



It's Time for Clockless Chips

By Claire Tristram
October 2001

Megahertz, shmegahertz. A few iconoclasts are building computer chips that dispense with the traditional clock. But they face big barriers in bringing their idea into the mainstream.

Digital anarchist: Karl Fant of Theseus Logic wants to take chip technology in a radical direction. A small cadre of researchers and companies are coming around to the idea.

"We're replacing dictatorship with anarchy!" Karl Fant tells me emphatically. Ponytailed and animated, the founder and chief technical officer of Theseus Logic fills the whiteboard with sweeping illustrative examples, kneeling down to use every bit of available writing space. He is in his socks. "Eventually every chip will be designed this way," he declares. "It's inevitable!"

Even in Silicon Valley, where company founders are known to indulge their nonconformist tendencies, Fant's Sunnyvale, CA, office comes as a surprise. His low desk is covered by a formless mass of memos and transcripts and other paper stuff, all mounding slightly toward the middle. There are no chairs — only pillows strewn artlessly about on the floor. If you happen to be me, you begin to regret wearing a dress and wonder where exactly you're meant to sit. But no: Fant leads you to a conventional conference room next door, where, thankfully, there is a chair. That's where he begins to evangelize about the coming revolution intended to wrest computer chips from the constraints of the past.

How? By throwing out the clock, the fundamental way that chips, since the dawn of the Computer Age, have organized and executed their work. Even those of us who know nothing about microprocessors know something about their clocks — Intel for years has used the clock speed of its microprocessors as a marketing tool, where faster is better. The number that dominates most computer ads, along with price, is a label like "1.3 GHz" (or gigahertz). That figure refers to the speed of the clock that governs the internal operation of the machine's microprocessor. Within every one-gigahertz microprocessor, for instance, there lies an oscillating crystal ticking one billion times a second. Engineers are trained to design chips where their first consideration is getting work done before the next clock-tick comes around. A chip without a clock would be about as useful as a page of text without any space between the letters. For most chip designers, throwing out the clock is difficult to imagine.

But not for Fant or his fellow iconoclasts working on clockless chips at startups, universities and corporate labs. It's a small group of ardent believers. Their annual conference attracts only a few hundred participants. Leaders in the field know one another well, and have one another's cell-phone numbers memorized. But while their methods and markets differ, they are united in their belief that clocked chips have run their course, and stand convinced that the advantages of their maverick approach, known alternatively as "asynchronous design" or "self-timed circuits," are so great that the chip industry will ultimately have no choice but to embrace it.

"Designers are realizing that distributing a clock across ever more complicated systems is becoming more and more difficult, and that sooner or later it won't work," says Alain Martin, a professor of computer science at



Caltech, who built the first clockless microprocessor in 1989. He points out that as chips get more complex, more and more of the power it takes to run them gets eaten up by the clock itself, which now needs to coordinate the work of millions of transistors.

Dispensing with this overhead confers large advantages on asynchronous chips. One is vastly improved electrical efficiency, which leads directly to prolonged battery life. The clockless technology also yields an edge in computing speed. In labs at Sun Microsystems, Intel and IBM, clockless chips have increased the pace at which high-end processors do their work. In 1997, Intel developed an asynchronous, Pentium-compatible test chip that ran three times as fast, on half the power, as its synchronous equivalent.

At Theseus, Fant has focused on still another benefit of asynchronous design. Because these chips give off no regularly timed signal, the way clocked circuits do, they can perform encryption in a way that is harder to identify and to crack. Improved encryption makes asynchronous circuits an obvious choice for smart cards — the chip-endowed plastic cards beginning to be used for such security-sensitive applications as storage of medical records, electronic funds exchange and personal identification.

Are Fant, Martin and other clockless champions right? Frankly, yes. And yet despite the technology's clear advantages, clockless chips remain more theory than practice. The Intel device, for instance, never made it out of the lab. The failure of clockless chips to gain ground, in fact, makes them a perfect case study of a development with overwhelming promise that nevertheless faces huge obstacles to market introduction — even in an industry known for continuous and rapid innovation.

