Intel's RDRAM Strategy a Sure Winner

CPU Maker Gains High-Bandwidth DRAM—One Way or Another

by Steven Przybylski, Consultant

Intel's adoption of Rambus's Direct RDRAM as its next-generation main-memory technology for PCs (see MPR 12/30/96, p. 4) has far-reaching implications for DRAM vendors, for the PC industry, and for non-PC system vendors as well. Intel's move was motivated by the incessant need to provide more system-level performance. Its strategy is working. By forcing the issue, Intel has ensured that its basic requirements for high-bandwidth memory will be met. Intel is clearly working in its own best interest; the question is how its move will affect everyone else.

If all goes according to Intel's plan, Direct RDRAMs will start appearing in high-end PCs in 1999 and mark the first time that a memory device not sanctioned by JEDEC—the semiconductor industry's primary standardizing body—will be used in mainstream PCs. Not since the introduction of MOSTEK's 4-Kbit DRAM in 1974 has a single company effectively determined mainstream DRAM architecture.

Why has Intel broken with tradition in such a big way? Simply put, bandwidth. Early last year, Intel quantified the bandwidth required to provide the compelling multimedia experience that is central to its long-term strategy of driving PCs into the home. Intel decided it needed an aggregate main-memory bandwidth of at least 1.6 Gbytes/s for high-quality 3D graphics and DVD processing at resolutions of up to $1024 \times 768 \times 16$. By early next century, system-level bandwidth on the order of 3 Gbytes/s will be needed. In contrast, the highest-performance memory systems in today's PCs generate a paltry 533 Mbytes/s.

The DRAM industry was not on track to deliver the bandwidth that Intel believes PCs need. In 1998, PCs will make the transition to 100-MHz memory-bus frequencies. At this frequency, SDRAM memory systems deliver 800 Mbytes/s—about half of Intel's new minimum requirement. Although DRAM vendors project SDRAM could hit up to 200 MHz, in real systems these devices will not deliver anything close to 1.6 Gbytes/s. The existing mainstream DRAM architectures cannot deliver the required bandwidth while simultaneously respecting the memory-size and bus-width constraints inherent in the PC platform. Some new high-bandwidth commodity DRAM is needed by 1999 for Intel to reach its larger goal.

Intel Had Several Options

So what could Intel do? There were basically five options.

1. Intel could wait and see. Intel had told the DRAM community of its requirements for future PCs, and presumably the community would respond. But left on its own, the DRAM industry moves slowly. For example, three years elapsed between the introduction of the first SDRAMs and any meaningful degree of compatibility among the vendors—and this was achieved only after significant effort on Intel's part. Time was short, however, and a passive approach would likely fail: Intel had to become proactive in determining the next generation of DRAM architecture.

2. Intel itself could develop an entirely new DRAM architecture and get the DRAM vendors to manufacture it. Unfortunately, this option would take too long and cost too much money, and the DRAM vendors would be reluctant to commit to another new architecture.

3. Intel could work with the DRAM vendors developing double-data rate (DDR) SDRAMs (now called SDRAM-II). These devices provide twice the bandwidth at any given frequency by transmitting new data on every phase of the clock. Of their several problems, the most important is Intel's belief that DDR SDRAMs will require a new DIMM socket to deliver 1.6 Gbytes/s in a PC. Furthermore, the outlook for data rates beyond about 2.4 Gbytes/s is poor. The prospect of a new socket with a short life span is not appealing. If a new socket is needed, it should ideally last for several generations of systems.

4. Intel could champion the work of the SyncLink Consortium, a different group of DRAM vendors and system companies working on a new very-high-bandwidth DRAM architecture now called SLDRAM. During the first half of 1996, however, the SyncLink Consortium was proceeding relatively slowly and had not decided on basic bus structure or framework. At the time, its work did not inspire confidence.

5. Intel could team with Rambus to ensure the new Rambus II protocol satisfies the needs of future PCs. This is the path Intel chose. Direct RDRAMs will be the result of the collaborative refinement of the original Rambus II specification.

To Intel, the choice between RDRAMs and SLDRAMs was clear: Rambus had a proven track record of delivering cheap, high-bandwidth systems; it was further along than the SyncLink Consortium; and it had existing relationships with many of the first- and second-tier DRAM vendors. Rambus's technology also has headroom, and its module solutions are compatible with the Intel/Microsoft initiative for sealed-case PCs. Nintendo had already demonstrated foolproof RDRAM expansion without the need for opening the box.

Intel's Win-Win Strategy

Fundamentally, the Rambus option is the option most likely to succeed and result in low-cost, high-bandwidth DRAMs that meet Intel's needs, be they Direct RDRAMs or one of the other alternatives. There is a chance that the DRAM vendors will manage, through their redoubled efforts, to make either DDR SDRAMs or SLDRAMs more attractive to Intel than Direct RDRAMs. If this happens, Intel still comes out ahead because it gets the bandwidth it needs. If, on the other hand, Direct RDRAMs are the lowest-cost widely available devices that meet Intel's bandwidth requirements, Intel again gets the bandwidth it needs and coincidentally further extends its control over PC system architecture. Only if none of these three options is available and cost-effective will Intel's multimedia objectives be in peril.

