THE EDITOR'S VIEW

What's Taking So Long?

High-Performance Processors Slip and Slide on Way to Production

In November of 1992, MIPS Technologies (MTI) announced the R4400, Intel made detailed disclosures about Pentium, Digital announced 200-MHz Alpha systems, HP revealed its PA7100LC, and the IBM/Motorola team presented the PowerPC 601. One year later, the first 601-based systems have just started shipping, and it's still tough to get a 66-MHz Pentium or 150-MHz R4400. The 200-MHz Alpha processor began volume shipments in August, but HP still has not announced availability of the 7100LC. What's going on?

Alert readers may catch the scent of vaporware here, but that is only part of the problem. While HP may have revealed the 7100LC before it had taped out, other vendors had functional silicon in hand before making their public statements. It appears that vendors are having problems coping with ever-more complex processors and the IC processes needed to support them.

First silicon of Pentium, for example, was received in mid-May of 1992 and quickly booted Windows. The company struggled with yields on its new BiCMOS process, however, and when Intel finally announced availability of Pentium last May, yields at the target clock rate of 66 MHz were still low. The vendor hopes to solve many of its problems by moving Pentium to an even newer 0.6-micron BiCMOS process, but this process will also have a lengthy learning curve.

When MTI announced the R4400, it rated the new chip at 150 MHz and claimed that parts would ship in 2Q93. We later discovered that these were separate claims; the parts shipping in 2Q93 could function at just 100 MHz, the same speed as the two-year-old R4000. The faster parts rely on a particularly tricky 0.6-micron process that NEC and Toshiba have just recently gotten the hang of; they are now delivering 150-MHz chips.

Digital announced availability of 200-MHz systems for 1Q93 but, to be fair, said that chips at this frequency would not ship in volume until 3Q93. The original 0.75-micron chips required a 3% voltage boost to hit 200 MHz. Digital shrunk the part to 0.68-micron CMOS to improve the yield at that speed. The company has a similarly long lead time for its newly announced 275-MHz parts: volume production is not expected until 3Q94. This delay is due entirely to the time needed to bring up and qualify a 0.5-micron process; the functional changes in the new parts are minor.

I would be remiss to neglect the tribulations of recent SPARC processors. In May of 1992, Sun announced its first SuperSPARC systems, while Ross Technology announced a 66-MHz hyperSPARC module. Ross' problems are documented on page 10; suffice it to say that the company had both functional problems and yield problems before recently announcing shipments.

Sun's 40-MHz SuperSPARC systems did not ship in volume until early 1993, and the promised 45-MHz systems were obsoleted last April, before they ever shipped in volume. The problems seem to have stemmed from Sun's inability to completely model the timing of the 3.1-million-transistor chip, and TI's problems with a new BiCMOS process. This story has a somewhat happy ending: TI now says it is shipping 60-MHz SuperSPARCs in volume using a 0.7-micron process.

Despite their best efforts, these vendors have found it difficult to shrink the time-honored gap of twelve months between first silicon and volume shipments. Intel built a complete model of Pentium using Quickturn's FPGA system and ran millions of line of code before the first parts were ever fabricated. Other companies have tried more efficient verification of the first parts and quick turnarounds to fix defects.

Such techniques improve time-to-market, but no one, it seems, has found a reliable way to meet clock-frequency goals. Since high-end processors typically use new IC processes that are developed in parallel with the CPU design, the exact process parameters are unknown when the circuits are being designed. Many vendors seem unable to predict clock rates in this uncertain environment. (HP and IBM do better in this regard.)

These problems don't affect less-aggressive designs. Sun, for example, was able to meet its clock rate targets for microSPARC and microSPARC-2 *(see* ***071501.PDF****)*. In the x86 world, vendors such as TI, AMD, and Cyrix quickly develop their 486 derivatives *(see* ***071504.PDF*** *and* ***0715MSB.PDF****)*. These chips generally take advantage of CPU cores and manufacturing processes that are known to be good.

Unless timing verification tools and process development times improve, speed and availability announcements of high-end processors should be not be taken as fact. In particular, unproven IC processes are just that; any glitches or miscalculations can delay availability or even reduce the rated clock speed until a next-generation process can take up the slack. *Caveat emptor.* ♦

Linley Gwennap

