

Asynchronous Design Shows Promise

But Lack of Tools and Skilled Designers Keeps Circuits Small and Simple

by Peter Song

Tick, tick, tick. The continuous beat of the ubiquitous clock signal is a comforting sound to digital designers, who know their signals have to be valid when the clock ticks but not between ticks. A growing number of designers, however, are making the scary leap to designs that do not use a clock. These asynchronous designs consume little power because transistors switch only when needed, not every cycle, and can provide performance benefits as well.

Asynchronous circuits are energy efficient—especially in CMOS technology, which consumes power only when switching—since only the transistors involved in the current computation are switched. They are well suited for algorithms that exhibit large differences between worst-case and average-case computing times. Without a clock to periodically generate huge voltage swings, asynchronous chips produce lower levels of electromagnetic interference. They are naturally immune from clock skew, a growing problem in clocked circuits as clock periods shrink and die sizes increase.

Asynchronous circuits are, unfortunately, much more difficult to design than clocked, or synchronous, circuits. Designing synchronous circuits is made easier by assuming that all state-holding elements (latches and registers) change together and in discontinuous steps. Assuming otherwise—that state-holding elements change independently or continuously—increases the need for highly accurate timing models, greatly complicates the design process, and requires highly inefficient test methods. The few designers who have tried it have found that the lack of simple and well under-

stood design methods and tools gets in the way of reaping the full benefits of being clock-free.

Clock-free chips are unlikely to displace clocked chips to any great extent in the foreseeable future, even in applications that demand the highest MIPS/W rating. Nevertheless, the pursuit of reliable and easy-to-use asynchronous design methods and tools continues, mostly in research environments. A few designs are starting to make the transition from research to products. Using an asynchronous ARM core, Amulet2e delivers 40 MIPS at 150mW and targets portable embedded applications. Sharp's DDMP (data-driven multimedia processor) uses eight CPUs—each clock-free CPU delivers 600 MOPS using only 40K gates and less than 60 mW—to compete against Chromatic's Mpact2, Philips's Trimedia, and TI's TMS320C6X. And Cogency (see MPR 10/6/97, p. 10) is now shipping a fully asynchronous DSP.

Eliminating Clock Has Benefits

Unlike many synchronous circuits that consume power every cycle everywhere, asynchronous circuits consume power only when and where it is needed. Although asynchronous circuits tend to use more transistors than synchronous circuits, the transistors not involved in computation do not switch. Philips designed a DCC (digital compact cassette) error-corrector chip that consumes 80% less power than a clocked design, despite using 70–100% more area for replacing a clock and registers with handshake components. Although many synchronous designs use power-saving techniques, such as not clocking idle functional units, these techniques are inherently part of asynchronous circuits.

Asynchronous circuits are well suited for implementing algorithms such as the ripple-carry addition, in which the worst-case computing delay—the carry rippling from the least significant bit of the sum to the most significant bit—rarely occurs. Although ripple-carry adders are smaller than other types of adders and produce fast results in most cases, they are not used in clocked designs in which results must be produced in the same number of cycles in all cases. Asynchronous designs can be optimized for average-case performance rather than for worst-case performance, resulting in simpler and smaller circuits without unduly compromising overall performance.

Clocked circuits must handle clock skew (the difference in arrival times of clock edges at different parts of the circuit) that bounds the minimum and maximum delay through combinational paths. Since wire delays do not scale linearly with decreasing clock periods, clock skew takes an increasingly larger portion of useful clock cycles as process geometries shrink. Furthermore, as more transistors are packed

Class	State change		Delay	
	Individual	System	Gate	Wire
Synchronous	Discrete	Synch	Bounded	Bounded
Micropipelines	Discrete	Async	Bounded datapath, unbounded control	Bounded datapath, unbounded control
Burst AFSM*	Discrete	Async	Bounded	Bounded
SIC† AFSM	Continuous	Async	Bounded	Bounded
Speed insensitive	Continuous	Async	Unbounded	Bounded (negligible)
Quasi-delay insensitive	Continuous	Async	Unbounded	Unbounded isochronic‡
Delay insensitive	Continuous	Async	Unbounded	Unbounded

Table 1. Sequential circuits can be classified by their assumptions about whether state changes are continual or discrete and synchronized or not and whether gate and wire delays are bounded (known) or unbounded (unknown). *AFSM = asynchronous finite-state machine. †isochronic = all wires originating from the same output have the same amount of delay. ‡SIC (single input change) = only one input is to change at a time. (Source: [1]).

into larger die, requiring larger clock-distribution trees or grids and stronger clock drivers, managing clock skew and power consumed by the clock becomes even more difficult. By replacing a global clock with handshake components that are local and operate independently, asynchronous circuits avoid these problems.

