

Xcelljournal

SOLUTIONS FOR A PROGRAMMABLE WORLD

FPGAs Command Center Stage in Next-Gen Wired Networks

INSIDE

Dorkbot uses Xilinx
Spartan-3E Starter Kit
to Make 19th-Century
Pipe Organ Rock!

Satellite-Based Computing
Flies Flexible Virtex Platform

Tools and Techniques to
Tame FPGA Power Budgets

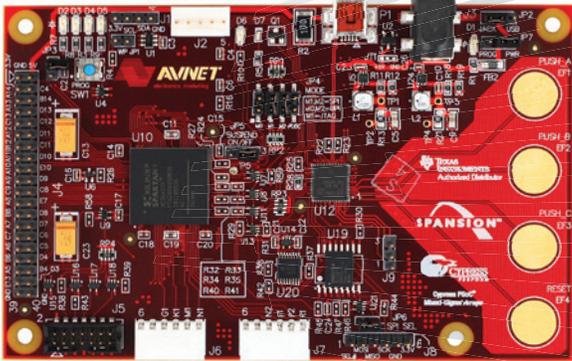


www.xilinx.com/xcell/

Xilinx® Spartan®-3A Evaluation Kit



DESIGNED BY AVNET



Target Applications

- » General FPGA prototyping
- » MicroBlaze™ systems
- » Configuration development
- » USB-powered controller
- » Cypress® PSoC® evaluation

Key Features

- » Xilinx XC3S400A-4FTG256C Spartan-3A FPGA
- » Four LEDs
- » Four CapSense switches
- » I²C temperature sensor
- » Two 6-pin expansion headers
- » 20 x 2, 0.1-inch user I/O header
- » 32 Mb Spansion® MirrorBit® NOR GL Parallel Flash
- » 128 Mb Spansion MirrorBit SPI FL Serial Flash
- » USB-UART bridge
- » I²C port
- » SPI and BPI configuration
- » Xilinx JTAG interface
- » FPGA configuration via PSoC®

The Xilinx® Spartan®-3A Evaluation Kit provides an easy-to-use, low-cost platform for experimenting and prototyping applications based on the Xilinx Spartan-3A FPGA family. Designed as an entry-level kit, first-time FPGA designers will find the board's functionality to be straightforward and practical, while advanced users will appreciate the board's unique features.



Get Behind the Wheel of the **Xilinx Spartan-3A Evaluation Kit** and take a quick video tour to see the kit in action (*Run time: 7 minutes*).

Ordering Information

Part Number	Hardware	Resale
AES-SP3A-EVAL400-G	Xilinx Spartan-3A Evaluation Kit	\$39.00* USD (Limit 5 per customer)

Take the quick video tour or purchase this kit at:
www.em.avnet.com/spartan3a-evl

Kit Includes

- » Xilinx Spartan-3A evaluation board
- » ISE® WebPACK™ 10.1 DVD
- » USB cable
- » Windows® programming application
- » Cypress MiniProg Programming Unit
- » Downloadable documentation and reference designs



Avnet Green Initiative



Accelerating Your Success™

1.800.332.8638

www.em.avnet.com

IS YOUR CURRENT FPGA DESIGN SOLUTION HOLDING YOU BACK?



FPGA Design | Ever feel tied down because your tools didn't support the FPGAs you needed? Ever spend your weekend learning yet another design tool? Maybe it's time you switch to a truly vendor independent FPGA design flow. One that enables you to create the best designs in any FPGA. Mentor's full-featured solution combines design creation, verification, and synthesis into a vendor-neutral, front-to-back FPGA design environment. Only Mentor can offer a comprehensive flow that improves productivity, reduces cost and allows for complete flexibility, enabling you to always choose the right technology for your design. To learn more go to mentor.com/techpapers or call us at 800.547.3000.

DESIGN FOR MANUFACTURING + INTEGRATED SYSTEM DESIGN
ELECTRONIC SYSTEM LEVEL DESIGN + FUNCTIONAL VERIFICATION

**Mentor
Graphics**[®]
THE EDA TECHNOLOGY LEADER

Xcell journal

PUBLISHER	Mike Santarini mike.santarini@xilinx.com 408-626-5981
EDITOR	Jacqueline Damian
ART DIRECTOR	Scott Blair
DESIGN/PRODUCTION	Teie, Gelwicks & Associates 1-800-493-5551
ADVERTISING SALES	Dan Teie 1-800-493-5551 xcelladsales@aol.com
INTERNATIONAL	Melissa Zhang, Asia Pacific melissa.zhang@xilinx.com Christelle Moraga, Europe/ Middle East/Africa christelle.moraga@xilinx.com Yumi Homura, Japan yumi.homura@xilinx.com
SUBSCRIPTIONS	All Inquiries www.xcellpublications.com
REPRINT ORDERS	1-800-493-5551



www.xilinx.com/xcell/

Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124-3400
Phone: 408-559-7778
FAX: 408-879-4780
www.xilinx.com/xcell/

© 2009 Xilinx, Inc. All rights reserved. XILINX, the Xilinx Logo, and other designated brands included herein are trademarks of Xilinx, Inc. All other trademarks are the property of their respective owners.

The articles, information, and other materials included in this issue are provided solely for the convenience of our readers. Xilinx makes no warranties, express, implied, statutory, or otherwise, and accepts no liability with respect to any such articles, information, or other materials or their use, and any use thereof is solely at the risk of the user. Any person or entity using such information in any way releases and waives any claim it might have against Xilinx for any loss, damage, or expense caused thereby.

Will FPGAs Lead Semiconductors out of the Economic Doldrums?

Programmable logic's 'forgiveness factor' should not be overlooked.

Over the last couple of months, I've read opinion pieces in the electronics trade press predicting FPGAs could very well lead the semiconductor industry's recovery from the economic downturn. I'm no financial guru, but I have to think this could well be the case.

In my 14 years covering the electronic-design space (EDA, ASIC, FPGA, PCB and memory industries), I've witnessed firsthand the maturation of the FPGA industry in its seemingly relentless quest to grab market share from the ASIC business. As FPGAs leveraged new process technologies to increase logic-cell counts in line with Moore's Law, they grew faster and lower in power with each new node and became ever more affordable. At the same time, FPGA vendors have added a slew of advanced functionality, hardwiring MPU cores into their devices and offering the latest soft IP.

Tools for FPGAs have kept pace with this complexity, allowing engineers to invent novel designs and implement them in FPGAs. What's more impressive is that most FPGAs are forgiving—if your design doesn't come out quite the way you want it, simply make a few tweaks and reprogram the FPGA until you have the functionality you need. If you need to accommodate a particular standard that is not quite fully defined, you can quickly add that new functionality once the standard has firmed up. You can even adjust your design and reprogram your FPGA after it has been deployed in the field. You don't have to be afraid to try something risky—something innovative.

This forgiveness factor, I believe, will become a much more powerful selling point for FPGAs over ASICs as system companies look at their budgets and examine the cost and risks of implementing designs in application-specific ICs. It certainly is no secret that the mask costs for designs in the latest process technologies have grown exponentially, making ASIC starts viable only for extremely high-volume applications. At the same time, the cost of ASIC tools to deal with new process complexity is skyrocketing.

Today, for example, most ASIC design teams must acquire design-for-manufacturing tools, sophisticated power-analysis tools to manage leakage and perhaps very soon (or already, in some cases) statistical analysis tools just to produce viable silicon. Chances are, to become proficient in these new tools will require a training ramp-up and a bit of troubleshooting. Then you have to deploy all these new tools optimally to make sure your multimillion-gate ASIC is right the first time and doesn't require a mask respin or 30. ASIC prospects are especially scary when you consider that during the last economic downturn, the dot-com bust, ASIC design starts plunged, from roughly 9,000 in 1999 to 4,000 in 2003, according to research firm Gartner Dataquest. And they have been steadily declining ever since. Considering that ASICs back in 2000 and 2003 were nowhere near as complex as they are today and mask costs were far less, one has to wonder how much lower ASIC starts may go.

Some may argue that application-specific standard products could step in to replace ASICs. You could indeed buy an ASSP that generally handles your specific task, and then modify the software running on it to do the job you want it to. But your competitors can buy the same part and implement similar, or even better, software in their product. In short, with an ASSP you no longer have the advantage of differentiating your hardware, just the software. Gartner Dataquest predicts a steady, albeit gentle, erosion in ASSP starts, from roughly 4,200 in 2008 to 4,000 in 2011. Meanwhile, FPGA starts are expected to grow from 90,000 to around 105,000 over that time frame.

Can you take innovative risks with ASICs and ASSPs? Can you afford to? These are questions that I think will resonate with greater frequency throughout the remainder of this economic downturn and likely afterward.

Does all this mean that FPGAs will, in fact, lead the semiconductor industry in the recovery? No one can say for sure, but the arguments ring true to me.



Mike Santarini
Publisher

LOWEST TOTAL COST... PERIOD.



GET UP TO 50% LOWER COST

- Integrated features and only two power rails minimize need for external components
- Save up to 70% in logic cell resources with dedicated DSP blocks
- Run cool with 11mW static power, 0 μ W in hibernate mode

In the highly competitive high-volume market, cost is king. Our latest Extended Spartan[®]-3A FPGAs deliver the integrated features, low static power, complex computation and embedded processing capabilities you need to achieve the absolute lowest total cost. Period.

Combine these advantages with the industry's largest selection of IP cores, reference designs and I/O standards and you have the most complete low-cost programmable solution available for your next high-volume design.

Visit us at www.xilinx.com to download our free ISE[®] WebPACK[™] design tools and start saving money today.



VIEWPOINTS



66

Letter from the Publisher

Will FPGAs Lead Semiconductors
out of the Economic Doldrums?...4

Xpectations Xilinx's Support Network:
Our Success Is Your Success...66

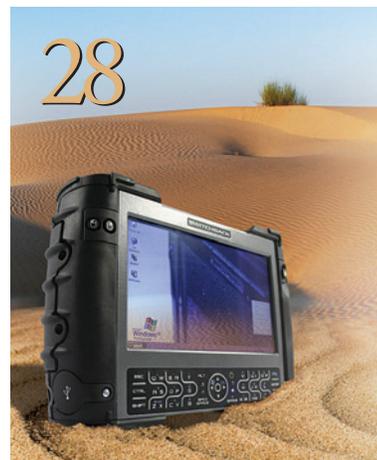
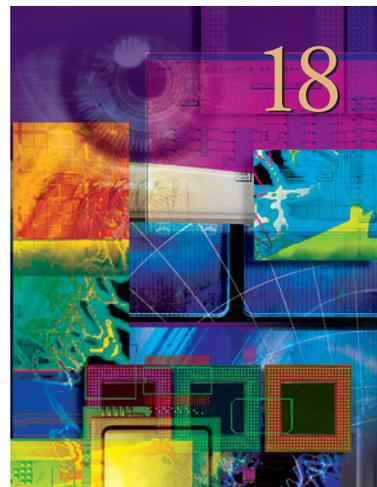
XCELLENCE BY DESIGN
APPLICATION FEATURES

**Xcellence in Wireless
Communications** Baseband
Development for 3GPP-LTE
Just Got Easier...14

Xcellence in Automotive & ISM
Updated Starter Kit Speeds
Video Development...18

Xcellence in Aerospace & Defense
A Flexible Platform for Satellite-Based
High-Performance Computing...22

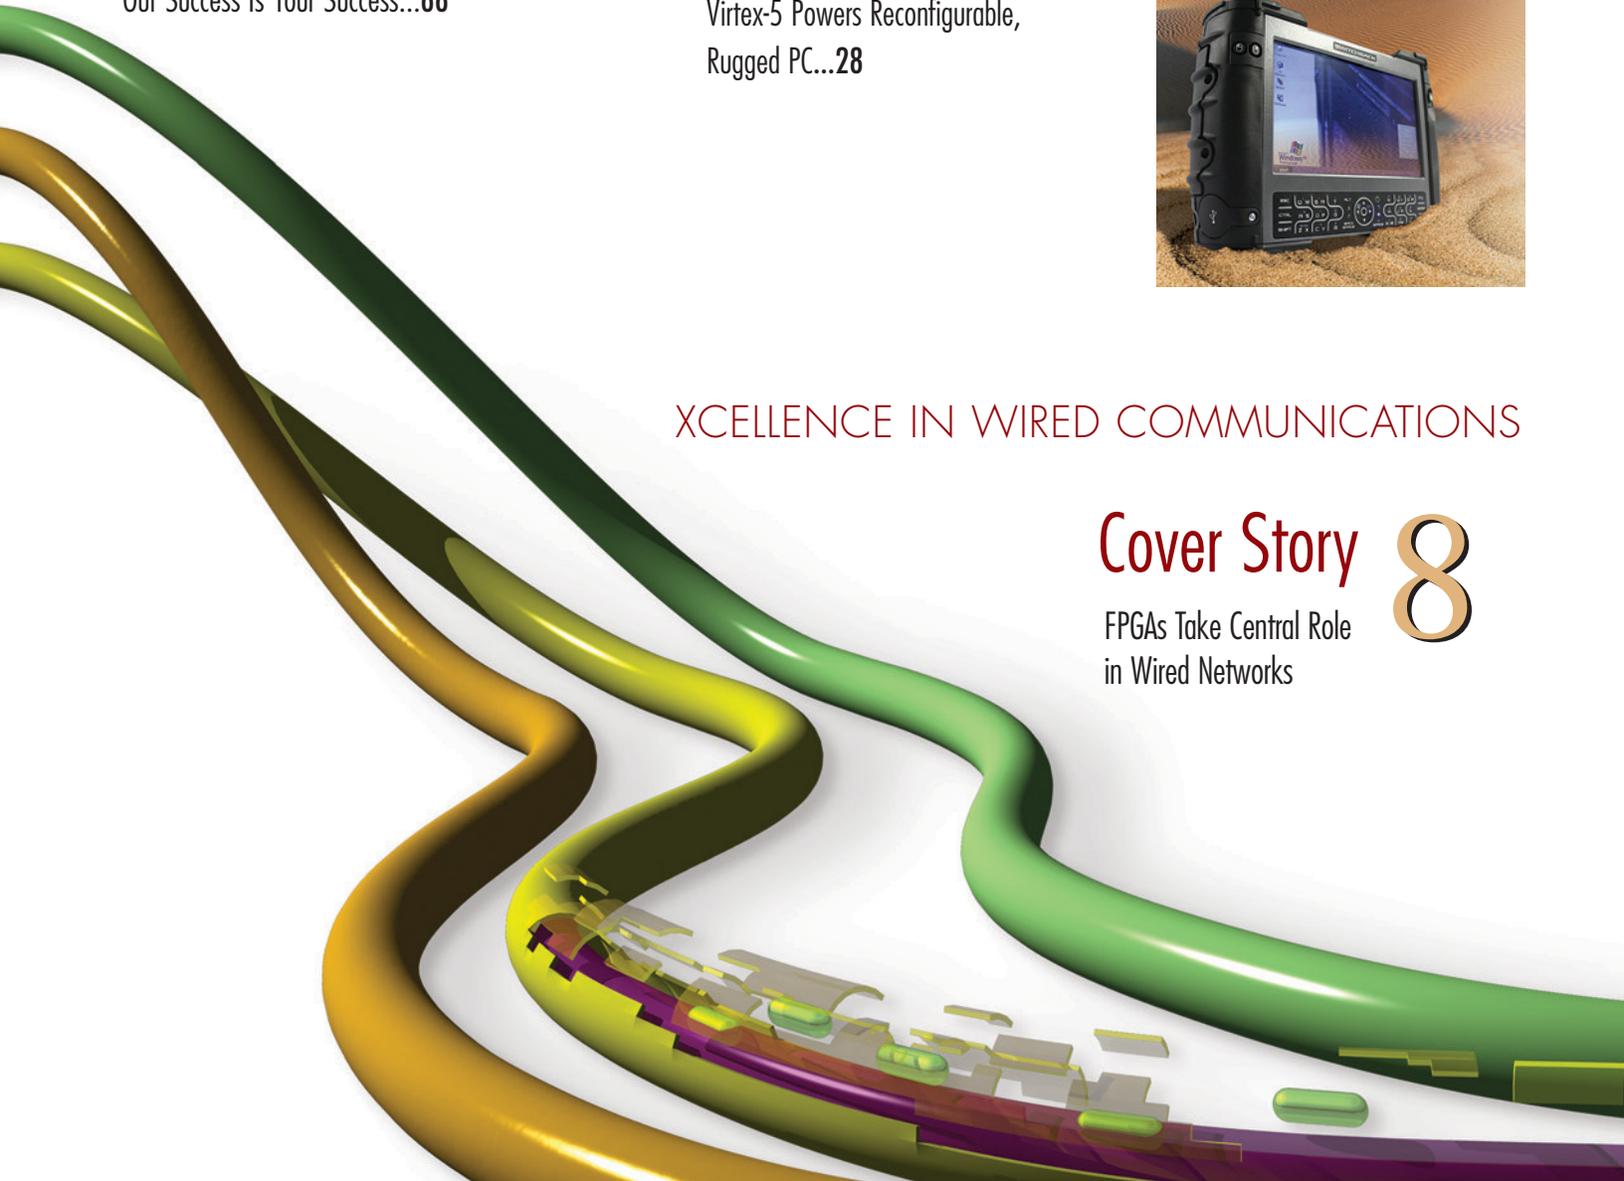
Virtex-5 Powers Reconfigurable,
Rugged PC...28



XCELLENCE IN WIRED COMMUNICATIONS

Cover Story 8

FPGAs Take Central Role
in Wired Networks



THE XILINX XPERIENCE FEATURES

Xperts Corner Floating Point: Have it Your Way with FPGA Embedded Processors...**32**

Xplanation: FPGA 101 Optimizing Xilinx FPGAs for Power...**36**

Xperiment Computer Interface Makes 19th-Century Pipe Organ Rock...**44**

Ask FAE-X Hidden in Plain View...**50**

Profiles of Xcellence Mark Moshayedi Drives STEC Enterprise Storage to Greener Pastures...**58**

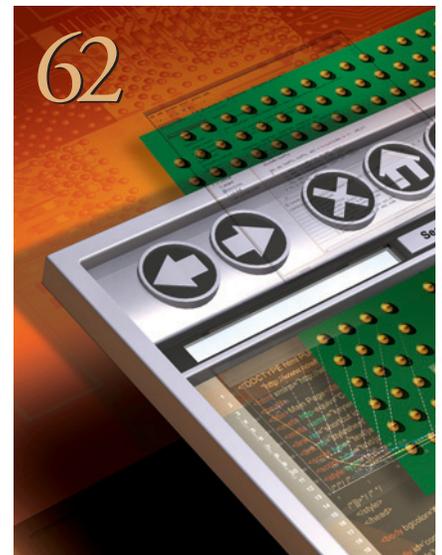


XTRA READING

Are You Xperienced?
Xilinx, Partners Offer Free Seminar on High-Performance System Design...**55**

Xamples A mix of new and popular application notes...**56**

Tools of Xcellence A bit of news about our partners and their latest offerings...**62**



FPGAs Take Central Role in Wired Communications

A demand for speed and the advent of multimedia fuel a need for advanced programmable devices in next-generation networks.

by Mike Santarini
Publisher, Xcell Journal
Xilinx, Inc.
mike.santarini@xilinx.com

The wired communications business has an insatiable need for speed. Fifteen years ago, data transport rates (aka bandwidth) were typically in the hundreds of thousands of bits per second (bps). Today's networks can hurl data across the globe at 10 Gbps, and at some points in the network transmission reaches terabit speeds. FPGAs have played a part in this evolution, and as FPGA technologies advance with Moore's Law they will likely take a more central role in next-generation wired networks.

In a bid to attract customers willing to pay more for new, high-bandwidth networks delivering multimedia content, telecommunication companies such as AT&T and Verizon are pressuring network equipment manufacturers to build faster systems that will speed the delivery of several types of data, not just voice. Steve Rago, principal analyst of broadband and Internet Protocol TV at market research firm iSuppli Corp., notes that telephone companies are urgently trying to transition their businesses from voice-only networks (see sidebar, page 12). Similarly, big corporations are also demanding faster network equipment that will allow employees to communicate more effectively worldwide. In the financial sector, for example, a speedy network can allow traders in far-flung locations to place trades quickly—here, faster communications translates literally into increased revenue.

Network equipment manufacturers such as Cisco Systems, Alcatel-Lucent, Nokia-Siemens Networks and Juniper Networks are among the many companies vying to be first to market with equipment that can offer carriers and enterprises 40-Gbps and 100-Gbps data transport speeds. To do so, they must first create a new generation of routers and switches powered by the latest generation of bleeding-edge ICs. What's more, they must accomplish this feat while standards for next-generation networks—namely, 40G and 100G—are still evolving.

Wired Network Basics

Today's wired communication networks are like a series of roads, highways and super-highways linking one destination to another. Each type of road has a speed limit, and the pokiest byways slow the overall traffic, increasing the time it takes for information to reach its destination.

When a typical user accesses an Internet site from a home PC and downloads a file, the request for data leaves the computer in a data packet at a maximum 1 Gbps—the copper wire connecting your PC to the carrier's access network limits the speed. The access network reads the data packet for, among other things, destination and size, and then forwards it to what's called a metro network—a faster series of electrical routers and switches that reads the packet and forwards the data to the next router along the line. The data ricochets from router to router in the metro network at 10 Gbps.

For long-distance routes, the metro network may ultimately connect to a data superhighway called the core, an optical network that shoots the information at the speed of light to a series of metro networks near the data server that contains the Internet file—be it Web page, video clip or music—you are trying to access. The data server then sends the requested data files back, sometimes the same way, through the network (Figure 1).

At each intersection or hub, a router must read the data packet for such information as destination and size, and determine the fastest route given network traffic conditions, before forwarding that data to its next stop. When negotiating longer routes onto the optical network, a router on the front of the optical network must translate that data from the digital signal suited for electronic routers to a pattern of light in the optical domain. Finally, at the end of the core network, another router must do the opposite, retranslating that data from light back to the format for an electrical packet. Then it must read the network traffic conditions for the fastest route and ship the data to the next electrical router or data server.

