



CommsDesign

Greening processor design

Michel Laurence, Octasic Inc.

Mar 22, 2010 (5:51 PM)

URL: <http://www.commsdesign.com/showArticle.jhtml?articleID=224100003>

Energy conservation has become a key social objective in all aspects of our lives. Due to the pervasiveness of electronics and computing, the processor industry accounts for an ever increasing share of world power consumption; be it for processors used in personal computers, domestic and industrial appliances, in massive server farms, or even in the growing number of mobile smart-phone devices.

For the engineers and designers creating these electronics, delivering a high-performance design, while keeping energy conservation in mind, remains a delicate balance. Fundamental advances at the silicon design level, however, have led to significant improvements in power efficiency for many of these electronics. To that effect in the last few years much of the attention of the industry has been focused on multicore architectures. While this architecture shift has certainly helped address some power concerns, a lesser known processor design methodology holds the promises for even more significant power savings. This is the adoption of self-clocking processing cores.

The power/performance/cost trade-off

A decade ago, migrating to new and smaller geometry process nodes was an exciting proposition: faster clock speeds, twice as many chips per wafer and of course much lower power consumption. As transistor geometries have scaled towards the most fundamental atomic dimensions, it has become increasingly difficult to get improvements in the three product fundamentals of performance, cost and power. Instead of obtaining improvements in all three, design teams have had to optimize for two of the three fundamentals; either optimize for power and cost with tradeoffs in performance, or optimize for performance gains and tradeoff higher power dissipation with marginal cost improvements. This has had to be done against a backdrop of much higher development costs and of course, increased design complexity and risk. These trends have in turn made the return on investment evaluations for complex SoCs much harder.

Until recently, the major performance gains in microprocessor evolution were done by using smaller and faster transistors made available in each technology node to increase clock frequency. Architectural gains and increased data path widths further improved performance while reduction in core operating voltages improved power consumption. At the 90nm technology node, transistor leakage current became a major challenge. Smaller transistors with lower threshold voltages could reduce chip area and increase performance, but brought a high cost in increased leakage current. Designers now face a difficult choice between increasing clock frequency to improve performance and paying a large penalty in power consumption, or reducing power with little gain in the performance per gate of the design and using more gates (silicon) for performance gains.

Multi-core architectures

The key strategy adopted by the industry has been the use of many slower processors within a single device and hence the well-known shift to multicore architectures. Reducing the clock frequency makes each processor more power efficient, and therefore greener.

Multicore processors have become increasingly popular, due largely to their ability to reduce power consumption and offer increases in system performance. Clearly two processors can do more useful work than one and use less power than simply clocking a single processor at twice the clock frequency. This architecture change allows greater power efficiency but causes two major side effects: reduced silicon efficiency and higher system complexity.

Since the multicore architecture requires multiple copies of the same core, it directly trades area to increase system performance. Any reductions in die size due to smaller transistor geometries are lost on the need to implement multiple cores.

Challenges with traditional processor design

As geometries scale below 40nm, leakage current continues to be an issue. Another added challenge is managing on-chip process variations. Variation in two key semiconductor process parameters, transistor threshold voltage and transistor effective length can have a substantial impact on the performance and power consumption of a design. As these variations in the manufacturing process increase, designers of SoCs at these nodes have to be much more conservative and slow down the frequency of operation in order to guarantee that the design will not only meet its performance specifications, but will also be manufacturable.

Another issue that has been well publicized for semiconductor evolution is heat generation and dissipation. Power dissipation is proportional to capacitance, and voltage squared. Prior to 90nm, as technology shrank the switching voltage was also reduced yielding a reduction in power. As the operational voltage is now close to the threshold voltage of a single transistor, it is not possible to get the same reductions in power dissipation to offset the shrinking area of the chip. As semiconductor process technology puts more gates in a given area without this commensurate, decrease in power consumed by the gates the heat generated for a given area increases. With a limited power dissipation capability of the package and heat sink, migration to smaller geometries with increased clock frequencies rapidly exceeds the ability to remove the heat thus limiting the performance achievable by a design.

The traditional digital logic design methodology combined with the ever greater challenges of migrating to smaller process nodes has forced the industry to look for a new solution to allow the combined gains of performance, power and cost improvements.

One way to address these core issues is by using a self-timed asynchronous architecture.

Processor design history & evolution

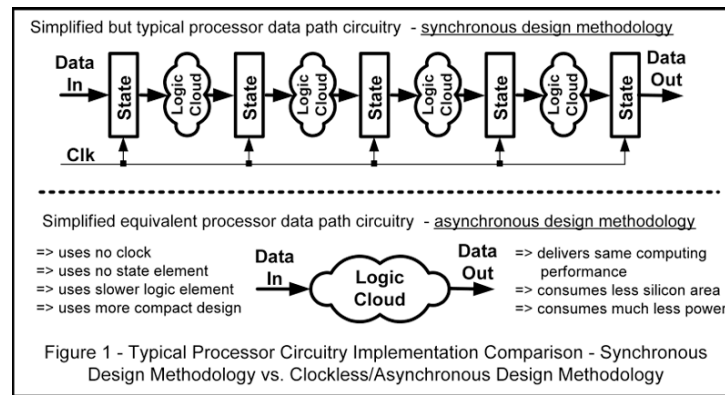
Employing an asynchronous design methodology over the more traditional synchronous methodology can be extremely advantageous from both a power consumption perspective and a silicon usage perspective. Until recently, however, asynchronous digital design has had an unfavorable reputation in the semiconductor industry, probably due in large part to traditional digital design education. Historically, most university electronics classes treated asynchronous design as an oddity that was more difficult to grasp and prone with the dangers of poor reliability due to race conditions and instability. By contrast, synchronous digital design classes dominated the curriculum and were taught as the safe and preferred design approach. Combined, these two factors make it easy to understand why asynchronous design has had limited success today.

