

Figuring out configurability

By Bridget Mintz Testa -- 10/1/2005

Electronic Business

<http://www.reed-electronics.com/eb-mag/article/CA6261842?pubdate=10%2F1%2F2005&&>



According to analysts, the market for configurable processors has a lot of headroom. That's why we were able to get (left to right) Chris Rowen, CEO of Tensilica; Gary Banta, CEO of Stretch; Carl Schlachte, CEO of ARC International; and Ronnie Vasishta, CEO of eASIC, all together for a photograph (John East, CEO of Actel, was traveling and unavailable).

Until 18 months ago, the most visible high point for configurable processors (CPs) was in 1996. In that year, Nintendo licensed a configurable IP core developed in 1993 by ARC International. The core became a graphics accelerator for Super Nintendo.

Since then electronics manufacturers have shown little interest in these processors. Their benefits weren't clear, and configurability itself was confusing. In the last 18 months, however, the number of licenses for configurable cores rose, indicating that manufacturers have finally grasped the benefits of these processors.

Even a short list of benefits is impressive. For example, a single CP can replace multiple processors in consumer electronics. When that happens, space and power demands and cost decrease and performance, efficiency and ease of use increase. CPs can also reduce or eliminate chip redesigns and respins. They can even make it possible to build ASICs in large quantities and can still be customized for various purposes.

To realize these and other benefits, the five competitors in this space—Actel, ARC International, eASIC, Stretch and Tensilica—must prove the worth of configurability vis-à-vis entrenched, familiar technologies such as RISC-based fixed-processor IP cores and DSPs.

Vendors of configurable processors—whether they are IP licensers such as ARC and Tensilica or fabless chip makers such as Actel, eASIC and Stretch—face potential competition from large, established, familiar, richer companies such as Altera, ARM, Intel, LSI Logic, MIPS Technologies, NEC, TI, Xilinx and, essentially, any other company that manufactures general-purpose or specialty processors. Although none of the established CP companies seem to view configurability as a vulnerability today, that could change if the potential challengers gain noticeable market share.

Many analysts think the CP companies have a shot at changing the landscape. After all, they've been around for a while—some as many as 20 years—frequently shifting business strategy as they've navigated the twists and turns of technology. It's a tough space. Leopard Logic, which we covered 18 months ago ("Leopard's Sweet Spot," February 2004) seems to be shriveled up and disappeared. But these five are too old to be startups and too experienced to be upstarts, and they haven't lasted this long without gaining some savvy.

To complicate matters, though, they have the challenge of trying to explain where they compete and where they complement each other, and that's not easy. Stretch is a Tensilica customer; Actel is an

ARM customer. The boundaries between them are undeniably fuzzy, but here's our roundup of the configurable landscape.

Pre- and postfab configurables

Conventional core processors can't be altered during design or after fabrication. "ARM or MIPS says, 'Here's our core; you fit it into your design,'" says Tony Massimini, chief of technology at Semico Research. CP vendors scoff at that inflexibility. Change can happen either while the processor is in the design phase or after it's been fabricated, according to Jeff Bier, founder and general manager of chip consultancy and analysis firm Berkeley Design Technology.

To change a processor before fabrication, designers write new instructions into its IP core. "The idea with this kind of CP is that you start with a general-purpose processor," says Tom Halfhill, senior analyst for In-Stat's Microprocessor Report (a division of EB 's parent company). "Then you add instructions to customize it for special purposes, such as compressing video or graphics. You start with plain vanilla and wind up with a hot fudge sundae."



"People have been restricted from customizing hardware, because it's so expensive. eAsic is reducing the cost of customizing hardware, so that it's affordable to the masses."

—Ronnie Vasishtha

Once the design is finalized, fabrication starts, and that flow is similar to that of an ASIC; this is where Tensilica and ARC play. Fabricated processors can be configured through software-programmable logic or other means; this is where Stretch's and Actel's products fit. This type of configurability is fast, because the chip is finished. Many analysts define configurability by the postfab method, yet it's limited in scope compared to the design-in method. eASIC claims to span both categories.

The players in configurable cores

Analysts envision success for configurable cores. "In 2005 the core-based market represents less than two billion units," says Semico's Massimini. "In 2009 that market will represent more than three billion units. Configurable cores will see an increasing market share during this period." The question remains whether ARC International and Tensilica can compete against the likes of ARM and MIPS (more about that later). Both have their advantages and stumbling blocks.

ARC International has followed a circuitous route to its current business model. When it designed that first configurable IP core back in 1993, it was a business unit of now defunct video game company Argonaut Software. ARC spun off into independence in 1998 and went public in the U.K. in September 2000.