Access or download typically occurs in a matter of minutes to a matter of seconds,

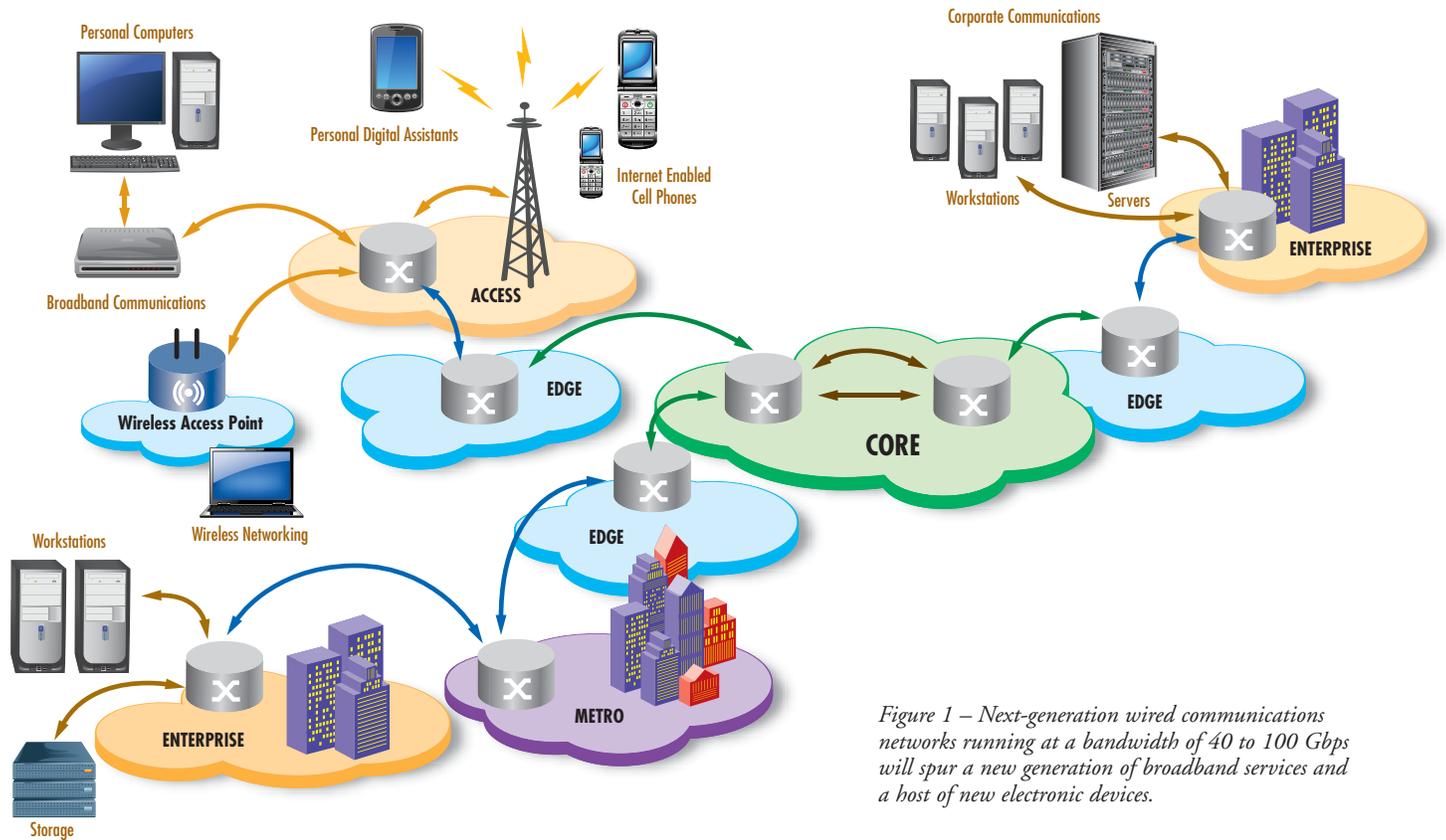


Figure 1 – Next-generation wired communications networks running at a bandwidth of 40 to 100 Gbps will spur a new generation of broadband services and a host of new electronic devices.

depending on the size and location of the file. But tomorrow's networks will hum along even faster, thanks in part to ever-advancing FPGA technology.

Telecom and Datacom Convergence

Today there are two types of wired networks—one for computing and the second for telecommunications. Traditionally, these networks have been separate, each with its own set of unique protocols, routing equipment, bandwidth requirements and rate of bandwidth growth. For example, the telecom industry has typically increased bandwidth in increments of roughly four (2.5 Gbps to 10 Gbps, now moving to 40 Gbps), while computer networking has done the job in leaps of 10x (100M, 10G, 100G). However, Xilinx distinguished engineer Gordon Brebner notes that during the last wired-network retooling a few years ago, a convergence of sorts took place at 10 Gbps, where physical signaling for Ethernet converged with signaling for telecom as both network types independently increased their top bandwidth rates. While

those networks still remain independent, the network industry has been doing its best over the past few years to merge them, specifically around Ethernet.

“Ethernet used to be the technology that simply connected you to your IT department,” said Brebner. “Now there is Ethernet everywhere, and it has developed into carrier Ethernet, in which the telecom industry is using this Ethernet technology internally in their network.” The 10-Gbps Ethernet (10GE) technology “has been standardized for a few years now,” Brebner said, “and currently, 40GE and 100GE are being drafted together as IEEE 802.03ba. The final standard is expected to be completed in late 2009.”

Brebner explained that 40GE would have been the next step for telecom, but the industry expects that rate will initially best suit enterprise networking. For longer transmissions, carriers will use the 100GE standard. It would not be surprising, however, if competition spurs some equipment companies to drive enterprise networking, too, to 100GE.

Inside the Router

To transport data at these speeds, network equipment makers will need to create very sophisticated equipment—routers, switches and transport systems—that employ extremely advanced circuitry.

For example, at the heart of a metro router is a series of line cards. Each line card receives data packets in a wide range of protocols, examines the packets for origin, size, destination and information regarding the rest of the network, and then forwards the packet to a switch. The switch, in its turn, shuttles the packet to its next destination along the network. The line card must accomplish all these computations in nanoseconds.

Traditionally, a line card consists of a CPU, a series of dedicated network processor units (NPUs) and a number of the highest-speed FPGAs available. As a packet enters a line card, an FPGA translates the raw data into formats that a given router can read. The processor coordinates the NPUs to read and route data, while the FPGAs facilitate some of the communication between the CPU and the NPUs.

Some equipment makers expect FPGAs to start playing a more-central role in the router, integrating the functionality of an NPU into the programmable logic fabric.

To handle packets properly, routers must understand multiple protocols. Indeed, said Brebner, they must support a variety of legacy and new protocols that are layered together in a single packet. Of course, if the entire world converged on one protocol or set of protocols, the network might be able to really speed up. But much of the differentiation that makes one network superior to its competitors lies in the protocols the routers use, Brebner noted. Carriers are not about to relinquish this competitive edge.

Next-generation wired networks will be transferring voice, Internet data and video simultaneously. This so-called triple play requires the development of new protocols and, inevitably, a series of refinements and modifications as carriers race to transfer this data more efficiently and safely.

That's why the ability to modify hardware and change functionality is becoming so important—it allows telecommunications equipment to take advantage of the new protocols and in so doing, delivers a huge advantage to OEMs. Many companies are shunning ASICs and ASSPs in their communications systems because those ICs offer the ability to modify only their own software. FPGAs, by contrast, let you modify the hardware, test software functionality in the software domain and then speed it up by creating a hardware implementation of algorithms in an FPGA.

Still other equipment makers expect FPGAs to start playing a more-central role in the router, integrating the functionality of an NPU into the programmable logic fabric. FPGA vendors tend to be the first silicon developers to use new IC processes. This trait has given them the full doubling-of-capacity benefits of Moore's Law, with the payoff of more real estate on each die for additional functionality. With each new generation of FPGA, the likelihood increases that customers can add functions

usually reserved for NPUs. Integrating the translation and interface functionality on a single chip ultimately speeds processing, reduces the overall bill of materials and lowers the power consumption of the router—and with it, the operating expenditures for the overall network. Moreover, because FPGAs offer hardware as well as software reconfigurability and can be modified in the field, network equipment vendors have an opportunity to upgrade their equipment while it's in use—indeed, even while it's still running.

With FPGAs evolving rapidly in ways that particularly suit wired communica-

tions applications, network equipment designers are making even greater use of these versatile devices in next-generation routers. Each new generation of FPGA technology includes a greater number of high-speed transceivers to match the increasing overall bandwidth of the network, even as the overall speed of each transceiver continues to climb. For example, the recently released Virtex®-5 TXT devices (Table 1) contain up to 48 RocketIO™ multirate transceivers running at 6.5 Gbps. They allow the device to deliver the 312 Gbps total bandwidth required for building network bridges.

Virtex-5 TXT FPGA Platform			
Part Number	XC5VTX150T	XC5VTX240T	
Slices	23,200	37,440	
Logic Cells	148,480	239,616	
CLB Flip-Flops	92,800	149,760	
Maximum Distributed RAM (kbits)	1,500	2,400	
Block RAM/FIFO w/ECC (36 kbits each)	228	324	
Total Block RAM (kbits)	8,208	11,664	
Digital Clock Manager (DCM)	12	12	
Phase-Locked Loop	6	6	
Maximum Single-Ended Pins (4)	680	680	
DSP48E Slices	80	96	
PCI Express Endpoint Blocks	1	1	
10/100/1000 Ethernet MAC Blocks	4	4	
RocketIO™ GTX High-Speed Transceivers	40	48	
Package (7,8)		Area	
FFA Packages (FF): flip-chip fine-pitch BGA (1.0 mm ball spacing)			
FF1156	35 x 35 mm	360 (40)	
FF1759	42.5 x 42.5 mm	680 (40)	680 (48)

Table 1 – Xilinx's serdes-heavy Virtex-5 TXT provides developers of next-generation wired communications equipment with a programmable platform for innovation.

In addition to the high-speed transceivers, the number of logic cells roughly doubles with each new generation of FPGA, in keeping with Moore's Law. These additional logic cells allow equipment manufacturers to place greater functionality within each FPGA, perhaps functionality that was previously assigned to NPUs.

Developing NPUs for each generation of equipment, or deciding which NPU is

right for the job, is one of the most trying issues equipment manufacturers confront, Brebner said. The choice is complicated by the fact that with each generation of router, a new batch of startups arises to build the NPUs to power them. "NPUs are about the most fractious area in the market," said Brebner. "Each NPU is designed in a particular way for various niches and functions." The vendor landscape is volatile:

NPU design companies have come and gone, leading some big vendors, such as Cisco Systems, to develop their own.

However, as FPGA technology advances in each generation, there is a greater opportunity for customers to integrate NPU intellectual property (IP) into FPGAs themselves. OEMs can also leverage the devices' reconfigurability so that, as data packets come in with different protocols, the FPGA can

Battle for the Broadband Bundle

Facing declining revenue as cable competitors poach their voice customers, telephone companies are rapidly bulking up their networks to offer multimedia services, potentially driving new growth in next-generation broadband equipment.

"For the last several years, traditional telephone companies have been losing their subscriber base at an alarming rate," said iSuppli analyst Steve Rago. "Roughly 4 to 10 percent of their subscribers are disappearing every year."

The erosion is occurring for several reasons. "Number one is that many folks are using mobile phones as their only phone line," said Rago. "The second reason is that the need for a second line for the Internet and in some cases even a fax is disappearing—with broadband you don't need a second line for the Internet." In addition, he said, cable multiple-service operators (MSOs) have been successfully snatching away traditional voice services, bundling voice with cable TV and the Internet. In last year's fourth quarter, said Rago, "cable added 1 million voice subscribers in the United States alone." The same transition is going on worldwide, he noted, even in mainland China, where the voice networks are only a dozen or so years old.

Telcos, for their part, are enjoying increased revenue from the broadband services they currently offer, Rago said. However, it hasn't offset the loss of revenue from voice. "The net result is that they are either holding steady [in terms of] revenue growth or growth is declining, which is not a very good position for Wall Street," he said. "If the telcos don't change the way they do business, they are going to become extinct—there won't be a need for them anymore."

Meanwhile, their competitors in the wired space, the MSOs, are not experiencing huge growth either, said Rago. "Actually, they are seeing revenues hold steady or even decline, with more competition from satellite companies and the telcos," he said.

To get back on a growth path, telephone companies worldwide have collectively decided to deliver video as a value-added service to voice, along with other offerings, Rago said. They are banking on one of these Internet-based services in particular: time-shifting TV, which will allow you to watch whatever program you want to see whenever you want to see it. "It will be a paradigm shift from the way you watch TV today," said Rago. "It's one of the advantages telecom companies have with IPTV over the MSOs."

Further, Rago said that instead of charging consumers for bits per second, the telcos have decided to do what MSOs do today: have users pay for the services they want. "You pay for your video service [IPTV, for example], you pay for your voice and your other services—you'll pay more depending on what services you add to your plan." Rago said that most telephone companies are already offering these new services or plan to soon do so.

Multibillion-Dollar Question

Meanwhile, MSOs will continue to attempt to lure traditional voice subscribers to their multimedia mix and will likely come up with their own value-added services.

But the multibillion-dollar question for all these companies, telephone and MSO alike, is not simply how to establish growth, but how to establish sustainable growth.

One challenge is "getting a fat enough pipe to the home to offer all of these new services," Rago said. DSL worldwide, and ADSL in particular, is still the biggest, he said, "but we've seen growth in broadband DSL and fiber to the home; or fiber near to the home and BDSL near to the curb. Fiber to the home is second now in terms of new services. It even surpasses cable modems."

Upgrading the access equipment to handle these new services will, of course, be key to making all this possible. "We're looking at equipment that can handle speeds anywhere from 30 Mbps to 100 Mbps," he said. New services such as time-shifting TV and video-on-demand will put tremendous pressure on bandwidth in the network. "It will drive a major need for innovations and enhancements in long-haul and the metro networking space," said Rago. It will also give rise to a new crop of high-speed-data consumer devices that will in turn shape the feature requirements for next-generation services.

Indeed, Rago said that "the battle for the broadband bundle" should result in novel technologies and services that ultimately drive innovations in related fields. But who will win that battle is anyone's guess at this point.

To read more about the competitive landscape, contact iSuppli for its latest report on consumer communications.

— Mike Santarini

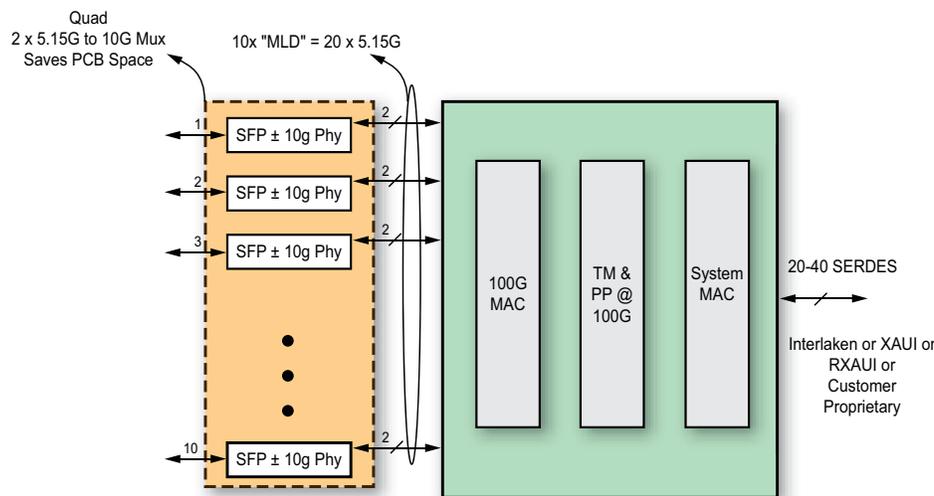


Figure 2 – Sarance Technologies' 100GE MAC solution implemented with Virtex-5 FPGAs

instantiate or implement an NPU architecture on the fly that is best suited to read the data, even run a security check in it, negotiate the fastest route to its destination and then forward the data there.

"FPGAs have traditionally performed embedded RISC and control plane functions," said Loring Wirbel, longtime communications watcher and moderator of *EDN Magazine's* new FPGA Gurus Web site. "Today's FPGAs can now handle a lot of datapath functions, so now seemingly a single FPGA—depending on how it's partitioned—can serve as an aggregation box in an enterprise or one of the blades in a big switching center. So you don't need coprocessing if you have partitioned things correctly. One of the stories of the death of the network processor is that slowly but surely, FPGAs started taking over the NPU's function."

Wirbel notes that traditionally, as each new generation of equipment has rolled out, there is a very brief opportunity for NPU vendors to field specialized packet-forwarding engines. But in the wink of an eye the opportunity passes, he said, and the engines get replaced with FPGAs.

"As we start moving to 40G and 100G networks, there will be a temporary place for very fast engines that just do the packet forwarding, just as there was for 1G and 10G [technology]," Wirbel said. "But the thing is, as you move to the new feature sizes, that is only going to be a narrow window and they may very likely just skip looking over

the ASSPs and go directly with an FPGA. Every generation has had that brief window where they use ASSPs, but with each generation that window is becoming more narrow—eventually it will stay closed."

Advancing FPGA Technologies for Wired Comms

Today, Xilinx's largest Virtex-5 TXT XC5VTX240T device contains 37,440 logic slices with a total of 239,616 logic cells. This architecture has afforded design teams and IP vendors great opportunities to innovate solutions that support XAUI, RXAUI, Interlaken, Sonet, ODN and many other wired standards with the most advanced FPGA silicon to date.

For example, Xilinx worked closely with Sarance Technologies to provide the industry's first 100GE media-access controller, a full-featured, IEEE 802.3ba-compliant solution implemented with Virtex-5 FPGAs (Figure 2).

Sarance announced in mid-2008 that its 100GE MAC solution was up and running on tier-one vendor hardware prototypes using two Virtex-5 FXT FPGAs, 10 external 10-Gbps physical-layer devices and a variety of system-side interfaces.

The 100GE MAC-to-Interlaken bridge solution that the new Virtex-5 TXT FPGA platform supports is a low-risk way to condense functionality into a single FPGA and three external quad serdes muxes. In this implementation, Xilinx's 64/66 and 64/67 encode/decode gearboxes are built

into the GTX transceiver, saving nearly one-fifth of the logic count and power consumption of the design.

In June 2008, telecommunications giant Comcast Corp. announced it had successfully completed a 100GE technology test over its existing backbone infrastructure between Philadelphia and McLean, Va., using the industry's first 100GE router interface. The system used the same Sarance Technologies' High Speed Ethernet IP Core (HSEC) running on a Virtex-5 FXT FPGA that is supported by the Virtex-5 TXT platform today.

The demonstrations follow up on early achievements in the 100GE domain. In November 2006, Xilinx FPGAs were the vehicle used to showcase the world's first successful 100GE transmission through a live production network demonstrated at the SC06 International Conference, the confab of high-performance computing, networking, storage and analysis.

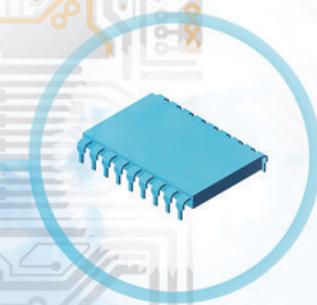
Finisar teamed with Level 3 Communications, Internet2 and the University of California at Santa Cruz to demonstrate the transmission of 100GE traffic over Level 3's DWDM network from the show site in Tampa, Fla., to Houston and back—a total of 4,000 miles.

The Xilinx FPGA electrically transmitted all ten signals to ten 10-Gbps XFP optical transceivers, which converted the signals into the optical domain. From there, the signals traveled to Infinera's commercially available DTN Switched WDM System, which handed them off to the Level 3 network.

Overall, FPGA technologies are advancing fast. With each turn of Moore's Law, FPGAs offer communications designers the ability to create higher-bandwidth, next-generation networks. In the not-so-distant future, network designers will give FPGAs a more-central place in their designs. How big that role turns out to be will depend not only on the silicon, but on the IP and hardware and software tools customers have at their disposal. Xilinx remains committed to creating innovations for the wired communications market, with an aim not just to maintain its leadership but to build on it with new programmable solutions. ●●●

Baseband Development for 3GPP-LTE Just Got Easier

With the release of LTE-Channel Encoder and Decoder, Xilinx helps customers speed Layer-1 subsystem development and address the performance and latency challenges of 4G wireless.



by David Nicklin
Senior Manager, Wireless Product Marketing
Xilinx, Inc.

The baseband processing signal chain presents both the greatest challenge and the best opportunity for innovation in the base transceiver station. No wonder, then, that it has become a key area for product differentiation among OEMs. Competition in baseband architecture designs has intensified with the realization that many of the techniques used for earlier 2G and 3G systems simply will not scale to meet the performance and latency requirements of the 3GPP Long Term Evolution (LTE) technology, wireless' fourth generation

Not only does the processing chain have to contend with far more processing than previously, but also, all the functions must be completed in much less time. Rounding out the set of challenges facing system architects is the need to develop a system that can meet the operators' aggressive capital- and operational-expenditure reduction targets. These major pressures on the baseband processing system design are illustrated in Figure 1.

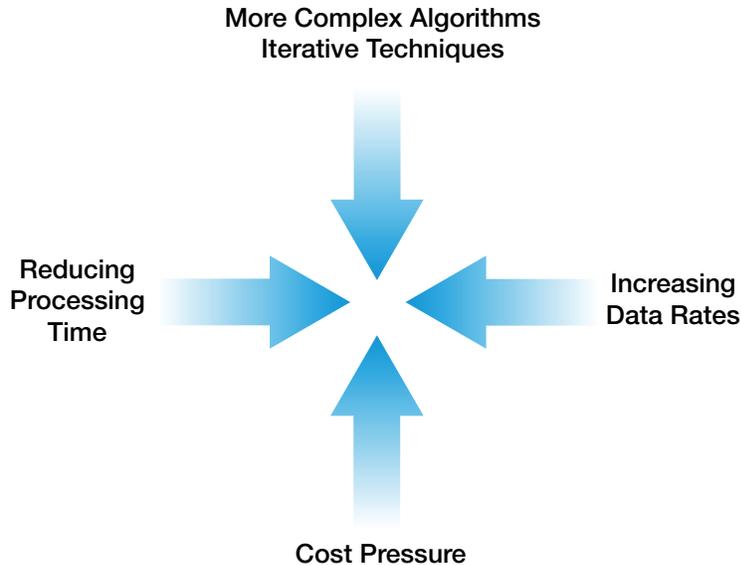


Figure 1 – Challenges in evolving baseband processing needs

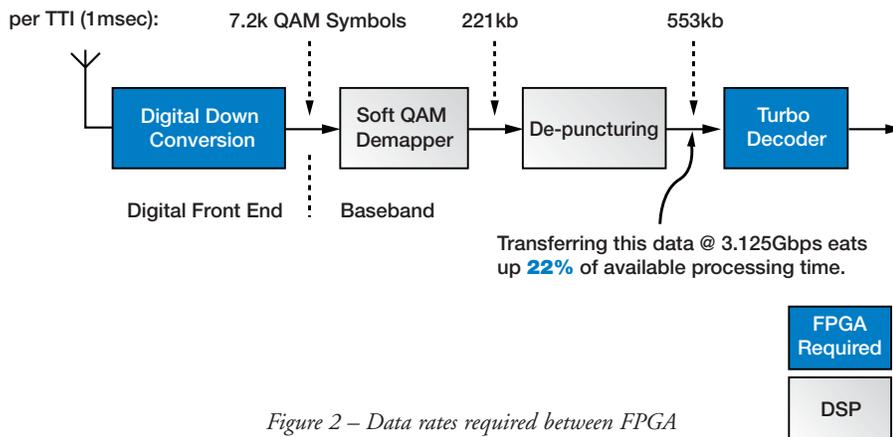


Figure 2 – Data rates required between FPGA and DSP in a typical LTE system

An FPGA-based solution can meet all of these demands, while sidestepping the usual performance issues and bottlenecks. Initiatives such as Xilinx's newly released LTE Uplink Channel Decoder and LTE Downlink Channel Encoder LogiCOREs™ seek to remove the barriers to FPGA adoption by incorporating many of the critical Layer-1 functions in a single IP solution.

Advances in silicon technology have been a key enabler in the success of wireless communications by facilitating the rollout of ever-more-complex algorithmic techniques from the research labs into products. One such example was the

deployment of iterative turbo error-correction techniques in 3G networks, a scheme that migrated from discovery to commercial release inside of 10 years. The pace of innovation has continued to accelerate, most notably with the exploitation of the spatial dimension in wireless communication through various multiple-input, multiple-output (MIMO) antenna techniques.

With the advent of 4G air interfaces however, the pressures have built to the point where traditional programmable DSP-centric channel card architectures are struggling to cope. The traditional partitioning between FPGA and DSP is constrained by performance bottlenecks that

come into play due to the enormous amount of data that has to pass to-and-fro between them.

So, how can we remove such bottlenecks? The key lies in simplifying the Layer-1 system architecture and eliminating all unnecessary chip-to-chip data transfers. This simplification process raises some uncomfortable questions about the scalability of architectures based on digital signal processors. Designers need a strong portfolio of intellectual property (IP), software and support to aid them with the transition to a Layer-1 system architecture in which most functions are implemented in programmable hardware rather than DSPs.

Simplifying Layer-1 Design

Let us look more closely at the problems that can occur when the FPGA is solely employed as a coprocessor to offload turbo decoding functions from the DSP processor. In analyzing the effectiveness of such a partitioning in a typical LTE baseband design (see Figure 2), Xilinx system architects have discovered that more than 20 percent of the available latency-timing budget can go into just shifting the data from a DSP processor to an FPGA and back again via SRIO connections. Shockingly, this is far from a worst-case scenario. Add to the mix data that is coded with higher-order modulation schemes, such as 64-QAM, two MIMO code words at a 1/3 code rate, over a full 20-MHz LTE band, and the numbers rapidly get much, much worse.

