

Guest Editors' Introduction: Clockless VLSI Systems

Soha Hassoun
Tufts University

Yong-Bin Kim and Fabrizio Lombardi
Northeastern University

■ **IT IS WITH GREAT PLEASURE** that we introduce the special issue on clockless VLSI design to the readership of *IEEE Design & Test*. This special issue consists of six articles selected to cover a wide spectrum of techniques and applications encountered in the design and manufacture of today's clockless VLSI systems. Authored by outstanding researchers, these articles cover experimental and speculative topics. As with all special issues, the topics that we cover here are just the highlights of the large volume of literature currently provided by the technical community.

Difficulties associated with traditional synchronous design have prompted many researchers to consider new alternatives. Three factors drive current research and development efforts in the design and test of clockless digital systems: low power, performance, and design for reuse.

Low power

The mainstream design style in use for today's processors is synchronous, that is, a clock regulates the internal timing. We measure how fast a computer can execute instructions by the number of clock cycles per second. Unfortunately, the clock consumes more power than most other components in the chip. The most disturbing aspect of this characteristic is that the clock only serves as a timer for computational tasks. It does not perform operations on data, it simply orchestrates the many computational parts of a digital system (whether chip, board, or computer).

New problems are also evident for power consumption at-large; that is, throughout the system in general. As the number of transistors on a chip increases, so does the power used by the clock. Therefore, in complicated chips, power consumption becomes an even more crucial topic. For mobile and portable electronics, chips must conserve power even more efficiently to maximize battery life. Low-power design is important

not only for portable systems but also for systems in which cooling and power consumption requirements are critical. Such systems include state-of-the-art microprocessors, which today have power dissipation in excess of 100 W. This emphasis in designing for low power and high performance makes designers interested in clockless design techniques, which offer advantages in these areas.

Performance

Today's computers operate in a synchronous manner. They maintain a main clock that controls the timing of the chips. As chip density and clock frequency increase, it is becoming increasingly difficult to develop schemes for distributing a global clock in new designs, such as systems on a chip. In addition, such clocks must have reduced skew to account for variations in manufacturing process, operating temperature, supply voltage, and clock-loading capacitance. A synchronous chip can only work as fast as its slowest component. Therefore, synchronous circuits exhibit worst-case behavior, that is, if one part of the chip is rather slow, it forces all other parts of the same chip to remain idle for some time to accommodate the slower operation. This results in a waste of computing that is obviously detrimental to the performance of the whole chip. Clockless systems exhibit average- or best-case behavior, because they permit faster operations to proceed without waiting for slower ones. The study of techniques and design methodologies for high performance and low power through alternative and innovative architectural considerations is imperative for designing next-generation VLSI chips.

Design for reuse

Current market trends require the use/reuse of cores in many complex designs. Design for reuse is a common integration practice for designing SoC implementation at relatively low production costs. Currently,

modularization of VLSI chips involves a traditional synchronous design methodology. But this approach is rather inflexible because it does not adapt well to new system contexts, such as novel process technologies and different operating clock frequencies. A clockless system, on the other hand, adapts more easily to technology advances because of its delay-insensitive property. It is also a very good candidate for SoC due to its modularity, ease of reusability, and robustness. Unfortunately, few designers are aware of the details necessary for clockless operation and design. High-performance VLSI systems for SoC can incorporate these design features into products while addressing design for reuse in IP block integration.

In light of today's innovations, clockless designs continue to gain importance and attention as difficulties in clock distribution, designing for low power, and noise management increase. Design-for-reuse issues in high-performance VLSI and the increasing delay sensitivity of conventional synchronous circuits have propelled clockless systems to the forefront of VLSI submicron technology. Clockless circuits hold many advantageous features over asynchronous circuits, such as at least partial delay-insensitive behavior, speed-independent design, and self-timed operation for pipelining. These circuits require innovative design techniques that depart from traditional practice.

The road to clockless

This special issue will provide a comprehensive venue for addressing some of the aforementioned topics. On the technology side, as transistor feature sizes move toward 0.1 micron and clock frequencies exceed the 1 GHz level, it is becoming increasingly difficult to develop schemes to distribute a global clock and analyze timing information. In today's VLSI systems, skew management of the clock network is extremely important. Skew is a function of load, network distribution, and device mismatch, as well as temperature and voltage gradients. Scaling down devices does not reduce clock skew, however. In fact, skew can increase significantly as a result of device and operational deviation stemming from the environment. As a result, skew is becoming one of the major obstacles in achieving high-frequency clock distribution in the design of sub-quarter micron CMOS ICs.

