System Architects Look to the Future

New Designs Needed to Meet Memory and I/O Bandwidth Challenges

by Michael Slater

Microprocessor performance continues to increase at a breakneck pace, but delivering full value at the system level is becoming ever more challenging. At October's Microprocessor Forum, a panel of six of the industry's leading system architects discussed the key challenges facing designers of future PCs.

Paul Rubinfeld, senior engineering manager at Digital Equipment, summed up the concerns voiced by many panel members: "One of the biggest problems is going to be bandwidth in and out of the chips. As the CPUs are executing instructions faster and faster, the ability to get data in and out, I think, is going to be one of the most severe limitations on future performance.

"Another problem that we're all facing is, as we turn up the clock rate, the circuit problems—how you actually get these high-speed circuits to work at 500 and 600 MHz is a big challenge, to say the least. To go along with that would be power dissipation, how you're going to keep these things cool, because there is a general trend for microprocessors to be hotter and hotter; you've got to manage that. How do you manage that in laptops?"

Bandwidth challenges are not limited to memory but also are showing up in I/O and may require significant changes in system design, as Lin Nease, a system architect at Hewlett-Packard, noted: "With the advent of multimedia, the home market seems to be dominating the technology advances we see in the types of systems we're talking about. Communications is going to be a breakthrough area, such as desktop videoconferencing.

"We're going to have to architect isochronous thinking from top to bottom in the system. The push to native (signal) processing exacerbates the problem, because now every one of these four or five paths goes through the CPU or will eventually. I believe it requires a culture that hasn't existed yet among the players in the industry."

Fred Pollack, who directs the group at Intel responsible for platform architecture and performance analysis, commented that "The amount of bandwidth we're looking at in terms of I/O could get to the same level that we need in terms of bandwidth from memory to feed the processor, and that's quite different."

EDO to Give Way to Sync DRAMs, Rambus

Pollack outlined his view of which memory technologies will address the bandwidth problems in PCs: "Today, in terms of PCs, it's EDO DRAM technology. What seems to be the next small step in the evolution is going to synchronous DRAMs that can do 1-1-1 bursts at 66 and 75 MHz. After that, it gets really interesting for two reasons. First, when we start talking about going to 64M DRAM technology, from a PC point of view, you're really interested in the granularity and the associated costs for that granularity. Second, you're also interested in getting the bandwidth needed for the next generation of processors. So we have to go beyond the 1-1-1 timing you get from synchronous DRAM, and there are probably a couple of contenders.

"One question is, what frequency you can take SDRAM up to? Can you take it above 100 MHz, 133 MHz? Or do you have to go to what amounts to less of an evolutionary step, more of a revolutionary step, and go to a different style of technology like Rambus? We don't have the answers to those things yet, but those are the two primary contenders at this point."

Pollack concluded, "It is not definite yet, but it is starting to look like PCs will move from EDO to SDRAMs toward the end of 1996 and early 1997. I think Burst EDO seems to be fading a little bit in terms of its popularity." Nease noted Intel's overwhelming influence on which technologies succeed: "This makes it easy for the rest of us—if Intel says it's sync DRAMs, we're done. Actually, I think the solution is going to come from packaging technologies that are going to allow significantly increased pin counts cheaply. In higher-end systems, there isn't as much pressure right now because the minimum increment (of memory size) isn't a big problem. The electrical issues cause you to have a certain amount of natural parallelism that sync DRAMs gives you without the core logic having to provide that. Maybe that's something that will evolve down (to PCs)."