One response is simply to add bigger "pipes" to move the data around by deploying more high-speed multigigabit-transceiver connections. While it is certainly possible to construct a system this way, it leads to an unnecessary increase in system power dissipation, as the relatively power-hungry high-speed serial connections transport the data back and forth and bridging functions are replicated, leading to the need for more hardware resources.

There is a better, more-optimal solution. By incorporating most of the Layer-1 functionality within the FPGA, a designer can free up this unnecessary overhead and

pass on the savings in improved system throughput and latency, while at the same time lowering power requirements. The power saving alone translates directly into improved reliability, reduced system cost and opex savings.

Such an architectural approach avoids the need for DSPs altogether—although they can be incorporated, if desired, to perform lower-rate functions. With this kind of a partitioning, the FPGA implements the entire Layer-1 baseband processing, leaving the other higher-layer functions, such as media-access control and HARQ processing, to a more cost-effective general-purpose processor or a network processor, which can also handle additional backhaul connectivity functions. Integrating all the high-performance and time-critical functions on a single platform FPGA effectively circumvents delay and bandwidth limitations, and partitioning becomes a much simpler task.

The key stumbling block to date in adopting such an approach has been the need to simplify the process, from design concept to hardware. Also, designers accustomed to a DSP-centric design flow need IP and development tools that make it easier to unlock the capabilities of the FPGA, and quickly and effectively develop baseband functionality within it.

Xilinx's LTE Uplink Channel Decoder and LTE Downlink Channel Encoder LogiCOREs remove such barriers to adoption by folding many of the critical Layer-1 functions in a single IP solution that can be customized via a GUI in the Xilinx CORE Generator™ tool. This design flow enables engineers with limited FPGA experience to concentrate on the wider system design, saving significant development and integration effort.

The Shape of Things to Come

The drive to ever-faster connectivity and low-latency connections is a key requirement of LTE and will remain so for future systems beyond 4G. As these newer data-centric wireless systems evolve, many companies that have adopted the traditional partitioning between DSP and FPGA will find the overhead involved in shifting data between separate chips has become unacceptable. An FPGA-based solution is now much more accessible to designers looking for that extra edge in their product design. Those who move past their ties to a legacy system design approach will deliver a product that surmounts the performance issues and bottlenecks their competition will continue to experience. ●●

Inside the LogiCOREs

Xilinx's recently released LTE Channel Encoder and Decoder LogiCOREs, designed for 3GPP rel8 E-UTRA eNB baseband processing, comply with specifications 3GPP TS 36.211 v8.2.0 and TS 36.212 v8.2.0 (2008-03). They support different configurations up to 20 MHz in bandwidth with normal (short) CP, 64-QAM modulation and two MIMO code words. Both can also handle FDD and TDD frame structures, suiting them for systems that have evolved from the TD-SCDMA standard.

Both the encoder and the decoder are supplied as standalone, parameterized IP blocks that can be easily incorporated into a customer's design using the Coregen software tool. Supporting them is a comprehensive set of tests, simulations and C models for system simulation in order to aid design integration.

More details on the new LogiCORES are available at the Xilinx IPcenter (www.xilinx.com/ipcenter). Designers can access additional wireless reference designs and solutions via the Wireless End-Market page (www.xilinx.com/esp/wireless).



Tis Better to Transmit than Receive

X5 tx

X5-TX, Virtex 5-based Transmitter Module with Integrated Wireless IP Cores!

Features

- (4) 500 MSPS or (2) 1 GSPS 16-bit DACs
- +/-1V, 50 ohm, DC or AC coupled inputs
- External or internal sample clock & trigger
- Xilinx Virtex5, SX95T or LX155T FPGA
- 512MB DDR2 DRAM
- 4MB QDR-II SRAM
- 8 Rocket IO private links, 2.5 Gbps each
- >1 GB/s, 8-lane PCI Express Host Interface
- Power Management features
- XMC Module (75x150 mm)
- PCI Express (VITA 42.3)

Applications

- Wireless Transmitter
- RADAR pulse generation
- High Speed Arbitrary Waveform Generation
- Electronic Warfare
- IP development

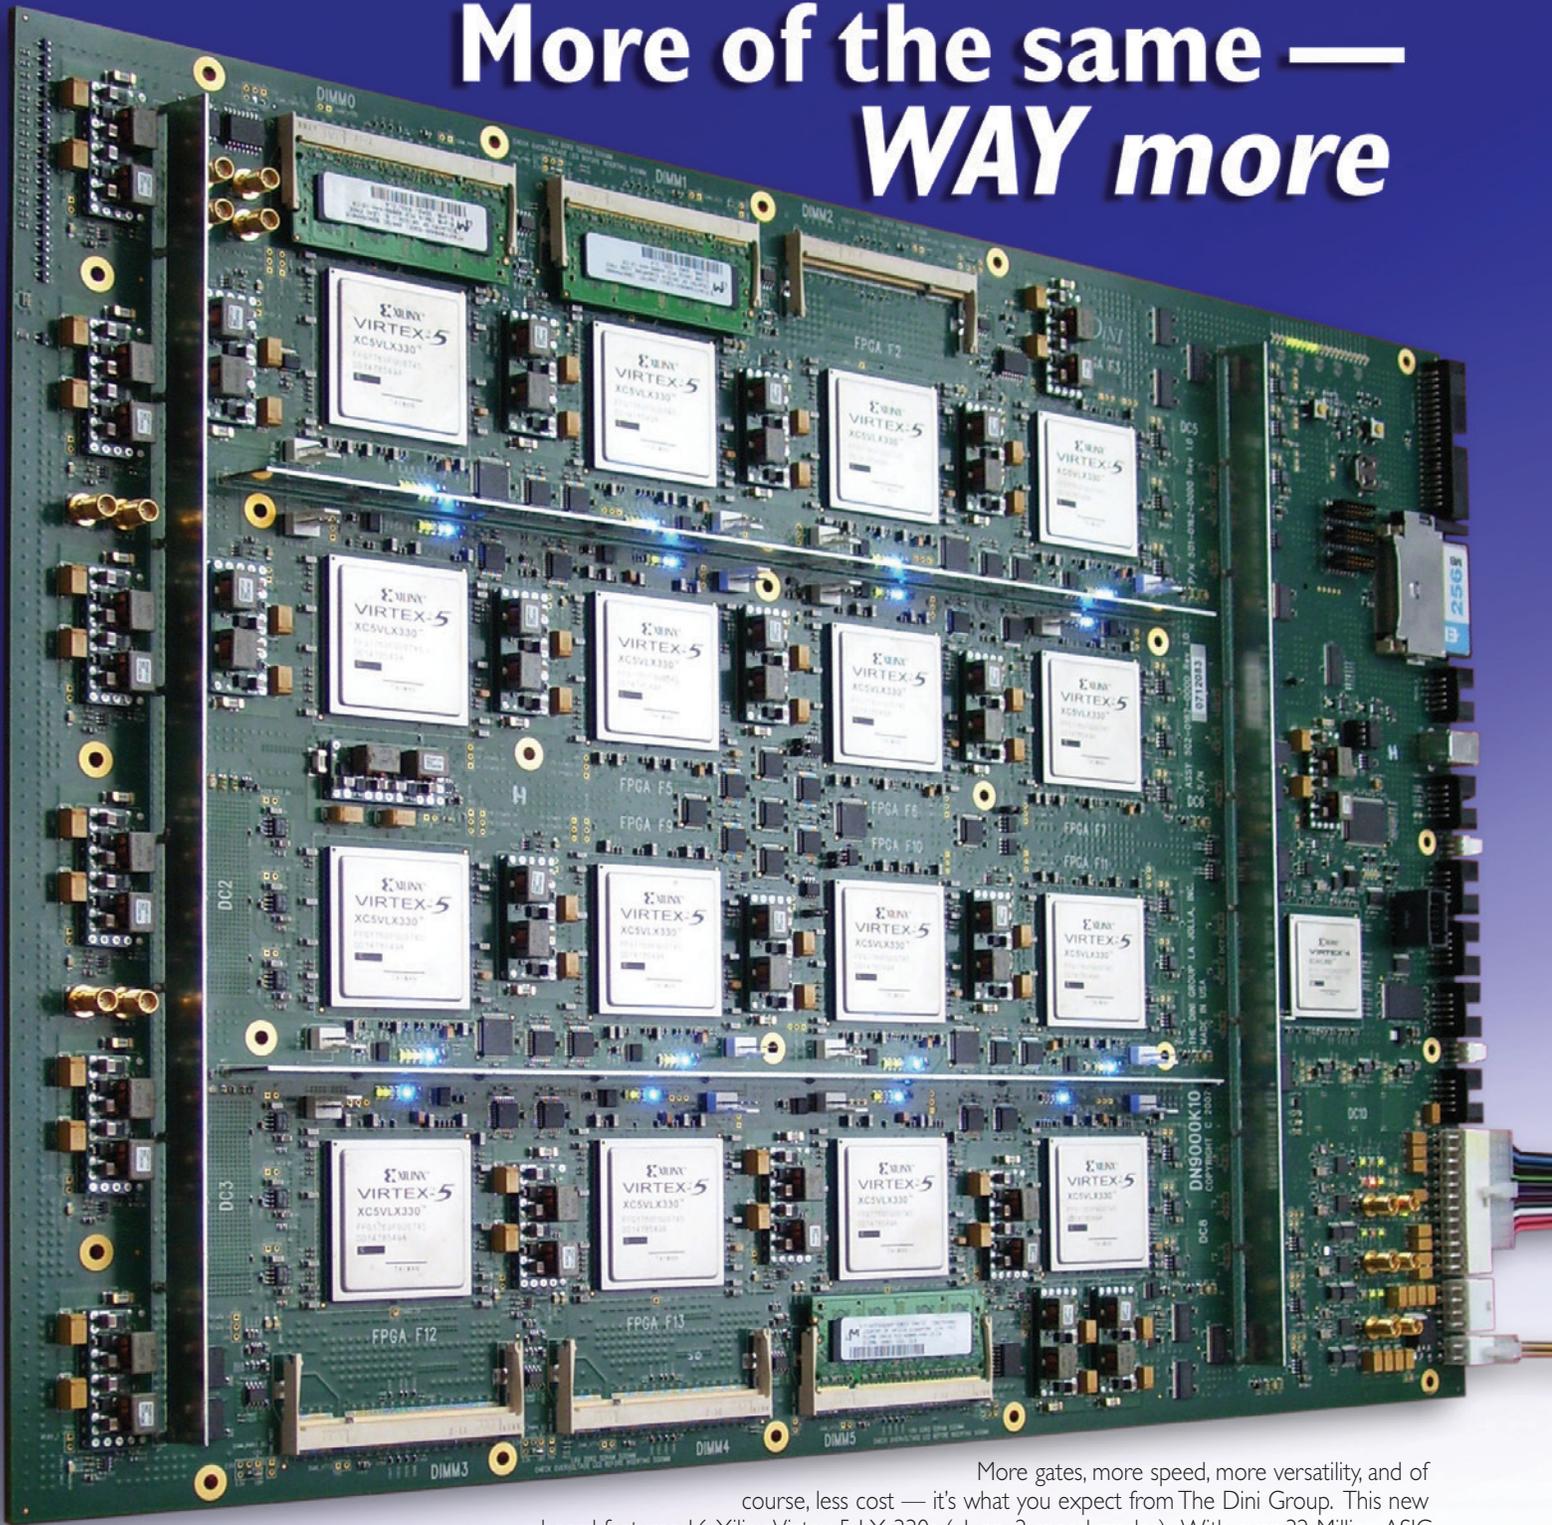
Download Data Sheets NOW!

wireless ip cores
R interface



Innovative Integration
... real time solutions!
805-578-4260 phone
www.innovative-dsp.com

More of the same — *WAY* more



More gates, more speed, more versatility, and of course, less cost — it's what you expect from The Dini Group. This new board features 16 Xilinx Virtex-5 LX 330s (-1 or -2 speed grades). With over 32 Million ASIC gates (not counting memories or multipliers) the DN9000K10 is the biggest, fastest, ASIC prototyping platform in production.

User-friendly features include:

- 9 clock networks, balanced and distributed to all FPGAs
- 6 DDR2 SODIMM modules with options for FLASH, SSRAM, QDR SSRAM, Mictor(s), DDR3, RLDRAM, and other memories
- USB and multiple RS 232 ports for user interface
- 1500 I/O pins for the most demanding expansion requirements

Software for board operation includes reference designs to get you up and running quickly. The board is available "off-the-shelf" with lead times of 2-3 weeks. For more gates and more speed, call The Dini Group and get your product to market faster.

The
DiNI
Group

Updated Starter Kit Speeds Video Development

The XtremeDSP Video Starter Kit, Spartan-3A DSP FPGA Edition 2, provides a high-performance development platform for complex high-definition systems.

by Joe Mallett
Senior Product Line Manager
Xilinx, Inc.
jmallett@xilinx.com

The video industry's shift to more-complex and integrated processing solutions along with demanding, next-generation video compression standards is driving system requirements for video performance that exceeds what standalone DSPs can deliver. No wonder, then, that many companies designing state-of-the-art video equipment are turning to FPGA platforms. In particular, many of them are joining companies in sensitive military, automotive, medical, consumer, industrial and security applications in choosing the Xilinx® Spartan®-3A DSP, which supplies more than 20 GMACs of DSP performance for less than \$30.

In addition to performance, this FPGA provides an integrated solution supporting DSP and embedded processing using MicroBlaze™ processors in the system. Those features enable OS support and drivers for specific market segments.

Xilinx recently updated the XtremeDSP™ Video Starter Kit—Spartan-3A Edition to help design groups get started with these kinds of advanced designs, and to accelerate development times. This kit can help you create advanced video systems for whatever application you are targeting.

Reference Design	Functionality Description
DVI Pass-Through	<ul style="list-style-type: none"> • Capture a video stream input port • Perform real-time image processing on the video stream • Display the processed video
DVI Frame Buffer	<ul style="list-style-type: none"> • Capture a video stream from a camera • Perform processing on the video stream • Buffer the video stream in external memory • Report memory bandwidth utilization data
Camera Frame Buffer	<ul style="list-style-type: none"> • Capture a video stream from a camera • Perform processing on the video stream • Buffer the video stream in external memory • Display the processed video at a different rate • Use a MicroBlaze to configure the video pipeline
Video Processing	<ul style="list-style-type: none"> • Capture video stream from an S-video interface • Support weave and de-weave access to video memory • Perform gamma processing and 2D FIR filtering
Hardware Co-SIM	<ul style="list-style-type: none"> • Initialize a point-to-point connection to the platform • Any design within System Generator can add a HW-CoSIM token • Validation acceleration with hardware in-the-loop

Video Starter Kit Version 2.0

The latest version of the XtremeDSP Video Starter Kit (VSK) for the Spartan-3A provides a comprehensive platform that accelerates the development of video applications on Xilinx FPGAs. Designed to leverage the cost/performance advantages of the Spartan-3A DSP FPGA device family, the kit offers an updated set of video reference designs based on an embedded design framework that allows customers to focus on their unique value-added development.

The VSK provides multiple reference designs that can accelerate the development of video applications running on Xilinx FPGAs. Each reference design is built upon a common framework and uses multiple interfaces for the I/O of video data to the FPGA. Table 1 lists each reference design and the video processing and connectivity capability it illustrates.

Table 1 – Reference designs

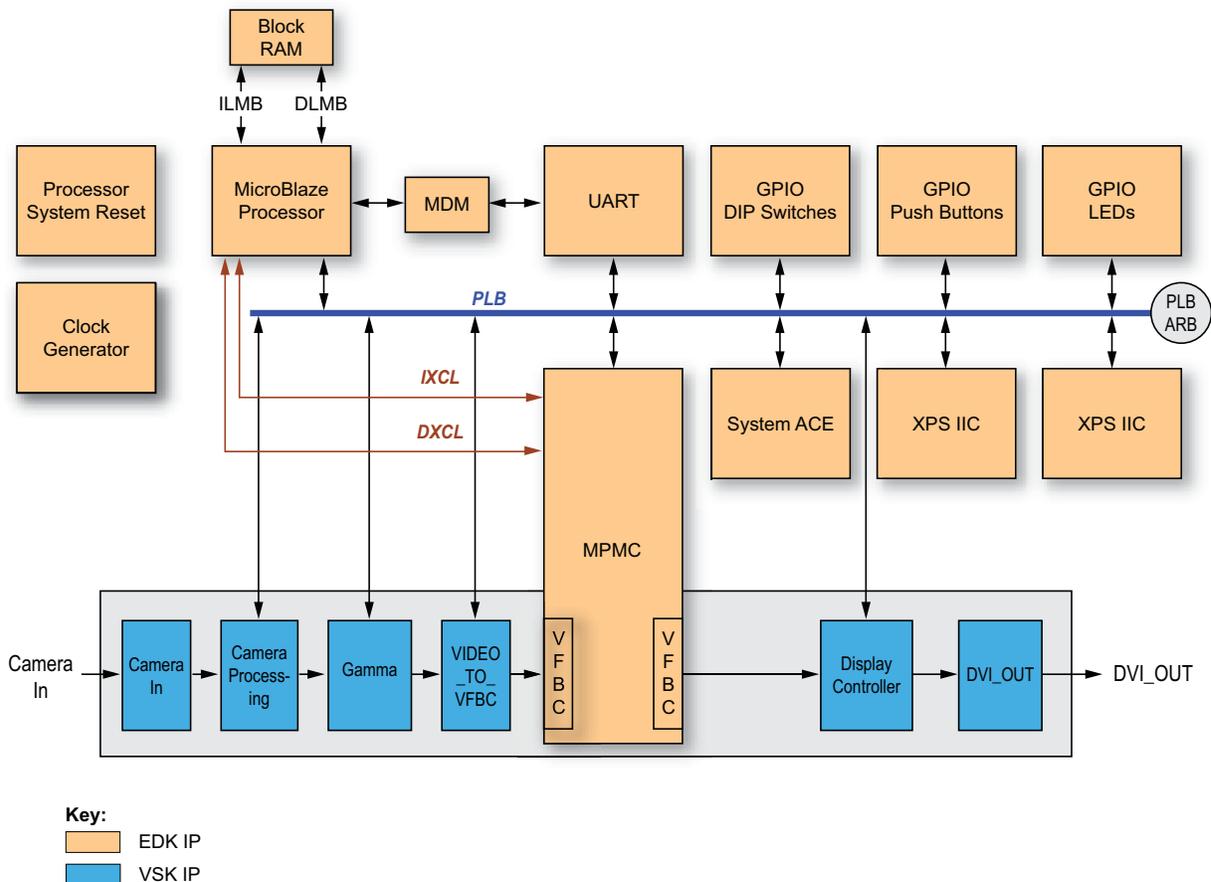


Figure 1 - Camera reference design

Starting with the camera reference design, a software developer can, from day one, begin implementing an operating system and programming the application layer using the EDK software development tools.

The reference designs help accelerate development by implementing specific data flows that are common to video systems. One such example is having a camera provide RAW image data to the FPGA for processing and display, as seen in the camera reference design shown in Figure 1.

The VSK provides all necessary source and project files for the reference designs, which developers use as a starting point. In the camera reference design, the camera processing block is a design developed in System Generator and integrated as a dedicated hardware peripheral with the EDK embedded system. This allows hardware designers to easily remove the example image processing, replace it with new or existing designs and integrate it within the

system without having to design the supporting hardware peripherals.

Embedded Processing

The migration to a complex hardware acceleration processing system is stepping up the need for embedded processing to handle all the real-time control, configuration and system interaction.

The tight integration means designers can convert DSP designs captured in System Generator into custom peripherals for Platform Studio, and connect them to the base system using the PLB bus. This allows a system designer to easily enable the system control and migration of existing system software with the adoption of a MicroBlaze v7 soft-core

processor. Designers gain performance and achieve system integration by exploiting the flexibility of the device to configure a hardware architecture optimized for a particular application. Starting with the camera reference design, for example, a software developer can, from day one, begin implementing an operating system and programming the application layer using the EDK software development tools (see Figure 2).

This flexibility adds a degree of freedom to the development process that also reduces the design complexity. The XtremeDSP Video Starter Kit gives a hardware or software developer a complete and easy-to-use design environment with example applications and full sup-

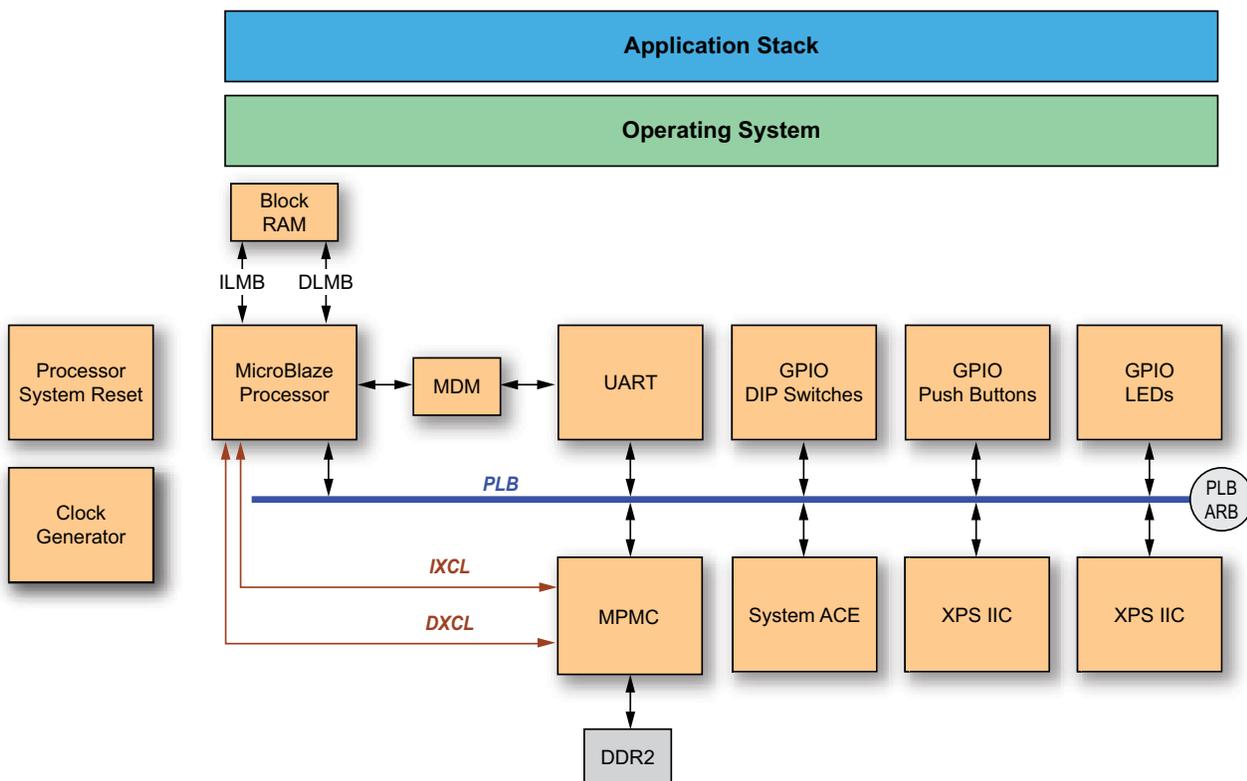


Figure 2 – Application programming on the VSK

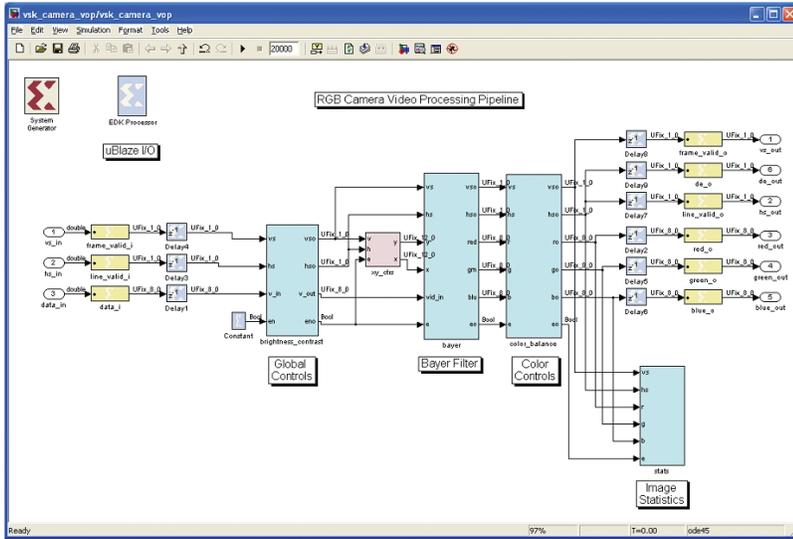


Figure 3 – System Generator diagram for a camera reference design

port for the standard Xilinx tool flows. That combination can help to accelerate the design process and allows for end product differentiation.

