

Is It Time for Clockless Chips?

David Geer

Vendors are revisiting an old concept—the clockless chip—as they look for new processor approaches to work with the growing number of cellular phones, PDAs, and other high-performance, battery-powered devices.

Clockless processors, also called asynchronous or self-timed, don't use the oscillating crystal that serves as the regularly "ticking" clock that paces the work done by traditional synchronous processors. Rather than waiting for a clock tick, clockless-chip elements hand off the results of their work as soon as they are finished.

Recent breakthroughs have boosted clockless chips' performance, removing an important obstacle to their wider use.

In addition to their efficient power use, a major advantage of clockless chips is the low electromagnetic interference (EMI) they generate. Both of these factors have increased the chips' reliability and robustness and have made them popular research subjects for applications such as pagers, smart cards, mobile devices, and cell phones.

Clockless chips have long been a subject of research at facilities such as the California Institute of Technology's Asynchronous VLSI Group (www.async.caltech.edu/) and the University of Manchester's Amulet project (www.cs.man.ac.uk/apt/projects/processors/amulet/).

Now, after a few small efforts and false starts in the 1990s, companies such as Fulcrum Microsystems, Handshake Solutions, Sun Microsystems, and Theseus Logic are again looking



to release commercial asynchronous chips, as the "A Wave of Clockless Chips" sidebar describes.

However, clockless chips still generate concerns—such as a lack of development tools and expertise as well as difficulties interfacing with synchronous chip technology—that proponents must address before their commercial use can be widespread.

PROBLEMS WITH CLOCKS

Clocked processors have dominated the computer industry since the 1960s because chip developers saw them as more reliable, capable of higher performance, and easier to design, test, and run than their clockless counterparts. The clock establishes a timing constraint within which all chip elements must work, and constraints can make design easier by reducing the number of potential decisions.

Clocked chips

The chip's clock is an oscillating crystal that vibrates at a regular frequency, depending on the voltage applied. This frequency is measured in gigahertz or megahertz. All the chip's work is synchronized via the clock, which sends its signals out along all circuits and controls the registers, the data

flow, and the order in which the processor performs the necessary tasks.

An advantage of synchronous chips is that the order in which signals arrive doesn't matter. Signals can arrive at different times, but the register waits until the next clock tick before capturing them. As long as they all arrive before the next tick, the system can process them in the proper order. Designers thus don't have to worry about related issues, such as wire lengths, when working on chips.

And it is easier to determine the maximum performance of a clocked system. With these systems, calculating performance simply involves counting the number of clock cycles needed to complete an operation. Calculating performance is less defined with asynchronous designs. This is an important marketing consideration.

The downside

Clocks lead to several types of inefficiencies, including those shown in Figure 1, particularly as chips get larger and faster.

Each tick must be long enough for signals to traverse even a chip's longest wires in one cycle. However, the tasks performed on parts of a chip that are close together finish well before a cycle but can't move on until the next tick.

As chips get bigger and more complex, it becomes more difficult for ticks to reach all elements, particularly as clocks get faster.

To cope, designers are using increasingly complicated and expensive approaches, such as hierarchies of buses and circuits that adjust clock readings at various components. This approach could, for example, delay the start of a clock tick so that it occurs when circuits are ready to pass and receive data.

Also, individual chip components can have their own clocks and communicate via buses, according to Ryan Jorgenson, Theseus's vice president of engineering. Clock ticks thus only have to cross individual components.

The clocks themselves consume

power and produce heat. In addition, in synchronous designs, registers use energy to switch so that they are ready to receive new data whenever the clock ticks, whether they have inputs to process or not. In asynchronous designs, gates switch only when they have inputs.

HOW CLOCKLESS CHIPS WORK

There are no purely asynchronous chips yet. Instead, today's clockless processors are actually clocked processors with asynchronous elements.

Clockless elements use *perfect clock gating*, in which circuits operate only when they have work to do, not whenever a clock ticks.

Instead of clock-based synchronization, local handshaking controls the passing of data between logic modules. The asynchronous processor places the location of the stored data it wants to read onto the address bus and issues a request for the information. The memory reads the address off the bus, finds the information, and places it on the data bus. The memory then acknowledges that it has read the data. Finally, the processor grabs the information from the data bus. Pipeline controls and FIFO sequencers move data and instructions around and keep them in the right order.

According to Jorgenson, "Data arrives at any rate and leaves at any rate. When the arrival rate exceeds the departure rate, the circuit stalls the input until the output catches up."

The many handshakes themselves require more power than a clock's operations. However, clockless systems more than offset this because, unlike synchronous chips, each circuit uses power only when it performs work.

