FULCRUM microsystems

Embedded.com

Asynchronous logic moves toward mainstream acceptance

By Bernard Cole, Embedded.com
March 6, 2006

The future of asynchronous logic design is looking a little bit brighter. As the semiconductor industry struggles with mounting problems trying to achieve significant yields, higher performances and lower power without significant increases in fabrication costs, developers are turning to asynchronous alternatives to solve these problems.

Asynchronous, or clockless, logic as a mainstream circuit logic alternative could be critical in many embedded designs in consumer electronics and mobile devices.

Several things have occurred to make this alternative more viable. First, companies active in developing asynchronous logic are shifting from selling a particular IP approach to becoming fabless IC companies, using their logic expertise to address segments of the market synchronous logic is having a hard time satisfying.

Second, the numerous variations – and names – of asynchronous logic are settling out to three or four – optimized for specific segments of the market. A third trend is the increasing use, even amongst the largest semiconductor companies, of asynchronous techniques to achieve the performance, power, and cost objectives the market demands.

And finally, efforts at universities and within the same asynchronous companies are increasingly focused on developing EDA tools and design flows that can be integrated into the custom and semi-custom methods now used by the industry for synchronous design.

Traditional synchronous design limits

Traditionally, most circuit designs in the mainstream are built with synchronous logic, small blocks of combinatorial logic separated by synchronously clocked registers. The biggest advantage of this approach is that that synchronous logic makes it easy to determine the maximum operating frequency of a design by finding and calculating the longest delay path between registers in a circuit.

But as devices move into the 90 nanometer range and below, it is it is becoming extraordinarily difficult to find and predict the critical path delays and to achieve the all-important timing closures. And as process technology works down to 45 nm and below, things only get worse, with shot noise, charge sharing, thermal effects, supply voltage noise and process variations all making calculations of delay more uncertain and difficult.

Because synchronous logic designs are always on, balancing power and performance becomes critical as integration levels increase. ‘Simple’ power consumption and power dissipation issues are not the only problem. In many millions of transistor SoC designs there are large clock current surges necessary, which tax a circuit's power distribution nets as well as the thermal stability of the circuit. There is also the growing inability to control noise and metal integrity. And in some of today's system-on-chip designs with millions of gates, the job of maintaining the global clocking across the area of a chip is becoming problematic.

While synchronous logic designers have been extraordinarily successful at squeezing every bit of performance out of their designs and at finding work-arounds to the myriad of design problems facing them in the nanometer range, it is becoming more expensive and it is taking longer to develop designs.

The advantages of clockless logic

At 90nm and below, asynchronous logic may be able to take advantage of the increasing process volatility. “Unlike synchronous designs where developers have to assume worst case values, asynchronous logic works with the average values and the average process,” said Peter Beerel, Associate Professor, Electrical Engineering-Systems Department at the University of Southern California, who heads the asynchronous logic research efforts there. “This is an enormous advantage in the face of the many process variations that must be dealt with in the current generation of 90nm and below designs. At 90nm, that can mean as much as a two fold improvement in performance.

“As every move down in geometries occurs, the problems get greater. The industry has been able to solve those problems, but it has not been cheap. Each year, as synchronous designs become more difficult to do and it becomes more expensive, asynchronous becomes more and more attractive. And each year the tools and methodologies the asynchronous community is developing get better.”

Unlike the familiar design that synchronous methodologies use, asynchronous circuits (also called self-timed, locally clocked, clockless and a number of other names descriptive of the different approaches) remove the need for a global synchronizing clock. Instead, the process of computation is controlled though local clocks and local handshaking and handoff between adjacent units. What this means for high performance and power-efficient design is that such local control permits resources to be used only when they are necessary.

Although asynchronous designs usually require more transitions on a computational path than synchronously designed CPUs, the transitions only occur in areas involved in the current computational task. Moreover, most forms of asynchronous logic are, to one degree or another, delay-insensitive because of their clockless derivation, making designs in the 90 nanometer and below regime easier, at least theoretically, because they depend on average clock skews, not worst case.

But, given the diversity of approaches, the difficulty in implementing the designs because of the lack of appropriate and familiar design tools and flows, and the perceived lack of sufficient performance improvements to justify the additional time and cost it takes, asynchronous logic has not had a warm welcome within most mainstream electronics companies.