Assumptions on Time Classify Sequential Circuits

The fundamental distinction between synchronous and asynchronous circuits is whether states change together (synchronously) or independently. As Table 1 shows, synchronous circuits are the most restrictive class of sequential circuits, designed under the assumption that states change together and in discrete steps, or in clock cycles. While some asynchronous circuits are designed under the same assumption, others are designed under the assumption that states can change continuously, thereby requiring signals to be valid at all times. Different sets of assumptions yield varying tradeoffs in design difficulties, circuit sizes and complexities, and reliability.

Asynchronous finite-state machines (AFSMs) are structurally different from *finite-state machines* (FSMs) in that the clocked state-holding elements are now replaced with the delay elements, which are needed to avoid race conditions that can cause the machines to change multiple times on a single input change. The clocked elements in FSMs contain snapshots of the combinational logic, allowing the combinational logic to change freely between the clock edges. The delay elements in AFSMs, in contrast, continually represent the state of the combinational logic.

AFSMs suffer from two major problems that make them unsuitable for datapath or deeply pipelined designs. The first is that they require the inputs to operate in fundamental mode—an input change should not occur until the machine settles into a stable state—severely limiting the degree of concurrency generally needed in datapath designs. The other is that they suffer from the additive-skew problem. That is, changes to the first machine's inputs are restricted to meet the fundamental-mode assumption of not only the first machine but all other machines connected to it as well. The first machine's inputs must not change faster than the time needed for it and the other machines to become stable.

Since each machine has skew—the difference between the minimum and the maximum propagation delays through the combinational logic—then for two machines in series, the first machine's inputs must not change faster than the sum of the first machine's skew and the propagation delay through the second machine. Meeting the additive-skew requirement results in extremely poor throughput in designs with many pipeline stages.

Circuits, including AFSMs, that assume bounded delays must be designed to function without timing problems in widely varying physical conditions. Designing such circuits requires a correct-by-construction design method and accurate timing models in all phases of the design process; the simple but practical synchronous design method of encap-

Why Clocked Design Is Easier

Assuming that states change only on clock boundaries allows designers to solve separately the functional and timing problems of synchronous circuits. Since the Boolean logic is precise and timeless but propagation delays are imprecise and variable, being able to meet the two requirements separately simplifies both the functional and the timing design processes and makes the required tools more efficient.

Designers can largely ignore the imprecise nature of propagation delays during the functional design process. Propagation delays are inherently imprecise and variable because they are affected by many physical and environmental factors, such as process variations, packaging, temperature, frequency, and voltage. They are difficult to model accurately without detailed physical data. The most accurate physical data is belatedly available from layout when the design is nearly completed.

Functional simulation is simpler when propagation delays are not modeled. To select the right set of features from many design options, designers generally simulate cycle-accurate or near-cycle-accurate behaviors of the features. Such simulators are much faster than event-driven simulators that model delays between events. Even when using event-driven simulators, most designers do not model the individual gate and wire delays but only the signal delays needed to chronologically order the events, resulting in faster simulation.

Timing analysis is also more efficient when logic functions are not modeled. Static timing-analysis tools calculate the longest delay from each input to each output of a module without evaluating the logic between the input and the output. Since only propagation delays and not logic functions are used, the modules can be arbitrarily large and can be built hierarchically using smaller modules.

Static timing analysis is so efficient that it can easily identify critical paths of multimillion-transistor circuits within hours. This is several orders of magnitude faster than dynamic timing-analysis tools, such as Spice, that use more detailed timing models. Although static timing tools are less accurate (most claim to be within 5–10% of the accuracy of Spice), their fast turnaround is indispensable for performing many iterations of logic changes and synthesis to meet timing goals of high-speed designs.

The penalty for this efficient design approach is that state changes can occur no faster than the longest combinational delay in the clocked circuit. If the clock is periodic—and it generally is, except during testing and debugging—its period must be longer than the longest combinational delay when computing the worst case. The penalty is an acceptable compromise for being able to design arbitrarily complex circuits.

