Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

ELIXENT IMPROVES D-FABRIX

D-Fabrix v2.0 Tweaks Parallel Architecture for Better Performance

By Tom R. Halfhill {6/27/05-02}

Elixent took the stage at **Spring Processor Forum 2005** to prove that listening to customers isn't a lost art. Using feedback from early adopters of its massively parallel configurable-processor core, Elixent has introduced D-Fabrix v2.0, which significantly boosts performance without increasing the overall gate count.

D-Fabrix embodies a concept that Elixent calls reconfigurable algorithm processing (RAP). It's an outgrowth of the Chess architecture described in an academic paper presented in 1999 at the ACM/SIGDA International Symposium on FPGAs. One of the paper's authors, Alan Marshall, is the chief technology officer of Elixent, a five-year-old U.K.-based spinoff from Hewlett-Packard Labs in Bristol, England. Marshall also presented D-Fabrix v2.0 at SPF.

The general concept of RAP is to build a chessboard-like array of four-bit ALUs, multiplexers, registers, local memories, and interconnecting switchboxes that work together as a massively parallel on-chip fabric. Using Elixent's proprietary tools, SoC developers can configure the fabric at design time to optimize it for specific algorithms. The D-Fabrix architecture lends itself to applications with inherent data parallelism. (See *MPR 2/9/04-15*, "Extreme CPUs Defy Conventions," and *MPR 7/21/03-01*, "Elixent Expands SoCs.")

Early D-Fabrix users include Matsushita Electric and Toshiba, which are interested in the architecture for consumer electronics and mobile products. Elixent says D-Fabrix can accelerate multimedia algorithms and other data-intensive tasks while requiring less silicon than conventional ASIC logic delivering the same performance. The first D-Fabrix chip, Toshiba's ET-1, uses Toshiba's Media Embedded Processor (MeP) as the on-chip controller for the fabric. MeP is a synthesizable 32-bit RISC core with a VLIW coprocessor. (See *MPR 6/10/02-02*, "New Processors for New Media.")

On the basis of user feedback and internal analysis, Elixent found that many of the four-bit ALUs in the fabric were being used as simple multiplexers instead of as compute

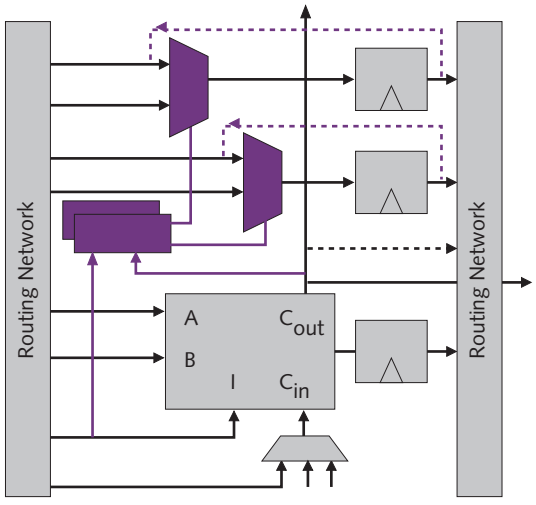

Figure 1. D-Fabrix v1.0 added dynamic multiplexers (shaded purple) to the switchboxes in the on-chip network fabric. The original Chess architecture that inspired D-Fabrix lacked these muxes, which forced the processor-configuration tools to use more than 50% of the available ALUs for that purpose. The muxes are controlled by the C_out signal or the bits received at the I input from the fabric's routing network.

Elixent Wins Japanese Design Award

Elixent has become the first company outside Japan to win the LSI IP Design Award. The British company won the trophy for its D-Fabrix technology at the recent SoC/SiP Developers Conference in Tokyo.

This seventh annual presentation of the award was sponsored by the LSI IP Design Award Committee and the Electric and Electronic Information Academic Promotion Foundation. The award was established in 1998 by the Development Bank of Japan, Matsushita, NEC Electronics, NikkeiBP, Toshiba, Renesas, Fujitsu, Rohm, Seiko Epson, and four Japanese universities.

The purpose of the award is to support new IP for SoCs and to promote the Japanese semiconductor industry. Although Elixent is British, the company has worked closely with Toshiba, and two prominent investors are Toshiba and Matsushita.

resources in the critical datapath. Elixent also found that programmers were writing more bit-manipulation code than expected and that the architecture could be improved for those operations. As a result, Elixent has introduced D-Fabrix v2.0, which addresses those issues.

Perplexed by Multiplexers

There's an old adage that when your only tool is a hammer, everything looks like a nail. Even before Elixent introduced D-Fabrix v1.0, the architects discovered that prototype processor cores generated by the configuration tools were using the arrays of ALUs for unanticipated purposes—especially for multiplexers. Indeed, more than 50% of the ALUs were used as muxes. This was inefficient, because a dedicated mux could be smaller and faster than an ALU. So in D-Fabrix v1.0, Elixent added dynamic muxes to the switchboxes that knit together the fabric's routing network. Figure 1 shows the role of the muxes in each switchbox.