System Generator supports hardware-in-the-loop co-simulation using the Spartan-3A DSP 3400A development platform, which can accelerate the performance of Simulink® simulations up to 100x. This acceleration enables video algorithm development and debug using real-time video streams read into Simulink using The Mathworks' Data Acquisition Toolbox™.

Hardware Acceleration

Now that the required processing bandwidth is outpacing the current capabilities of standalone DSP processors, hardware acceleration is becoming a necessity in many video applications. FPGAs enable this hardware acceleration while delivering the additional benefits of system integration and architecture repartitioning.

The migration from a standalone system processor to the integration of a coprocessor requires some design exploration as the hardware designer is looking at the various functions to accelerate. The first challenge is the need to have a variety of design flows that enable abstract model programming using MATLAB™ and Simulink, and the easy integration of any existing VHDL/Verilog designs.

Designers can first implement video algorithm designs as a MATLAB or Simulink model using the optional Video and Image Processing Blockset. As the development moves to the next stage, hardware implementation is easily enabled through System Generator for DSP, which provides a rich set of DSP building blocks, optimized for Xilinx devices, for use within the Simulink modeling environment (see Figure 3).

Once the hardware design is completed in System Generator, the HW-CoSIM functionality accelerates the validation time by placing hardware in the loop. For complex systems, this can greatly improve the testing run-time and thereby increase the number of iterations completed within a given amount of time.

The XtremeDSP Video Starter Kit–Spartan-3A DSP Edition coupled with the integrated reference designs provides the ideal platform for video developers by enabling various modes for processing data from streaming to frame buffer based. Video developers can quickly design and validate hardware peripherals (HW-CoSim) using System Generator. Integration of hardware peripherals and embedded processing accelerates the development of complex video systems within industrial imaging, broadcast, consumer, medical and automotive applications. ●

Supporting Your Future
HUNT ENGINEERING
 www.hunteng.co.uk

USB connected Programmable FPGA systems

V-II Pro PowerPC®

- Virtex®-II Pro XC2VP7
- 256 Mbytes DDR Memory
- Configurable digital I/Os
- PowerPC® boot FLASH
- USB 2 or Standalone

Software Defined Radio

- Virtex®-II FPGA 1M gates
- 2 ch 125MSPS A/D and D/A
- TI C6203 DSP
- 32Mbytes SDRAM
- Configurable Digital I/O
- USB 2 or Standalone

Imaging with Virtex®-4FX

- Virtex®-4 FX12 FPGA
- 128Mbytes DDR Memory
- CameraLink connection
- VHDL Imaging Library
- USB 2 or Standalone

Programmable hardware with cables, device drivers, loading tools, examples and Power Supply. Systems can be used connected to a PC using USB, or can function standalone (without USB) using the initialisation PROMs.

sales@hunteng.co.uk
 +44 (0)1278 760188
www.hunt-rtg.com

A Flexible Platform for Satellite-Based High-Performance Computing

Space-grade Virtex FPGAs and a reconfigurable system architecture satisfy demanding size, weight and power requirements and accelerate design cycles.

by Ian Troxel
Future Systems Architect
SEAKR Engineering, Inc.
Ian.Troxel@seakr.com

Greg Lara
Marketing Manager
Xilinx, Inc.
greg.lara@xilinx.com

Developers of space-based electronic systems face increasing pressure to deliver higher levels of performance while working within more-aggressive project schedules and tighter budgets. Many space-based systems now under development call for advanced video equipment that will capture images with extremely high resolution and then relay those images instantaneously back to earth. To do this, new systems need to include advanced processing circuitry, greater storage capacity and the ability to transfer large data files quickly over long distances. However space-based systems have a unique set of size, weight and power (SWAP) constraints that can prove taxing for designers.

Requirements to do more for less are driving the adoption of commercial, off-the-shelf (COTS) devices, such as FPGAs. The flexibility inherent in reconfigurable FPGAs offers tremendous benefits for developers of space-based systems in terms of SWAP constraints, cost and productivity.

One way to get the maximum leverage from available engineering and budget resources is to create a flexible payload that can be deployed in multiple missions. Our company, SEAKR Engineering, Inc., employed reconfigurable Xilinx® Virtex® FPGAs to create a flexible high-performance computing platform that serves as the heart of a variety of space-based systems. This reconfigurable computing (RCC) methodology has enabled our engineers to achieve demanding performance targets within SWAP, cost and time constraints for a number of missions, most notably our onboard processor for Raytheon's Advanced Responsive Tactically Effective Military Imaging Spectrometer (Artemis), our Programmable Space Transceiver, our Programmable Space IP Modem and the Orion Vision Processing Unit, currently under development.

Application-Independent Processor Architecture

This new platform we developed, called the Application Independent Processor (AIP), comprises a mix of scalar processors and RCCs in a flexible, scalable architecture that supports open standards (Figure 1). Its flexible I/O architecture allows us to mix and match boards to create different configurations that best suit our application requirements, delivering what we call mission-unique functionality. The AIP leverages the unique capabilities of Xilinx's SRAM-based FPGAs to enable in-orbit reconfigurability for additional flexibility and SWAP benefits. The AIP also supports a variety of single-event effect (SEE) radiation-mitigation techniques to ensure reliable operation in different orbits.

The heart of the AIP system architecture is a reconfigurable computer board hosting a trio of Virtex®-4 FPGAs (Figure 2). Our investigation of available components concluded that Virtex FPGAs were the only devices that could achieve our performance targets and provide the characteristics required for spaceflight. For the most demanding applications, Xilinx offers Virtex-4QV space-grade devices. These FPGAs incorporate the same architecture as their commercial-grade counterparts (enabling low-cost system development,

prototyping and evaluation), but undergo special processing and screening to Class-Q and Class-V requirements.

Working in concert with sequential processors, the Virtex-4 FPGAs serve as a coprocessor for accelerating key processing-intensive tasks. The three-FPGA board architecture offers flexibility for addressing the unique requirements of different missions. In some applications we use the three FPGAs for SEE mitigation techniques that require component-level redundancy. In others, we partition a large coprocessor across multiple devices and use the ring bus connecting the three FPGAs via an LVDS interface for high-speed communication among the devices. Employing an extended 6U form factor, the card has two connectors for interboard communication: one for a CompactPCI backplane and another for a high-speed serial network.

Each FPGA has direct access to dedicated high-speed memory on the RCC board and to a connector supporting expansion and customization via a high-speed mezzanine card. This architecture enables us to expand the capabilities of the RCC board with mission-specific I/O, memory, analog circuitry and even

additional logic. The mezzanine card, which also plays a part in SEE mitigation in certain applications, joins the RCC board through a set of three connectors, each providing 170 LVDS I/Os.

Moving mission-specific functionality to the mezzanine card has enabled us to use the same FPGA-based processing card for multiple unique applications. The common architecture has reduced project risk, costs and schedules.

Mitigating Radiation Effects in FPGAs

Flying reconfigurable FPGA-based systems in space requires special considerations to ensure reliable operation in a high-radiation environment, because the SRAM-based configuration circuitry is susceptible to radiation-induced upset. The first consideration is the choice of component. In addition to industrial and military temperature-grade options, Xilinx offers V-grade Virtex-4 and Virtex-5 FPGAs that undergo special processing to make them immune to radiation-induced latch-up and deliver guaranteed performance for total-dose effects. These devices also undergo extensive characterization in neutron and proton beams to

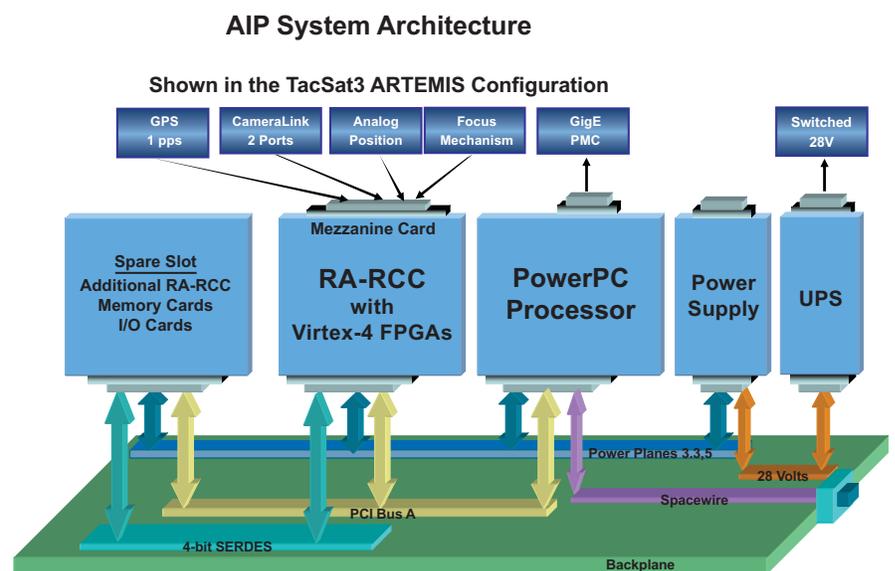


Figure 1 – AIP architecture, showing combination of RCC boards and other system boards in Artemis

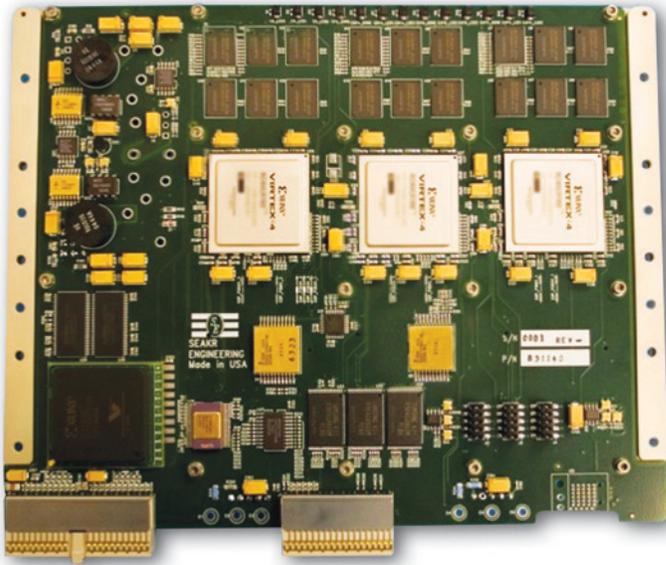
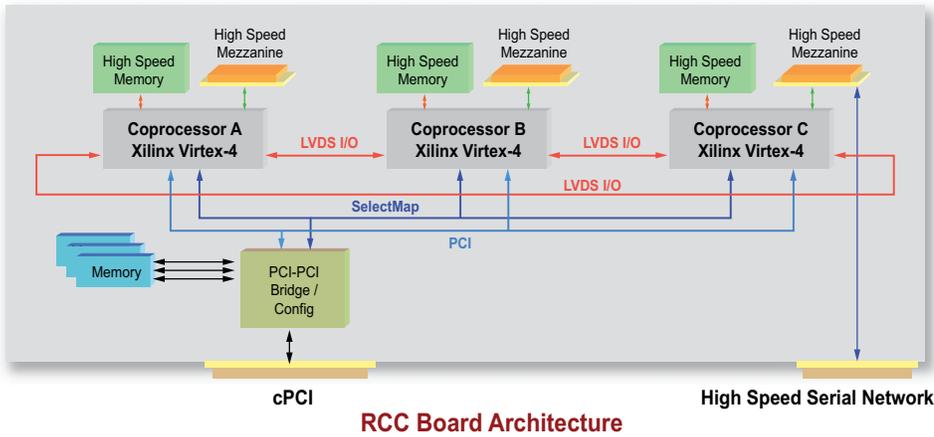


Figure 2 – RCC board architecture diagram and photo

generate reliable predictions of single-event upset (SEU) and single-event functional interrupt rates for particular orbits. This data guides engineers in selecting an upset-mitigation scheme appropriate for the application and orbit.

Upset mitigation for reconfigurable FPGAs generally involves some combination of hardware triple redundancy and configuration memory scrubbing. Hardware triple redundancy involves tripling critical circuits to ensure uninterrupted operation even if one element experiences a radiation-induced upset. It also adds a voter circuit that compares signals arriving from the triplicated logic branches and rejects an invalid signal resulting from an upset.

Designers can choose from among a number of approaches to satisfy the performance and availability requirements of their system. One technique involves redundant FPGAs and an external radiation-hardened voter circuit. Another approach is device-level mitigation, which involves triplicating mission-critical logic inside a single FPGA, along with the associated voter circuits. Traditionally, engineers have tackled this triple-modular redundancy (TMR) design technique by hand. Xilinx now offers special design tools that automate TMR implementation within an FPGA. Factors influencing the choice of mitigation scheme include the size of the target circuit, the level of radi-

tion in the selected orbit and uptime requirements of the circuit.

The basic concept behind memory scrubbing is to rewrite the configuration memory more frequently than upsets accumulate. Designers can choose from among a number of memory-scrubbing techniques to suit different upset rates and uptime requirements. At one extreme, the simplest technique involves reloading a complete bitstream into the configuration memory. This method involves low overhead but requires that the circuit be inactive for a period at least as long as a configuration cycle. More-advanced techniques exist for applications with stricter uptime requirements, higher upset rates or both. Leveraging the partial-reconfiguration capability of Virtex FPGAs, these methods involve circuits that detect memory upsets and then initiate reconfiguration of only a selected subset of the memory array.

AIP in Action

We have already used the AIP architecture in four separate missions. The combination of the FPGA-based RCC board and flexible mezzanine card has enabled our engineers to build a variety of processing and communications systems rapidly and to implement mitigation schemes suitable for the unique requirements of each mission.

The first incarnation of the AIP was Artemis—the Advanced Responsive Tactically Effective Military Imaging Spectrometer—which will be flying on the TacSat-3 satellite, scheduled for launch in the second quarter. Designed to provide situational awareness on the battlefield, Artemis performs advanced image processing on data the satellite collects and delivers it to soldiers in the field via a narrowband downlink. Our engineers realized that an RCC approach would be required to meet the spacecraft's size, weight and power goals: dimensions of 7.8 x 11.41 x 10 inches; a weight of 18 pounds; and power of 40 watts (with a hard limit of 50 W).

Two Virtex-4 FPGAs perform sensor data acquisition along with preprocessing functions such as calibration. An embedded processing system based on the MicroBlaze™ soft-processor core coordinates memory

access and processor coordination, while a PowerPC® single-board computer handles image generation and target cueing. Figure 1 shows the Artemis system architecture.

Because the image data path is not mission critical, configuration memory scrubbing provides suitable mitigation for Artemis. That is, the designers were able to satisfy the availability requirements without resorting to logic triplication or redundant devices. Furthermore, we determined that we could use commercial-grade FPGAs to

flash memory; each bitstream configures the FPGAs to process a specific waveform and frequency. In this way the system is able to support multiple waveforms using a minimum amount of hardware (see Table 1).

The flexibility of the RCC board provides benefits that start with initial system development. The delay between requesting and receiving a spectrum slot can be greater than one year. Reconfigurable hardware allows designers to initiate development before receiving

Internet in Space

Packet-based networking in space promises to provide the same flexibility and robustness available in terrestrial networks. Long a mainstay of wireline networking equipment, reconfigurable FPGAs lend the same benefits of performance, flexibility and design acceleration to space-based applications, as demonstrated by our Programmable Satellite Internet Protocol Modem. PSIM extracts Ethernet frames from standard satellite communications waveforms and facilitates IP routing on the spacecraft. Packet-based satellite communication enables beam- and waveform-independent routing of data through virtual circuits. Compared with standard bent-pipe satellite communication channels, packet-based networking improves scalability and throughput, enables decentralized multicast and is flexible enough to offer fine-grained quality-of-service.

The PSIM comprises 12 Virtex-4V FPGAs mounted on four RCC cards, along with two sequential processors and an analog switch card in a ruggedized chassis (Figure 3). The FPGAs perform waveform processing, while the sequential processors provide Ethernet interfaces and packet switching.

The availability requirements for this mission called for a more-aggressive mitigation scheme than the ones we used in Artemis or the PST. The system must provide uninterrupted end-to-end control, as recovery from an error would take too much time and reduce availability below target requirements. As a result, SEAKR engineers implemented a mitigation scheme that enables on-the-fly correction of errors while providing uninterrupted service.

We triplicated the FPGA logic in three devices on each RCC board. A radiation-hardened logic device on the mezzanine card serves as a majority voter. Memory scrubbing takes place in the background, transparent to the operation of the network. The mezzanine card also provides the physical interfaces to the router.

The mission on which the PSIM is to fly is scheduled for launch in the second quarter.

Receiver/Uplink	L-Band: 1,760 to 1,840 MHz
	S-Band: 2,025 to 2,120 MHz
Transmitter/Downlink	S-Band: 2,200 to 3,200 MHz
Space Ground Link System	FSK-AM Command Uplink (1 kbps, 2 kbps)
	Subcarrier BPSK Telemetry Downlink (256 kbps)
Universal S-Band	Subcarrier BPSK Command Uplink (<= 4 kbps)
	Subcarrier BPSK Telemetry Downlink (256 kbps)

Table 1 – Programmable Satellite Transceiver communications details

reduce the system cost. Two Virtex-4 FPGAs share the image-processing workload, while the third socket remains unpopulated to minimize power.

The AIP methodology delivered huge productivity dividends after its first application, enabling us to save roughly one year of development time on each subsequent project thanks to a substantial reduction in nonrecurring-engineering costs.

RCC Enables a Flexible Transceiver

The second mission for the AIP was in the Programmable Satellite Transceiver. The PST system provides frequency-agile satellite communications on multiple radio bands. SEAKR engineers concluded that even high-end PowerPC processors could not provide the necessary heavy lifting within the SWAP requirements of 3.86 x 6.85 x 7 inches, 10 pounds and Rx power of 10 W (Tx, 45 W).

To meet these requirements, our designers exploited the in-system reconfiguration capability of Virtex FPGAs. The system stores multiple configuration bitstreams in

the spectrum assignment and then implement the required frequency later. This capability also enables developers to adapt the system to the requirements of subsequent missions. SEAKR is currently developing additional waveforms for future development.

The nature of the PST mission simplified the radiation-mitigation requirements. The communication system maintains end-to-end control of the channel and is tolerant of data errors: in the event of corrupted data, the system responds by retransmitting the affected packets. This inherent error tolerance means that configuration memory scrubbing provides suitable SEU mitigation for the control path. To protect intermediate processing results, we triplicated the memory on the mezzanine card.

To complete the system, the AIP board joins RF modules and a power module in an extended 6U form factor chassis that's designed to withstand shock and other stresses of launch.

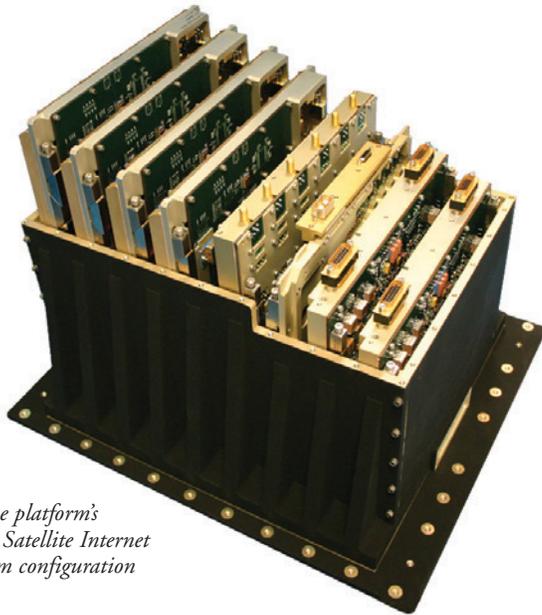


Figure 3 – The platform's Programmable Satellite Internet Protocol Modem configuration

High-Performance Video for Manned Spaceflight

The most recent application of the AIP architecture is the Visual Processing Unit (VPU) for the Orion Crew Exploration Vehicle. The VPU provides a reconfigurable platform for processing image algorithms driving pose estimation, optical navigation and compression/decompression. The system receives image data from a variety of sensors: star tracker, vision navigation sensor, docking camera and situational-awareness camera.

To process this volume of data takes a combination of sequential processors and FPGA-based RCC cards. Virtex-4 FPGAs implement video-processing algorithms such as feature recognition, graphical overlay, tiling and video compression. They also incorporate a MicroBlaze soft-processor core to coordinate algorithm cores and processor communication. A LEON fault-tolerant processor-based single-board computer is dedicated to system coordination, error handling, RCC configuration and oversight, and interconnect control.

A mezzanine card provides the sensor interfaces, implementing LVDS links to all three FPGAs for maximum flexibility in video stream selection and mitigation schemes.

Because of the “human-critical” nature of the tasks the VPU carries out, SEAKR engineers selected Virtex-4QV space-grade FPGAs and implemented an aggres-

sive mitigation scheme. Combining TMR methodology and configuration memory scrubbing ensures transparent correction of control path corruptions.

In conclusion, leveraging the capabilities of Virtex FPGAs, SEAKR engineers have developed an application-independent processor for space applications and demonstrated its flexibility on several missions. The RCC serves as a key component in satellite-based image processing and communications, flexible radio communications, space-based networking and navigation for human spaceflight.

Space-grade Virtex FPGAs are COTS components that offer the performance required for demanding data-processing and communications systems. These reconfigurable FPGAs enable a flexible, scalable architecture that reduces development cost and accelerates design cycles. In addition to supporting rapid development and flexible manufacturing on the ground, Virtex FPGAs offer the ability to reconfigure on-orbit for additional, significant SWAP benefits.

Future generations of V-grade reconfigurable FPGAs promise to offer greater size, weight and power benefits by delivering higher logic capacity, greater integration of hardened IP blocks, higher performance and lower power consumption. Radiation-hardened reconfigurable Virtex FPGAs will simplify the designers’ task and extend SWAP benefits further by eliminating the need to implement logic- or device-level redundancy. ●●

X5 com

We Need to Talk...

X5-COM PCIe IO Module featuring 4 Ethernet/SRIO/Gigabit Serial Ports

APPLICATIONS

- ✦ Ethernet packet processing
- ✦ Communications test equipment
- ✦ FPGA computing node
- ✦ Wireless Remote Radio Head IF Processor for OBSAI
- ✦ Real-time data encryption for Ethernet
- ✦ Remote IO interfacing
- ✦ High Speed Data Recording and Playback
- ✦ IP development

FEATURES

- ✦ Four communications ports: Gigabit Ethernet, Aurora, Infiniband, Serial Rapid IO (requires supporting IP in FPGA)
- ✦ Xilinx Virtex5 SX95T or FX100T FPGA
- ✦ SX95T: 640 DSP MACs
- ✦ FX100T: 2 PowerPC+ 256 DSP MACs
- ✦ Industry-standard SFP modules support up to 4.125 Gbps over Copper or Fiber Optic cables (rates depend on FPGA speed and cable)
- ✦ 512MB DDR2 DRAM supporting 4 GB/s transfer rates
- ✦ 4MB QDR-II SRAM for computations
- ✦ 4 Rocket IO links for data plane using P16
- ✦ >1 GB/s, 8-lane PCI Express Host Interface
- ✦ Conduction cooling and thermal monitoring
- ✦ XMC Module (75x150 mm)
- ✦ Adapters for Desktop, CompactPCI, and Cabled PCI Express

wireless ip cores

R-interface

Windows, Linux, Mac OS, Solaris, FreeBSD

Innovative Integration
... real time solutions!