From providing IP to targeting niches

Rather than trying to fight the institutional momentum behind synchronous logic, a number of asynchronous logic-based companies – such as Achronix, Fulcrum, Handshake Solutions, Silistix, and Theseus - have stopped trying to sell their particular approaches as IP to an unconvinced design community. Instead they are using the special advantages their particular asynchronous design methods give them to target specific market niches where this methodology will give them an edge.

Rajit Manohar, founder and chief technology officer of Achronix, said the company has already demonstrated a 650-plus megahertz FPGA built with relatively conservative 180 nanometer design rules, even looser than conventional slower synchronous-based FPGAs.

Later this year the company will introduce an Ultra line of FPGAs fabricated with 90 nm CMOS that will operate at clock frequencies in the 700 MHz to 1.2 GHz range, based on a synchronous interface, an asynchronous core and a set of software tools to convert synchronous design flows to asynchronous logic. It is targeting many applications where ASICs have predominated and were previously inaccessible to FPGA manufacturers.

Fulcrum is targeting networking processing applications with its asynchronous logic technology and EDA tools, initially in a crossbar switch it incorporated that into its PivotPoint SPI-4 switch chip, now used by about 15 vendors, said Mike Zeile, vice president of marketing. The same crossbar core is also used in its most recent family of 10-Gbit Ethernet switches — called FocalPoint — supporting up to 24 ports.

The asynchronous paths allow the devices to offer full 10-Gbit speed and 200-nanosecond total latency through the chip in a conservative, low-leakage, 130-nm process. The registers and Ethernet ports are traditional, synchronous logic, while the sections that determine the performance and power consumption - the crossbar and its SRAM — are asynchronous.

Handshake Solutions has focused on the low power advantages of clockless logic and has been successful in the smart card market, where they have sold millions of units of eight bit asynchronous MCUs. The company is taking advantage of the fact that its parent company Philips is well positioned in automotive electronics, where it has worked with ARM Ltd. on the introduction of a low-power, 32-bit ARM designed for that market.

Silistix has targeted the on-chip interconnect segment of the market, where existing globally synchronously-clocked shared-bus topologies are proving inadequate. According to David Fritz, vice president of marketing, it has developed a globally asynchronous, locally synchronous interconnect fabric it calls Chain, supported by EDA tools and libraries called Chainworks to allow designers to generate asynchronous and delay-insensitive links between traditional synchronous logic blocks in existing designs. It supports multiple local bus protocols, including AHB, APB, and AXI as used by ARM Holdings plc, enabling existing IP blocks to be used without modification.

Theseus Logic uses a low-power optimized asynchronous logic methodology as part of its new product and service strategy aimed at developing low cost, highly integrated mixed signal system-on-chip devices for wireless sensor nodes.

Companies and engineers who want to look under the hood to see how these performance and power consumption advances are achieved can examine a handful that research and experience have shown to have advantages in particular segments of the market. Achronix and Fulcrum, for example, base their approaches on the quasi-delay insensitive (QDI) logic pioneered at the California Institute of Technology, and in the former case, refined at Cornell University. A QDI circuit does not use any assumption of or knowledge of delays in operators and wires and are the most conservative asynchronous circuits in terms of the use of delays. But they are also the most robust to variations in physical parameters because the circuit's dependence on delays is minimal.

Handshake Solutions uses a combination of regular synchronous and asynchronous logic in a clockless variation it calls handshake logic, which it implements in several variations, as two-phase or four-phase protocols, and as single-rail or double-rail data encoding. In most recent designs it has opted for a four-phase single rail implementation because it does not require dedicated 'asynchronous' standard cells, allowing the use of standard EDA tools and reuse of standard data path blocks.

Silistix's asynchronous approach is based on research at the University of Manchester involving the use of self-timed packet-based networks to solve timing closure problems in complex system on chip designs. Theseus also uses a delay insensitive clockless logic, but bases it on a class of NULL Convention Logic (NCL) circuit techniques developed by the company's founder, Karl Fant, which integrates data transformation and control into a single expression, thus yielding inherently clockless, delay insensitive circuit.

Synchronous edges closer to asynchronous

On their side of the logic divide, traditional semiconductor companies have not been blind to the perceived advantages of asynchronous logic and have maintained on-going research efforts, and in some cases they have come up with variations that are essentially clocked synchronous circuits, but that owe a lot to asynchronous methodologies, including mesochronous, plesio-synchronous and adiabatic logic.

