THE EDITOR'S VIEW

What Happened to MIPS?

By Michael Slater

Last month's acquisition of MIPS Computer Systems by Silicon Graphics was the end result of a long series of difficulties that have plagued MIPS. While the merger appears to be a positive development, the very fact that being absorbed by one of its customers seems positive is a sign of MIPS' problems.

The ACE initiative, which remains key to the company's future, also caused the company much grief. ACE tried to be too many things to too many companies, leaving itself vulnerable to shifting plans at many other companies and open to attack on many fronts.

ACE set up the MIPS architecture as the challenger to the Intel x86 architecture, and doing so was like waving a red flag in front of a very powerful bull. The heart of the ACE/MIPS story is that R4000-based systems will provide two to three times the performance of Intel-based systems at a comparable price. In response, Intel turned up the hype on its future P5 processor, claiming at numerous conferences that the P5 was only slightly behind the R4000 in schedule and would out-perform it on both integer and floating-point programs. This claim, even though it neglects entirely the issue of price and the fact that the R4000 will be shipping in a faster version by the time the P5 is in production, largely took the wind out of ACE's sails.

The emergence of 386-compatible processors from AMD, C&T, and Cyrix has also hurt MIPS. The now-competitive x86 market is forcing prices down, making it harder for MIPS to maintain a compelling price/performance edge and taking away MIPS' advantage as a multiple-sourced processor.

Since ACE also includes a UNIX component, it prompted attacks from Sun, HP, and IBM. Combined with Intel's attack, concerns about MIPS' financial status, DEC's promotion of Alpha, and troubles at Compaq, ACE was under siege. The press, in a striking example of pack journalism, turned vicious. One MIPS executive commented that "It became clear that it simply wouldn't matter what we did; anything good mostly got ignored, and anything bad got amplified. There were people who wrote very visible articles where they selected the most damaging single thing from a three-minute interview to be quoted, and when they were called on it, basically said they weren't interested in objectivity but in generating a good story." Press reports frequently mentioned financial difficulties at MIPS, but few noted that the company turned a $2.5 million profit in the fourth quarter of 1991—despite all the attacks and the softness in the computer industry—and has $50 million in the bank. While it is true that MIPS' systems business was struggling, the company was not on the verge of bankruptcy, as much of the press—and most of MIPS' competitors—might lead you to believe.

MIPS made some strategic errors in its microprocessor development that resulted from a shifting application focus and limited resources. The R4000 was designed to serve a variety of markets, from high-end multiprocessor systems to mainstream PCs, with one piece of silicon in different packages. For companies like SGI and DEC, however, the R4000 didn't have an aggressive enough performance target. For the PC market, the chip was burdened with unnecessary features.

Completion of the R4000 was delayed by two key factors: the decision to implement a full 64-bit architecture (which, by MIPS' own reckoning, cost six months) and the support for multilevel caches, numerous bus protocols, and multiprocessor systems. The multiprocessor support was particularly complex and troublesome to design and debug, as MIPS tried to satisfy the disparate requirements of several major customers. These features are unimportant to the R4000's success as a competitor to the x86 in mainstream PC markets, yet they delayed the product substantially. If MIPS had had the resources to split its development project in two (one high end and one low end), it could have pursued a 32-bit, no secondary cache part with limited (or even no) multiprocessor support, and this chip could have been in production by the end of 1990. This would have put MIPS in a far better competitive position. The VRX R4000 derivative (see p. 11) now being developed by MIPS and NEC is the type of device that has a real chance to compete in the PC market, but consider how much more potent it would have been two years earlier.

For all of its troubles, the MIPS architecture remains one of the strongest long-term challenges to the x86 architecture's dominance in the PC market. Two of the largest semiconductor makers in the world (NEC and Toshiba) are building R4000 chips; several projects are underway to develop low-cost, low-power derivatives; and Microsoft remains committed to a MIPS version of Windows NT. If MIPS is able to achieve its goal of two to three times x86 performance at comparable pricing, if Windows NT is popular, and if many application vendors provide MIPS versions of their Windows applications, then it has a real chance to penetrate the Intel-dominated world of personal computers. If not, SGI may wish it hadn't inherited ACE in the process of taking control of the MIPS chip-design team.♦