After its IPO, ARC pictured itself as a one-stop shop for embedded-system designers. In 2004, though, after two years of losses and falling revenue, the company began to focus on the company's IP and associated design tools. ARC's losses went from approximately \$41 million in 2003 to just \$9.2 million last year, and revenue climbed from \$17.6 million in 2003 to \$22.2 million in 2004. It's doing something right, says Rich Wawrzyniak, senior analyst for ASICs and SoCs at Semico, because its number of licensees has doubled: Of ARC's 103 licensees for its 32-bit RISC-based ARC 600 and ARC 700 CP cores, 50 deals were negotiated in the last 18 months." Customers include BroadCom, Fujitsu, Infineon, LeapFrog, SanDisk and Sony.

One possible reason for the uptick: a "hedge your bets" technology strategy. "Customers like to have the potential for configurability, but they may not use it," says Will Strauss, principal and founder of the DSP market research and consultancy firm. "Everybody likes having it, just as everyone likes having DSP capability."

Carl Schlachte, ARC's CEO, cites very different reasons: "We have a slew of customers that know what they want to build but can't get it from existing companies. When you buy Intel, Motorola, ARM or MIPS chips, you can't change them. For many applications, this is fine. But when your applications are very specific, you need a microprocessor you can customize. This is why people call us." Schlachte says customers know that the optimized cores work and don't require respins, so they can get products to market faster and in higher volumes.

ARC expects to announce a new video-optimized core before the end of 2005. It's a follow-on to the company's core optimized for audio.

Competing with ARC

Tensilica, founded in 1997, is ARC's sole competitor in the configurable IP core space. "They are running head to head," says Marcus Levy, president of the Embedded Microprocessor Benchmark Consortium. Even so, "low-key" best describes the competition between the two companies. "Today competition is between configurable processors and fixed processors rather than between CP companies," says Wawrzyniak. "There's plenty of room for growth right now,"

Chris Rowen, Tensilica's founder and CEO, worked with RISC microprocessor pioneer John Hennessy as a graduate student at Stanford University (Hennessy is now the university's president). He followed Hennessy when he founded MIPS Computer Systems (now MIPS Technologies), but Rowen later went to Synopsys. "In 1997 I decided that the world needed a more automated approach to microprocessor design that wasn't being advanced by anyone," says Rowen. "We really democratized the idea of advanced microprocessor design, so you don't need a Ph.D. to do design."

With \$64 million raised in four venture capital rounds—the last in April 2001—Tensilica developed two 32-bit RISC-based cores: Xtensa V and Xtensa LX. There's also the Xtensa HiFi Audio Engine, a core optimized for embedded audio applications. The cores come with all the necessary tools and equipment for chip and system designers. More than 75 companies—including Agilent, Broadcom, Cisco Systems, Fujitsu, Hughes, NEC, NTT and Sony—have licensed Tensilica's technology.

Financial analysts have been expecting the company to launch an IPO since 2004, because it appears to be financially healthy. "Tensilica must be breaking even or making money, because it isn't burning through it," says Halfhill. "It hasn't had to go back and get more."

The company still hasn't announced a time frame for an IPO, possibly because of some potential patent problems. In 2004 an anonymous party challenged claims in the company's key technology patent. The situation worsened when the U.S. Patent Office examined Tensilica's claims and rejected them. Tensilica appealed successfully, and according to a company spokeswoman, the Patent Office will officially reinstate all claims sometime before the end of the year.

Although ARC and Tensilica are the only two IP companies in the nascent CP market, there's other competition. The biggest threat they face comes from fixed general-purpose processor (GPP) IP companies, specifically ARM and MIPS. "The battle has always been ARM's standard, legacy instruction set with lots of tools versus Tensilica's greater optimization," Levy says. Designers must weigh the known and familiar against the innovative and risky.

In fact, both ARM, with the OptimoDE, and MIPS, with the CorExtend CP, have developed their own configurable processors. Neither company is really pushing its CPs—yet—but that may well change if CPs suddenly become very popular. "The larger companies can't be ignored," Bier says.

Might ARM or MIPS—especially ARM, which has recently been on a buying spree—gobble up the CP IP companies? "It doesn't look like they're interested in ARC or Tensilica right now," says Halfhill. "They'd just get an incompatible architecture, and there's no need to eliminate the 'competition,' because they don't see ARC and Tensilica succeeding."

The players in software-programmable logic

On the postfabrication side, the boundaries between the competitors are even fuzzier. Take fabless FPGA manufacturer Actel, founded in 1985. Because it doesn't use a configurable core at all, some analysts don't think Actel belongs in the configurable mix.