CLOCKLESS ADVANTAGES

In synchronous designs, the data moves on every clock edge, causing voltage spikes. In clockless chips, data doesn't all move at the same time, which spreads out current flow, thereby minimizing the strength and frequency of spikes and emitting less

A Wave of Clockless Chips

In the near future, Handshake Solutions and ARM, a chip-design firm, plan to release a commercial asynchronous ARM core for use in devices such as smart cards, consumer electronics, and automotive applications, according to Handshake chief technical officer Ad Peeters.

Sun Microsystems is building a supercomputer with at least 100,000 processors, some using asynchronous circuits, noted Sun Fellow Jim Mitchell.

Sun's UltraSPARC IIIi processor for servers and workstations also features asynchronous circuits, said Sun Distinguished Engineer Jo Ebergen.

Fulcrum Microsystems offers an asynchronous PivotPoint high-performance switch chip for multigigabit networking and storage devices, according to Mike Zeile, the company's vice president of marketing. The company has also developed clockless cores for use with embedded systems, he noted.

"Theseus Logic developed a clockless version of Motorola's 8-bit microcontroller with lower power consumption and reduced noise," said vice president of engineering Ryan Jorgenson. Theseus designed the device for use in battery-powered or signal-processing applications.

"Also, Theseus and [medical-equipment provider] Medtronic have worked on a [clockless] chip for defibrillators and pacemakers," Jorgenson said.

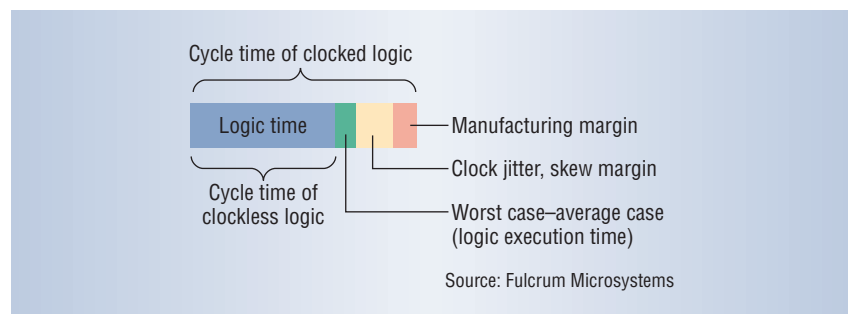


Figure 1. Clockless chips offer an advantage over their synchronous counterparts because they efficiently use cycle times. Synchronous processors must make sure they can complete each part of a computation in one clock tick. Thus, in addition to running their logic, the chips must add cycle time to compensate for how much longer it takes to run some operations than to run average operations (worst case – average case), variations in clock operations (jitter and skew), and manufacturing and environmental irregularities.

EMI. Less EMI reduces both noise-related errors within circuits and interference with nearby devices.

Power efficiency, responsiveness, robustness

Because asynchronous chips have no clock and each circuit powers up only when used, asynchronous processors use less energy than synchronous chips by providing only the voltage necessary for a particular operation.

According to Jorgenson, clockless chips are particularly energy-efficient for running video, audio, and other streaming applications—data-intensive programs that frequently cause synchronous processors to use considerable power. Streaming data applications have frequent periods of dead time—such as when there is no sound or when video frames change very little from their immediate predecessors—and little need for running

error-correction logic. During this inactive time, asynchronous processors don't use much power.

Clockless processors activate only the circuits needed to handle data, thus they leave unused circuits ready to respond quickly to other demands. Asynchronous chips run cooler and have fewer and lower voltage spikes. Therefore, they are less likely to experience temperature-related problems and are more robust.

Because they use handshaking, clockless chips give data time to arrive and stabilize before circuits pass it on. This contributes to reliability because it avoids the rushed data handling that central clocks sometimes necessitate, according to University of Manchester Professor Steve Furber, who runs the Amulet project.

Simple, efficient design

Companies can develop logic modules without regard to compatibility with a central clock frequency, which makes the design process easier, according to Furber.

Also, because asynchronous processors don't need specially designed modules that all work at the same clock frequency, they can use standard components. This enables simpler, faster design and assembly.

RECENT ADVANCES BOOST PERFORMANCE

Traditionally, asynchronous designs have had lackluster performance, even though their circuitry can handle data without waiting for clock ticks.

According to Fulcrum cofounder Andrew Lines, most clockless chips have used combinational logic, an early, uncomplicated form of logic based on simple state recognition. However, combinational logic uses the larger and slower p-type transistors. This has typically led to large feature sizes and slow performance, particularly for complex clockless chips.

However, the recent use of both *domino logic* and the *delay-insensitive mode* in asynchronous processors has

created a fast approach known as *integrated pipelines mode*.

Domino logic improves performance because a system can evaluate several lines of data at a time in one cycle, as opposed to the typical approach of handing one line in each cycle. Domino logic is also efficient because it acts only on data that has changed during processing, rather than acting on all data throughout the process.