805-578-4260 phone
www.innovative-dsp.com

Save days to weeks on your next FPGA debug cycle



Agilent Logic Analyzers

Up to 1.2 GHz timing, 667 MHz state, and 256 M deep memory

OR



Agilent Mixed Signal Oscilloscopes

4 scope channels + 16 timing channels

+

Agilent FPGA Dynamic Probe

Application software to increase visibility inside your FPGA

=

Fastest FPGA Debug Available

- Perform real-time functional and parametric debug that time-correlates internal FPGA activity with the surrounding system
- Change internal FPGA probe points in seconds without design changes
- Get fast serial bus decode for I²C, SPI, CAN/LIN and RS-232/UART

Start saving time now. Download FREE application information. www.agilent.com/find/fpgatools



Agilent Technologies

Virtex-5 Powers Reconfigurable, Rugged PC

RMT's SwitchBack uses Xilinx Virtex-5 FPGA in a PC that users can customize in the field and upgrade on the fly.

by Shane Lewis
 Director of Technology Development
 RMT, Inc.
slewis@ropermobile.com

The U.S. military and companies in heavy industries such as mining, transportation, warehousing, logistics and public safety all have strict requirements for personal computers. First and foremost, their PCs must be ruggedized to endure physical abuse, extreme heat and cold, and exposure to moisture, even submersion. At the same time, these rugged PCs need computing functionality that's on par with the latest commercial PCs, but surpasses them in security and global communications capabilities. Customers looking for this type of computer also require very specific peripheral features targeted at mission-critical tasks. But up until recently, these buyers have been forced to use standard, off-the-shelf PCs that often don't fully meet their needs.

This deficiency presented our design group at RMT, Inc., with the enormous challenge of designing a modular and customizable computing solution—a “common platform” that truly executed the customers' strictest requirements and met their expectations. Savvy R&D teams know the pitfalls of trying to build a “one-size-fits-all” device, which too often results in a kluge of difficult compromises. Our design team set out to defy the odds and engineer a truly adaptable computer platform that would be field-reconfigurable and customizable at the circuit level, while at the same time remaining an elegant, rugged and user-friendly system.



The result of this effort is the SwitchBack. This computer is truly different at every level, and redefines PC architecture through its innovative use of the Xilinx® Virtex®-5 FPGA. The difference between the traditional PC and the SwitchBack architecture is quite extraordinary.

Traditional Open PC Architecture

The base of any modern legacy PC is an x86 processor and associated chip set for either Linux or Windows, namely Windows XP or Vista. It's the legacy support behind this code set that has enabled it to dominate the PC world and subsequently, to constrain the embedded computer space, where operating systems and processor technologies tend to be more fragmented. If you open any desktop, laptop or tablet PC designed to run Windows XP or Vista, you will find a chip set/CPU circuit topology.

This architecture, which has been the de facto standard for all PCs, both rugged and commercial, for many years, has one fundamental limitation. Task execution must either be written for a specific processor or must be plugged into one of the many available expansion ports as external hardware. The boundaries of what can be done are closely defined around the wiring of the dedicated ASICs that make up the chip set itself.

SwitchBack Architecture

To get a different outcome, the RMT team knew that we had to rethink this design. We crafted SwitchBack's patent-pending architecture to be both field-reconfigurable as well as compatible with Windows-based applications. The concept of field reconfiguration is widely utilized today in embedded computing, a trend that FPGAs have enabled.

While programmable logic may play various supporting roles on some x86 PC-based motherboards, the FPGA is the hero in the SwitchBack, directing the computer's functions (Figure 1). The Virtex-5 is the primary controller of all major subsystems. From the moment the user presses the power button, the FPGA controls all peripherals, including the display itself, as well as the flow of most data. This scheme

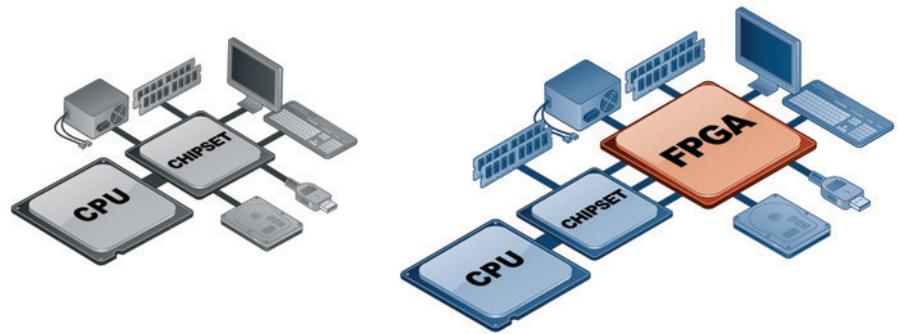


Figure 1 – The difference between a traditional PC architecture (left) and the revolutionary SwitchBack is that in the latter, the FPGA is the primary controller. This improves processing time and makes the SwitchBack reconfigurable and customizable.

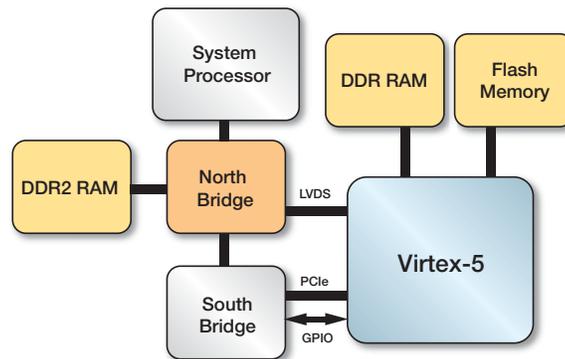


Figure 2 – In the SwitchBack architecture, primary system control is a function of the Virtex-5 FPGA, not the main x86 CPU processor.

makes it possible to access and view data without booting the Windows operating system—a feat that's virtually impossible with a traditional PC. The SwitchBack takes it a step further by allowing users to access and control custom functions and peripherals without the assistance of the processor or operating system. These are programmed into the BackPack, a modular system that attaches to the rear of the SwitchBack and allows full control of any peripherals without waiting for slow-moving processors or operating systems.

The Heart of the SwitchBack

The amazing flexibility and control of the SwitchBack is made possible through an architectural design based on the Xilinx Virtex-5 LX30T. The RMT team selected

the Virtex-5 for its vast array of internal resources as well as its RocketI/O™ and PCI Express® Endpoint Block. We then set the PCI Express interface as the major connection point between the embedded system and the Intel x86 system. Although there are other interface bridges between the components, this one is the major data pipeline for processing and control.

Hardware Design

The fundamental layout of SwitchBack's hardware design is essentially two systems that can operate independently. The Virtex-5 FPGA, not the x86 CPU processor, is the primary control system (Figure 2). It actually configures immediately upon system boot-up, before allowing the secondary system to begin booting.

Additionally, the FPGA has its own RAM and flash memory for both configuration and program storage, should users need to program secondary operations into the FPGA. This revolutionary architecture and its FPGA-defined algorithms provide several key capabilities that ordinary PCs cannot accomplish. They include control of all system resources allowed by the main processor or its base chip set, and independent and autonomous control of peripherals, including those programmed into the SwitchBack's BackPacks. The FPGA also includes processor-independent functions to assist the main processor.

Control of system resources and peripherals is straightforward and well-understood in the embedded world. But the use of processor-independent functions raises new possibilities for realizing the SwitchBack's full potential. The options include reconfigurable hardware, open FPGA for open architecture, additional customization via BackPacks and a Backpack modular development kit.

Reconfigurable Hardware

The Virtex-5's internal resource array—including the ExpressFabric Architecture, block RAM, 1.25-Gbit/second Select I/O and DSP48E slices—offers a multitude of possibilities when creating functions and processes that would normally be done in software on the main processor. With these chip features, we were able to implement many capabilities into the FPGA to offload a portion of the main processor's duties, or create from scratch entirely new subsystems that behave like physical expansion cards in a regular PC.

By using the PCI Express bus between the two systems, we can connect the new hardware expressed in the FPGA to the main processor as though it were physical hardware plugged into an expansion port (Figure 3). This allows the construction of new types of devices that operate independently of the main x86 processor. These devices or functions may include data format conversion, custom logic interfaces, hardware emulation, independ-

ent microprocessors, communications devices, arithmetic cores and autonomous Backpack control.

In fact, we call these functions “virtual devices” because they appear to the operating system as separate hardware, yet there is no physical circuit card to implement them. One or more devices can be implemented at the same time, expanding the platform far beyond the features expressed by the onboard chip set.

Open FPGA for Open Architecture

Part of our vision was to preserve the PC's open architecture platform and extend that flexibility further to include the FPGA itself. The capacity of the Virtex-5 FPGA is greater than the SwitchBack requires for its main control functions. Our team purposely chose a larger FPGA to allow customers to add their own programming and functionality to SwitchBack, enabling them to develop a truly custom, precision tool that can accomplish the same results as multiple independent subsystems attempting to work in parallel.



Figure 3 – Any reconfigurable hardware expressed in the FPGA can connect to the main processor, which can access it as though it were physical hardware plugged into an expansion port.

Thanks to the innovative use of the Virtex-5 in the SwitchBack, customers need no longer settle for a suboptimal computer and an assortment of dissimilar peripherals connected by a mess of cables.

We also intentionally underutilized the Virtex-5 FPGA as the system master in SwitchBack. For example, of the Virtex-5 LX30T FPGA's 32 DSP48E slices, the SwitchBack uses only one, leaving the remaining 31 available for the end user. Depending on the buyer's resource needs and the configuration of the SwitchBack at purchase, available resources such as registers, LUTs, BRAM, DCM and PLLs may be available for customer-specific use. In general, the SwitchBack uses less than half of the FPGA for system management and general housekeeping, leaving the majority of the remaining resources for the user to define.

A custom update interface tool allows users to easily update the flash memory inside the SwitchBack. This simple software update requires no JTAG or specialized equipment to modify the FPGA's configuration file. Installing new hardware is quick and easy. Users simply reboot the SwitchBack system and watch the new hardware appear in the Hardware Manager, ready for immediate use.

We also packaged the SwitchBack's requirements for the FPGA into an easy-to-implement core, which allows the end customer to quickly add logic, registers and data buses to the remaining FPGA space available.

Our Firmware Development Kit (FDK) allows customers to modify and reconfigure the SwitchBack to meet their own specific needs, as the mission changes. The SwitchBack can quickly adapt to the in-field situation with a unique module upgrade, custom logic changes to the FPGA or both. This scheme effectively redesigns the system in the field. By providing this capability in an FDK, customers with FPGA experience and the proper place-and-route tools can create and shape the SwitchBack to fit their exact needs.

Additional Customization via BackPacks

The SwitchBack is capable of further customization through its Backpack technology. BackPacks are customer-specific modules that users can securely attach to the back of the SwitchBack. RMT initially created the Backpack to eliminate the need for external peripherals and to add multiple ports. BackPacks can take on an infinite array of sizes, shapes and complexity—handling additional processing capabilities, for example—endowing this personal computer with the computational clout of a supercomputer. In this way, the Backpack is a field-customizable system that can morph SwitchBack into a highly integrated, precision tool.

The RMT team routed the GPIO from the FPGA directly to the Backpack port so that logic in the FPGA could gain access to the attached Backpack, making it possible to control any type of device without ever involving the main system processor. The SwitchBack's promise can now be fully realized when combining the reprogrammable FPGA with the unlimited potential of external adaptability.

Modular Development Kit

Building upon the success of the Backpack, we found that FPGA-savvy customers could develop and program their own BackPacks if we provided them with the right tool kit. The SwitchBack's Modular Development Kit (MDK) allows them to design custom BackPacks that will make use of SwitchBack's unique architecture. In many cases, customers have electronic devices or circuit cards they wish to integrate for rapid testing and deployment. The MDK makes doing so a snap. It provides a functional circuit board, cables, schematics and mechanical CAD data, allowing a customer to build a Backpack in a matter of days.

By using SwitchBack's onboard FPGA and Backpack technology, users can build entire subsystems of specialized functions

and tie them to the computer rapidly. Although a Backpack is part of the computer itself, it can operate completely autonomously, sharing only processed data rather than bogging down the main processor with extra duties, as can happen when attaching devices through traditional methods (such as USB). Thus, customers can implement BackPacks with equal or greater processing responsibility without burdening the main processor.

The Xilinx Virtex-5 with DSP48E slices can be put into service to provide additional signal processing at an enhanced and improved efficiency level. A few examples of possible adaptations include image processing, software radios, cryptography, network security and analog modems.

Users can increase the performance of particular functions when they build the analog topology of these functions into a Backpack while letting SwitchBack handle the processing duties onboard.

Thanks to our design group's innovative use of the Virtex-5 in the SwitchBack, customers no longer have to settle for a suboptimal computer and an assortment of dissimilar peripherals connected by a mess of cables. In this way, the SwitchBack enables rapid system design and deployment for mission-critical applications.

SwitchBack is the next evolutionary step in computer technology. Essentially, it's a reconfigurable PC platform that customers can customize in the field and upgrade on the fly. Since it can run mission-critical applications without interruption, it is especially suited for operations in key markets such as the U.S. military, mining, transportation, warehousing and logistics, public safety and many others.

To learn more about the SwitchBack and its revolutionary architecture, visit www.ropermobile.com/products/switchback or e-mail or phone the author at RMT, Inc., at slewis@ropermobile.com or (480) 705-4200, ext. 306. 

Floating Point: Have it Your Way with FPGA Embedded Processors

Implementing an FPU for the Xilinx PowerPC 440 is easy.

by Glenn Steiner
Senior Manager
Xilinx, Inc.
glenn.steiner@xilinx.com

Ben Jones
Senior DSP Design Engineer
Xilinx, Inc.
ben.jones@xilinx.com

Peter Alfke
Distinguished Engineer
Xilinx, Inc.
peter.alfke@xilinx.com

When creating embedded applications employing numerical processing, it's important to keep the arithmetic operations simple, generally by using an integer or fixed-point representation. This helps to minimize the cost and power and to maximize the speed of an implementation in hardware.

FPGAs are well suited to performing fixed-point operations, and offer the ability to construct highly parallel data path solutions in logic or soft- or hard-processor-based implementations. The Xilinx® PowerPC® 440, the latest hard processor within the FXT series of the Virtex®-5 FPGA family, offers a superscalar capability, allowing users to program the device to perform one or two fixed-point operations in parallel at clock rates of up to 550 MHz.

While users can program the device to perform most calculations using integer or fixed-point arithmetic, they often must rearrange calculations and insert scaling operations to calculate results to sufficient accuracy. For complex algorithms, this can be time-consuming and often

results in programs that are application-specific and not reusable. The alternative is to adopt a standard floating-point representation to provide a high dynamic range that's adaptable to many applications. The reward is that you need not modify the algorithm to obtain a fixed-point implementation for any particular application or operating environment, nor extensively modify the code for subsequent projects and applications.

While Xilinx offers a very efficient emulated floating-point solution for the PowerPC 440 processor based upon the IBM floating-point performance libraries, the core still requires tens of cycles to perform each operation. Hardware acceleration of floating-point operations in the form of a

floating-point unit (FPU) reduces this cycle time. The PowerPC 440 processor in the Virtex-5 FXT provides an efficient interface for connecting a hardware accelerator, such as the Xilinx soft FPU, to the processor core. This scheme bridges the 128-bit auxiliary processor unit (APU) interface provided on PowerPC 440 processors to coprocessors via the fabric coprocessor bus (FCB). One such coprocessor, the Xilinx LogiCORE™ IP Virtex-5 APU-FPU, gives Virtex-5 FXT users floating-point on the PowerPC “their way,” using either software emulation or a dedicated soft-logic FPU. Figure 1 shows a typical implementation of a PowerPC 440 processor connected to the Virtex-5 APU-FPU via the FCB.

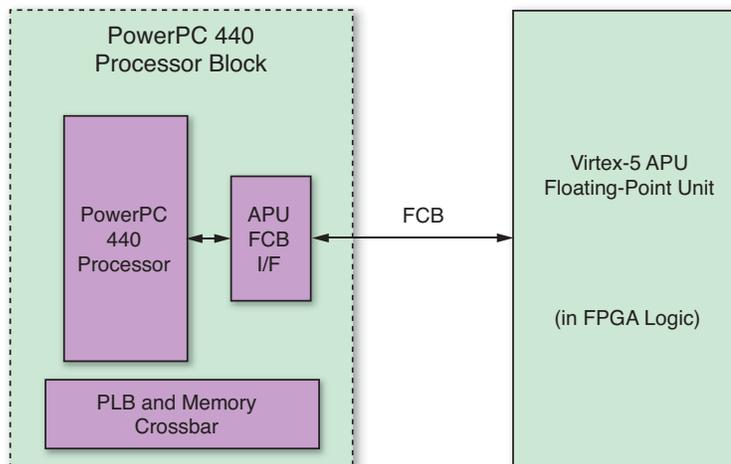


Figure 1 – Embedded processor system containing an APU-FPU core

About the PowerPC 440 FPU

Xilinx designed the APU-FPU specifically for the PowerPC 440 processor embedded in the Virtex-5 FXT FPGA. The tight coupling of the FPU to the processor through the APU interface lets the floating-point unit directly execute native PowerPC floating-point instructions to achieve typically 6x acceleration over software emulation.

The Xilinx PowerPC FPU complies with the IEEE-754 standard for single- and double-precision floating-point arithmetic, with minor exceptions. Xilinx offers variants optimized for 2:1 and 3:1 APU-to-CPU clock ratios, allowing the PowerPC processor to operate at maximum frequency. Autonomous instruction issuing hides arithmetic latency and decreases the cycles per instruction. What’s more, these optimized implementations leverage the device’s high-performance DSP features to reduce operator latency and trim logic count and power consumption. Xilinx supports the APU-FPU flow in the Xilinx Embedded Development Kit (EDK).

Figure 2 provides an overview of the FPU architecture. The APU-FPU comprises execution units, register file, bus interface and all the control logic necessary to manage the execution of floating-point instructions.

The FPU comes in two variants. The double-precision version executes all of the floating-point instructions, including sin-

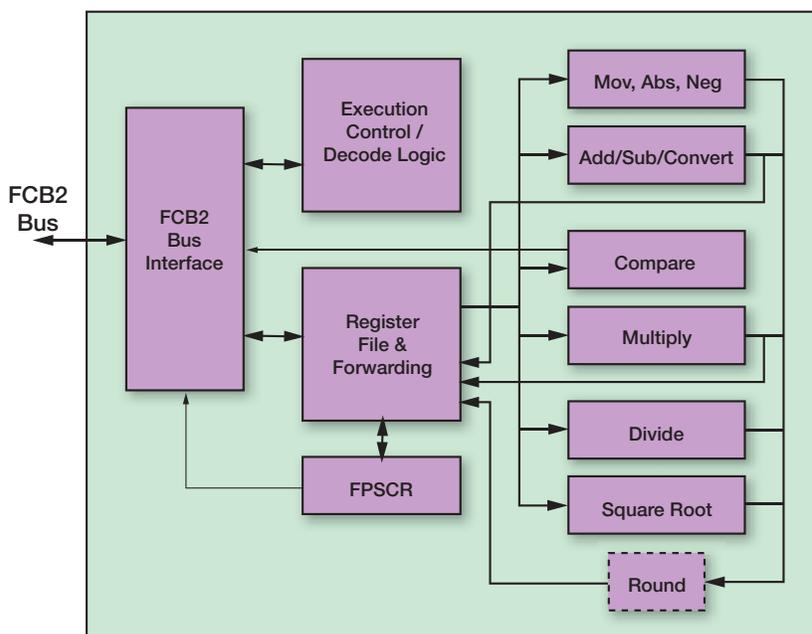


Figure 2 – Virtex-5 FXT PowerPC 440 floating-point coprocessor architecture

gle-precision, with the exception of the PowerPC ISA graphics subset (fsel, fres and frsqtrte). This means you can use the FPU with a range of commercial compilers and operating systems, listed at www.xilinx.com/is/embedded/partners/listing.htm.

A single-precision variant of the APU-FPU, which the Xilinx compiler supports, uses fewer resources. When this FPU is employed, double-precision operations are performed using software emulation.

Tying APU-FPU to PowerPC 440

There are two ways to attach the APU-FPU to the PowerPC 440 processor. The first method is to use the Base System Builder (BSB) wizard within the Xilinx Platform Studio design tool. The second method is to simply attach the APU-FPU unit to an existing design.

With the BSB wizard, you specify a target board and desired processor, either PowerPC or MicroBlaze™. Then, via a

On average, the soft FPU is six times faster than software emulation. The single-precision FPU is typically 13 percent faster than the double-precision version.

series of check-boxes and drop-downs, you select the IP you want to include in the design. The BSB wizard makes it easy to quickly assemble and test a basic processor system. Connection of the APU-FPU is as easy as checking the box indicating that you wish to include an FPU (see Figure 3,

top screen). The wizard implements a double-precision FPU optimized to operate at one-third of the processor clock frequency. You can further customize the FPU for a higher clock rate or for single precision.

If you don't want to use the wizard, the other method is to simply drag the APU-

FPU IP from the IP Catalog to the System Assembly View, and then configure the FPU. The bottom screen in Figure 3 shows the IP Catalog on the left and a newly added FPU in the System Assembly View. By right-clicking on the FPU and selecting Configure IP, you can pick the desired precision (single or double) and choose whether you want the FPU optimized for low latency (one-third clock rate) or high speed (one-half clock rate). Finally, connect the FPU to the FCB and link the FPU/FCB clock to the appropriate clock (typically one-half or one-third of the processor clock rate).

Have Floating Point Your Way

Provided free of charge with Platform Studio, the Virtex-5 APU-FPU delivers customized floating-point support. You may choose to implement the single-precision FPU using approximately 2,500 LUT-register pairs or the double-precision FPU using roughly 4,900 LUT-register pairs. Alternatively, you can run your software application with floating-point emulation and no additional FPGA logic.

You can choose your performance level up front: either select the appropriate FPU, or implement your design and determine if software emulation is meeting your requirements. If it's not, you can upgrade to a soft FPU.

Clearly, if the performance you're getting from software emulation is sufficient, then you don't need an FPU. But for higher performance, you can use an APU-FPU. Use the double-precision FPU if your application requires it or you are using a partner compiler. If your application needs only single-precision arithmetic and you are using the Xilinx-provided GNU compiler, then the single-precision FPU will reduce your logic requirements. Remember, if you choose the double-precision FPU, it will execute single-precision arithmetic by rounding the result to provide single-precision accuracy.

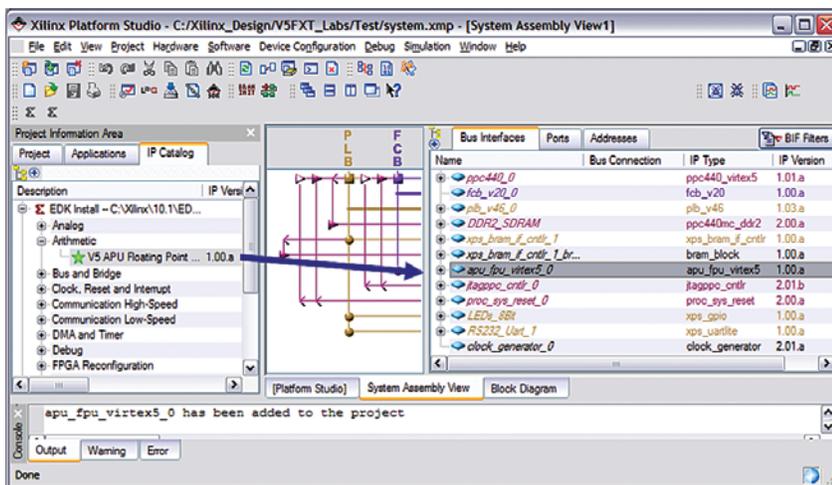
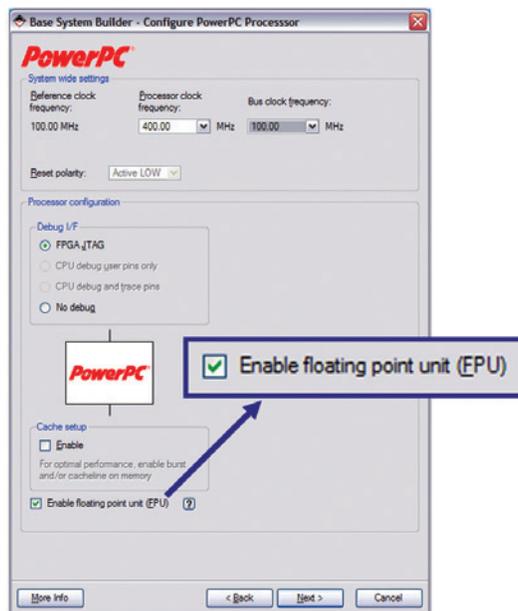


Figure 3 – Adding an FPU to an existing PowerPC processor design via BSB wizard (top) and via System Assembly View

Typical Performance Gains

When evaluating the need for a hard or a soft FPU, you should determine how floating-point-intensive your code is. Frequently, code contains a mix of floating-point, integer, memory and logic operations. Thus, while benchmarks can be good indicators of potential performance improvement, there is nothing better than running your own code.