This last outcome is extremely unlikely, because almost everyone in the DRAM community is financially motivated to have at least one of these options succeed, and no one of consequence is motivated to have none of them succeed. Rambus's licensing strategy involves cutting a unique deal with each DRAM vendor. Since each license involves slightly different royalty rates, the DRAM vendors are naturally at odds with one another. Those with marginally better arrangements would benefit from the widespread adoption of RDRAMs: when DRAM prices are depressed, a half-percent difference in cost is significant.

This arrangement makes it difficult for DRAM vendors to work in unison toward some other alternative. Also, within each of Rambus's Far East licensees, there are important individuals who would lose face if the company were to back away from Rambus. The beauty of Rambus's strategy is that companies much larger than itself are financially and emotionally motivated to work toward its success. The beauty of Intel's strategy is that even if Rambus fails, Intel succeeds.

The DRAM Industry Is Losing Control

The most likely outcome is that by the year 2000 or 2001, Direct RDRAMs will be the lowest-cost high-bandwidth DRAMs on the market. What will this mean for the DRAM industry? First of all, only niche DRAM vendors will be able to survive without a Rambus license. Thus, it is not surprising that most of the first- and second-tier vendors that have yet to negotiate licenses are now lining up at Rambus's door. Already, Mitsubishi and Micron have joined Rambus's seven existing DRAM partners in committing to produce Direct RDRAMs. For the remaining DRAM vendors, the relatively short amount of time between now and 1999, and the possibility that the last one in will get the worst deal, makes it dangerous to delay further.

Once in, companies that are still in the DDR SDRAM and especially the SLDRAM camps will have tough resource-allocation problems. Given tight development budgets (due to depressed DRAM prices), limited design teams, the constant need for EDO and SDRAM shrinks, and contractual obligations to develop Direct RDRAMs, these companies may not have enough resources available to actively develop additional new devices.

Many within the DRAM community are concerned that moving control of the mainstream DRAM architecture behind Intel's and Rambus's doors will be bad for the DRAM and system industries alike. They argue that there will be less innovation and progress when the DRAM vendors are no longer able to innovate architecturally. However, the primary DRAM vendors informally decide as a group what the architecture and speed of the mainstream devices will be. Deviation from the herd is not tolerated by the marketplace. Not since the 1970s have individual DRAM vendors had the power to innovate architecturally. In a sense, Rambus and Intel are more able to innovate than the DRAM vendors because they will be able to dictate adoption of new ideas.

More plausibly, the most important requirement for DRAMs is that they be cheap. The DRAM vendors know more than anyone else how to continue making DRAMs as inexpensively as possible. Some argue that moving control of the architecture out of the hands of the DRAM community jeopardizes this long-term trend.

Rambus's royalties are an emotional issue for many in the DRAM industry, yet these royalty relationships are commonplace in the DRAM industry. Texas Instruments, for example, currently derives more income from its DRAM patent portfolio than Rambus can reasonably expect to generate within the next decade. The aggravating issue is not so much royalties *per se,* but new and blatantly aboveboard royalties. Also, because Rambus is an intellectual-property company, its licensing relationships do not have the same sense of reciprocity and *quid pro quo* as do other licensing arrangements in the industry.

System Vendors Will End Up Paying

Intel has made sure Rambus will not get too greedy. As part of its deal with Intel, Rambus agreed to limit its royalties from

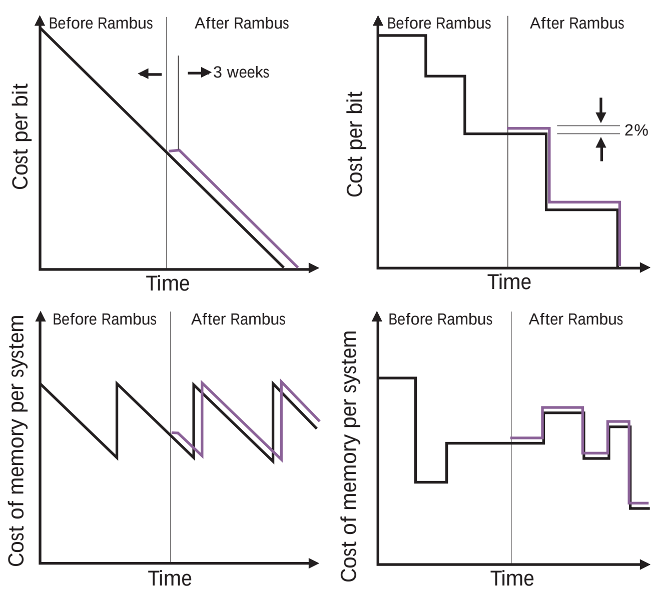

Figure 1. If the price of memory declines steadily, the Rambus royalty effectively translates into a short delay in the transition from one main memory size to another. When the price of memory falls in cascades, the royalty translates into an increase in the average cost of memory in a system.

any RDRAM licensees exceeding some shipment threshold to 2% of RDRAM revenues. Intel argues that in light of an average 35% per annum decline in the price of memory, the 2% royalty on RDRAMs is equivalent to a three-week pause in that decline. In theory, after the transition, everything is as if the royalty were not there. This world view is shown on the left side of Figure 1.

