W H I T E P A P E R

FPGAs: Fast Track to DSP

Mentor Graphics®

Implementing DSP on FPGA

DSP, for many engineers in the field, is a proprietary device used to implement digital-signal processing functions. These conventional pieces of silicon have been used for years to quickly and efficiently digitize analog signals, arithmetically alter that digital information in a desired way, then convert the digital back into something sensory that we humans can typically see, hear, or feel. These digital signal processors, or DSPs, are found throughout the world of electronic equipment, from cell phones to flat-screen TVs.

Dedicated DSP processors and cores are still a widely popular means of achieving digital signal processing needs. In little more than a few days time, a programmer can take a DSP chip and write an algorithm that will efficiently perform the digital signal processing task required. But as CPUs have incorporated arithmetic coprocessors and extensions that optimize digital signal processing tasks, and as other silicon alternatives have emerged, the necessity of a separate DSP chip is no longer viewed as inevitable. In fact, there seems to be increasing evidence that the negatives of discrete DSP chips might outweigh their benefits.

First, there is the issue of fixed costs. All DSP chips are viewed as IP (intellectual property), and as such, have per-use costs associated with them. While such considerations aren't much of an impediment to a company producing a limited number of end products, many digital signal processing needs occur in high-volume products where steep discounting of retail prices over a short period of time requires the lowest per-unit cost possible.

Next, there is the issue of performance. Compared to either ASIC or FPGA hardware solutions, a DSP chip is by far the slowest option. ASIC chips run as fast as the technology allows. Similarly, an FPGA that has been optimized to perform a digital-signal-processing task, will run anywhere from 10 times to more than 1000 times faster than a single DSP chip. Whereas a DSP processor typically employs serial processing, the parallel capacities inherent to either ASIC or FPGA architecture will always give them both a significant edge over DSPs.

Finally, the power drain of most DSP chips can quickly become an issue, especially if the end product will rely on battery power to increase its portability. The explosion of battery-driven technologies such as Wireless LANs brings this power issue to the forefront of engineering considerations. While DSP chips are often quick to program into your system during the development phase, their battery-hungry ways make them less attractive options for digital signal processing with significant portability needs. Carter Horney, a WLAN researcher associated with Forward Concepts, a DSP research company, puts it bluntly: "DSPs just draw too much power for portable applications."

To see how DSP solutions come up short in wireless scenarios only requires a quick look at the numbers. The digital signal processing needs of an analog, dial-up modem runs at 8,000 samples per second, whereas a wireless LAN application often calls for 80-160 million samples per second. Even the industry leader of DSPs, Texas Instruments, doesn't use a DSP for WLAN, according to its director of business development, Yoram Solomon.

Historically, when the price, performance, or power downsides of DSP chips have become too much of a deterrent, most designers have turned to ASIC hardware solutions. A hard-wired, cell-

based custom chip, ASIC chips can be equipped to run as fast and as powerfully as the upper limits of chip technology allow. Once designed, ASIC chips are infinitely cheaper than DSP chips. First, the cost is in direct proportion to the amount of silicon used to achieve the chip's function. The more efficient the architecture, the lower the eventual price of the ASIC chip. For end products manufactured in large enough volumes, the cost savings of ASICs are even more impressive. Whereas the per-use IP costs on DSP chips are fixed, the cost of ASIC hardware continually diminishes over longer manufacturing runs. Seen exclusively from a post-design perspective, an ASIC solution will be faster, more efficient, and cheaper than its DSP chip alternative.

But as wonderfully as ASICs stack up to DSP chips in bench comparisons, the low cost of an ASIC chip completely ignores the additional time and NRE (non-recurring engineering) costs that are required to get you there. Proponents of ASICs all but admit that their hardware solution works best when high volumes of chip manufacture are anticipated, the product's time-to-market isn't critical, and a design team's prowess at creating a first-silicon functional chip that eliminates post-sale upgrades.

But in today's fast-paced environment of semiconductor innovation, in which a cutthroat sense of urgency is only magnified by a chip's all too imminent obsolescence in the marketplace, long lead times and guarantees of high-volume production are often in short supply. And as much as any engineer hopes that his or her ASIC architecture has been perfectly tailored to the needs of the engineering task, the real world of ASIC design is littered with costly return trips to the fab.

DSP vendors have been painfully aware of the lost market share that more powerful general-purpose CPUs and ASIC hardware solutions have both inflicted on their software DSP chip solution. Their response has been to create hardware-accelerated engines to increase performance gain. While engines such as dedicated Viterbi decoders and matrix multipliers enhance a DSP chip's parallelism, the performance gain is less than 100% per additional engine and the DSP chip gets more expensive as it becomes more application specific.

An increasing number of designers are turning to programmable logic devices, or FPGAs, as a way to navigate the software-hardware extremes of DSP or ASIC design solutions. An FPGA that's been enhanced for digital-signal-processing gives you unlimited customizing options in a chip without all the silicon physical-design work required for an ASIC solution.

How does an FPGA compare to DSPs and ASICs in terms of the performance and power? According to recent benchmarking efforts by BDTI (Berkeley Design Technology), an Altera Stratix 1S20-6 chip supported more than a dozen channels of BDTI's Communications Benchmark, while the 300-MHz Motorola MSC8101 supported only about one-fifth of one channel. (It should be noted that, due to FPGA's flexibility to trade off parallelism against resources used, and the resulting necessity on the designers' part to optimize applications as a whole, BDTI had to create a Communications Benchmark that offered a comparison at the application level instead of at the algorithm kernel level as had been adequate for comparisons between DSP chips.)