References and Resources

- [1] Hauck, S. "Asynchronous Design Methodologies: An Overview." In *Proceedings of the IEEE, Vol. 83, No. 1*, pp. 69–93, January 1995.
- [2] Sutherland, I. "Micropipelines." In *Communications of the ACM, Vol. 32, No. 6*, pp. 720–738, June 1989.
- [*] Davis, A., S. Nowick, eds. *Asynchronous Digital Circuit Design*, Springer-Verlag, 1994.
- [*] Hulgaard, H., et al. "Testing Asynchronous Circuits: A Survey." University of Washington, Dept. of CSE International Report, TR # 94-03-06, 1991.
- [*] www.cs.man.ac.uk/amulet/async
- [*] www.cs.utah.edu/projects/acs

run out). Although some precision timing-analysis tools, such as Spice or EPIC's TimeMill, do treat functional and timing attributes as one indivisible entity, they are utterly inefficient as functional simulators that must deliver millions of simulation cycles for overnight regression.

The inability to suspend and advance state machines at will in asynchronous chips makes debugging and testing them difficult. Whereas the entire state of synchronous chips can be frozen, advanced, and even modified via scan operations, the state of asynchronous chips cannot. Without clocked registers, single-step operations are not possible: with asynchronous chips, a change in input pins propagates until the entire chip reaches a stable state. Asynchronous chips either run at full speed or remain halted; it is not possible to stop them until they halt naturally.

Existing manufacturing-test equipment is inadequate for testing asynchronous circuits. For circuits that assume bounded gate and wire delays, the widely used stuck-at fault model is inadequate and therefore requires path-delay tests for detecting delay faults. Path-delay tests involve the application of test vectors and the timely sampling of responses—each sample is taken after a time that allows a path to produce a correct output if it is defect-free but an incorrect output if it contains a delay fault. These tests are significantly more expensive to perform than stuck-at fault tests. In addition, adding extra logic to eliminate hazards during design directly conflicts with the need to eliminate redundant logic for testing, as it makes detecting faults difficult.

Circuits that use the continuous-time model are also more susceptible to noise than those that use the discrete-time model. Whereas state changes are designed to take place after noise has time to subside in the discrete-time model, circuits must be stable at all times in the continuous-time model. Noise and timing hazards can be dealt with by slowing down the clock in synchronous chips, but they generally cause malfunctions that cannot be fixed in asynchronous chips. Asynchronous circuits, however, generate lower levels of noise and EMI than clocked circuits.

Clock-Free Chips Won't Displace Clocked Chips

Clock-free circuits can have low power consumption and low levels of electromagnetic interference. They can be optimized for average-case, instead of the infrequent worst-case, computing times to simplify designs or yield smaller sizes without compromising overall performance. Due to difficulties in designing and testing asynchronous circuits, however, many of these potential benefits have yet to be turned into competitive advantages over synchronous circuits.

Asynchronous circuits cover a wide range of circuits and design methods. Although delay-insensitive circuits are less than practical, they introduce the ideas of using transition-based signaling for control and using either dual-rail encoding or the bundled-data scheme for passing data. Closest to synchronous circuits, and therefore not taking full advantage of being clock-free, micropipelines offer the most practical way to design asynchronous circuits. Micropipelines avoid additive skew, as well as hazards, by isolating adjacent stages with pipeline registers. Although not enough is known about applying micropipelines to more general circuit structures, large-scale design efforts, such as those undertaken by Cogency and the Amulet group, are providing more insight.

All the asynchronous design methods, however, are likely to be more complicated than the more familiar synchronous design methods. Although most designers learn about asynchronous design techniques in their introduction to logic design, there are no products from which to gain practical experience and further develop skill. This situation is unlikely to change in the foreseeable future unless a non-trivial asynchronous design can clearly demonstrate technical superiority over synchronous designs.

There is a circle of dependence that hinders wider acceptance of asynchronous circuits and more aggressive development of their design methods and tools. Given the enormous mismatch between the states of synchronous and asynchronous designs, asynchronous designs are unlikely to deliver higher performance than synchronous designs, even at comparable power-consumption levels. Without demonstrably clear market advantages, few may commit to asynchronous designs when easier solutions can be found. Many portable applications that are highly sensitive to EMI may find clock-free designs attractive, however.

Asynchronous circuits are more likely to be used within synchronous chips than by themselves. By themselves, they are less likely to have the levels of complexity needed to make them complete solutions in most applications. Due to accelerating advances in semiconductor processes, the minimum level of integration in a single chip is likely to grow faster than the complexity of the asynchronous chips that can be designed. By implementing clock-free only the modules that are well suited for it—datapaths that tend to have simple and repetitive logic, modules that have few interfaces and sporadic inputs, and algorithms that exhibit large computing delays between the worst case and average case—the difficulties of designing clock-free chips can be largely mitigated. 