To get a feel for how well the APU-FPU performs with code that is floating-point-intensive, Table 1 shows selected benchmark data for a Virtex-5 FXT PowerPC 440 processor operating at 400 MHz, with software emulation and with the processor connected to the double-precision APU-FPU running at 200 MHz.

The listed numbers are a subset of a larger suite of benchmarks that Xilinx has run to evaluate the performance of the processor floating-point units. On the average, the soft FPU is six times faster than software emulation and the single-precision FPU is typically 13 percent faster than the double-precision FPU.

In situations where floating-point is dominant, you can boost the performance of the soft FPU by optimizing the code to take full advantage of the FPU pipeline. The FIR filter benchmark is a good example of the potential performance gains. The nonoptimized code was typical “textbook code”—though easy for humans to read, it tends to execute inefficiently with most FPUs. However, by implementing loop unrolling, maximizing the retention of constants in the FPU registers and interleaving other code between floating-point instructions, you can obtain significant performance improvements for your design. In this example, the optimized filter code was 3.8 times faster than the nonoptimized code and 30 times faster than software emulation.

Overall, the Virtex-5 FXT with its PowerPC 440 processor offers numerous options for embedded applications. You can implement your design with or without an FPU, trading off the higher performance of the FPU vs. software emulation, to tailor the processing-power resources of the Virtex-5 FXT to best suit your design requirements and thus “have it your way.”

Benchmark*	Units	Software Emulation	Double-Precision FPU	FPU Speedup Over Software
1k fast Fourier transform	Iterations/sec	83	637	7.6
FIR filter (nonoptimized)	MFLOPS	6.5	51.4	7.9
FIR filter (optimized code)	MFLOPS	6.5	194	30
Whetstone	MFLOPS	6.2	34.8	5.6
Bytemark LU decomposition	Iterations/sec	8.1	43.6	5.4
Bytemark neural net	Iterations/sec	0.255	1.42	5.6
SPEC PID	Iterations/sec	1.07	3.826	3.56

* EDK 10.1 SP2 GNU compiler with flags: -O3 -funroll-loops

Table 1 – Typical floating-point performance, 400-MHz processor and 200-MHz FPU



One Board to Rule them All!

X5 GSPS

Lord of RF Signal Capture

Features

- Two 1.5 GSPS, 8-bit A/Ds (Nat ADC08D1500)
- +/-1V, 50 ohm, SMA Inputs
- Xilinx Virtex5, SX95T FPGA
- 512 MB DDR2 DRAM
- 4MB QDR-II SRAM
- 8 Rocket IO Private Links, 2.5 Gbps each
- >1 GB/s, 8-lane PCI Express Host Interface
- Power Management Features
- XMC Module (75x150 mm)
- PCI Express (VITA 42.3)



Perfect for

- Wireless Receiver
- WLAN, WCDMA, WiMAX front end
- RADAR
- Electronic Warfare
- Electronic Counter Measures (ECM)
- High Speed Data Recording
- Electronic Surveillance
- Spectral Analysis
- IP Development

wireless ip cores

R interface



Innovative Integration
... real time solutions!

805-578-4260 phone

www.innovative-dsp.com

Optimizing Xilinx FPGAs for Power

Designers can rely on a multitude of tools and techniques to tame their power budgets.

by Matt Klein
Principal Engineer, Technical Marketing
Xilinx, Inc.
matt.klein@xilinx.com

As IC processes have advanced over the last half dozen years from the 130-nanometer to the 90-nm and now the 65-nm node, at each step power management has grown in importance. It was at the 130-nm node that manufacturers started noticing that transistors leaked power, even in standby mode. At 90 nm, the operating voltage of ICs decreased, but leakage continued to rise, wasting a greater percentage of the device's power. At 65 nm, both these trends continue. Indeed, leakage at the 65-nm node is so pronounced that many designers consider managing power as important as meeting performance specifications.

Because FPGA vendors traditionally design for a broad range of applications and endow their devices with a plethora of high-speed transistors, FPGAs have not been the most power-conservative devices. Like other silicon designed in the most advanced processes, they use transistors that leak. However, designers can leverage an FPGA's programmability and use related tools to accurately estimate power and then employ optimization techniques to make their FPGA designs and the PCBs that contain them much more power efficient.

There are two primary types of power consumption in an FPGA: static and dynamic. Static power consumption is caused by leaking transistors—those that leak even when they are not doing tasks in a design. Dynamic power is the power the device consumes when it is running a task—toggling nodes as a function of voltage, frequency and capacitance. It is important to understand both power types and how each varies under different operating conditions so that you can properly optimize them to meet your design's power budget.

Static and Dynamic Power and Variation

Leakage current becomes fairly significant for both ASICs and FPGAs at 90 nm and even more challenging at 65 nm. To obtain higher performance from the transistor, its threshold voltage needs to be lowered, but that also increases leakage. Xilinx has done many things to minimize leakage, but nonetheless, the variation in static power from leakage is about two-to-one between worst case and typical process. Leakage power is also strongly influenced by core voltage (V_{CCINT}), varying with the cube of V_{CCINT} . Static power rises by approximately 15 percent for only a 5 percent increase in V_{CCINT} . Lastly, leakage is strongly influenced by junction (or die) temperature.

Figures 1 and 2 illustrate the variation of static power from leakage with voltage and temperature.

Other sources of static power consumption in the FPGA are DC currents from operating circuits, but for the most part they are notably process and temperature invariant. Examples include I/O DC currents (such as I/O termination voltages on terminated standards like HSTL, SSTL and LVDS) as well as DC currents in current driver I/O types like LVDS. Some FPGA analog blocks also are sources of static power consumption, likewise process and temperature invariant. Among them are the digital clock manager (DCM), a clock control element in Xilinx® FPGAs; phase-locked loops (PLLs), which are available in the Xilinx Virtex®-5 FPGA; and IODELAY, an element used to select programmable delays on input and output signals in Xilinx FPGAs.

Dynamic power is the power consumed during switching events in the core or I/O of an FPGA. To calculate dynamic power, we must know the number of toggling transistors and traces, capacitance and toggling frequency. Transistors are used for logic and programmable interconnects between metal traces in the FPGA. The capacitance consists of transistor parasitic capacitance and metal interconnect capacitance. The formula for dynamic power is:

$PD_{DYNAMIC} = nCV2f$, where n = the number of toggling nodes, C = capacitance, V = voltage swing, f = toggle frequency.

Tighter logic packing (through internal FPGA architectural changes) reduces the

number of switching transistors. Using smaller transistors trims routing lengths between them, which reduces dynamic power. So the 65-nm transistors in the Virtex-5 FPGA have lower gate capacitance and shorter interconnect traces, a combination that drops node capacitance by about 15 to 20 percent. That in turn lowers the dynamic power.

Voltage also has an effect on dynamic power. Moving from the 90-nm to the 65-nm process node reduces dynamic power in Virtex-5 FPGA designs by approximately 30 percent, simply by decreasing V_{CCINT} from 1.2 volts to 1 V. That plus architectural enhancements allowed a net dynamic power reduction of 40 to 50 percent compared

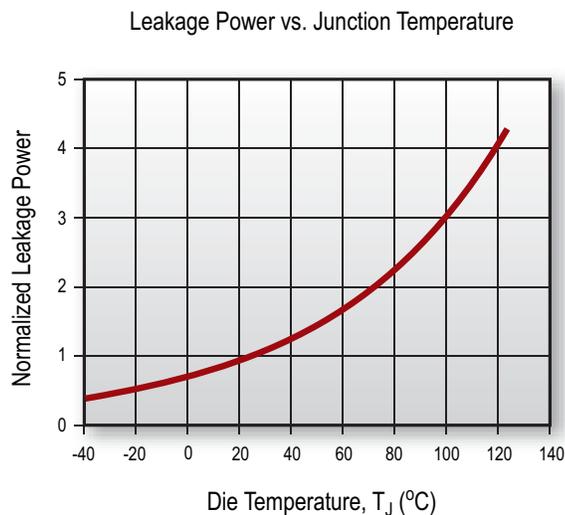


Figure 1 – Leakage power variation with die temperature

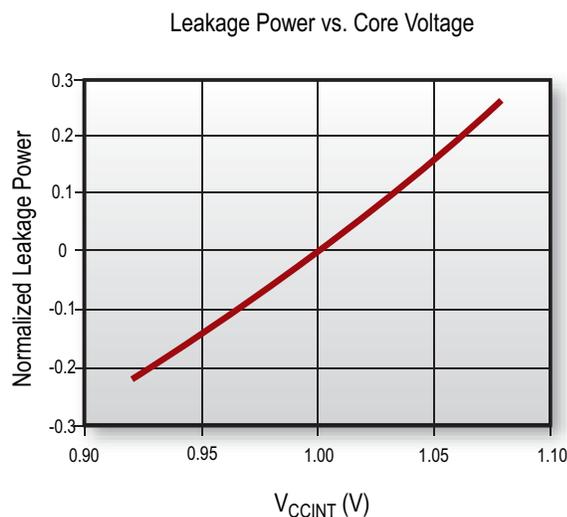


Figure 2 – Leakage power variation with core voltage (V_{CCINT})

with 90-nm technology. (Note: While dynamic power varies with the square of V_{CCINT} , it is largely temperature and process invariant for the core of the FPGA.)

FPGA Power Analysis Tools

Xilinx has two types of power analysis tools. We designed the first, the XPower Estimator (XPE) spreadsheet tool, for use before a designer employs implementation tools. Use the second tool, XPower Analyzer, after you have implemented your design to check how the changes you've made affect power consumption.

The XPower Estimator gives you a quick power estimation based on user descriptions of resource utilization in the FPGA, toggle rates, loading, etc. in a spreadsheet environment. This is the tool to use for your initial power evaluation, selection of power supplies and regulator, as well as any cooling solutions for the system (heat sinks, fans and the like).

With this Microsoft Excel-based tool, system architects can make device-, design- and system-oriented power decisions. You simply enter the estimated design parameters, such as resource utilization, operating environment, and clock and toggle rates. XPE then calculates estimated power for a given design and reports total power and maximum junction temperature as well as rail-based and block-based power.

In setting up the estimation run, the tool's Process function is an important feature. It allows you to see typical or worst-case power consumption by various blocks. Primarily, the static power from leakage on the V_{CCINT} supply is very process dependent. Further, the Voltage Source Summary lets you quickly see the effect on power consumption when voltages are varied. That is especially important to understand relative to V_{CCINT} , which is one of the power supplies representing all the core logic. Both the process variation and voltage variation selections in the XPE tool ensure that you can determine proper worst-case power supply sizing.

One other valuable feature of XPE is the Thermal Information/Summary, which allows you to specify heat sink, PCB properties and temperature information. This

ensures that the design will also meet the thermal specifications for commercial-grade or industrial-grade devices. The Block Summary, meanwhile, shows the power from each block and the Power Summary displays the sum of the quiescent and dynamic power.

Each of the tabs in the XPE tool allows you to enter utilization and toggle rates for a given type of resource, such as clocks, logic, I/O, block RAM (BRAM), PLLs, DSP and so on.

Finally, XPE's Graphs tab/sheet gives you a graphical look at power by function, process, voltage and temperature varia-

Change Dump (VCD) and Switching Activity Interchange Format (SAIF) files.

If you are using either the VCD or SAIF formats, you need to create representative simulation vectors so the tool can record the toggle rates of nodes in the system, which in turn allows you to access the data later. In the absence of these simulation files, the user can have the XPower Analyzer tool perform a vectorless simulation. This type of simulation uses mathematical and statistical modeling to propagate starting toggle rates through the actual design logic. It then generates a result containing toggle rates of each node in the design.

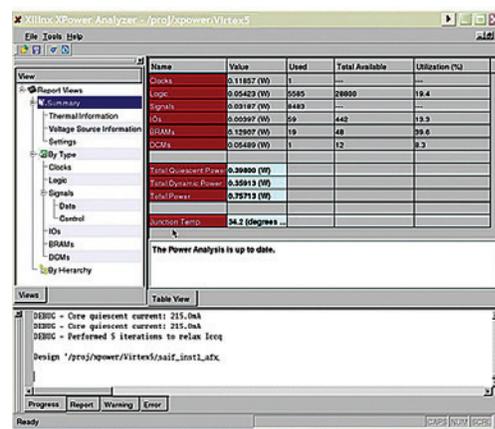


Figure 3 – Xilinx XPower Analyzer summary page

tion. The Power by Function graphic, in particular, lists each feature and shows its power consumption, allowing you to identify features that could best benefit from optimization.

The second Xilinx power analysis tool, XPower Analyzer, provides an even more accurate view of the power breakdown based on exact resource information it extracts during the implementation. You can supply the tool with test and simulation vectors, or perform vectorless power estimation. This tool uses characterized capacitance data for physical resources in the FPGA design.

XPower Analyzer is tied into the Xilinx Integrated Software Environment (ISE®), and it accepts post-place-and-route information from several internal Xilinx file formats. It also accepts industry-standard Value

With both the vector-based (from VCD and SAIF) files and the vectorless variety, XPower considers physical connectivity of the placed-and-routed design and exact resource usage. The tool cross-references the activity or toggle rate at each node with characterized capacitance data for physical resources and individual dynamic power consumption of each block at given toggle rates. The result, shown in Figure 3, represents total power and maximum junction temperature, and contains rail-based, block-based and hierarchically based power reporting.

XPower gives you a detailed look at where your design is consuming power and lets you do “what if” analysis to make more-informed choices for which blocks could most benefit from optimizations, ranging from simple ones through rearchi-

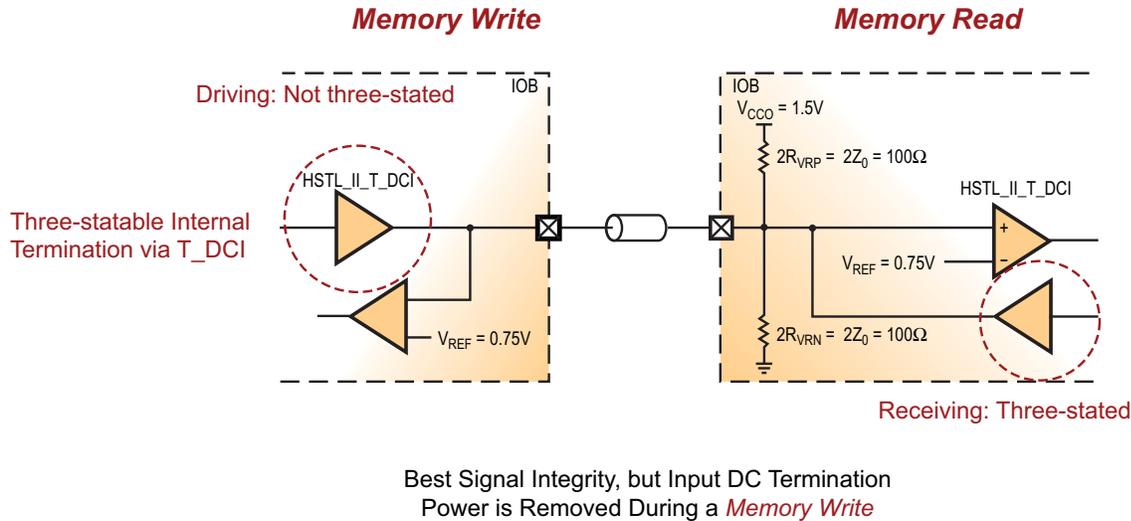


Figure 4 – FPGA pin shown during memory read and memory write using T_DCI

testing. Additionally, you can use XPower to document the actual power specs for a given design and pass that information to the board level.

Reducing Power with FPGA Design Techniques

While the process shrink to 65 nm gives the Virtex-5 an inherent dynamic power reduction, you can also employ new tools, tricks and techniques to further reduce the juice.

One way to attack power consumption is by selecting the right FPGA for your design and then leveraging its programmability to further optimize your design’s power consumption. Your design choices can affect both static and dynamic power consumption.

Static power from leakage is proportional to the amount of logic and, hence, the number of transistors used to construct a given FPGA. Thus, if you reduce the number of FPGA resources you are using, you can probably implement your design in a smaller device, which lowers your leakage power. The effect of moving to the next-smaller device is shown in Figure 5.

To reduce the size of your design, you can employ several techniques, starting with time slicing of logic functions. That is, if two sets of circuitry perform a linear set of functions in the FPGA and are copies of each other, one may use one set of that circuitry, run it at twice the rate and multiplex

	Static Power Reduction Going to Smaller Device
330k → 220k	-33%
220k → 110k	-51%
110k → 85k	-24%
85k → 50k	-46%
50k → 30k	-33%

Figure 5 – Static power reduction with part-size reduction

the data going into a single instance of that circuitry. This uses half the logic.

Another way to reduce logic size is to use Xilinx’s unique partial-reconfiguration feature to replace sections of circuitry with new sections when only one section is needed at a time.

You can also move functions to nonlimiting available resources—state machines to BRAM, for example, or counters to DSP48 (Xilinx multiply, add, DSP block), registers to shift register logic and BRAM to lookuptable RAM (LUTRAM). You can also make sure that you are not overconstraining the timing for your design, which can duplicate logic/registers.

Also, you should take full advantage of the hard IP blocks (BRAM, DSP, FIFO, Ethernet MAC, PCI Express) we’ve implemented in the FPGA architecture.

Another way to reduce static power is to carefully audit your design and eliminate redundant DC consumers. Often, your design may employ blocks with extraneous or hidden DCMs or PLLs. That can happen if you are redesigning blocks and forget to remove them, or if you’re building a next-generation product with a bit of legacy code. Extracting the DCM or PLL to the top level of your design, allowing blocks to share the resources, will further reduce the size of the design and DC power.

Making wise use of memory blocks will also help reduce your FPGA design’s dynamic power consumption, and in turn its overall power consumption. Since dynamic power is a function of capacitance (area or length) and frequency, you should examine the way your design accesses block memory and identify areas where you can optimize capacitance and frequency.

Xilinx FPGAs include two types of memory arrays. BRAM, which we provide in 18k or 36k bit size, is optimized for large memory blocks. LUTRAM is optimized for small granularity and is based on the lookup table in the FPGA. LUTRAM comes in units of 64 bits in Xilinx Virtex-5 FPGAs.

Of these two types, BRAM typically consumes more power. Its enable rate is typically the largest source of BRAM power consumption, while its toggle also contributes, but is secondary. Designers can

Xilinx FPGAs have some interesting features in the area of clock gating. For example, you can use the BUFGMUX clock buffer to have the FPGA shut off a global clock or dynamically select a slower one.

take a few actions to minimize BRAM power consumption. For example, you can enable the BRAM only during an active read or write cycle. You can also make sure to use LUTRAM instead of BRAM for small memory blocks, reserving the BRAM for larger ones. Additionally, you can try to use BRAM for multiple large blocks.

Another technique is to arrange the memory arrays to minimize area and maximize performance or to minimize power consumption. Figure 6 shows a 2k x 36-bit storage array optimized for speed and area. We formed it by using four 2k x 9-bit

The other half of Figure 6 shows Xilinx's Block Memory Generator, which allows you to build arbitrarily sized memory arrays and optimize them for speed or power. Figure 7 shows the Xilinx Power Estimator for this case, comparing the power consumption between N blocks running with a given enable rate to N blocks with an enable rate of N/4. The results show a 75 percent reduction in dynamic power.

Xilinx tools will help you pick the right memory array for the job. Consider two sets of memory storage areas needed for a

85 percent power savings over implementing them in BRAM. That's because we had a lot of wasted space in the BRAM and inefficiently used sixteen 18-kbit blocks to get 16 very small (64 x 32-bit) memories.

When we look at the power comparisons for the second case of 16 sets of 18k bit array, the XPE tool shows the opposite for large memory arrays (Figure 9). Implementing them in BRAM instead of in LUTRAM gives a 28 percent power savings, attributed to the many small-granularity objects that need to be turned on and interconnected.

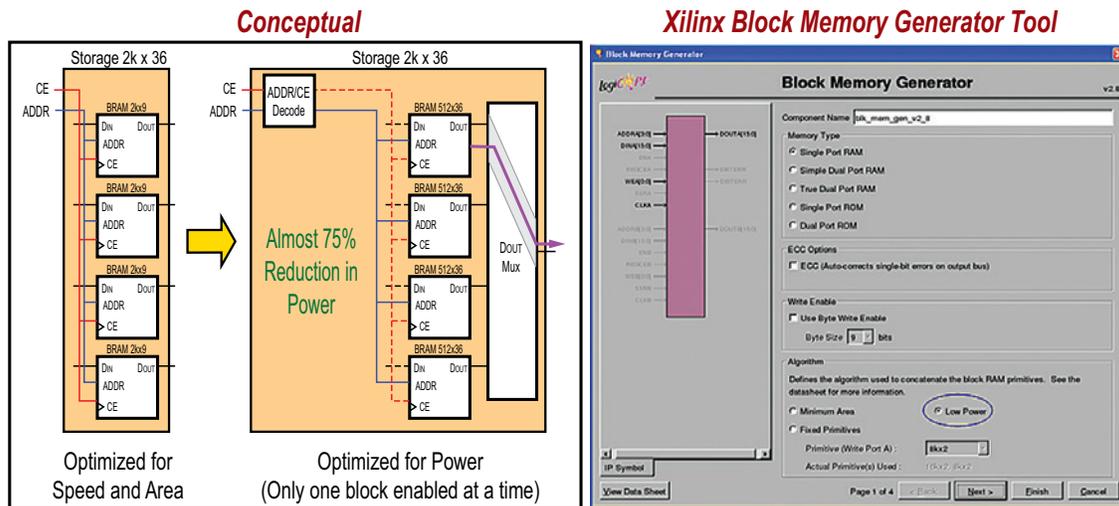


Figure 6 – Speed and area vs. power-optimized memory array (left) and Xilinx Block Memory Generator with power vs. area selections

blocks in parallel and always enabling all four blocks when a new value is needed. You can also employ another arrangement of 2k x 36 bits by constructing four 512 x 36-bit blocks, but decoding the lower two address bits to select which 512 x 36-bit block is being accessed. In the latter case, accessing no more than one memory block at a time reduces the power consumption by 75 percent compared with the first case.

design. In one case we need 16 sets of 64 x 32-bit memory structures (total bits = 32k) running at 300 MHz. In the other case we need 16 sets of 512 x 36-bit memory structures (total bits = 294k).

If we look at power comparisons for the 16 sets of 64 x 32-bit memory structures, what the XPE tool shows us (Figure 8) is that small memory arrays are best to implement in LUTRAM. Doing so delivers an

Xilinx FPGAs also have some interesting features in the area of clock gating. For example, you can use the BUFGMUX clock buffer to have the FPGA shut off a global clock or dynamically select a slower clock. You can also use the BUFGCE clock buffer to perform cycle-by-cycle clock gating in a manner very similar to the cycle-gating technique designers use for ASIC design.

You should consider using both features.