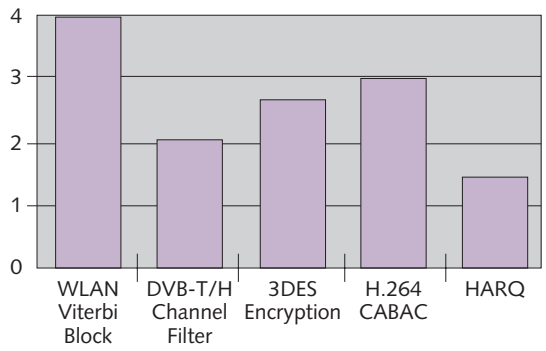

Figure 2. This chart shows some results of Elixent's internal benchmarking. These tests compare the simulated performance of a D-Fabrix v2.0 processor with the baseline performance of a D-Fabrix v1.0 processor of equal size. Elixent says a customer has achieved similar results—more than 2× improvement in performance per area—with a different internal benchmark suite. (WLAN: wireless local-area network; DVB-T/H: digital video broadcasting, terrestrial/handheld; 3DES: Triple Data Encryption Standard; H.264 CABAC: context-adaptive binary arithmetic coding for the H.264 video codec; HARQ: hybrid automatic repeat request, a retransmission protocol for communications.)

As customers began using D-Fabrix v1.0 to create application-specific processor cores, further analysis revealed that the dedicated muxes helped, but didn't entirely solve, the problem. Overall, the processors were still using 27% of the ALUs as muxes, even while some of the new muxes went unused. Some applications were worse in this respect than others: for instance, a processor configured for running decryption algorithms used 50% of the ALUs as muxes. Moreover, Elixent found that 15% of the ALUs were used for bit extraction, and 5% were used solely for executing bitwise logic. Most of the bit extractions consisted of removing a single bit from a nybble. These results sent Elixent back to the drawing board (or the CAD station).

After some investigation, Elixent discovered that the automated configuration tools were still using too many ALUs as muxes because the routing was more efficient that way. In the densely routed fabric, the wiring around the ALUs was better than the wiring around the muxes, so the tools found it more efficient to stack multiple four-bit ALUs when constructing a mux for words longer than four bits. In addition, combining an ALU and a dedicated mux avoided the routing congestion of using two muxes in a single switchbox.

The first step toward a solution was to improve the wiring around the dedicated muxes, so the smart configuration tools would be less tempted to use ALUs instead. At the same time, Elixent reassigned some switch positions in the local wiring of the ALUs to improve their connections with the switchbox muxes. This reduced the routing congestion when using both muxes in a switchbox. Finally, Elixent added some new logic outside the ALUs to support bit manipulation. Adding that logic to the ALUs themselves wasn't practical, because the necessary signals aren't available inside the ALUs.

One result of these modifications is that the silicon area required for an ALU and a switchbox increases by about 10% in D-Fabrix v2.0. However, the increase is more than offset by the reduced utilization of ALUs as muxes. Better yet, critical paths become shorter, because now the processor needs fewer ALUs, and their local routing in D-Fabrix v2.0 is more efficient. Elixent predicted that a D-Fabrix v2.0 processor core would deliver about 66% more performance than a D-Fabrix v1.0 core of the same size, but an actual implementation delivers about 100% more performance, on average. Of course, a customer could make the opposite trade-off by designing a D-Fabrix v2.0 processor that delivers the same level of performance as a D-Fabrix v1.0 processor, but with fewer gates. Figure 2 shows Elixent's

benchmark results of a D-Fabrix v2.0 processor on a cycle-accurate simulator.

D-Fabrix Chips Coming Soon

Real-world performance will be easier to evaluate when actual chips based on the D-Fabrix architecture become available. That day is coming soon, says Elixent. Customers are now taping out multiple SoC designs, and Elixent says the chips will be manufactured using "widely variant" fabrication processes. As with all massively parallel architectures, the main challenge is writing software that exploits the processor's potential throughput. In this case, it's not an insurmountable challenge, because the D-Fabrix designs are intended for multimedia and communications applications brimming with data parallelism.

Certainly it doesn't hurt to have early adopters like Matsushita and Toshiba in Elixent's camp. Both companies are also strategic investors in Elixent (Matsushita through its venture-capital arm, the Panasonic Digital Concepts Center). With two industrial/consumer-electronics giants hovering in the background, and the first chip designs nearing production, Elixent is approaching the day when it can demonstrate the value of its unique architecture. ◇

Price & Availability

The D-Fabrix v2.0 Reconfigurable Algorithm Processor core is available now from Toshiba as part of the TC280 ASIC library (0.13 micron) and is being ported to other proprietary deep-submicron processes. For more information about D-Fabrix and licensing fees, contact Elixent at *www.elixent.com*.