The traditional design approach used in industry has been tied to synchronous logic, in which timing throughout an entire chip must be carefully analyzed and controlled to ensure its correct operation. But according to *The International Technology Roadmap for*

Semiconductors, published by the Semiconductor Industry Association, the industry will soon be able to manufacture chips so complex that timing analysis and clock distribution will become completely intractable problems, thus requiring a shift to novel clockless architectures. Because of the rapid growth in portable equipment use (such as cellular phones and GPS units), and the continuing trend in processors with stringent requirements for power dissipation, energy efficiency has become an important attribute of VLSI designs. In a system with a global clock, latches and registers operate and consume dynamic energy during each clock pulse, despite the fact that many of these latches and registers are inactive (that is, they have no new data to store). Such overhead and performance issues associated with a clock-imposed control mechanism are rather burdensome as fabrication innovations allow for faster and larger designs.

Clockless implementations are good alternatives for meeting performance, power, and design-for-reuse criteria. At circuit level, clockless operation differs considerably from asynchronous behavior. Clockless circuits are at least quasi-delay-insensitive with speed-independent behavior readily applicable to different applications, such as self-timed pipeline designs. Even though the design of clockless chips is still in its infancy, new innovations in this arena have already started to appear in many commercial products. Companies have announced microprocessor prototypes and support circuitry that incorporate elements of clockless technology and are planning to gradually integrate an "island" of clockless logic into future-generation chips. Some companies are already marketing clockless chips that, for example, give pagers up to twice the battery life of their competitors' products. Furthermore, designers are already investigating the use of clockless chips for mobile devices and smart cards. These innovations are based on the so-called null convention logic (NCL) as a way of letting clockless chips know when an operation is complete.

This special issue

The convergence of all of these issues makes clockless VLSI design a relevant, yet extremely difficult, objective to attain in digital systems. This special issue provides the reader with a timely account of state-of-the-art research on this topic. We can broadly divide the articles in this special issue into three parts: The first set, consisting of two articles, provides an historical introduction to this topical area; the second set has two articles that deal with the NCL; and the third set contains two articles that deal with system-level topics.

In the first article, Alain J. Martin, Mika Nyström, and Catherine G. Wong present the evolutionary development of microprocessors with asynchronous features. These microprocessors exploit so-called quasi delay-insensitive principles to conservatively design robust circuits with a small dependency on timing. Encoding information about signal validity within the signal itself removes most of the timing assumptions, resulting in a novel operational mode. Caltech has used this design model in a family of microprocessors spanning the past 15 years, from the 1988 Caltech Asynchronous Microprocessor (CAM, a 16 bit RISC machine), to the architectural features and timing issues of the 1998 Caltech MiniMIPS (for high-performance computing), to its most recent prototype, the 2003 Litanium 8051 microcontroller. Although commercial CAD tools and design methodologies for synchronous industrial designs outpace those available for asynchronous designs, these three generations of Caltech asynchronous microprocessors have repeatedly proven the feasibility of asynchronous designs.

In the second article, Stephen Unger of Columbia University deals with designing (static) CMOS logic circuits for asynchronous operation to reduce transistor channels at no operational cost. This technological feature systematically increases speed (with possible testability improvements) while reducing both chip area and energy dissipation. Unger shows that, for dynamic operation, it is possible to eliminate a transistor altogether without any negative effects on the asynchronous circuit's behavior.

The following two articles discuss NCL and the related logic paradigm, providing the basis for timing analysis of delay-insensitive and self-timed operation. First, Satish K. Bandapati, Scott C. Smith, and Minsu Choi consider a variety of unsigned multipliers designed using NCL. These circuits have substantially different architectures, thus resulting in a large variance in circuit performance in terms of measures such as power, area, and speed. It is significant that meeting different measures requires different architectures: for example, a parallel pipelined dual-rail multiplier is the best choice for attaining high performance. On the other hand, using a parallel nonpipelined quad-rail multiplier best reduces power consumption.