Mesochronous involves logic designs in which various portions of an SoC design are not synched independently of the logic signals, but where the clock signals accompany the data. In plesio-synchronous logic, multiple clock domains distributed throughout a chip share the same clock, but with timing from a separate parallel clock signal distribution system. In adiabatic logic, instead of supplying a constant voltage to a chip and then clocking signals through, circuits have a periodic, sinusoidal power system that activates logic gates as needed.

Also included in such alternatives are techniques such as Intel's self-resetting logic, which shares many characteristics with asynchronous logic, including the ability to use the same custom and semi-custom EDA tools and design flows.

Building asynchronous logic with synchronous tools

USC's Beerel said that clockless logic companies and researchers are moving toward similar EDA tools and design flows, reflecting the same trends that are occurring in synchronous logic design. Currently, there are essentially two main approaches in mainstream circuit design: full custom design flows, which may use EDA tools but in which there is a high degree of so-called “handcrafting”, that is, optimizing the design for performance and semi-custom, fully automated methods.

With the latter, using tools provided by firms such Artisan, Cadence, Magma, Mentor Graphics, Synopsis and Virage, the strict CAD tool flow imposes limitations on what a circuit designer can do in terms of optimizing performance, power, or some other important parameter. To develop a repeatable and predictable design flow and a consistent set and interface, restrictions are placed on the kinds of circuits that can be used and how they can be used.

As a result, where an Intel can build a full custom circuit with a 2-3 GHz clock rate, users of the more mainstream, fully automated, semicustom approach must settle for 500 MHz in a typical design. “But in exchange for that the developer gets a mature, well-oiled set of CAD tools that allows the creation of chips with much less effort and expense than the alternative,” said Beerel, “with a much faster turnaround and with greater assurance of a working design at the end. The tools are also well supported and upgraded each year with incremental improvements in terms of power and performance. “

A similar dichotomy is emerging in the asynchronous community, but in reverse. Because they are focused on competing with synchronous, many of the leading clockless-based logic companies such as Achronix, Fulcrum and Silistix have developed custom design flows based on a combination of standard EDA tools and methodologies, but they in many cases depend primarily on careful handcrafting of the final design to optimize their circuits for the targeted markets.

Theseus - which pioneered the use of standard EDA tools in the design of clockless circuits - and Handshake out of Philips have what may be the most automated clockless circuit design flows, with many traditional EDA tools. The only question in the case of Handshake is its use of a C-like hardware description language for asynchronous logic called HASTE, rather than the industry standard Verilog or VHDL. Despite its simplicity compared to Verilog and VHDL used with most EDA tools, HASTE is a non-standard approach and requires a learning curve to use, become proficient with it and use it effectively, said Beerel.

“While it is a serious hurdle to Handshake’s acceptance, it can be overcome,” he said. “And they have been in this for quite a while and have created a niche for themselves in the low power segment of the market, in smart cards and low power 8-bit microcontrollers. They also have some penetration of the automotive market. The big question is whether they can make HASTE as mainstream as VHDL or Verilog.”

The search for the holy grail

The holy grail of the clockless design community, whether at the companies or at universities such as Caltech, Cambridge University, Columbia University, Leeds University, USC, and University of Utah, is to create an EDA tool flow that is as close to the traditional one for synchronous designs as possible. Same languages, simulation, and modeling tools, with asynchronous rules inserted when they are the only thing that that can accomplish a specific circuit level task.

“If it is necessary to develop the tools and building blocks for every aspect of a circuit design for asynchronous logic from scratch,” said Beerel, “there would not be any possibility of asynchronous logic moving into the mainstream of integrated circuit design. What we are working on here at USC is identifying those points in the design flow that require tools that are unique to clockless/asynchronous logic and everywhere else use traditional tools.”

Beerel and his coworkers have taken a two-phase asynchronous handshake approach -- developed by clockless pioneer Ivan Sutherland at Sun Microsystems for use in control circuits - and extended it with a proposed single-track full-buffer circuit family for use in both the control and data path.

“Using our single-track family, we have developed a prototype library of cells in 0.25 micron technology and were able to build a 600,000 transistor chip that achieved between 1.2 to 1.4 GHz measured performance,” he said. “The key point is that this chip was designed using a push-button fully automated place-and-route design flow.” They are now working with Fulcrum on extending these fully automated techniques to commercial devices, under a joint research grant to USC.

“I think we are getting very close to a crossover point this year where the industry will see applications and tool flows where the only practical solution will be asynchronous logic,” said Beerel. “We may be on the eve of a revolution. That may be my enthusiasm talking, but what I see in the lab and on the production line tell me that there is a tipping point in the near future.”