The Path Not Taken

The founders of modern computer technology contemplated asynchronous design as early as 1946. But these early computer engineers chose instead to go with a clock. "At the time, it was the right choice," says Jo Ebergen, a senior staff engineer at Sun who works in an asynchronous research group headed by Sun fellow and vice president Ivan Sutherland. (In 1989, Sutherland, best known as a pioneer in computer graphics, wrote a paper that nearly single-handedly reignited interest in clockless-chip technology.) "The circumstances in which they had to design, using vacuum tubes and relay circuits, meant that they really couldn't build a reliable computer without a clock governing the whole thing," he adds. By using a clock, engineers could build in fail-safe measures that made computers reliable even when the parts they were made from weren't.

From that first choice came the steamroller effect of Moore's Law, wherein nearly all research, development and production in the semiconductor industry has focused on clocked chips. By the 1960s, the notion of clockless chips had all but disappeared — kept alive only by an esoteric paper or two coming out of universities. In today's chips, therefore, the clock remains the key part of the action. As a microprocessor performs a given operation, electronic signals travel along microscopic strips of metal — forking, intersecting again, encountering logic gates — until they finally deposit the results of the computation in a temporary memory bank called a register. Let's say you want to multiply 4 by 6. If you could slow down the chip and peek into the register as this calculation was being completed, you might see the value changing many times, say, from 4 to 12 to 8, before finally settling down into the correct answer. That's because the signals transmitted to perform the operation travel along many different paths before arriving at the register; only after all signals have



completed their journey is the correct value assured. The role of the clock is to guarantee that the answer will be ready at a given time. The chip is designed so that even the slowest path through the circuit — the path with the longest wires and the most gates — is guaranteed to reach the register within a single clock-tick.

With a central timepiece governing the action, engineers don't have to worry about the varying lengths of millions of infinitesimally small wires; signals can arrive at the register in any order, as long as they all settle in before the clock next ticks. Teams of hundreds of engineers can coordinate their work around the unifying principle of the clock. And we all benefit: the discipline of clock-based design has enabled the magic of exponential growth in chip performance to endure for more than 30 years. "The clock has to go down as one of the most brilliant ideas in design," says Kevin Normoyle, a Distinguished Engineer at Sun who works on the design of Sun's Sparc microprocessors. "It's so simple, and yet it's an approach that has scaled up and now works for millions of transistors."

But after a point, cranking up the clock speed becomes an exercise in diminishing returns. That's why a one-gigahertz chip doesn't run twice as fast as a 500-megahertz chip. The clock, through the work it must do to coordinate millions of transistors on a chip, generates its own overhead. The faster the clock, the greater the overhead becomes. The clock in a state-of-the-art microprocessor can consume up to 30 percent of the chip's computing capability, with that percentage increasing at an ever faster rate as clock speeds increase. It's as if a factory became overrun with stopwatch-wielding supervisors who improved efficiency but also took up more and more space held by workers and machines.

Clocked chips are becoming serious power hogs, too: the job of coordinating tens of millions of transistors at a billion ticks per second requires the consumption of a lot of energy, most of which ends up as heat. Patrick Gelsinger, chief technology officer at Intel, referred to the problem in his keynote speech at the International Solid-State Circuits Conference last February. Gelsinger was only half-joking when he said that if microprocessors continue to be run by ever faster clocks, then by 2005 a chip will run as hot as a nuclear reactor.

Perhaps the most pressing problem with conventional microprocessors, though, is that you can only speed up the chip's clock so much before banging into some inconvenient physical realities. In today's one-gigahertz chips, electronic pulses signifying binary ones and zeroes can — just barely — make it across the chip within a single beat of the clock. But in the two-gigahertz chips expected to arrive in the next couple of years that will no longer be true. The role the clock plays now, synchronizing all the work on a chip, will begin to break down.

