THE EDITORIAL VIEW

Is It Soup Yet?

Sometimes It's Hard to Pin Down Whether a Chip Is "Available"

Most processors follow a familiar product life cycle: first silicon is built and tested internally, a revised version is sampled to a few key customers, and then a final version is completed. At that time, the clouds part, the sun shines, and a new processor is available in volume to any company that wants to buy it. In the life of a processor, this event is like graduating from college and heading off to a first job.

Some processors, it seems, head for graduate school instead. These chips have a big announcement ceremony but don't appear in systems for a while. Or perhaps only a few customers can get parts. Another common trick is to put the chip only in very expensive systems that not many people will actually buy.

A recent example of this phenomenon is Sun's UltraSparc-2 processor. Last June, Sun Microelectronics (SME) said volume availability of the 250-MHz UltraSparc-2 would begin in about a month; one representative even said sister company Sun Microsystems (SMCC) was shipping the 250-MHz part in a few high-end servers and would soon follow with high-volume workstations. The processor existed in this nebulous state for months, with SME constantly claiming that something big would happen soon.

In fact, UltraSparc-2 had run into some nasty manufacturing problems. Texas Instruments was using the chip to bring up its 0.29-micron CMOS process and had problems getting the chip to yield well. For periods of weeks, the chip did not yield at all. After months of trying, the process is finally producing "thousands" of working parts.

During 1996, SMCC shipped no servers with UltraSparc-2 processors. The system house was redesigning its products for the higher bus speeds needed to get the best performance out of the 250-MHz part; this rework took longer than expected. Last month, SMCC began shipping its first UltraSparc-2 servers, with prices starting at $46,000. Higher-volume workstations are planned to follow "soon."

This sequence of events caused us to mistakenly report the chip was shipping last year, although it is now clear it was not. Even today, shipments are relatively low, but we now consider UltraSparc-2 to be shipping in volume. Given Sun's track record, it remains unclear when the 300-MHz parts originally promised for the end of 1996 will surface. We wish processor vendors would be more forthcoming about the state of their products.

Pinning down a volume shipment date is important because, for the most part, we try to compare processors that are currently available against each other, not against some vendor's vaporware. When we compare products that aren't shipping, we try to choose chips that we expect will reach volume availability at the same time. When a chip is in a nebulous state, it's hard to make fair comparisons.

A less extreme example is the R10000. This chip was announced at 200 MHz in January 1996. To its credit, Silicon Graphics (SGI) quickly rolled out a slew of servers using the part; the entry price for these R10000 systems was $89,000, and boxes with a 200-MHz CPU were priced well into six figures. In addition, NEC said it was charging an outrageous $3,000 for the 200-MHz CPU chip alone, and Toshiba refused to even quote a price.

Like UltraSparc-2, the R10000 had encountered manufacturing problems, although not as severe. The initial design had low yields but supplied enough chips for the relatively high priced servers. Yields at 200 MHz, however, were vanishingly small. SGI eventually made shipments at 195 MHz but not quite 200. Our dilemma: do we describe the chip as a 200-MHz design, as specified, or a 195-MHz design, as manufactured? We chose the latter.

After about six months, yields on the R10000 improved significantly, and SGI was able to announce a line of workstations using the R10000, although again the 200-MHz version did not appear anywhere in SGI's lineup. Apparently, NEC and Toshiba have managed to squeeze out enough 200-MHz chips for Siemens-Nixdorf to announce a high-end server product, and we recently published benchmark results for that system.

In contrast, a vendor like Intel samples more chips than a company like SGI ships even in its highest-volume products. Intel would claim that, by its standards, none of SGI's chips is in volume production. We try to hold Intel to a different standard than a RISC processor maker; in either case, we consider a chip to be in volume production when it is widely available to the vendor's customers.

To be more specific, we don't count systems from tiny companies (e.g., DeskStation) or systems with six-digit price tags as volume production. Occasionally, we may penalize a processor because the lead system vendor slipped its schedule. But these rules of thumb are useful in separating the chips that are really shipping from the ones that just think they are. Potential buyers should look for real chips instead of accepting hype. M

Linley Gwennap