VIEWPOINT

Die Like a Man

The Risk in RISC for Semiconductor Makers

By Nick Tredennick

In a couple of articles last year, (μPR 5/29/91 p. 12 and μPR 6/26/91 p. 12), I picked on the workstation companies (and the trade press) for their continued haughty attitude and blind optimism in the face of my predictions for their imminent troubles as the PC-compatible computers eat into their markets. I've decided to take a new tack and pick on the companies making CPU chips for the workstation companies.

One of the most memorable recordings I ever heard was of radio traffic taped during a dogfight in the air war over Vietnam. It began with an American fighter pilot screaming on the radio: "Help! Help! There's a MiG on my tail! I can't get away!" This monologue continued for some time until one of the other American fighter pilots keyed his microphone and replied: "Shut up and die like a man." I don't know the outcome of the dogfight or the fate of either pilot, but the recording made a lasting impression on me. Fighter pilots have an extremely macho code and blubbering about your imminent death won't be tolerated by your peers.

The CPU chip makers serving the workstation companies seem a lot like the macho fighter pilot who is willing to "die like a man." I don't hear them blubbering and complaining: they seem to be doing everything they can to survive even though they probably know the cause is hopeless.

So who's signed up to die? The four companies making SPARC chips (LSI, Cypress, TI, and Fujitsu) and the six companies making MIPS chips (LSI, IDT, Performance, NEC, Siemens, and Toshiba). I'm not predicting the companies themselves are going to die (Fujitsu could probably give us all horses and still be profitable), but just that they aren't going to make much money on their CPU chip businesses. I think the situation was hopeless when these companies signed up. Could they have known from the beginning the situation was hopeless? I think so. Some of them probably knew there was no business case when they signed up to do the chips. Signing up was a combination of macho culture and wishful thinking. Macho culture: A semiconductor company isn't a real semiconductor company unless it makes a CPU. Wishful thinking: There is so much money in the PC-compatible CPU market something has to happen which will allow someone other than Intel to reap some of the benefit.

Some executive decision-makers probably bought the palaver about the inherent superiority of RISC being great enough to attract users on performance differential alone. They may have overlooked software issues. Software for the x86 is cheap and sells in high volume; software for workstations is expensive and sells in low volume. Software helps drive and sustain the x86 market; software restricts the workstation market. None of the new RISCs will run x86 software directly. Stories about the RISCs being able to emulate the x86 with better performance are apocryphal, but may have influenced the decision nevertheless.

More than 20 million PCs are shipping every year. Intel is selling the bulk of the CPUs for those systems and is probably grossing about $70 per socket for the 10 million 386s and about $360 per socket for the 2 million 486s. That translates to an annual gross revenue of about $1.4 billion. Other semiconductor companies would like to get some of that money. If your company could get just 7% of the CPU sales, it might be a $100 million business. You would be willing to spend a lot of R&D dollars to capture a few percent of that market.

Strategists from some of the semiconductor companies probably thought it was just a high-end microprocessor market and not exclusively an x86 market. Not so. As Intel's Dave House has pointed out, PC-compatible systems using the Intel x86 CPU are shipping more than 21 million a year, Macintosh systems using the Motorola 680x0 CPU are shipping about 1.8 million a year, Sun workstations using the SPARC CPU are shipping about 230,000 a year, and MIPS workstations are shipping about 80,000 a year. Shipments of SPARC and MIPS CPUs are two orders of magnitude below shipments of x86.

The entire workstation market for this year will probably be about half a million units. Sun gets a little less than half of the workstation market and everyone else shares the rest. If you are one of the four semiconductor companies making SPARC, that may seem like a good deal, so let's look at what it means in sales. If one of the semiconductor companies gets half of the sales to Sun and the other three split the rest of the market, that means one company gets to ship 115,000 SPARC CPUs and the other three ship fewer than 40,000 each. The six companies making MIPS processors are in worse shape. If one of the MIPS "semiconductor partners" has half of the market and the other five split the rest of the market, that means one company gets to ship 40,000 MIPS CPUs and the other five ship about 8,000 each.

Research and development costs for one of these microprocessors have been estimated to be between $30 million and $100 million. Let's assume the RISCs are cheaper at $30 million and the x86 development cost is about $100 million. Let's also assume the semiconductor technology and die sizes for the current generation RISC and x86 chips are roughly equivalent, so the number of chips per wafer and the wafer processing cost for each CPU style is about the same. [The 1991 MIPS R4000 die is 468 × 611 mils. The old (1989 original) Intel 486 die is 414 × 619 mils; the current version is half that size. The original R4000 and original 486 designs use 1-micron technology; both are now in production using 0.8-micron versions.]