Clockless to the Rescue

By throwing out the clock, chip makers will be able to escape from this bind. Clockless chips draw power only when there is useful work to do, enabling a huge savings in battery-driven devices; an asynchronous-chip-based pager marketed by Philips Electronics, for example, runs almost twice as long as competitors' products, which use conventional clocked chips.

Like a team of horses that can only run as fast as its slowest member, a clocked chip can run no faster than its most slothful piece of logic; the answer isn't guaranteed until every part completes its work. By contrast, the



transistors on an asynchronous chip can swap information independently, without needing to wait for everything else. The result? Instead of the entire chip running at the speed of its slowest components, it can run at the average speed of all components. At both Intel and Sun, this approach has led to prototype chips that run two to three times faster than comparable products using conventional circuitry.

"Look at it this way," says Sun's Ebergen. "You give me a folder, I work on it, I give it back to you, and the fact that I give it back indicates I'm done. We don't have to communicate every five seconds. We might do the job much faster by agreeing between the two of us when to get things started and when to get things done and not worry about synchronizing our work every step along the way."

Another advantage of clockless chips is that they give off very low levels of electromagnetic noise. The faster the clock, the more difficult it is to prevent a device from interfering with other devices; dispensing with the clock all but eliminates this problem. The combination of low noise and low power consumption makes asynchronous chips a natural choice for mobile devices. "The low-hanging fruit for clockless chips will be in communications devices," starting with cell phones, says Yobie Benjamin, a technology strategist for the consulting firm Ernst and Young. So convinced is Benjamin of the technology's promise that he has personally invested in Fulcrum Microsystems, a clockless startup out of Caltech.

Two other new firms, Theseus and Manchester, England-based Self-Timed Solutions, are focusing on clockless chips for smart cards. Fant maintains that a key problem holding back smart cards is that conventional chips make it easy to crack the chip's security codes by watching the signals. "The clock is like a big signal that says, 'Okay, look now,'" says Fant. "It's like looking for someone in a marching band. Asynchronous is more like a milling crowd. There's no clear signal to watch. Potential hackers don't know where to begin."

Speed, energy efficiency and stealth sound like important goals for any chip, not just those used in a few niche applications. But while Sun, IBM and Intel all have small research groups working on asynchronous designs for specialty applications, neither they nor anyone else has announced work on a general-purpose clockless microprocessor. This seems an odd oversight. An industry that considers the improvement of processor speed to be an almost sacred goal has forsaken one of the most promising avenues for making chips go faster. You just have to ask why.

Why, for example, did Intel scrap its asynchronous chip? The answer is that although the chip ran three times as fast and used half the electrical power as clocked counterparts, that wasn't enough of an improvement to justify a shift to a radical technology. An asynchronous chip in the lab might be years ahead of any synchronous design, but the design, testing and manufacturing systems that support conventional microprocessor production still have about a 20-year head start on anything that supports asynchronous production. Anyone planning to develop a clockless chip will need to find a way to short-circuit that lead.

"If you get three times the power going with an asynchronous design, but it takes you five times as long to get to the market — well, you lose," says Intel senior scientist Ken Stevens, who worked on the 1997 asynchronous project. "It's not enough to be a visionary, or to say how great this technology is. It all comes back to whether you can make it fast enough, and cheaply enough, and whether you can keep doing it year after year."



Philips's asynchronous chip has given the company's pagers the ability to last almost twice as long, on the same battery power, as clocked alternatives. But its debut in 1998 followed a decade of dedicated research. Asynchronous researchers from the beginning understood that their task wasn't just to build another chip, but rather to build a way to design, test and manufacture that chip. And that wasn't easy.

Playing Catch-Up

The first huge barrier to bringing clockless chips to market is the lack of automated tools to accelerate their design. Twenty years ago, a handful of engineers could lay out a chip's circuitry on paper. Today, hundreds of engineers work in teams, and the only hope of coordinating their actions is to use sophisticated computer-aided tools. But asynchronous designers face a chicken-and-egg problem: if there is no mass market for asynchronous chips, there's little incentive to create tools to build them; if there are no tools, no chips get produced. The same problem applies to the development of chip-testing technologies. Without any significant quantity of asynchronous circuits to test, there is no market for third-party testing tools.