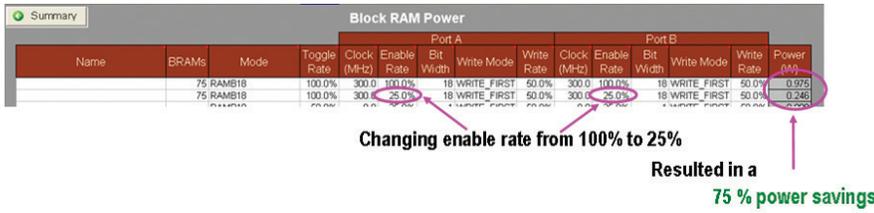


Figure 7 – XPE results of power-optimized array

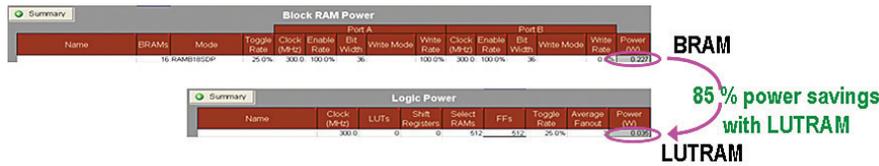


Figure 8 – Small memory power estimation using block RAM or LUTRAM

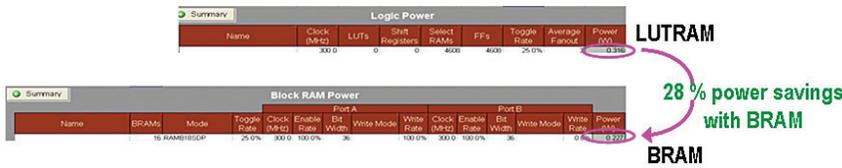


Figure 9 – Large memory power estimation using LUTRAM vs. block RAM

They are especially helpful in designs where certain blocks are not in use but contribute to power consumption. In those cases, you can turn off a very large clock domain with thousands of clock loads either on a clock-cycle-by-clock-cycle basis or for many, many clock cycles.

You can also rein in dynamic power consumption by reducing glitch energy. In designs that contain combinatorial logic and registers, occasionally various inputs to a block of combinatorial logic will arrive at slightly different times, generating short-duration glitches that can propagate to other structures and waste power (see Figure 10). By using more pipelining between layers of logic, you can block a glitch from propagating to other structures, which will reduce dynamic power.

Reducing Power at the Board Level

The PCB designer, mechanical engineer and system architect have several things to consider at the board level to reduce power consumption in the FPGA. Both the core voltage and the junction temperature of the FPGA have a strong influence on various components of power consumption.

Keeping control over V_{CCINT} core voltage is one way to reduce power consumption at the board level. Static power from leakage and dynamic power are both highly dependent on the core voltage of the FPGA.

Thus, one way to reduce leakage is to set the core voltage close to nominal (1 volt) rather than at the high end of the Virtex-5's operating range (1.05 V = +5 percent). With modern switching regulators you can achieve a voltage tolerance of ±1.5 percent

vs. the ±5 percent specification. Keeping the core voltage at the 1-V nominal rather than the 1.05-V maximum setting can reduce static power from leakage by 15 percent and dynamic power by 10 percent.

You can also reduce power consumption by keeping the junction temperature under control. The thermal properties of the FPGA, PCB, heat sink, ambient temperature, airflow and FPGA power for a given design all influence what the junction temperature of the FPGA will be.

One simple and somewhat obvious way to reduce the FPGA's junction temperature is simply to use a more thermally efficient PCB or heat sink. Then, any changes the FPGA designer can make to reduce power will be a bonus. At elevated junction temperature, such as 100°C, a reduction of 15°C will reduce static power consumption from leakage by 20 percent.

Another way to reduce power consumption is to monitor the temperature and voltage in an FPGA. The Virtex-5 FPGA includes an analog block called System Monitor, which monitors external or internal analog voltages and die temperature. System Monitor is wrapped around a 10-bit A/D converter, which can provide accurate and reliable results over a temperature range of -40°C to +125°C. The A/D converter digitizes the output of on-chip sensors; you can use it to monitor up to 17 external analog inputs to check environmental aspects of the system performance.

The block contains configurable thresholds and warning levels, and it stores the results of its measurements in configurable registers that easily interface to user logic or a microprocessor. Additionally, you can read the values via the JTAG port or even at power-up, prior to the FPGA being configured.

As power from the core voltage is driven down through transistor improvements, reduced capacitance and lower voltage, I/O power becomes another important consideration that you need to be aware of to balance power and performance. You can also reduce overall power consumption by making smart I/O choices. To make an informed choice, you have to consider each FPGA design's I/O interface requirements.

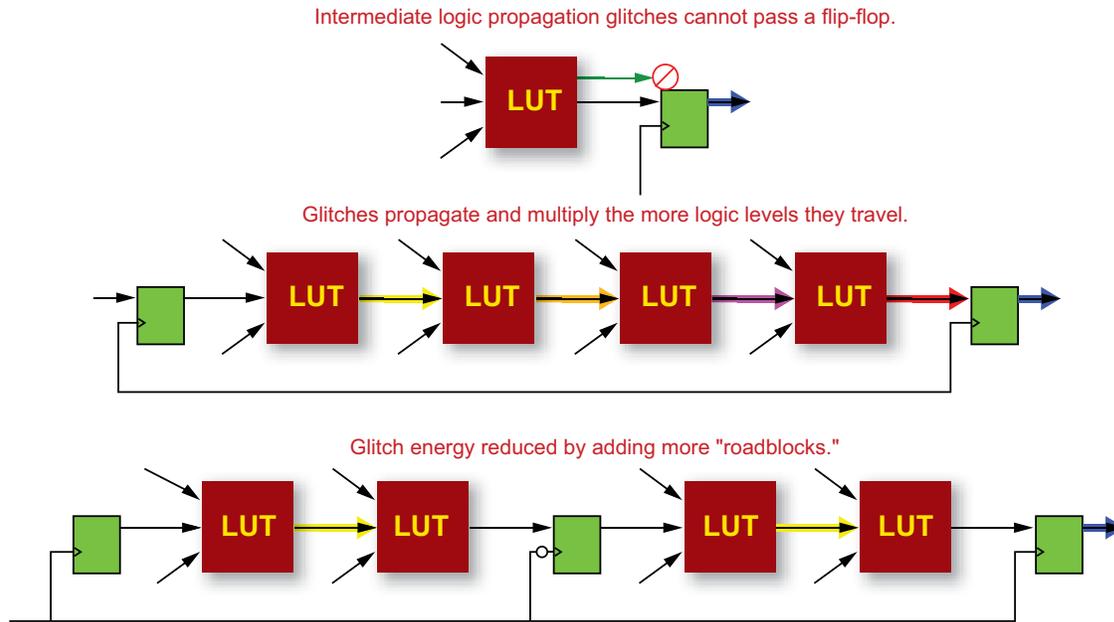


Figure 10 – Glitch propagation and blocking with inserted flip-flops

For example, interfacing to a memory (DDR2, QDR, RLDRAM, etc.), may require termination inside the FPGA for signal integrity, but it will typically consume more power and raise junction temperature.

Meanwhile, if you are interfacing an FPGA to an ASIC/ASSP, you must select the interface based on what the ASIC/ASSP target specifies (LVDS, HSTL, etc.). And if you are interfacing one FPGA to another FPGA, you may choose the interface based on the design's performance needs and may be able to more readily find opportunities for power optimization.

While both inputs and output consume power, the reference standards like LVDS, HSTL and SSTL consume the most. For outputs, the higher-drive-strength standards consume the most power and the power varies linearly with output enable rates and toggle rates. However, LVDS is an exception in that it is based on a fixed current source, which is independent of toggle rate.

For inputs, the referenced standards consume a lot of power because their receive structure incorporates a differential receiver and also because they include a selectable internal termination. Both consume DC power.

A feature called T_DCI (dynamically three-statable digitally controlled imped-

ance) in the Virtex-5 allows the user to dynamically remove the termination when a given I/O pad is used as an output. This is useful for the data bus or memory interface, and depending on the read vs. write ratio, it can rein in a fair amount of power (see Figure 4).

When selecting I/O interfaces, wise choices that balance performance and power are important. You should use interfaces like LVDS when your design needs absolute maximum performance and minimum noise or the target device requires the I/O standard.

Because termination generally consumes a large amount of power, you need to use it wisely and take into account the balance of power and performance. Schemes that use external termination or no termination can greatly reduce power.

Yesterday, Today and Tomorrow

Ever since power management began to loom as a big issue, Xilinx has been diligently building power-optimization technologies into tools throughout our ISE suite. For example, in addition to launching XPE and XPower Analyzer, a few years ago we gave ISE a power-optimized router that works based on known capacitance of routing resources inside the FPGA.

Also, you can configure ISE's power-optimized synthesis engine to automatically locate small arrays in the source code and synthesize them into LUTRAM. At your command, the engine will locate large arrays (in a size you specify) and synthesize them into block RAM. If it finds a large counter, it can implement it in a DSP48 block. It can also make smart choices when replicating logic to ensure it implements only the optimal amount.

More recently, Xilinx has introduced an optimized placer that will group functions together to minimize routing distance and, hence, capacitance. A related set of tools called PlanAhead™ allows you to take hierarchical groups of logic and physically place them in rough areas inside the FPGA. This helps reduce capacitance and speed up routing time as well.

As Xilinx continues to pioneer technology on the latest process nodes, we anticipate that dynamic and static power will continue to pose challenges. But at the same time, we are diligently working not only to optimize our tools and methods for power management, but also to make a concerted effort to nip power problems in the bud—in silicon.

For more information on Xilinx power management, visit www.xilinx.com/power.



inrevium

by Tokyo Electron Device Ltd.



More FPGA Needs in More Applications

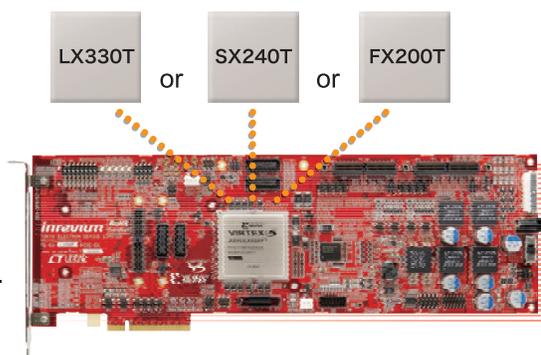
NEW! Virtex[®]-5 High Density PCI Express Platform

Tokyo Electron Device Limited has released three inrevium Virtex-5 High-Density PCI Express Platforms.

These PCI Express Gen 1 & 2 capable platforms utilize Xilinx Virtex-5 LX330T, SX240T and FX200T FPGAs, the highest density FPGAs available.

Expansion I/O connectors enable a wide variety of interfaces by connecting various optional boards.

In addition, the large-scale ASIC prototype development can be realized by a cable connecting of multiple FPGA boards. Available now throughout North America, Europe and Asia.



Jump-start your next FPGA design with the inrevium platform, visit at
<http://www.inrevium.jp/eng/x-fpga-board/>



TOKYO ELECTRON DEVICE LIMITED

World Headquarters

Yokohama East Square 1-4, Kinko-cho, Kanagawa-ku, Yokohama City, Kanagawa, 221-0056 JAPAN
Tel.+81-45-443-4016 E-mail.psd-sales@teldevice.co.jp

US office

2953 Bunker Hill Lane, Suite 300 Santa Clara, CA 95054, USA
Tel.+1-408-919-4772

Inrevium boards are available through:



Nu Horizons Electronics Corp.

Phone: +1-888-747-NUHO (6846)
URL <http://www.nuhorizons.com/x-fpga-board>



HiTech Global Design & Distribution, LLC

Phone: +1-408-781-8043
URL <http://www.hitechglobal.com>

Computer Interface Makes 19th-Century Pipe Organ Rock

How a group of engineers in Edinburgh, Scotland, used a Xilinx Spartan-3E Starter Kit to create a robotic organist.

by Gareth Edwards
Design Manager
Xilinx, Inc.
gareth.edwards@xilinx.com

It all started, as these things often do, with a conversation in the pub.

“Do you know the pipe organ upstairs in the Forest Café?”

“Yes.”

“We should build a robotic organist to play it.”

“Of course we should!”

And with that casual exchange, so began Project Waldflöte.

My day job is as a design manager in the IP group at Xilinx Scotland, but in my spare time, I’m also part of an informal movement called “dorkbot,” which promotes grassroots collaborations between the engineering-and-scientific community and the artistic community; its tongue-in-cheek motto is “people doing strange things with electricity.” I belong to the Edinburgh-based chapter (named either “dorkbot alba” or “dorkbot Edinburgh,” depending on who you are talking to). Members have, in the past, built a pixel-mapped LED top hat, a self-propelled toothbrush, a persistence-of-vision poi juggling device, a barely electromagnetic screwdriver and various noise-generating boxes. Injuries are surprisingly rare.

The dorkbot Edinburgh group meets every other Tuesday in the Forest Café, a volunteer-run nonprofit gathering place near the University of Edinburgh. I had been attending dorkbot workshops in the café for a few weeks when one night, I ventured upstairs to repair some stage lighting and, surprisingly, found myself in a church, complete with pulpit, balcony and, most important, a 16-foot pipe organ (Figure 1).

It turned out that the building the café occupies was once the meeting place of the Edinburgh Congregational Church—hence the organ. But this wasn't the instrument's original home. Initially installed in Dublin Castle in Ireland in the late 19th century by Gray and Davison, the famed London-based pipe organ builders, the organ, for reasons unknown, was moved to Edinburgh in 1900. There it has stayed, in varying states of repair, ever since.

So after the conversation in the pub, we leaped into inaction. Over the course of seven months' worth of Tuesday evenings, we pondered, poked, prodded and prototyped several ways of driving the organ's keyboard.

As for the name, we settled on “Project Waldflöte” for a stop on the organ called *Waldflöte*. It means “forest flute” in German, and since the organ is in the Forest Café, it seemed poetically apt.

Getting the Mechanics Right

It became obvious very early on in the development that we could partition the problem into a mechanical part and an electronic part. Once we had a solution to the mechanical problem, we could proceed with the construction of both parts relatively independently.

One of the main constraints was financial—we didn't have any real money to spend, just whatever our group of roughly a half-dozen core people was willing to throw into the pot. Scouring the surplus market, we found some solenoids that looked like they could fit the bill. We could get a hundred of them for about a pound each (roughly \$1.50), so we ordered half a dozen to play with on the organ itself.



PHOTO: MARTIN LING

Figure 1 - The Forest Café pipe organ in Edinburgh, Scotland, was built by the London firm of Gray and Davison in the late 19th century.

What we found was that the size of the solenoids was ideal, but the travel of the core was a bit less than we would need for consistent triggering of the organ's white keys. Although we could drive the black keys directly with the solenoid core, we would need some kind of lever for the white ones.

You can see the first prototype of the solenoid assembly in Figure 2 and a dia-

gram of how it works in Figure 3. For the white keys, the top plywood lever is hinged at the back with a piece of duct tape, and is pulled down when the solenoid is energized. When the solenoid is released, the organ key itself provides the upward force—there is no need for an additional spring. For the black keys, a small pin protruding from the bottom of



Figure 2 - Prototype solenoid assembly

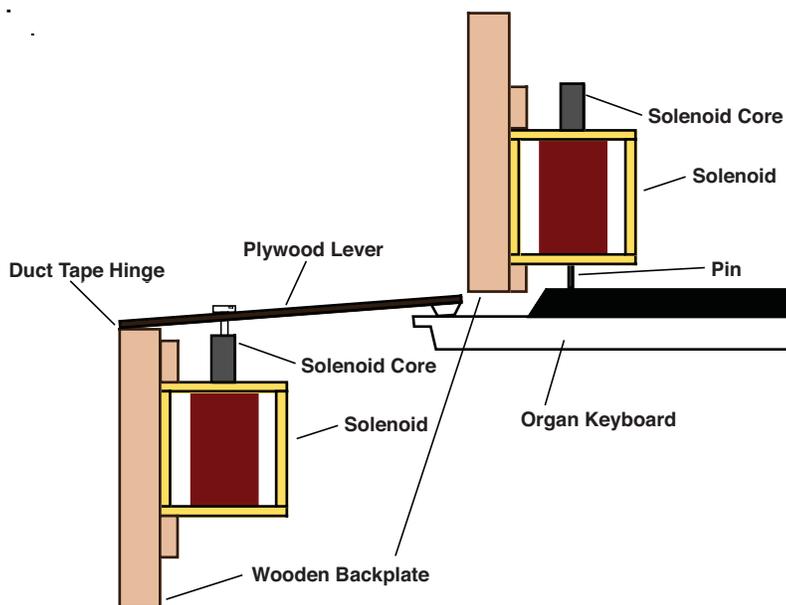


Figure 3 - Mechanical layout

the solenoid pushes down directly on the key with sufficient force and travel to sound the note.

Testing with this assembly showed that the keys were indeed successfully pressed. It also demonstrated that I am incapable of dividing by seven—the spacing on which I had placed the solenoids was not even close to matching the actual spacing of the octaves of the keyboard, so we could only test one key at a time. Still, we had proven the principle was sound, so we went ahead and ordered the parts for a full-length key rig and then started on the design of the electronics.

Designing the Electronics

At that point we sat down and roughed out the architecture; that basic diagram can be seen in Figure 4. On the left, MIDI messages arrive from the outside world (I'll talk more on the MIDI protocol below). On the right is a shift register chain; the controller toggles the “clock” signal while driving the appropriate “data” value to fill the shift register chain, then asserts the “strobe” signal to deliver the contents of the chain in parallel onto the solenoid driver inputs.

We implemented the shift register/driver chain with 74HC595 shift register ICs. However, the experiments with the sole-

noids had demonstrated that each one would require around a 350-milliamp drive from a 15-volt supply—way beyond what a CMOS output stage can deliver. To meet this requirement, we added a ULN2803A Darlington output stage to each shift register IC. This chip also has an integral protection diode to shunt the high flyback voltage that a solenoid can generate when the current is turned off, which saves adding a discrete diode to the layout. We built a few prototype driver boards on stripboard, each capable of driving 16 solenoids.

The Controller Design

Although there were a number of ways we could have implemented the controller (including on an Arduino platform or using some other microcontroller), we chose to do it on a Xilinx® Spartan®-3E Starter Kit, since, in my day job for Xilinx, I had access to the board and I knew the tool set inside out. In particular, I knew my way around the debug tools such as the Platform Studio SDK and ChipScope™, and since this was likely to be a project in which we debugged live, that would save time down the line. We used the Xilinx Embedded Development Kit to create the MicroBlaze™ subsystem that would be the heart of the design (Figure 5).

In addition to the MIDI interface and shift register interfaces, we chose to add a serial RS-232 console to help us debug the system. The RS-232 protocol might seem a bit old-school, but in this kind of project its presence can be invaluable. We also added some GPIO ports to drive LEDs and read switches and pushbuttons, to allow some interactivity without having to use the console.

Writing the MicroBlaze Firmware

We had decided that the best input interface for the system would be a MIDI port. The Musical Instrument Digital Interface has been, since the 1980s, the standard way to connect digitally controlled musical instruments such as synthesizers to other instruments or to a controlling computer, and it was therefore obvious that we should use it too. MIDI would give us maximum flexibility in the devices we could connect to the organ.

The MicroBlaze maintains an internal map of the state of the entire keyboard and which keys the system is pressing—that is, which solenoids the system is energizing.

MIDI is a unidirectional low-speed serial protocol, operating at 31250 baud. It consists of a variety of message types, but for our purposes, the only important ones are NOTE ON and NOTE OFF. Each NOTE ON message consists of three bytes.

Byte 1 is 0x9n, where n is the channel number.

Byte 2 is the note number from 0 to 127, where middle C is No. 60.

Byte 3 is the velocity value from 0 to 127.

NOTE OFF is very similar, except the first byte is 0x8n.

In our implementation, we decided to listen on all channels simultaneously

(known as “omni” operation). And since a pipe organ keyboard is not velocity-sensitive, we can safely ignore all the velocity bytes.

The EDK UART IP core receives the MIDI message bytes and presents them to the MicroBlaze processor one at a time through a FIFO. The MicroBlaze maintains an internal map of the state of the entire keyboard and which keys the system is currently pressing (that is, which solenoids the system is energizing). The firmware uses a static lookup table to figure out which solenoid is associated with the musical note in the event and uses this as an index into the map; the arrival of a

NOTE ON message sets the corresponding map entry to “1” and a NOTE OFF message sets the entry to “0.”

After the map is updated, the solenoid registers are refreshed with the entire contents of the map; by bit-banging on the GPIO port, the MicroBlaze processor writes the map one bit at a time onto the data input to the shift register and toggles the clock signal to move the shift register along by one position. Once the entire shift register has been updated with the map contents, the MicroBlaze then writes a rising edge onto the STROBE line, which copies the values in the shift registers into the output registers, energizing or de-energizing the correct solenoids to make beautiful music.

We implemented the firmware as a software state machine; for an embedded application without a real-time operating system, this can provide some of the capabilities of a multithreaded application but without the overhead of an actual thread implementation. A static array of structs describes, for each current state, what action the system should take for a particular event:

```
const midi_state_table_entry_t
MIDI_STATE_TABLE[] =
{
    {INHIBITED,PANIC,
     MidiSM_Panic,INHIBITED},
    {ANY_STATE,PANIC,
     MidiSM_Panic,INIT},
    {ANY_STATE,INHIBIT,
     MidiSM_DoNothing,INHIBITED},
    {ANY_STATE,OTHER_STATUS_RECEIVED,
     MidiSM_ClearMessage,INIT},
    {INIT,NOTE_ON_OR_OFF_RECEIVED,
     MidiSM_StoreStatusByte,NOTE_ON_OR_OFF},
    {INIT,DATA_RECEIVED,
     MidiSM_DoNothing,INIT},
    {NOTE_ON_OR_OFF,NOTE_ON_OR_OFF_RECEIVED,MidiSM_
```

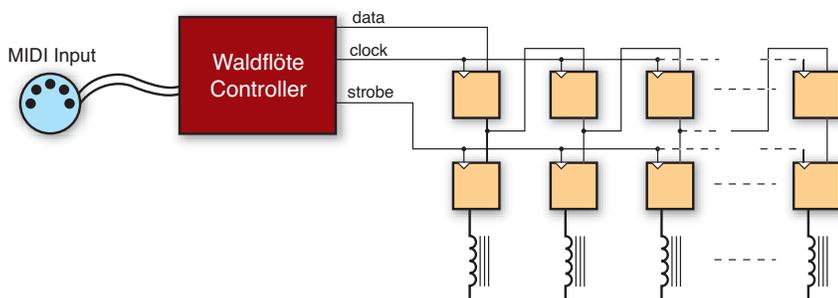


Figure 4 - Electronics architecture

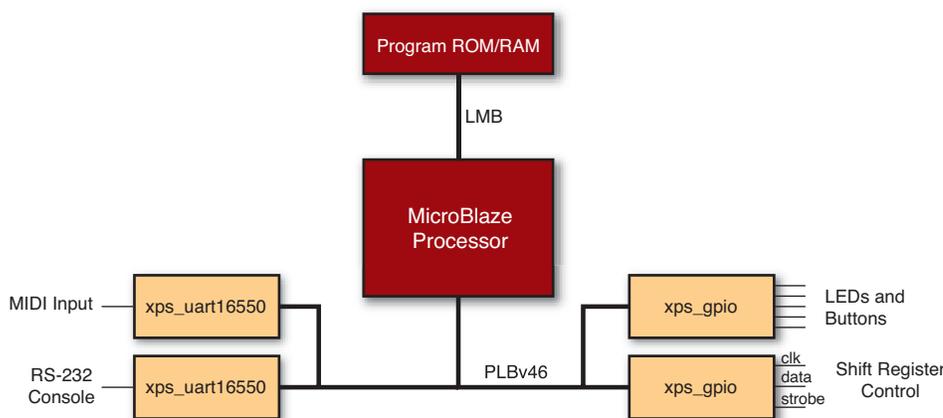


Figure 5 - MicroBlaze subsystem

We've successfully played some very complex and fast pieces of music, from classical to rock; there doesn't seem to be any serious limitation in the speed of the solenoids and drivers.

```

StoreStatusByte,NOTE_ON_OR_OFF},
{NOTE_ON_OR_OFF,DATA_RECEIVED,
 MidiSM_StoreNoteNumber,NOTE_ON_OR
_OFF_NUMBER},
{NOTE_ON_OR_OFF_NUMBER,
 NOTE_ON_OR_OFF_RECEIVED,MidiSM_St
oreStatusByte, NOTE_ON_OR_OFF},
{NOTE_ON_OR_OFF_NUMBER,
 DATA_RECEIVED,MidiSM
_NoteOnOrOffComplete,
NOTE_ON_OR_OFF},
{INHIBITED,ENABLE,
 MidiSM_DoNothing,INIT},
{LAST_STATE, LAST_EVENT, 0,
LAST_STATE},
};

The first entry in the struct is the cur-
rent state; the second entry, the event that
has arrived; the third entry, the state transi-
tion function needed to handle the event;
and the fourth entry, the next state.

The code that implements the business
end of the state machine looks something
like this:

XStatus MidiSM_
DoStateTransition
(midi_state_machine_t *pInstance,
u8 event)
{
    const midi_state_table_
    entry_t *pTable = pInstance-
    >pStateTable;

    // Search for a match in the
    state table
    do {
        if ((event == pTable-
        >received_event)
            && ((pInstance-
            >current_state == pTable-
            >state)
            || (pTable->state ==
            ANY_STATE)))
        {
            (*pTable-
            >transition_function)((v
            oid *)pInstance);
            pInstance->current_state
            = pTable->next_state;
            return XST_
            SUCCESS;
        }
        pTable++;
    } while (pTable->state !=
    LAST_STATE);

    // Aaargh, something bad happened
    - should never get here
    XASSERT_NONVOID_ALWAYS();
}

```

The event loop supplies an event as an argu- ment to this function and, depending on the current state and the event, some action is taken and the state of the system changed. The types of events include the arrival of bytes on the MIDI interface, the arrival of characters on the console and presses of the panic button. All experienced MIDI hackers know a panic button is a must-have feature to save your ears and your power supplies—it unconditionally turns off all the solenoids and returns the system to a known-safe state.

