



LSI Logic's Leverage RapidChip Heads to 90nm

What has over two million real ASIC gates, runs at over 200MHz with twenty levels of logic, burns about the power of a cell-based ASIC, and is economically feasible to deploy even in mid-to-small volume production? The answer is "Not an FPGA". While these specs may sound close to the marketing picture painted by programmable logic vendors, there is a vast gulf between the brochure and real-world performance in an average application. Structured and platform ASICs, however, can realistically reach these goals in your average application, and getting design teams to understand that fact is one of the biggest challenges faced by structured ASIC marketers in working with design teams jaded by years of spin-laden specsmanship from the FPGA industry.

LSI logic announced this week that their increasingly successful RapidChip platform/structured ASIC family is headed to 90nm. While the step to the next process node may come as no surprise, the implications merit serious consideration, particularly for high-end FPGA customers who are either pushing the limits of FPGA technology or going into volumes where FPGA unit costs are prohibitive. It also means that platform ASICs are taking another giant bite out of the cell-based market space. With their 90nm announcement, LSI rolled out densities up to 10 million "these-are-not-the-same-as-system" ASIC gates. This upward leap in density clearly blows any lingering theory that the company might be protecting its cell-based offering by throttling back on structured ASIC.

Structured and platform ASIC has now transcended the "will they catch on?" question and gone straight into "how big will they get?" At this juncture, the answer appears to be "really big." Structured ASIC now could join low-cost FPGA in the lucky category of devices with almost unbounded growth potential.

We can split the custom digital logic market into bins almost an unlimited number of ways, but consider the span of devices available from “simple” CPLDs through cell-based ASICs. If we limit ourselves to medium and higher volumes (let’s pick 4-digit numbers and up for convenience), we can then segment roughly by performance/density points. At the low end, we have essentially one bin that contains CPLDs and low-cost FPGAs. These devices all have small 2-digit or 1-digit price tags (US dollars) and can handle probably up to the range of 100K ASIC gates of complexity (or a little better, depending on your gate math). Low-cost FPGAs are cost effective in their density range from a volume of one up through millions of units. No analyst would come under suspicion for predicting huge growth potential for this market segment.

Moving up the density ladder one notch, we get to the problematic plateau of high-density FPGAs. High-density FPGAs face problems such as power, price, performance, and now platform ASIC that threaten to marginalize them into the role of “specialty” devices. Even today, to effectively cross into medium-volume effectiveness, high-density FPGAs depend on cost-reduction strategies such as Altera’s HardCopy (structured ASICs directly from FPGA) and Xilinx’s EasyPath (FPGAs cost-reduced by application-specific testing). Both of these strategies produce devices that are no longer fully reprogrammable FPGAs and that are density limited to the size of the original FPGA.

Carrying on to the next density level (the single digit millions of ASIC gates), we enter a domain that will probably become the exclusive property of structured and platform ASIC. With 90nm offerings moving this range up to 10-million ASIC gates with copious amounts of on-chip memory and robust IP libraries, there is no technology that threatens serious competition with these devices. FPGAs don’t match their density, performance, power consumption, or unit cost. Full-blown ASICs cost ten to a hundred times more to develop because of mask, NRE, tool, and design team costs.

Finally, in the rare air at the very top of the density/performance range, we have cell-based ASICs and custom silicon based on customer-owned tooling (COT) methodologies. With every process node, the design challenge for these devices becomes tougher, the mask and NRE costs rise higher, and the number of applications that require their unique performance or density capabilities becomes smaller. Already a good number of these designs are going into second-order applications where they are re-sold as ASSPs, FPGAs, platform ASICs, or other broadly useful devices that can take advantage of economy of scale to offset their massive development costs.

If this picture continues with no unexpected technology discontinuities, we get a scene with almost unbounded growth in structured/platform ASICs and low-cost FPGAs and marginal futures for higher-density FPGAs and ASICs. With each generation, the two more accessible technologies will press their density and performance numbers upward, stealing market share away from their less fortunate high-end brethren. LSI Logic is now proving that it has no compunction about pushing RapidChip straight toward the heart of their historical bread-and-butter product offerings. This means that at least one company sees enormous potential in platform ASIC.

LSI Logic is already on to the next step of the game. Many players will participate in the growing structured ASIC market, and everyone will have access to comparable levels of silicon technology. The axis most likely to determine success and failure in this space is application-appropriate IP. When a design team evaluates alternatives in structured ASIC, the deciding factor is most likely to be the IP collection that backs up each offering. LSI Logic has a long history of success at the IP- and application-focused marketing strategy. They come into platform ASIC with a wealth of experience helping their customers succeed by understanding the customer's design problems in advance and offering a combination of IP and services that accelerate and simplify the design process.

LSI is taking RapidChip into the communications, storage, consumer, and mil-aero markets where there are numerous mid-volume, capability-hungry applications waiting to take advantage of the low-NRE, high-performance combination offered by structured ASIC. In each area, LSI has a focused marketing effort, a comprehensive IP portfolio, partnerships with key industry players, and a growing list of blue-chip customers doing designs with RapidChip. Each market segment ideally demands slightly different mixes of hard-IP on their device, and one of the early criticisms of RapidChip was that almost every customer engagement created the need for a new "slice" (LSI's term for a particular platform's configuration of IP). With a couple of years of platform experience behind them, however, LSI now has a large selection of slices already designed, and the need to customize continues to dwindle.

One often overlooked advantage of the structured ASIC approach is the simplification of the design tool flow compared with cell-based ASIC and COT. In the case of RapidChip, the design process is supported by what LSI calls RapidWorx, which is a tool set that strongly resembles the design kits that have been used by FPGA designers for the past decade or so, right down to the Synplicity-based synthesis. Synplicity was the first to jump on the

structured ASIC wagon with both feet, and they stand to gain a commanding position if the market truly takes off. With the customers' tool expectations set much closer to the FPGA end of the dial, Synplicity had the infrastructure and experience to tackle the problem without much adjustment, and they got a big jump on the rest of the EDA industry in supporting structured-ASIC's needs.

From the customer perspective, there is no need to mix-and-match a point-tool based flow. Everything needed to design successfully is included in the design kit. Aftermarket tools are not ruled out, but they are most likely to show up in the front-end of the design flow in design creation and management, IP integration, and (we're still not sure exactly what it means) ESL. While the FPGA and structured ASIC boom means a mandatory adjustment for semiconductor companies, it carries strong implications for EDA as well.

Since these are "platform" devices, one would also expect a robust infrastructure for the development and deployment of embedded software and for the process of hardware/software co-design. To date, much of the focus from LSI seems to be on the hardware side of the equation, with the software side coming mainly as windfall from partnerships such as ARM and the supporting cast they bring with them. As more customers take advantage of the true embedded computing capabilities of structured ASIC platforms, however, watch for the software development component to push toward center stage.

Today, however, LSI's focus is on deepening the penetration with the current 130nm offering with capabilities like SerDes and increased memory offering, while driving toward their 90nm introduction in order to support emerging IP and interface standards not available at the 130nm process point. For designers of high-performance systems that surpass FPGA's capability but don't warrant ASIC's investment, progress can't come fast enough.

Kevin Morris, FPGA and Programmable Logic Journal

July 5, 2005