A few exceptions to this rule exist. One such exception began back in the early 1980s with a group led by Professor Steve Furber at Manchester University in England, who developed a low power asynchronous ARM processor. Another example is with southern California-based Fulcrum Microsystems, which has been successfully selling for quite a while Ethernet interconnect chips based on its internally developed asynchronous design methodology. More recently, Achronix Semiconductor in San Jose, California has released a family of very high-performance FPGAs based on its asynchronous technology. These asynchronous products are not widely known or understood by the industry, which has formed similar opinions as to those of the students receiving an education in digital design, which is that asynchronous designs are difficult, unreliable and have poor performance. However, a fresh look at the issues of asynchronous design has proven those doubts to be unfounded given the right architecture and rigorous design methodology.

Asynchronous design advantages

Asynchronous circuits can deliver a variety of advantages. At one extreme, the technology can be used to deliver the highest speed circuits. At the other end of the spectrum, the technology can be used to deliver the lowest possible power consumption devices. Asynchronous, or clockless circuits can also be readily used to develop high-performance, low-power, and silicon- efficient processors.

To understand the advantages of asynchronous circuits, it's important to first compare the main design techniques against universally used synchronous approaches. **Figure 1** illustrates graphically at a high level the typical combined effects of moving from a synchronous to an asynchronous implementation methodology in a typical highly-pipelined processor design.



[Click on image to enlarge.](#)

Elimination of clock trees - Synchronous high-speed processors require large clock trees to keep sequential blocks synchronized. These clock trees require high-power buffers to drive complex high-capacitance networks that cover most of the chip. Clocks change state twice per cycle, consuming power on both positive and negative edges. These clock trees do not perform any information processing, thus provide no useful computing work, yet they consume a significant portion of the total power. Eliminating the clock trees alone can reduce power consumption by as much as 40 percent in a high-performance processor.

Elimination of pipeline state elements - Modern synchronous high-performance processors rely heavily on pipeline design techniques. Those pipelines require the use of a very large number of inter-stage flip-flops and state elements to support at high clock frequency operation. However, these inter-stage flip-flops and state elements also don't contribute to the actual data processing and computing tasks performed by the processor. In an asynchronous design these storage elements are discarded, saving the silicon space they occupy and the large amount of power they consume.

Elimination of lost margin timing - These inter-stage flip-flops require set-up and hold times which force a significant portion of the time between clock edges to be unusable for computation in high-frequency synchronous designs. Moreover given that each sub-micron technology shrink tends to increase path timing uncertainty, this further shortens the active period that can be used to achieve useful work between clock edges. This also means that in a synchronous design, the inter-stage circuit logic needs to be designed to operate increasingly faster than a single clock period to perform the same work. This requires the increased use of larger, higher power buffers in the datapath. In an asynchronous processor design, the logic does not have to deal with such small time steps. They can be built using slower, smaller and lower power circuits, while still delivering the same level of overall performance. Because the gates can be slower, it allows more use of low-leakage high-threshold voltage (HVT) gates, which drastically reduces leakage and further reduces power consumption and die area.

Reducing wire interconnect length - The silicon area savings discussed above translate into even more power savings, because wires connecting two elements get shorter as the circuits between these elements shrink. Shorter wires have less capacitance, thus switching them requires less power by using smaller buffers.

By applying asynchronous methodologies to processor implementation, one can eliminate clock trees and sequential elements, which do not contribute to processing and computing tasks. They also reduce timing constraints on the circuits that do useful processing work, reducing the area and power consumption. This means lower power can be achieved using less silicon to do an equivalent function, thereby substantially reducing die area and real product costs. In many applications, the power savings can be by far the most significant, given the challenges facing a designer today.

A bright future for evolved processor designs

In summary, the elimination of all clock circuits results in large power savings; while the elimination of all inter-stage storage and state elements saves both silicon area and power. Additionally, by using slower, smaller and less power-hungry logic elements in the remaining circuitry, designers can obtain the same performance, while saving on the silicon and power, which are even further reduced by contracting the total size of the circuit.

Another important implication of using asynchronous cores is power dissipation because it is the limiting factor for the total number of cores that can be placed in a single package. A device based on an asynchronous core can pack over five times as much processing in a single device greatly reducing total system cost and bringing new possibilities to SoC applications.

In an era where power, performance and cost are the subjects of such prevalent considerations, an asynchronous design methodology has too much to offer the processor industry to be ignored. Reducing power can be a singular benefit to those manufacturers that not only want to reduce cost and heighten performance, but also lessen their impact on global energy consumption. The semiconductor vendors who will master the technology and apply it to their processor and product designs will deliver an unparalleled level of power efficiency and cost performance that should leave their less opportunistic contemporaries' offerings years behind.

About the author:

Michel Laurence is co-founder and executive chairman of Octasic Inc.

In 1998, Michel Laurence co-founded Octasic. He was CEO and Chairman of the Board of Directors from 2001 to 2009, overseeing R&D and the company's strategic direction. As of May 2009, Laurence is Executive Chairman. Prior to Octasic, Laurence co-founded InnoMediaLogic (IML) in 1996, and within a span of four years as Chairman and President, he developed the company into a thriving multi-million dollar VoP solutions vendor. NMS Communications purchased IML in 2000 and Laurence was appointed VP and General Manager of NMS Communications' Network Access Business Unit.



[Copyright © 2003 CMP Media, LLC](#) | [Privacy Statement](#)