Don North, who manages a system architecture research group in Apple's Advanced Technology Group, further explained the opportunity for memory devices: "We have left most of the bandwidth that is available in memory chips inside them and still haven't been able to export the tremendous parallel bandwidth that is available. Rambus has gone probably the farthest of any in being able to stream data in and out, but there's still an extensive amount of internal bandwidth that's just being dropped on the floor. ...When you're building workstations and higher-end systems, it's not so much of an issue. You'll pay more to get extra performance. It's really not only a memory vendors' problem but also a system vendors' problem. We're asking all the memory vendors to compete equally with parts that are plug compatible

The Microprocessor Forum's panel of system architects featured, from left to right, Paul Rubinfeld (Digital), Lin Nease (Hewlett-Packard), Don North (Apple), Richard Oehler (IBM), Fred Pollack (Intel), George White (Corollary), and moderator Michael Slater.

and not really giving them any incentive to innovate, because we won't pay them any extra for the innovation. The system vendors are going to have to bite the bullet to get the extra performance out of the memories."

Pollack commented on Rambus' prospects for becoming more significant in the 64M generation: "There's a lot to say for Rambus because it solves the bandwidth problems, it solves the granularity problems, and those are real pluses in its favor. The disadvantage is that it is not evolutionary. You can talk about designing a chip set and a board today that support both synchronous and EDO DRAM, but it's very difficult do a transition where a chip set would support both synchronous DRAM and Rambus. Even if Rambus is the right answer for the 64M generation—and it might be, I don't know—then the question is how do we evolve to that."

Host Multimedia Processing to Spread

The panel was almost universal in its support of instruction set extensions for multimedia. Richard Oehler, who represented IBM in the group that revised the POWER architecture to create PowerPC, responded to the moderator's question about why that group rejected multimedia extensions: "They were certainly considered and they haven't been rejected. You saw that we've fixed some of the little-endian byte-alignment issues, and we certainly are moving the architecture over time. You'll see an evolution of this architecture; I just can't tell you when it will happen."

According to Rubinfeld, Digital will also join in adding multimedia instructions, saying "It's been considered, and there will be multimedia support in future Alpha chips."

Pollack clarified Intel's divergent strategies for I/O processing in desktops and servers: "When we're talking about servers, it is very important to have intelligent I/O. You don't want the processor having to go out and do reads and writes of control registers. You want much more intelligent I/O and a scalable I/O system.

"In the desktop realm, I think it's a price/performance issue, and you'll see a scale of things. When you see PCs below $2,000, where cost is really at a premium and designers are trying to save a dollar here or a dollar there, host-based processing makes a lot of sense; you've got the CPU there, and it allows you to bring new capabilities that you wouldn't otherwise have. At the same time, that gives you a wide range of platforms that have these new base-level capabilities, which attracts ISVs to write to those base-level capabilities and also attracts users. Then people want even more capabilities than you can see at the low end, and that creates a market for a higher-end system, like a better videoconferencing system. The market is very diverse, and there is room for dedicated processors as well as NSP."

Apple's North generally agreed but offered a caveat: "In doing all the multimedia functions, you can add extensions to the instruction set, as shown by UltraSparc, but it may be a false economy to drag all that multimedia data into the CPU off a network connection and then send it to the display, when ideally you'd like to go directly from a network through the appropriate interconnect logic and not have to drag that data through the main CPU bus, using up all the bus bandwidth that you didn't have enough of in the first place."

When asked about the prospects for unified memory architecture (UMA), Pollack commented "If you're going to talk low cost in terms of a PC, it really means you have to be talking about 8M of DRAM vs. 16M. If you take that extra 1M for the frame buffer, now you have 7M for Windows 95, and that has a significant performance penalty just taking away that 1M—never mind the fact that you then have to share the bandwidth as well. The concern that I have with UMA at the low end is that it might take away too much performance, even as a cost tradeoff. On the other hand, there's a lot that people trade off when you start talking about PCs that are below $1,500; there have been far worse sins committed in the past in terms of leaving performance on the table."

This concern has also been raised by Microsoft since the Forum, and it appears that realistic Windows 95 systems with UMA will need 12–16M RAM.