The delay-insensitive mode allows

Clockless chips offer power efficiency, robustness, and reliability.

an arbitrary time delay for logic blocks. "Registers communicate at their fastest common speed. If one block is slow, the blocks that it communicates with slow down," said Jorgenson. This gives a system time to handle and validate data before passing it along, thereby reducing errors.

CLOCKLESS CHALLENGES

Asynchronous chips face a couple of important challenges.

Integrating clockless and clocked solutions

In today's clockless chips, asynchronous and synchronous circuitry must interface.

Unlike synchronous processors, asynchronous chips don't complete instructions at times set by a clock. This variability can cause problems interfacing with synchronous systems, particularly with their memory and bus systems.

Clocked components require that data bits be valid and arrive by each clock tick, whereas asynchronous components allow validation and arrival to occur at their own pace. This requires special circuits to align the asynchronous information with the synchronous system's clock, explained Mike Zeile, Fulcrum's vice president of marketing.

In some cases, asynchronous systems

can try to mesh with synchronous systems by working with a clock. However, because the two systems are so different, this approach can fail.

Lack of tools and expertise

Because most chips use synchronous technology, there is a shortage of expertise, as well as coding and design tools, for clockless processors.

According to Jorgenson, this forces clockless designers to either invent their own tools or adapt existing clocked tools, a potentially expensive and time-consuming process.

Although manufacturers can use typical silicon-based fabrication to build asynchronous chips, the lack of design tools makes producing clockless processors more expensive, explained Intel Fellow Shekhar Borkar.

However, companies involved in asynchronous-processor design are beginning to release more tools. For example, to build clockless chips, Handshake uses its proprietary Haste programming language, as well as the Tangram compiler developed at Philips Research Laboratories.

The University of Manchester has produced the Balsa Asynchronous Synthesis System, and Silistix Ltd. is commercializing clockless-design tools.

"We have developed a complete suite of tools," said Professor Alain Martin, who heads Caltech's Asynchronous VLSI Group. "We are considering commercializing the tools through a startup (Situs Logic)."

There is also a shortage of asynchronous design expertise. Not only is there little opportunity for developers to gain experience with clockless chips, but colleges have fewer asynchronous design courses.

A HYBRID FUTURE

No company is likely to release a completely asynchronous chip in the near future. Thus, chip systems could feature clockless islands tied together by a main clock design that ticks only for data that passes between the sections. This adds the benefits of asyn-

chronous design to synchronous chips.

On the other hand, University of Utah Professor Chris Myers contended, the industry will move gradually toward chip designs that are “globally asynchronous, locally synchronous.” Synchronous islands would operate at different clock speeds using handshaking to communicate through an asynchronous buffer or fabric.

According to Myers, distributing a clock signal across an entire processor is becoming difficult, so clocking would be used only to distribute the signal across smaller chip sections that communicate asynchronously.

Experts say synchronous chips’ performance will continue to improve. Therefore, said Fulcrum’s Lines, there may not be much demand for asynchronous chips to enhance performance. Furber, on the other hand, contended there will be demand for clockless chips because of their many advantages.

“Most of the research problems are resolved,” Myers said. “We’re left with development work. [We require] more design examples that prove the need for asynchronous design.”

Said Intel’s Borkar, “I’m not shy about using asynchronous chips. I’m here to serve the engineering community. But someone please prove their benefit to me.”

Added Will Strauss, principal analyst at Forward Concepts, a market research firm, “I’ve yet to see a commercially successful clockless logic chip shipping in volume. It requires thinking outside the box to find volume applications that benefit from the clockless approach at a reasonable cost.” ■

David Geer is a freelance technology writer based in Ashtabula, Ohio. Contact him at david@geercom.com.

Editor: Lee Garber, *Computer*,
l.garber@computer.org

How to Reach *Computer*

Writers

We welcome submissions. For detailed information, visit www.computer.org/computer/author.htm.

News Ideas

Contact Lee Garber at lgarber@computer.org with ideas for news features or news briefs.

Products and Books

Send product announcements to products@computer.org. Contact computer-ma@computer.org with book announcements.

Letters to the Editor

Please provide an e-mail address with your letter. Send letters to computer@computer.org.

On the Web

Explore www.computer.org/computer/ for free articles and general information about *Computer* magazine.

Magazine Change of Address

Send change-of-address requests for magazine subscriptions to address.change@ieee.org. Make sure to specify *Computer*.

Missing or Damaged Copies

If you are missing an issue or received a damaged copy, contact membership@computer.org.

Reprint Permission

To obtain permission to reprint an article, contact William Hagen, IEEE Copyrights and Trademarks Manager, at whagen@ieee.org. To buy a reprint, send a query to computer@computer.org.

Innovative Technology for Computer Professionals
Computer