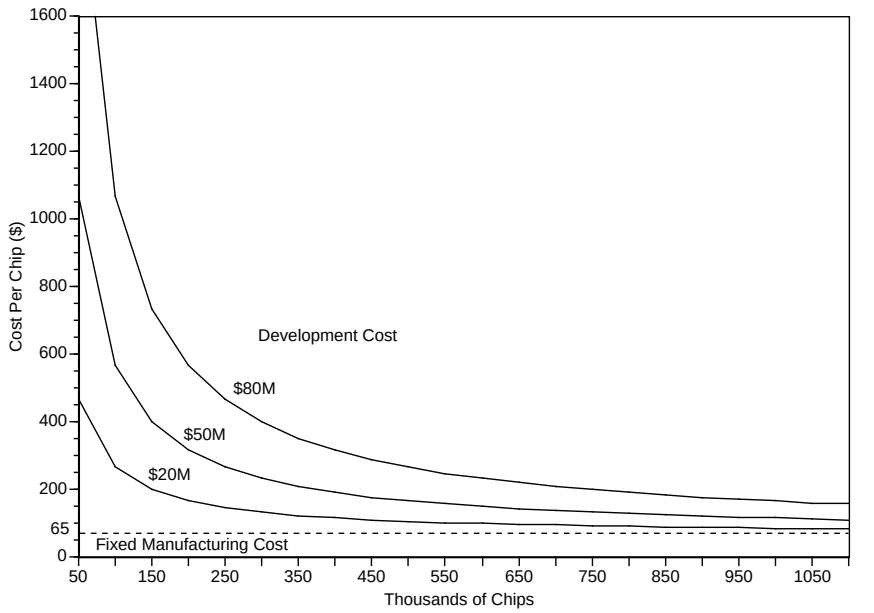

Figure 1. Cost per chip to amortize various development costs, based on a $65 fixed manufacturing cost.

At the original chip sizes, there are about forty chips per wafer. Assuming that everyone is running six-inch wafers at a cost of $600 per wafer and everyone gets about the same yield for a wafer, I'll take a wild guess and suggest all the manufacturers get 12 to 14 good chips per wafer. The processing cost for each of the advanced CPUs is therefore about $45 per chip. To this processing cost we should add about $10 for packaging (more dollars for packages with more pins) and perhaps another $10 for testing. Fixed manufacturing cost for a packaged and tested top-of-the-line CPU is probably about $65. (Reducing the chip sizes reduces manufacturing cost and further polarizes the following analysis.)

That takes care of the fixed cost for the chip; now we have to amortize the development cost across the chip sales to see what the lower bound for pricing will be. For estimating amortization, I'll assume the life of a high-end CPU design to be two years. While the life of a high-end design may be well over two years (as evidenced by the 386, which is only now reaching peak volumes), companies wishing to compete in the high-end CPU market will have to produce a new design about every two years and will, therefore, have to amortize the cost of each development over two years to stay even.

In two years, Intel will probably ship over five million 486 CPUs, so amortizing the development cost adds $20 to each chip, giving Intel a lower-bound price of $85 for the 486 CPU. (This neglects the fact that Intel is already producing the 486 in a second-generation smaller size.) Next in line is the SPARC developer with half of the Sun market. In two years this company might ship about 250,000 SPARC CPUs, so amortizing the development cost adds $120 to each chip, giving this company a lower bound of $185 for its SPARC CPU. The remaining three SPARC developers have lower-bound costs of about $440 for their SPARC CPUs. I don't know how the MIPS semiconductor partner deal works. If we assume MIPS allows their semiconductor partners to share the development cost equally (and without profit to MIPS), each will pay a $5 million share of the R4000 development cost. The MIPS partner with half of the market will ship 80,000 CPUs in two years, so they will have a lower bound of about $125 for their MIPS CPU. The other five MIPS partners will each ship 16,000 CPUs in two years, so they will have a lower bound of about $375 for their MIPS CPUs. If we assume MIPS gives the design to the partners and only collects a royalty on the part, the lower bound for cost for all the partners would be about the same and be somewhat over $250.