Desktop Multiprocessors Controversial

George White, president of Corollary—which has built a business around PC-compatible multiprocessor systems—was surprisingly pessimistic about MP on the desktop: "I don't think multiprocessor desktops make a lot of sense. If you were selling CPUs and a lot of the desktops had two CPUs in them instead of one, you could make a lot more money. But I'm not sure that it does the end users very much good. Our opinion is that at least in the mainstream market—not an engineering workstation but a normal PC—the fastest chip Intel makes is fast enough for what 99.9% of the people want to do, so putting two of them in the box would only appeal to a very small percentage of the users—and then only if the OS knew what to do with the second processor.

"I saw someone running a demonstration once of a dual-processor Pentium system. I walked up to him and said that it would look a lot better if he had a way of turning off the second processor so you could see the difference. He said, 'Actually, it would look a lot better if I had a way of turning *on* the second processor.' "

Pollack was more optimistic: "I think MP is for the desktop as well as servers. It's important to understand it is for a segment of the desktop. In my view, for the next couple years, it's for about the top 10%. This I think is the size of the market, not what you would realize in actual sales. There's a number of things that do scale well; a good example is Photoshop, which has been multithreaded on NT. For those kinds of applications, for software development, for financial services, all of these are useful to do in an MP context."

North agreed that MP would be valuable for a certain segment of users that is close to Apple's heart: "The 10% of the market that Fred was referring to, the Photoshop and desktop publishing area, is 2% bigger than Apple's market share. We know for a fact that applications like Photoshop accelerate tremendously, and scale almost linearly, with additional processors. The incremental cost of adding another processor is relatively inexpensive compared to the rest of the system."

Pollack concurred, pointing out that "I think it's very attractive if you buy a system for say $5,000 or $6,000 and then you have the ability for another $1,000 to add another processor. That's a small fraction of the system price—even when you get down to a $3,000 desktop, with the ability to add another processor for $500. You don't need to get a lot of benefit out of it to motivate that kind of sale."

As for getting application developers to make the necessary changes to their applications to support MP, Pollack was again optimistic: "With people coding to the Win95 interface, you have multiple tasks you can code to within a process. Increasingly, there will be more and more applications, but initially it is a limited segment."

Others were more skeptical. HP's Nease noted that "It's a simple economic issue for a lot of the software suppliers. They're really not in business to allow us to do experiments with hardware. They're in business to compete with one another on functionality and feature sets. ...They can't delay a release for three months for something that does not constitute 99% of the market."

White also was skeptical of getting MP support in mainstream applications: "If you multithread your application and debug it on a uniprocessor, it will likely be broken when you run it on a multiprocessor. ...You've got to have a computationally intensive application to go through that work (of testing and debugging in MP configurations). The database guys are going to do it, and it is relevant on servers, but I don't think Excel is going to be rearchitected."

Backside Cache Buses Attractive

Most high-end processors today have private L2 cache buses, and most of the panelists agreed that this will become more widespread. When asked whether this feature will move down from the PowerPC 620, IBM's Oehler commented that "On the 620, we call it the 6xx bus, as opposed to the 60x bus. I think you will see that bus not just in 620-type systems but also in 604 MP systems, because it is independent of the processor."

Apple's North agreed that isolating cache traffic from the system bus is important but noted that this can be done in various ways: "The 60x processors are extremely sensitive to memory bandwidth in terms of the effect on performance. In the future, given the right packaging technology, you could get extra pins to provide a dedicated bus. Or you could use some form of inline cache that would isolate the processor and provide the same kind of capability. There's at least a couple ways to do it."

Digital's Alpha is the one high-end processor family not to have a separate external cache bus; Rubinfeld noted that "Digital is also seriously looking at, in the future, having a private cache bus."

Demanding Times Ahead

As this discussion illustrates, bandwidth issues are at the center of system architects' concerns for next-generation systems. Not only memory bandwidth but also I/O bandwidth are becoming more challenging. On the memory side, Rambus holds a glimmer of a solution, but the industry must find a way to make a revolutionary leap to a new approach. This is probably not the only such leap the industry will have to make to gain the full benefit of the high-performance processors we will have available at the end of the decade. ♦

This article includes only a few highlights from the full panel discussion. Audio tapes of the entire Microprocessor Forum are available from MicroDesign Resources.