Next, Theseus Logic's Steve Masteller and Lief Sorenson address an analyzer tool currently under development. It promises to yield an automated environment for clockless design. This tool incorporates static timing analysis, orphan checking, ATPG, and

possibly synthesis. Masteller and Sorenson outline the NCL Analyzer with respect to cycle decomposition, by which a gate-level netlist characterization provides an accurate identification of cycles. This results in an elegant technique that improves over previous works, namely, on specific features such as presence of reset capabilities or register and inverter cells.

The fifth and sixth articles present system-level topics that are emerging for clockless operation. In the fifth article, Juha Plosila, Tiberiu Seceleanu, and Pasi Liljeberg propose an asynchronous structure for implementation in SoC. Segmentation of buses is a widely employed method for permitting modular and dynamic operation in the communication infrastructure. This method arranges segments to permit concurrency to satisfy different operational constraints in the SoC. The authors show that this bus' asynchronous mode of operation also permits self-timing that supports improvement in either execution speed or power consumption, depending on the application. Remarkably, this technique preserves parallelization through an intersegment topological arrangement that efficiently organizes communication with high signaling speed (through a so-called central arbiter).

Finally, Woo Jin Kim and Yong-Bin Kim address the design of high-performance circuits without the insertion of storage elements. They propose synthesis and delay-balancing scripts as part of a tool to automate the design of these circuits. The authors describe techniques that account for process, voltage, and temperature variations to accurately control the clocking process. Although frequency adjustment can accurately control the clocking process in conventional pipelining, wave pipelining requires a more comprehensive technique that must account for additional parameters such as skew. The authors present different solutions to this problem, including a localized scheme that offers many advantages for CMOS integration.

WE SINCERELY HOPE that this special issue will be a reference publication for future research. These articles cover topics that are timely and important, and the authors have done an excellent job of presenting the material. We extend our sincere thanks to all the authors and reviewers. We also thank Rajesh Gupta, editor in chief of *IEEE Design & Test*, for allowing us to organize this special issue. Finally, a special thanks is due to the editorial staff for editing and assembling this issue. Please feel free to contact us if you have questions or comments. ■



Soha Hassoun is an assistant professor in the Department of Computer Science at Tufts University. Her research interests include CAD, VLSI design, and computer architecture.

Hassoun has a BS in electrical engineering from South Dakota State University, an MS in electrical engineering from the Massachusetts Institute of Technology, and a PhD in computer science and engineering from the University of Washington, Seattle. She is a member of the ACM, a senior member of the IEEE, and a Fellow of Tau Beta Pi.



Yong-Bin Kim is a faculty member and holder of the Zraket Endowed Professorship in the Department of Electrical and Computer Engineering at Northeastern University. His research

interests include testing and design of digital systems, high-speed digital/analog IC design, clocking schemes for high-performance VLSI systems (including on-chip clock skew analysis and clock distribution), high-speed IC signal integrity and physical CAD tool development, low-power and high-speed circuit design methodology and technology, deep-submicron device phenomena, high-speed system integration for signal processing and communication applications, innovative circuits and system applications, and merged DRAM logic technology. Kim has a BS in electronic engineering from Sogang University, Seoul, South Korea; an MS in electrical engineering from the New Jersey Institute of Technology; and a PhD in computer engineering from Colorado State University. He is a senior member of the IEEE.



Fabrizio Lombardi chairs the Department of Electrical and Computer Engineering at Northeastern University, Boston, where he also holds the International Test Conference

Endowed Professorship. His research interests include testing and design of digital systems, ATE systems, configurable and network computing, defect tolerance, and CAD for VLSI. Lombardi has a BSc in electronic engineering from the University of Essex, UK; and an MS in microwaves and modern optics and a Diploma in microwave engineering, both from University College London. He is the associate editor in chief of the *IEEE Transactions on Computers*, an associate editor for *IEEE Design & Test*, and the founding general chair of the IEEE Symposium on Network Computing and Applications.

■ Direct questions and comments about this special issue to Yong-Bin Kim, Department of Electrical and Computer Engineering, 327 Dana Research Building, Northeastern University, 360 Huntington Ave., Boston, MA 02115; ybk@ece.neu.edu.

For further information on this or any other computing topic, visit our Digital Library at <http://computer.org/publications/dlib>.



Take your e-mail address with you
Get a free e-mail alias
from the
Computer Society
and
DON'T GET CUT OFF
you@computer.org
Sign up today at
computer.org/WebAccounts/alias.htm