NGIO, impose less latency, and provide a smoother transition for existing PCI designs. Intel expects Future I/O to be more expensive to implement than NGIO and says NGIO is better suited to the needs of the server market. We hope that Intel and the Future I/O supporters can find a way to cooperate, reconcile their differences quickly, and avoid a standards war in the PC industry.

Direct RDRAM Yield Problems Delay Camino

Though the eventual adoption of Rambus's Direct RDRAM memory technology across Intel's core-logic product line is inevitable, the two companies have recently been forced to delay the debut of this technology on consumer desktops. None of the 15 memory-chip vendors that have signed up to make DRDRAMs was able to meet Intel's original schedule for volume production.

Intel put the best possible face on the situation, noting that "all major Direct RDRAM milestones [have been] achieved." This refers only to sampling and testing, however. The ultimate milestone—volume production of the 400-MHz devices required by Intel's 820 (Camino) chip set in time for its previously planned June introduction—has already been missed. Intel has put off the Camino introduction to September, by which time enough 400-MHz DRDRAMs should be ready.

The DRAM vendors blame the aggressive timing requirements of the 400-MHz Rambus interface, which yields 800 Mbits/s per pin and 1.6 GBytes/s (peak) for each 16-bit device. Yield, they say, would be fine at 300 MHz, and indeed many vendors can produce such parts today. Intel and its OEM customers were apprised of this situation some time ago but rejected the option to introduce Camino with the slower DRDRAM. Instead, Intel and the OEMs decided to wait until the promised 400-MHz speeds could be delivered by the DRAM vendors.

Many DRAM vendors have already decided to skip the 64/72-Mbit generation and are now focused on delivering 128- and 144-Mbit chips as quickly as possible. Intel expects that only six DRAM vendors will offer 72-Mbit parts, and just two will produce 64-Mbit devices. The limited availability of the 64-Mbit variety, expected to be the most popular for mainstream desktop systems, will further reduce the availability and boost the cost of Direct RDRAM for initial Camino-based systems. By the first half of 2000, however, Intel expects 128- and 144-Mbit parts will be available from 11 of the 15 announced DRDRAM makers.

Graphics chips, Intel says, will begin to use Direct RDRAM for local storage in 1H00, providing a secondary market for these parts. What Intel hasn't attempted to explain is what will happen to the large quantities of 300-MHz Direct RDRAMs that will be manufactured this year. These chips cannot be used on desktop PCs and are unlikely to be used on graphics chips that are even more performance hungry. The slower DRDRAMs might eventually find a home in notebook PCs, where performance demands and power budgets are lower—but the market for DRDRAM in mobile systems isn't likely to develop until mid-2000.

The transition from SDRAM to DRDRAM is of great concern to major PC OEMs. Intel described a chip currently under development that will interface a DRDRAM chip set to SDRAM memory devices, allowing the use of less-expensive and more widely available SDRAMs on Camino systems. Intel considered and rejected the use of such a chip on the motherboard—taking the place of the third Rambus module (RIMM) socket—because of potential signal-integrity problems.

In place of this solution, Intel offers the S-RIMM (a modified RIMM with the interface chip plus SDRAMs) and a DIMM riser, which fits in a RIMM socket and is equipped with the interface chip plus two SDRAM DIMM sockets. Only the latter option provides compatibility with existing SDRAM modules, but Intel acknowledges that this plan has other problems. A single DIMM riser can accommodate only two DIMMs, and only one DIMM riser can be used in each system. This capacity limit, combined with the physical fragility of the resulting assembly, may make the DIMM riser option impractical for OEM use.

IDF Communicates Intel's Desire for Control

The February IDF did its usual excellent job of communicating the details of Intel's contributions to the PC platform. The show also reinforced another important message: Intel intends to control every key element of the PC platform. To accomplish this goal, the company will pursue its own proprietary technology, even when comparable or superior technology already exists as an open industry standard.

In some cases, Intel's pursuit of control will hurt its own customers by forcing them to discard work they've already done, to accept less powerful or less flexible technology, or to spend extra money on features, such as 1394, that Intel should include in its own products.

Several of Intel's key OEM customers have already broken ranks, especially in the profitable and hotly contested server market. Intel was not a key participant in the server business until recently, but the company is taking steps to gain control of this market too. The PCI-X and Future I/O initiatives, however, demonstrate that major server OEMs have significant objections to Intel's plans.

Other OEMs surely have similar objections to NGIO and other recent Intel proposals but cannot afford to publicly break ranks with their most crucial supplier in the absence of viable alternatives. Where such alternatives exist—particularly PCI-X, Future I/O, and the IEEE FireWire and OpenBoot standards—the coming year may find more OEMs choosing to meet their customers' needs rather than to serve Intel's needs. M