Intel is a special case, since they are also shipping about 10 million 386 CPUs per year. I feel comfortable assuming they amortized the design cost of the 386 in the first two years it was shipping. Since the die for the 386 is much smaller than the 486, Intel is probably making that CPU for less than $10. If they amortized the 486 development cost into the cost of the 386 once the development of the 386 was fully amortized, it would amount to perhaps $3 per chip. (They probably didn't have to do this since the cost of 386 development was probably covered by extra money from the 286 design, and so on.) If the 386 is selling for more than $20, there is some extra money going somewhere. It is. Price and cost are vastly different things. The current, 1000-piece price of a 33-MHz 386DX from Intel is $125. Intel would still be making money if they gave away a

free 486 with every purchase of a 386DX. At their current volumes, there is no way the SPARC and MIPS chip makers can compete with Intel on cost. Costs are lower for Intel.

Figure 1 illustrates the dilemma faced by the SPARC and MIPS chip vendors. The figure plots the cost per chip against chip shipments for various development costs. Figure 1 assumes the manufacturing cost of a good die is $65. If the development cost for the chip is $40 million and you expect to ship 800,000 units, then the cost per chip is $115 ($50 of amortized development cost plus $65 manufacturing cost). Here's the dilemma: Intel, with its x86, is always operating well off the right side of this chart (millions of units), while most of the SPARC and MIPS chip vendors are operating well off the left side of this chart (fewer than 50,000 units). It's not a happy situation for anyone except Intel. Wishful thinking takes over in predictions for workstation CPU volumes off the right side of Figure 1.

Cypress is attempting a novel solution to this problem—they plan to offer Pinnacle prototypes for $10,000 each. If they could sell just 3000 samples at this rate, Cypress could amortize the development cost and price production parts to compete with the 486. (This assertion rests on the risky assumption that the cost to produce the module is well below $10,000.) This strategy might work for Intel, since there probably are 3000 customers wanting early 586s that badly. For Cypress, the strategy may amount to no more than the ability to announce the availability of parts without the risk of having to ship any.

I divide the CPU market into four segments: mid-range x86, low power, high-end x86, and high-end any-architecture. I listed them in what I presume is decreasing market size. The mid-range x86 market is the PC-compatible market. The high-end x86 market is the high end of the PC-compatible market. These markets are open only to x86 CPUs and are not open to encroachment by any other architecture. The low power market is primarily an x86 market, but it is still early enough, particularly with regard to pen-based computers, to be open to any-architecture. The high-end any-architecture market is open to any-architecture (including the x86).

So far, the SPARC and MIPS implementations have been designed for the high-end any-architecture market. A couple of years ago, most of the nine companies supplying the high-end any-architecture market noticed that the volumes were too low to support nine companies. (Intel ships almost as many 32-bit CPUs per week as all the SPARC and MIPS vendors together ship all year.) At this point, executives embarked on several not-necessarily-exclusive strategies:

- Slow or stop development of high-end any-architecture CPUs.
- Continue developing high-end any-architecture CPUs and hope for the best.
- Form or join strategic alliances to influence the future of the workstation market.
- Develop derivative products based on the current design to help amortize development cost.
- Increase emphasis on software or other product lines.

For the five large companies (Siemens, TI, Fujitsu, NEC, and Toshiba), it doesn't matter whether signing up for SPARC or MIPS was just a bad decision or a deliberate hedge, and their subsequent decisions may likewise be unaffected by the economic logic of the situation. Hedging is high-technology lottery. "There's a 0% chance Intel will license high-end x86 for us to produce at a favorable rate. There's a 0.1% chance ACE will succeed and the MIPS architecture will displace the PC in high-volume shipments. Let's sign up to sell MIPS chips just in case ACE is successful. It'll only cost us a few million, it won't even show up on the bottom line, and it could be a big win."

For the smaller companies, the choices are difficult. The best strategy is probably to dump the hardware and sell specialized software for the PC-compatible market. Sun and Silicon Graphics seem to be embarking on this path, though their internal strategic planners may or may not be aware of it. (Strategic planners probably find the enormous volumes and cheap hardware in the PC-compatible market financially attractive for their niche software products, irrespective of whether they believe PC-compatibles will eventually displace their own workstation hardware in the market.) The strategy might work for a workstation company with proprietary software and an installed base. The chip vendors have neither. Strategic alliances are not likely to help the chip vendors much either. Alliances like SPARC International and ACE are designed to benefit the workstation companies; any benefit to the chip vendors is incidental. And even though these alliances are designed to benefit the workstation companies, their ultimate effect will be to benefit the PC-compatible business. (For rationale, see µPR 6/26/91 p. 12.)