While even FPGA vendors admit that FPGAs don't outperform ASICs in either the performance or power categories currently, the differences here are not nearly as substantial as those between FPGAs and DSPs. And when it comes to speed, an optimized FPGA chip can run neck and neck with its ASIC counterpart.

Price has long been the deciding factor for designers in their preference for creating ASIC chips over an FPGA. And it's true — once fabricated, ASICs are significantly cheaper than FPGAs off

the shelf. But such comparisons completely ignore both the upfront costs of ASIC in terms of their initial development, fabrication, and verification, and the post-market expenses required if an ASIC chip requires upgrades for any reason.

Fabrication of an ASIC chip, which often costs upward of $500,000, is a particularly high-stakes moment in the development process. If the design has hidden glitches that only manifest themselves after fabrication, you've got a very expensive tutorial chip on your hands, not to mention the ever-growing number of man-hours that will be required to debug your faulty chip.

Given this all-or-nothing reality of ASIC development, it is easy to see how design teams working on ASIC hardware solutions could, quite reasonably, move more cautiously — and therefore, less nimbly — through their chip's development. But in a fast-changing marketplace, agility becomes a premium commodity, and herein lays the vast potential of FPGA design.

While the higher price tag associated with FPGAs (as compared to ASICs) has been perceived as a major obstacle in the past, more design teams are taking into account NRE costs involved in ASIC development and coming to the conclusion that the flexibility and agility of FPGAs represent a major shift in the economics of hardware architecture.

Vendors promoting FPGAs have developed a number of features that fortify the chip's arithmetic capabilities. Initially, FPGAs were equipped with low-latency carry-chain-routing lines to speed addition and subtraction operations across multiple logic blocks. Next, embedded AND gates accelerated array multiplication operations through their successful integration with block-to-block diagonal routing lines.

Today, LUTs (look-up-table) are combined with fast carry chains to provide highly efficient addition and subtraction operations, while embedded dedicated multiplier-function blocks on-chip and complete MACs handle multiplication and division functions. FPGAs like Altera's Stratix EP1S and Xilinx's Virtex-II and Virtex-II Pro tout their ability to integrate dedicated-arithmetic units.

The use of nonstandard 18-bit data inputs in Altera and Xilinx's arithmetic structures provides a keen advantage in programming digital-signal-processing, as these functions rarely call for exact 16-, 24-, or 32-bit precision. FPGAs with LUT and register structures give you the data precision required in the logic blocks. Several FPGAs also allow you to tap into LUT for both logic functions and small RAM arrays, thereby eliminating the concern that these chips—with their predefined bus widths on embedded-memory blocks — might otherwise pose with regard to memory. This density and bus-width flexibility within FPGA's LUT structures allows a designer to augment or swap out dedicated RAM arrays to enhance advanced signal processing found in such applications as wireless LANs.

Of great concern to any designer contemplating a switch to FPGAs as a means of achieving more sophisticated digital signal processing is the fear of adequate design-software support. Embedded circuitry is all fine and good, but how do you know if FPGAs are giving your signal-processing algorithms all the power, performance and efficiency that you're paying for both in terms of money or disruption of your design team's normal workflow?

To account for such parameters, FPGA vendors often provide a third-party synthesis compiler together with their own place and route software. In combination, these tools can situate multipliers and MACs via your HDL code automatically. Additional assistance with coding styles helps guide the design software to create the desired result.

Higher-level functions, such as filters and transforms, are made simple through pushbutton utilities. Creation of a fully serial or fully parallel circuits are possible, and some vendors' core generator tools are robust enough to make a hybrid approach that is both fast and easy on the power.

Ultimately, this versatility might be the real key to the rise of FPGAs in the signal-processing realm. Key tasks could be assigned to hardware to take advantage of optimum speed, power consumption, and per-unit costs, while other tasks are performed by software to speed time-to-market and ease legacy compatibility. With digital signal-processing continuing to be a critical component in the evolution of wireless networks and in multimedia, FPGAs will grow in their ability to deliver state-of-the art signal processing in a fraction of the time that an ASIC can be brought to market, and at a price (once factors other than the mere price of the actual chip) that will only become more affordable than ASIC design over time.

Visit our web site at www.mentor.com for the latest product news.

Corporate Headquarters
Mentor Graphics Corporation
8005 SW Boeckman Road
Wilsonville, OR 97070-7777
Phone: 503-685-7000

Sales and Product Information
Phone: 800-547-3000
503-685-8000

Silicon Valley Headquarters
Mentor Graphics Corporation
1001 Ridder Park Drive
San Jose, California 95131 USA
Phone: 408-436-1500
Fax: 408-436-1501

North American Support Center
Phone: 800-547-4303

Europe Headquarters
Mentor Graphics Corporation
Immeuble le Pasteur
13/15, rue Jeanne Braconnier
92360 Meudon La Forêt
France
Phone: 33 (0) 1-40-94-74-74
Fax: 33 (0) 1-46-01-91-73

Pacific Rim Headquarters
Mentor Graphics (Taiwan)
Room 1603, 16F
International Trade Building
No. 333, Section 1, Keelung Road
Taipei, Taiwan, ROC
Phone: 886-2-87252000
Fax: 886-2-27576027

Japan Headquarters
Mentor Graphics Japan Co., Ltd.
Gotenyama Hills
7-35, Kita-Shinagawa 4-chome
Shinagawa-Ku, Tokyo 140
Japan
Phone: 81-3-5488-3030
Fax: 81-3-5488-3021

Printed on Recycled Paper

1-03-WCI