Bier disagrees: "You can put a processor on an FPGA and use the FPGA to augment the processor capability, and that's a CP," he says.

That's exactly what Actel did. It licensed a soft version of the fixed 32-bit ARM7 Thumb processor core for its FPGAs. The configurability isn't in the ARM core—it's in the FPGA's programmable logic.

Actel's choice raises questions. Why pick a fixed processor core instead of a configurable one? And how did the company talk ARM into this arrangement when, according to Halfhill, fear of IP theft meant that ARM prohibited the use of its cores on FPGAs other than for prototyping?



"Not every application requires configurable processors, but more and more do. They will grow in marketplace acceptance."
—Rich Wawrzyniak, Semico Research

The answer to the first question is that because the FPGA's programmable logic handled configurability, Actel didn't need a CP core. "Why not get the world's most accepted embedded processor and use that?" reasons John East, Actel's CEO and president.

Winning the deal with ARM took more effort. Actel dealt with ARM's fear of theft of IP from FPGAs by encrypting the ARM7 core's instruction set. Actel's customers can use the core, but they can't see it. What they can't see, they—or their customers—can't steal.

To persuade the IP giant that licensing was a good idea, East pointed out that small startup companies can't afford the millions of dollars in licensing fees and royalties ARM charges. Yet if ARM has no traction with these small companies, it can't dominate the IP market in the future. "We developed a model for them," East says.

It's a straightforward model. ARM licenses IP to Actel, which pays all the licensing fees and royalties. Those don't pass through to customers. Instead, Actel sells its ARM-enhanced FPGAs at high unit prices that are still much less than the customers would have to pay directly for ARM IP. The bulk of the FPGA sales money goes directly to ARM, but "Actel makes a decent profit," East says. "ARM gets royalties, and Actel gets new customers. ARM hopes the Actel customers succeed and make full ASICs, using ARM IP in high volume. Actel will take the 99 percent that don't go on to high volume."

Actel's configurable-FPGA software was still in the alpha testing phase as of this writing, but East expects to ship silicon by the end of 2005. In the meantime, the company can rely on its traditional FPGA revenue, which was \$165.5 million in 2004.

As a traditional FPGA vendor, Actel vies directly with FPGA makers Altera and Xilinx, which have each developed "processor cores designed specifically for putting on their FPGAs," says Halfhill. "The advantage is that it costs only \$500 to \$1,000 to license one of their cores. A disadvantage: The cores are proprietary. A second disadvantage: The cores aren't familiar, so they're not as easy to use as ARM or MIPS cores."

Actel's ARM license provides a competitive advantage over Altera's and Xilinx's proprietary, unfamiliar products. Halfhill believes that, for low-volume applications, the idea of an FPGA plus a processor core will develop more traction. According to Bier, neither Altera nor Xilinx has publicly cited its IP-enhanced FPGAs as keys to its business. For all these reasons, Actel has little real competition for its enhanced FPGA. But if Altera and Xilinx crank up sales on their core/FPGA combos, Actel could find configurability success a lot harder to achieve.

Even more than Actel, privately owned eASIC, which has venture funding of at least \$12.5 million, blazed its own path to "configurability." The fabless company, founded in 1999, builds both structured-ASIC cores and the chips themselves. Its business strategy centers on eliminating the costs and turnaround times associated with standard and structured ASICs. "Our strategy is to invigorate customization in the industry," says Ronnie Vasishta, eASIC's CEO. "People have been restricted from customizing hardware, because it's so expensive. eASIC is reducing the cost of customizing hardware, so that it's affordable to the masses."

Reducing that cost lies in the way eASIC builds chips and in its version of configurability. Ordinary structured ASICs are made in much the same way ASICs are, except that whereas every ASIC chip layer is customized, structured ASICs start with a stack of standardized layers and add just a few custom layers on top.

eASIC takes standardization up to the very top layer of the chip—only the connecting via layer between the top two metal layers of its eASICore chip is customized, and only about 1 percent of that. Prefab configurability is possible because customers provide the data for that last, top layer. eASIC finishes the chip in the foundry. It achieves postfab "soft" configurability by downloading a bit-stream that customizes the chip's logic. "Anyone doing FPGAs could use our chip fabric, but the result is closer to an ASIC in density," says Vasishta.

eASIC's fabrication technique can produce even a single chip cost-effectively. The company uses electron beams to directly "write" chips on wafers instead of using more-traditional and much more expensive mask technology. Writing an entire wafer costs about a thousand dollars with the electron beam, Vasishta estimates, whereas creating a one-layer mask runs between \$40,000 and \$150,000, depending on the chip layer involved and its complexity. And masks, of course, may not come out right the first time.