In the case of its pager chips, Philips decided the only way out of this trap was to itself invest in developing the tools it needed. "After 13 years of research, we are now close to an effective and efficient test approach for asynchronous circuits," says Philips research fellow Kees van Berkel, who has worked on the Dutch giant's asynchronous team since the early 1980s. And Philips is not alone in this quest. In an effort to create momentum for asynchronous chips, two computer scientists — Steven Nowick at Columbia University and Steve Furber at the University of Manchester — have each developed design tools that they are giving away as shareware. "Tools are now the show stoppers," says Nowick. "If you don't have tools you can't do things in portable ways, and you can't train people to become experts."

Beyond a new generation of design-and-testing equipment, successful development of clockless chips requires people who understand asynchronous design. Such talent is scarce, as asynchronous principles fly in the face of the way almost every university teaches its engineering students. Conventional chips can have values arrive at a register incorrectly and out of sequence; but in a clockless chip, the values that arrive in registers must be correct the first time. One way to achieve this goal is to pay close attention to such details as the lengths of the wires and the number of logic gates connected to a given register, thereby assuring that signals travel to the register in the proper logical sequence. But that means being far more meticulous about the physical design than synchronous designers have been trained to be.

An alternative, used by Theseus and others, is to open up a separate communication channel on the chip. Clocked chips represent ones and zeroes using low and high voltages on a single wire; "dual-rail" circuits, on the other hand, use two wires, giving the chip communications pathways, not only to send bits, but also to send "handshake" signals to indicate when work has been completed. Fant additionally proposes replacing the conventional system of digital logic with what he calls "null convention logic," a scheme that identifies not only "yes" and "no," but also "no answer yet" — a convenient way for clockless chips to recognize when an operation has not yet been completed. All of these ideas and approaches are different enough that executing them could confound the mind of an engineer trained to design to the beat of a clock. It's no surprise that the two newest asynchronous startups, Fulcrum Microsystems and Self-Timed Solutions, are populated by students



coming out of Caltech and the University of Manchester, where clockless-chip research has been going on the longest.

For a chip to be successful, all three elements — design tools, manufacturing efficiency and experienced designers — need to come together. The asynchronous cadre has "very promising ideas," says Max Baron, microprocessor analyst and editor of the industry newsletter *Microprocessor Report*. "But they don't have the actual machine, and they haven't proven they know how to build it."

Though it will take far longer for clockless chips to go mainstream, we're already seeing the beginnings of that transition as well. Intel, which shelved its asynchronous-chips project in 1997, incorporated elements of its clockless technology into the Pentium 4 chip that it released this year. "We're introducing asynchronous design from the bottom up, designing in some pieces of unlocked logic in a chip that is still of conventional design," says Stevens. "At this point, if we can do something asynchronously, and it's better in terms of power consumption, then we will do it."

So what of Karl Fant's flamboyantly predicted revolution? In an industry as mature as chip making, there's no replacing dictatorship with anarchy overnight. But over time, the balance will probably shift toward clockless design; enough articles will be written, enough tools built, enough engineers educated that it will no longer be unrealistic to imagine marketing such a chip even outside of specialized niches. "Once people understand how to do this easily, it will become more natural to think about asynchronous," says Sun engineer Normoyle. "People won't do it because it's interesting. We'll do it because it's easier than something else. Our only goal is to be better than the other guys. The switch will come when synchronous is no longer good enough."

The winners in this next wave of innovation will be the companies that choose the right time to jump off the curve. Clockless chips have the promise of revolutionizing the industry, of rapidly accelerating the relentless drive toward faster and cheaper chips that we've come to expect from Moore's Law. Who is to say what might be possible? Why not an all-asynchronous chip compatible with Intel products?

"If someone does that, they will have a serious competitive advantage for a number of years," says Intel's Stevens. Translation? "So yeah, we're worried."

Let the anarchy begin.