The best-looking alternative for the chip vendors appears to be the development of derivative products. So far I have been talking about the high-end any-architecture segment of the CPU market. High end microprocessors used as a CPU in a system are a tiny fraction of the microprocessor market. The rest of the microprocessors (well over 98%) go into embedded control applications. The embedded control market is over a billion units per year (predicted to exceed 1.7 billion units in 1992). Strategic planners for the SPARC and MIPS chip vendors probably eye this market with a certain infatuation. Any market with a billion-unit volume must have some room for a few more vendors.

Viewing the embedded control market as a single large market open to any microprocessor would be the same mistake as viewing the microprocessor-as-a-CPU market as a single entity. The embedded control market may be open to any-architecture, but it isn't open to any microprocessor. For the convenience of this discussion, I'll divide the embedded control market into two segments: commodity and boutique. The commodity market segment is characterized by very low margins, very low average selling prices, and very high volumes. The commodity market is over a billion units. To compete in the commodity market, you must have an excellent custom design and high-volume manufacturing capability. The boutique market segment is everything else and represents approximately zero percent of the embedded control market.

The current SPARC and MIPS designs, even in their embedded control costumes, don't qualify to compete in the commodity segment of the embedded control market. Still, the boutique segment of the embedded control market is millions of units and has more opportunity than the high-end any-architecture segment of the CPU market. But here the SPARC and MIPS chips, designed for the high-end any-architecture CPU market where performance is king and cost and power are secondary, are being adapted to compete in the embedded control market, where cost, function, and power consumption are often more important than raw performance. Motorola, National, AMD, Intel, and others design microprocessors for specific segments of the embedded control market. I doubt they view competition from adapted SPARC and MIPS designs with much trepidation. Further, prices for embedded control microprocessors don't have the advantage of the artificially inflated price floor characteristic of the CPU market (supported by the high profit margins Intel is able to command) to help amortize development costs. It's a tough market for SPARC and MIPS; even if they sell in the boutique segment of the embedded control market, these designs will never make the transition to the commodity segment.

There's one more market opportunity to look at: the low-power segment of the CPU market. The pen-based notebook computer market may be substantial someday, but the architecture window is probably closing with the three candidates being ARM, Hobbit, and the x86. For this market, power dissipation is king and everything else is secondary. The current SPARC and MIPS designs don't qualify to enter this segment, but at least they are not yet locked out.

The situation looks bleak, but I'm impressed: these companies are trying everything they can and I don't hear any blubbering and complaining. At least they are going to die like men in the grand tradition of fighter pilots. ♦

MIPS ARCsystems

Continued from page 9

Conclusions

These first-generation ARC systems should be competitive in the workstation market, but other than serving as software development vehicles, they have little to do with the PC business because of their high prices and lack of appropriate application software. It is a shame that MIPS didn't price the systems more aggressively, but the company can't afford to buy market share.

In the next two years , MIPS should finally be able to reach mainstream PC price points while offering significantly higher performance. By then, of course, there will be a half dozen or more makers of x86-compatible microprocessors, and it is hard to know just what the QED and VRX chips will be up against. At a minimum, they will have to compete with 486-based systems at the low end and P5-based systems at the high end—not to mention SPARC-based systems and the other RISC competitors.

So far, only Acer, Olivetti, and a start-up called Carrera Computers (Laguna Hills, CA) have been announced as ARCSystem licensees. Of the 90 system vendors that ACE claims to have in its camp, few have demonstrated any level of commitment to building ARC systems. Just as with SPARC-based systems, there is a promise of numerous PC makers manufacturing ARC systems, creating a competitive market for binary-compatible RISC workstations. So far, more companies have failed than have succeeded trying to do this with SPARCstations, and few have achieved a volume of even a thousand units a month—tiny by the standards of the PC market.

At the moment, there are few companies with a demonstrated, serious commitment to making ARC systems, and it remains to be seen whether or not it will reach critical mass as a standard. Compaq's recent decision to shelve its ACE plans leaves the initiative without a first-tier PC maker, and DEC's religious pursuit of Alpha has raised serious questions about its commitment to MIPS-based systems. SGI's proposed purchase of MIPS further clouds the issue, though it seems that this is the least of ACE's problems.

While it seems very unlikely that ARC systems will ever become a dominant standard, it is too early to count it out as one of the top few survivors of the early 1990s. While ACE has largely collapsed, its core—a standard for MIPS-based systems that will run Windows NT as well as UNIX—is progressing reasonably. There won't be any quick victories in the battles to take some of Intel's market share, but there will probably be some slow, hard-fought successes, and MIPS, SPARC, and PowerPC remain the top contenders. ♦