Few of eASIC's customized chip customers want their names revealed, but Vasishta says the company—which started chip production just this year—is shipping millions of units. The one customer that isn't shy, STMicroelectronics, licensed eASIC's structured-ASIC core for its SPEAR chip. The company publicly announced that it took only 24 hours to design its chip with eASIC's core. According to Vasishta, turnaround time from design completion to final silicon may take only three to four weeks and nonrecurring engineering costs are zero.

Like ARC and Tensilica, eASIC is facing competition from much larger, well-established competitors, namely LSI Logic and NEC. But Jim Turley, owner and principal analyst of Silicon Insider, still thinks the company's future is bright. "Because structured ASICs are new, customers aren't quite sure how to evaluate them," he says. "There's a lot of room in the structured-ASIC space for capabilities such as ease of use, price and performance. eASIC is doing well, because the structured-ASIC market is growing and the company fills a nice slot between FPGAs and ASICs."

Still more configurability

Just as eASIC and Actel defy categorization, so too does the version of configurability of Stretch, with its S5000 family of processors. Some analysts compare the processors to FPGAs; others compare them to structured ASICs. Some call them hybrids. Forward Concepts' Strauss says, "Stretch is making a DSP, although you can also use it for traditional processor purposes."

Founded in 2002, it achieves configurability by embedding programmable logic in a Tensilica IP core architecture. With Stretch's software support tools, chip developers can program and configure the processor after it's fabricated, using ordinary C/C++ code. "Stretch processors can stay configurable forever," says Gary Banta, the company's founder and CEO. "You can add instructions to the architecture at any time."

According to Banta, Stretch has no intention of customizing its processor. "It's hard to build a horizontal business," Banta says, "but it results in a healthier, profitable business." Banta

acknowledges that this form of growth is slow. "The design wins aren't big or fast," he says, "but once they get momentum, they keep going."

For now, Stretch can keep going for another year on whatever is left of its original \$62 million in venture funds. By then, Banta expects to be selling products. Stretch has announced two small startup customers: biometric security company A4Vision and wireless technology company PureWave Networks. Banta says, however, that Stretch "is working with about a hundred customers that are deeply engaged in using our tools and working on development or evaluation."

Notwithstanding the confusion about Stretch's chip, most analysts view the company favorably. "In some ways, I think, it has the best of both worlds," Levy says. "It has configurability, but the product is already built. So the customer buys the product off the shelf and then designs configurability."

The next stage

After their somewhat wayward paths, CPs are finally gaining some acceptance. With design wins, vendors are answering questions about how the technology works and how the processors can be used. The future looks promising. "Not every application requires configurable processors, but more and more do," says Wawrzyniak. "They will grow in marketplace acceptance."

Customers are still figuring out just when and where it makes sense to take risks on the new technology. According to Jordan Selburn, a principal analyst with iSuppli, the CP vendors must still prove to reluctant customers that there are financial and technological advantages to configurability over use of traditional FPGAs and ASICs. Although most are still private, Selburn says that "their revenue streams indicate that they are finding success."

It's entirely possible, then, that the fuzzy boundaries of the configurable vendors will eventually play to their—and the industry's—advantage. Even if side-by-side comparisons are difficult, Wawrzyniak, for one, thinks the CP vendors are clear on where they fit and what they offer. As a result, he sees system designers seeking the optimum mix of configurable and traditional processors. "When they figure this out," he says, "the configurable processor market will reach the next stage of its growth."

Will configurable processors ever be part of the mainstream, or will they always be a niche option? Send your thoughts to feedback@eb-reedbusiness.com.

THE CONFIGURABLE-PROCESSOR LANDSCAPE

| Company | Public or private | CP name | Type of configurability |
|--------------------|-------------------|------------------|--------------------------------|
| Actel | Public | Not yet named | FPGA + ARM7 IP |
| Altera | Public | Nios II | FPGA + Altera IP |
| ARC International | Public | ARC 600, ARC 700 | IP core |
| ARM | Public | OptimoDE | IP core |
| eASIC | Private | eASICore | Structured ASIC + PL* |
| MIPS | Public | CorExtend | IP core |
| STMicroelectronics | Public | SPEAR | SoC + PL* + ARM IP |
| Stretch | Private | S5000 | PL* embedded in Tensilica core |
| Tensilica | Private | Xtensa | IP core |
| Xilinx | Public | MicroBlaze | FPGA + Xilinx IP core |

*PL= programmable logic

SOURCE: ELECTRONIC BUSINESS

Bridget Mintz Testa is a regular contributor to Electronic Business.

© 2005, Reed Business Information, a division of Reed Elsevier Inc. All Rights Reserved.