August 22, 2002

Faster Chips That March to Their Own Improvised Beat

By ANNE EISENBERG

THERE are no jazz riffs in the rhythms of electronic circuitry. Instead, all the operations of a microprocessor are governed by a metronomic central clock - a crystal oscillator that beats out a single lockstep rhythm more relentlessly than Toscanini on a bad day.

The calculations done by a chip are timed by this tiny clock, so that answers can be produced at a specific instant and permit the next logical step to proceed.

If the clock distribution system sends the timing signal to the circuits at a rate of 1.5 gigahertz, that means the clock is ticking a billion and a half times a second. All the tens of millions of additions and multiplications, as well as all the other operations, will be synchronized to this beat, whether each operation needs that sliver of time or not.

In the last few decades, though, a small number of circuit designers and other researchers have decided that they want their circuits to march to a different, clock-free drummer that lets them proceed at their own best speed.

To that end, these researchers have developed self-timing, or asynchronous, circuits. This ability, they argue, will lead to improved computer performance, providing faster operations and reduced power consumption, particularly in the ever smaller, ever faster circuits of the future.

At Sun Microsystems Laboratories in Mountain View, Calif., a dozen or so engineers are focusing on designing asynchronous circuits, said Jo Ebergen, an engineer who leads the group with Ivan Sutherland.

Dr. Ebergen invoked the metaphor of a bucket brigade to explain the advantages of asynchronous systems. "In a brigade, each person has to take and send along buckets, but the rhythm can't be faster than the slowest person with the heaviest bucket," he said. In synchronous computers, the clock's basic pace must be slow enough to accommodate the slowest action within the group.

An asynchronous circuit, in contrast, is like a bucket brigade in which each person acts based on local circumstances, passing a bucket to the next person as soon as the person is ready, or taking a rest when no buckets are needed. Instead of being driven by the inflexibility of a clock, an asynchronous system uses local coordination circuits to exchange signals that a job is done.

Advances in the speed and complexity of integrated circuits will soon force designers to learn the newer asynchronous techniques and their potential, said Steven Nowick, an associate professor of computer science at Columbia University.

"The number of transistors is going through the roof," he said. Within a few years, 100 million to one billion transistors are expected to reside on a single chip. "In another 5 to 10 years, there's not going to be any other way to handle systems this complex," he said.

Asynchronous designs are not a new idea in circuitry: the mathematician Alan Turing and other computing pioneers experimented with them. But advances in the technology, along with the pressure to deal with a future of more complex chips, have stirred greater interest in the subject, Dr. Nowick said. As an example, he cited a workshop convened earlier this month by the Defense Advanced Research Projects Agency to consider increases in financing for asynchronous computing.

Some asynchronous circuits are already used in Sun's UltraSparc IIIi processor chip, Dr. Ebergen said. He said his team was working with other product groups within Sun to use asynchronous technology in applications, similar to UltraSparc, where circuits help in communication between the memory and the memory controller.

"As technology improves," he said, "more and more asynchronous techniques will have to be applied." All the clock domains on a chip have to communicate, but domains sometimes have different rates. "Whenever there is an interface, asynchronous circuit techniques will have to be used," he said, "and this will only increase as more clock domains appear on a chip."

While the Sun group is concentrating on the potential speed of asynchronous circuits, the asynchronous circuits group at Philips Research in Eindhoven, the Netherlands, is exploiting other potential benefits, including the possibility of lower power consumption and reduced electromagnetic emission.

Ad Peeters, a senior scientist at Philips Research who leads the group, said that several popular Philips products on the market already use asynchronous circuits, including a microcontroller used in pagers and integrated circuits for smart cards.

He said that asynchronous circuits had recently begun gaining acceptance among skeptics trained in synchronous logic. "At first, we had to be careful not to scare customers by the 'A-word,' " he said. "But Philips's successes in the smart-card market make it easier to come out and be proud of our achievements."

Erik Brunvand, an associate professor of computer science at the University of Utah, predicted that asynchronous systems would gradually come into their own as they provide a clear benefit. The Philips pager with its asynchronous microcontroller is an example, he said, because it eliminated the heavy beat of clock signals that produces electromagnetic interference, interrupting radio frequency communication.

While he would prefer systems that are completely asynchronous, Dr. Brunvand believes that asynchronous circuits would initially find a use as handmaidens to the dominant synchronous technology. "One of the ways industry will latch on is to use asynchronous techniques to get synchronous domains connected and working reliably," he said.

In one approach, local domains are all synchronous but communicate with one another asynchronously. In a variation, a single clock signal coordinates an entire system, but local pieces signal their completion asynchronously. "Either way, there's an interesting mixing of these different approaches," he said.

Jan Rabaey, a professor of computer science at the University of California at Berkeley, said that the globally asynchronous, locally synchronous patterns were starting to appear on chips. "As transistors get smaller, down in the 30- to 50-nanometer range, asynchronous connects will play a bigger and bigger role," he said. "In 5 to 10 years it's going to be the right thing to do."