Unfortunately, the price of memory doesn't decline smoothly, but in stair-step fashion, as shown in the right side of Figure 1. When prices are essentially constant, there is a royalty-induced difference between the cost of RDRAMs and that of a truly royalty-free alternative. As a result, the 2% royalty on RDRAMs and the up-to-5% royalties on chip sets and other devices containing RDRAM controllers would add directly to the system cost.

The fiction in this analysis is that there is no truly royalty-free alternative. SDRAMs are thoroughly patented, and when all is said and done, SLDRAMs are likely to come burdened with some royalties as well, even to Consortium members. Today's EDO DRAMs could be considered royalty free in some sense, but they cannot generate the required bandwidth. Furthermore, in the DRAM world, there is a very strong relationship between volume and cost. Even in the presence of a small die-size penalty or royalty, the dominant architecture will become the lowest-cost and lowest-price alternative. Intel expects that as Direct RDRAMs migrate from the high end to the full spectrum of PCs, any initial price premium will entirely dissipate.

To put this in perspective, the cost of memory for a sub-$2,000 PC has traditionally fluctuated between $50 and $200. A real 2% increase in the price of memory translates to between $1 and $4 of added cost to the system vendors, on top of the hidden premium for the memory controller. Given that the price points for PCs are essentially fixed by the marketplace, some portion of these new costs would, at least initially, have to come from either the system feature set or the vendor's already wafer-thin margins.

The designers of high-end consumer products are generally pleased by Intel's adoption of Rambus. The success of Nintendo 64, a high-volume low-cost system based on RDRAMs, is very encouraging. Of the existing memory technologies, Rambus is the best at providing high bandwidth from a small amount of memory.

The developers of workstations and large-scale computer systems are less comfortable. Very wide EDO and SDRAM memory systems have been adequate for meeting the bandwidth, latency, and memory-size requirements of these systems. This community—especially those planning to use high-end Intel processors—is rightfully concerned that Intel is pushing Direct RDRAMs to be too PC-centric. Already suffering from extreme cost and performance pressure from PCs, these companies are worried about the possibility of eventually having to use a suboptimal memory technology because the PC industry has driven its price below that of the alternatives. The prospect of having to directly or indirectly pay royalties for the right to build a memory controller adds insult to injury.

Rambus Set for 1997 IPO

And what about Rambus? The company obviously has a right to be ecstatic. Intel has validated its technology and its years of hard work. For once, the DRAM vendors are knocking on its door instead of the other way around. Yet there is risk, too. Intel has at times been a fickle partner. It could plausibly change its mind again, especially if DDR SDRAMs and/or SLDRAMs prove viable. Also, Rambus is presently focused on its Concurrent RDRAMs, which are making good headway in graphics and multimedia subsystems. There is a danger that in all the commotion over the PC main-memory marketplace, Rambus will lose sight of the current marketplace.

To date, Rambus has derived most of its income from development contracts. In the next few years, as RDRAMs penetrate the graphics marketplace, Rambus will likely accrue as much royalty income from memory controllers as from RDRAMs. If, hypothetically, by 2002 or 2003 Direct RDRAMs account for half of the PC industry's 70% share of an $80 billion DRAM marketplace, at an aggregate DRAM-plus-controller royalty rate of 1.5%, Rambus would gross $420 million in revenue with next to no cost of goods. So, having turned a quarterly profit for the first time in the fourth quarter of 1996, Rambus has set the wheels in motion for an IPO in 1997. With its novel business model and Intel more than just looking over its shoulder, Rambus's IPO promises to be the one of most noteworthy since Netscape's.

Intel Wins Either Way

So what do we have? Intel needed more bandwidth and it had no confidence that the DRAM industry was going to provide it in time. It had to become more proactive in setting direction. Of the alternatives, RDRAM was the most attractive for several reasons.

There are really only two possible outcomes. Either Direct RDRAMs will be cheap and widely used by 1999–2000, or the DRAM vendors will manage to make SLDRAMs more attractive to Intel than Direct RDRAMs. In the latter case, DDR SDRAMs may have a role to play as an intermediate step.

There are several factors—having to do with the structure of the DRAM industry, the state and structure of the Direct RDRAM effort, and Rambus's licenses—that give Direct RDRAMs a big leg up. But even if Intel changes its mind and ultimately uses something other than Direct RDRAMs, it will have succeeded in getting the DRAM industry to provide the bandwidth needed to continue making the purchase of new PCs compelling. M

Steven Przybylski is the principal consultant for the Verdande Group in San Jose, Calif. (www.verdande.com). *He will present the seminar, "New DRAM Architectures for Main Memory and Graphics," at MDR's PC Tech Forum in May.*