Waldflöte in Action

Figure 6 is a photograph of the organ with the contraption in place. Hiding the keys at the bottom are the solenoids' wooden back- planes—each plank has 30 or more sole- noids mounted on it, along with some recycled can capacitors to provide the ener- gy reservoirs for the solenoids. We C- clamped the whole driver assembly to the organ. At the top you can see the Spartan- 3E Starter Kit board and the interface strip-

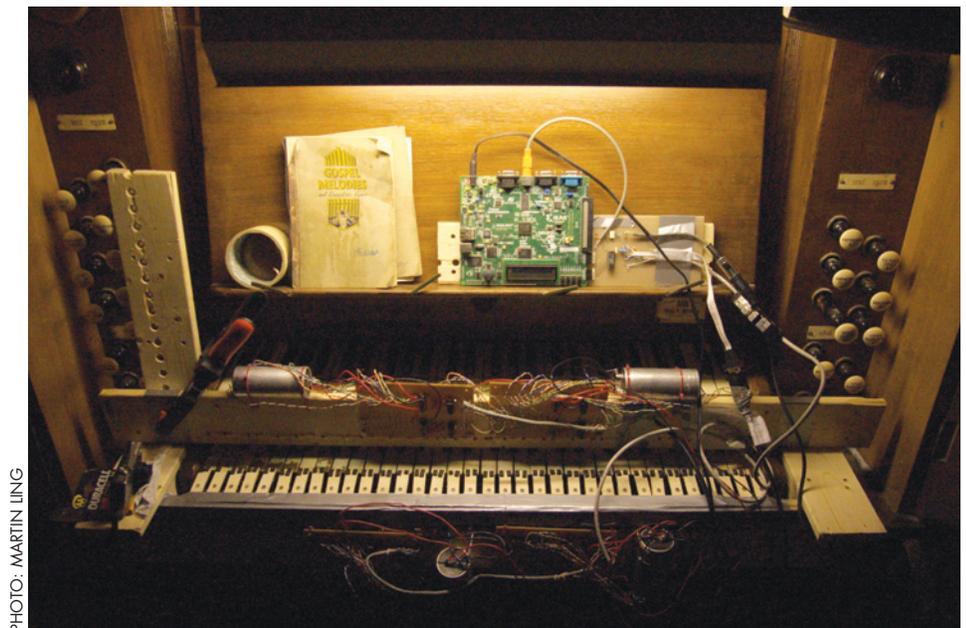


PHOTO: MARTIN LING

Figure 6 - The finished controller in place; the robotic organist plays everything from rhapsodies to rock.

board on its right; we wired these to the driver assemblies using recycled CAT5 cable.

It's hard to describe the operation of the organ in print, so I encourage you to follow the Internet link at the end of the article and take a look at the videos we've uploaded. The main thing you will notice is the click-clacking sound of the solenoids as the robot plays, be it "Moonlight Sonata" or "Jump"—this is the sound of the solenoid cores bottoming out in the coil, not of the actual levers hitting off the key. However, when you are down in the main hall of the café rather than standing up on the balcony where the organ is situated, the solenoid noise is much less noticeable. Only the majestic sound of the pipes dominates.

We've successfully played some very complex and fast pieces of music, from classical to rock, using the system; there doesn't seem to be any serious limitation in the speed of the solenoids and drivers. The solenoid power supply generally draws less than 4 A from the 15-V supply even in the most demanding pieces. And even though we are overdriving the solenoids slightly, there is no perceptible heating in the solenoid coils. All in all, we are very, very pleased with the system and proud to have been involved in its creation.

So what's next for Waldflöte? Well, we've informally invited some musicians to create compositions for the new instrument (particularly, composers who are excited by the idea of a 53-fingered performer that never tires), and we are thinking about staging a recital. Another possibility is to mechanize the operation of the stops of the organ, so we can have changes in volume and timbre during a performance under software control. We're also pondering ways of driving the bass pedals of the organ, which will bring the longest and lowest pipes into play. Finally, and possibly the most achievable option, we are considering putting a service onto the Internet that would allow the public to upload their own MIDI files to the system and have the audio from the organ streamed back to them in real time.

But then again, we might just go back down the pub.

To see and hear Waldflöte in action, visit <http://dorkbot.noodlefactory.co.uk/wiki/WaldFlote>. 

GET PUBLISHED



WOULD YOU LIKE TO WRITE FOR XCELL PUBLICATIONS?

It's easier than you think!

Submit an article draft for our Web-based or printed publications and we will assign an editor and a graphic artist to work with you to make your work look as good as possible.

For more information on this exciting and highly rewarding program, please contact:

Mike Santarini
 Publisher, Xcell Publications
xcell@xilinx.com



See all the new publications on our website.

www.xilinx.com/xcell

Hidden in Plain View

Xilinx offers a wealth of resources to simplify the prep work for your next design.



by Barrie Timpe
Field Applications Engineer, Southern Ohio Valley
Xilinx, Inc.

Designers are under the gun, pressured to do more in less time. All too often, customers and design management want it all, and they want it now. At the same time, silicon is growing more complex as each new device family piles on features. The FPGA is no longer a relatively simple array of logic blocks, a peripheral I/O ring and a centralized clock tree. As densities have grown, hardened silicon resources (block memories, DSP slices, advanced I/O blocks, multigigabit transceivers, PPC405/440 and the like) have added powerful features and performance capabilities to programmable logic devices.

All of these factors put an increased burden on designers. Luckily, Xilinx users have access to a wealth of resources that can reduce overall design effort and risk. It's all too easy to overlook some of the more important ones during these high-pressure development cycles, starting with the most basic: documentation.

User Guides Abounding

With previous generations of products (for example, the Virtex®-II), the datasheet was often the definitive introduction to the silicon capabilities, outlining detailed functional and performance characteristics associated with the various device elements. Due to the increasing complexity of newer families, however, the datasheets are now largely reserved for the associated DC and switching characteristics. It's up to user guides to fill the role of introducing the features and functions of the devices.

A dozen or more separate user guides address specific topics of interest in the current flagship Virtex®-5 FPGA, including configuration, system monitor, trimode Ethernet MAC, gigabit transceivers, PCI Express hard block, packaging/pinout, printed-circuit board design and so on. The more cost-optimized Spartan®-3 products have two user guides, one focused on the silicon capabilities and the other for device configuration. A third, optional guide explains the DSP48A slice unique to Spartan-3A DSP devices.

User guides are often the best place to start to familiarize yourself with a new device family. Don't overlook the bookmark index and search features; they will help you quickly navigate to specific areas of interest.

Deviations from the datasheet may give rise to errata. These typically apply to engineering-sample revisions, but may also affect production-step device revisions. As a designer, it is important to become familiar with the details and associated considerations outlined in the respective errata, including devices affected and how to identify them. Errata are posted on the Web site's Customer Notices page; users can also receive e-mail updates via MySupport Alerts. It is important to stay on top of the errata, since they

The user guides and schematics can be a great resource for peripheral interfacing, power supply reference or other features you want to include in an end-product design.

may be updated during the life cycle of the product. Meanwhile, Xilinx Change Notices detail updates to device availability, package or test locations, package materials, revised specifications and the like.

Application Notes

Xilinx Application Notes outline issues of a more specific slant in terms of actual design considerations. More than 1,000 of these documents now exist, with new ones coming out regularly and the existing documents getting occasional updates. Check the various categories and the associated titles to see if there's one you could leverage on a current or future design. Many application notes also have accompanying design files, such as RTL source or constraint files. Here are a few examples of the most popular application notes.

- XAPP058, XAPP424 and XAPP502 describe techniques for updating devices in the field.
- XAPP802 is an overview of approaches for memory interfacing, by device and memory technology.
- XAPP1052 and XAPP859 offer DMA/initiator example designs for the Virtex-5 PCI Express Block Plus core.
- XAPP485 and XAPP486 cover Spartan-3E LVDS serializer/deserializer designs. Versions are also included with the Spartan-3A Starter Kit reference designs.
- XAPP866, XAPP856, XAPP873, XAPP855 and XAPP860 describe Virtex-5 LVDS interfacing.
- XAPP514 is a collection of audio/video interfacing application notes. Contact your FAE for updated versions for the Virtex-5.

Application notes can be useful in providing information on how something is possible—often you know it is, but you may not know where to get started or simply don't have a handle on some of the considerations. Even if the application note does not fit your exact requirement, it can often serve as a baseline or may spark derivative ideas. The idea is to incorporate application notes and extend them, as appropriate, for a user's specific design requirements.

Development Boards

Development boards and kits are another crucial resource. They are most often used in the early stages of a design cycle with a new device family, but can be appropriate at other times as well. These boards and kits can function as a platform to evaluate project feasibility. You can use them as an initial development platform during your own hardware design or to gain familiarity with introductory designs and tool flows. However, they can also be helpful even when you have no apparent need for an evaluation board. For example, some of the "getting started" guides help with common tasks like downloading a design to the FPGA/PROM and generating or loading designs onto a System ACE™ CF card.

The user guides and schematics can also be a great resource for peripheral interfacing, power supply reference or other features you want to include in an end-product design. They come with reference designs that can bring designers up to speed on IP usage, configuration flexibility and so on.

A detailed review of reference designs can turn up a wealth of practical design nuggets. For example, the EDK TFT controller, mutex and mailbox cores first launched on

Virtex-4 reference designs (ML405 and ML410) before they made their way into the embedded development kit.

Designers can lean on other user guides, characterization reports and white papers as well. For those with rusty skills, the Programmable Logic Quick Start Guide (*ug500.pdf*) can be a good place to get started. Xilinx design tools (for example, ISE®, EDK, AccelDSP™ and System Generator) also include a number of example designs or tutorials with the installation directory. While there is no substitute for practical experience, the existing documentation provides a framework for novice users to understand the tool flow and options.

We also recommended reviewing the release notes and answer records for the tools and IP. A search of the Xilinx Web site with appropriate search criteria (for example, “MIG v2.3” or “MIG release notes” for our Memory Interface Generator tool) will yield related answer records of known issues, questions and recommendations. A number of useful “index” answer records can provide summary information (search for “master answer”). Of course, nothing stands still, especially technology. So, even after becoming familiar with the documentation necessary for your design, these resources will continue to change as new documents are introduced and existing ones are updated. But you won't need to comprehensively crawl through the public Xilinx site to stay current.

Fortunately, a convenient mechanism exists to notify users of changes. MySupport Alerts act almost like a Web-based “diff” utility in bringing important changes to your notice. Once you have an account on the Xilinx Web site, log in and then update your profile to include the alerts you would like to receive. The site offers filtering by type of alert, device family, board/kit and IP. As updates are published, you will get an e-mail with a list of changes, URLs and brief overviews of each change. You can drill down into the respective documents, each with its own revision history. Xilinx batches these alerts and typically issues them weekly.

Only a Mouse Click Away

The Xilinx User Community Forums are a Web-based arena for the exchange of ideas and questions relating to Xilinx silicon, tools and designs. Since its launch in July 2007, the site has drawn more than 6,500 registered users and likely many more “lurking” as readers.

This is a powerful and easy-to-use resource, organized by functional areas into “boards.” It is not a place to advertise products or services, however. Nor is it a replacement for closed-loop technical support (like Xilinx WebCases, a formal mechanism for receiving Xilinx technical support from customer application engineers). Although a number of Xilinx employees read these forums and actively participate (often after work), there is no guarantee you will get an answer. But by following a number of guidelines, you can increase the site's utility to you and the wider community.

The first tip is to review the datasheets and user guides before posting. Designers could find the answers to many of the questions posted by simply reading the appropriate documentation. In addition, search before posting. It's likely that others have addressed many of the same topics before. Also, provide the appropriate information so someone can help you. A rundown of what have you tried, what tools and IP you are using and the specific error message is much better than “it doesn't work.” In other words, keep your question specific; vague ones are hard to answer.

It's generally not realistic to expect a specific example with your exact configuration on a particular board. Many times an answer may guide you in the right direction, but it will still require some work to address your specific scenario. Be realistic. Don't expect someone to do your professional homework (or university assignment), or to download your design, review a lot of code and find an obscure bug. Post in the appropriate forum and do not cross-post the same question to multiple forums. Start a new thread if you are switching subjects rather than “hijacking” an existing thread. Make your post succinct and easy to access; nobody wants to download a word-pro-

cessing document or design archive if it isn't necessary. And if someone's response has been helpful, say “thank you.”

More important, post your own ideas and share your successes. Post follow-up resolutions to aid others who face similar issues. The idea is for the forum to become a place to share concepts with the community—not merely to serve as a help line.

Eric Steven Raymond, the open-source advocate and spokesman known to many simply as “ESR,” has his own useful guidelines in his article “How to Ask Questions the Smart Way.” While his suggestions were initially intended for usenet and IRC discussions, this content is appropriate to Web-based forums as well. Warning: Some audiences may view some of the content as politically incorrect. Nor does Xilinx as a corporation officially endorse it.

Help from Freeware

ADEPT—the acronym stands for Advanced Debugging/Editing/Planning Tool—is a freeware application that enjoys a cultlike following with many users and Xilinx technical employees because of its specific capabilities and ease of use. Xilinx itself does not officially distribute or support ADEPT.

Developed by Xilinx strategic applications engineer Jim Wu, ADEPT gets frequent updates to address new features and to support updated versions of ISE as they are released. Users must install the ISE Foundation or WebPACK™, since ADEPT accesses those databases to provide its functionality.

ADEPT supports Virtex-4, Virtex-5 and Spartan-3 Generation devices. It provides important insight into aspects such as footprint, clocking resources, transceiver location and attributes, and clock skew. It also offers useful insight into your design, including logic utilization (by hierarchical decomposition) and control set usage. ADEPT can help generate Orcad and Viewlogic (now Mentor) format schematic symbols, which is especially useful with our larger packages approaching 2,000 pins. After downloading the zip archive, simply install this 32-bit Windows application by extracting the zip archive and setting its environment

A Hidden Gem: PicoBlaze IP

The PicoBlaze 8-bit sequencer is a little bit like the bicycle stashed in your garage behind the sedan and the minivan. The capabilities and targeted application are different, but this venerable technology is an important resource to consider depending on where you are going and what baggage you are taking along.

PicoBlaze is the marketing umbrella for a family of 8-bit programmable sequencers developed by Ken Chapman of Xilinx UK. The latest version, kcpsm3, is targeted for Spartan-3 Generation, Virtex-4 and (unofficially) Virtex-5 FPGAs. There are other versions of PicoBlaze (Virtex/Spartan-II, Virtex-II and CoolRunner-II) as well. While there are slight differences in features and targeted device family support among them, all leverage a common lineage and share many similarities and design concepts.

In truth, PicoBlaze isn't often the focus of much marketing promotion. It has been around for a while, and its capabilities are often overshadowed by the performance and flexibility of its bigger siblings, the MicroBlaze and the PowerPCs, Xilinx's 32-bit embedded processing offerings. Yet, even for designs deploying the MicroBlaze or PowerPC, PicoBlaze can be a natural complement for specific areas of the design decomposition.

Technically, PicoBlaze is a full 8-bit microcontroller. However, that nomenclature usually sets the expectation that a graphical integrated development environment and C compiler are available. While third parties do supply such tools for PicoBlaze, they are not part of the reference design release; nor does Xilinx support them. "Programmable sequencer" better positions how most engineers see the PicoBlaze.

The PicoBlaze/kcpsm3 consists of an ALU, program counter stack (for nested subroutines), sixteen 8-bit registers, 64 bytes of scratchpad memory, program counter and control, and interrupt support circuitry. The control of the sequencer is captured in assembly language, and the assembler (a 32-bit Windows executable appropriately named kcpsm3.exe) will produce an RTL netlist of the sequencer instantiation and block RAM initialization. The design files include an HDL example, netlist (for schematic flows), UART and FIFO, testbench and JTAG boot loader. The documentation also includes a detailed user manual with an example design for an alarm clock. The Spartan-3E and Spartan-3A Starter Kit reference designs contain many more. System Generator also supports PicoBlaze.

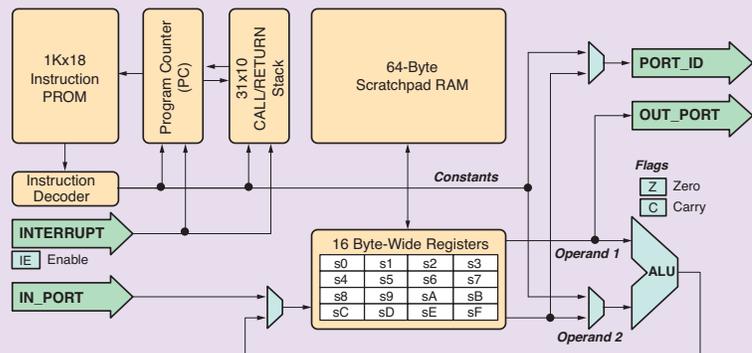


Figure 1 – PicoBlaze Embedded Microcontroller Block Diagram

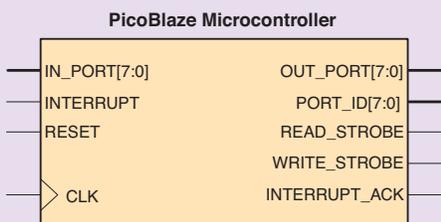


Figure 2 – PicoBlaze Interface Connections

As an embedded sequencer, the device is best suited for applications in which the control flexibility quickly outgrows expression in traditional HDL. PicoBlaze is a programmable state machine; its useful features include a fixed resource footprint (one 18-kbyte BRAM and 96 Spartan-3 slices) and a deterministic two clock cycles of execution time per instruction. PicoBlaze implements the instructions you would expect to find in an 8-bit microcontroller, including bit operations (compare, add, subtract, rotate/shift, test and/or/xor,) data movement and program execution (call, jump, return).

PicoBlaze includes an I/O space of 256 inputs and 256 outputs with an 8-bit address, 8-bit data input and 8-bit data bus. Since PicoBlaze is embedded in the FPGA, this I/O port is a key to interfacing with other logic in the FPGA and peripheral components. One common application is interfacing to user I/O, including LEDs (PWM), 2x16-character LCDs, switch inputs, rotary encoders or keyboard scanners. It is also well-suited for low-speed communications interfacing such as a serial port, SPI and I²C.

Other interfaced examples include the Maxim/Dallas DS2432 secure EEPROM. It is also possible to interface PicoBlaze to a fast simplex link as an intelligent coprocessor for a MicroBlaze soft processor. You could use PicoBlaze for basic calculations, such as an ADSR (attack, decay, sustain, release) envelope for an audio synthesizer.

In short, the PicoBlaze sequencer is a nice complement to the capabilities of traditional HDL design for FPGAs. The PicoBlaze user guide (*ug129.pdf*) provides guidance on the trade-offs with traditional FPGA design capture; it can be a helpful resource when decomposing a larger design and deciding on the applicable capture technique of particular functions or logical blocks. The PicoBlaze download pages (www.xilinx.com/picoblaze) are a great place to start. — *Barrie Timpe*

Clearly, the key to successful design is not just to understand the tasks, challenges and possible risks, but to be able to draw appropriately from a variety of resources.

variable to point to its location. A typical usage would be as follows:

- Launch the application.
- Select the appropriate family, device and package from the top subwindow.
- Load the device (this may take a few minutes, so be patient).
- Enable the appropriate display options via the View menu (this may also take a while).
- Optionally, load an existing ucf or ncd file.
- Sort by appropriate field (pin number, bank, type, etc.).

You can also export many of the views to a spreadsheet for design documentation and exchanges with other engineers.

Although there are other ways to perform this task, I often use ADEPT to investigate the resource utilization by module. I also find it helpful when correlating the number and locations of required IDELAYCTRLs on designs that utilize Virtex-4 IDELAYs or Virtex-5 IODELAYs, such as DDR2 memory interfaces.

For an introduction to this tool, Jim Wu's Web site (<http://home.comcast.net/~jimwu88/tools/adept/>) supplies the user guide, link to the download files, release notes history and descriptions of other possible functions (for example, how to port a PCI pinout). However, please do not contact Xilinx technical support for questions on ADEPT.

The Power Issue

Not long ago, power was a concern only for certain classes of applications, such as handheld or battery-powered designs. Now it is nearly always at the top of the list of design criteria, even with (or above) issues like performance, functionality and cost. Exacerbating factors include lower-

voltage supplies with higher associated current, larger number of supply rails for peripheral components, increased energy costs and continued trends to higher-efficiency distribution systems.

Few FPGA and board designers are experts in power supply systems. But trends like gate leakage in deep-submicron geometries now make it critical to understand the contributions of both static and dynamic power. Static power is the cover charge—the amount the device will dissipate even if your design is doing nothing. It is a function of device family, device size and junction temperature. Dynamic power is design-dependent, based on resources used, resource configuration and toggle rates. Understanding both types is important if you hope to gauge a realistic estimate of your power budget.

XPower is a Xilinx utility that can estimate power by extracting resource utilization from an existing design. However, it is most useful when the design nears its final form. The XPower Estimator spreadsheets (distributed per family) are a more appropriate resource earlier in the design cycle. You enter your estimated resource utilization (logic, clocking, memory, I/O, etc.) and see the resulting power. The XPower Estimator also automates thermal calculations based on the package, calculated power dissipation, ambient temperature and airflow. It's easy to quickly play "what-if" scenarios with various implementation approaches.

The Power Solutions page (www.xilinx.com/power) includes links to these resources, supporting documentation, Webcasts and application notes. It also has links to popular third-party power solutions providers, many of which have targeted literature and reference designs to aid in the power supply design for the board.

Other Learning Options

For those who need to go further, Xilinx Educational Services offers free online training and Webcasts to provide introductory

information on tools and devices. A number of other solutions are available for a fee. For example, Xilinx offers one- and two-day classes via an authorized training network. They provide more-detailed formal training on technology, devices and tools.

The company's Titanium Dedicated program, meanwhile, offers access to an applications engineer for a period of a week or more (either on-site or remote) to address specific areas of concern in your design. QuickStart is another popular option. It bundles two days of on-site training with three days of on-site consulting from a Titanium applications engineer.

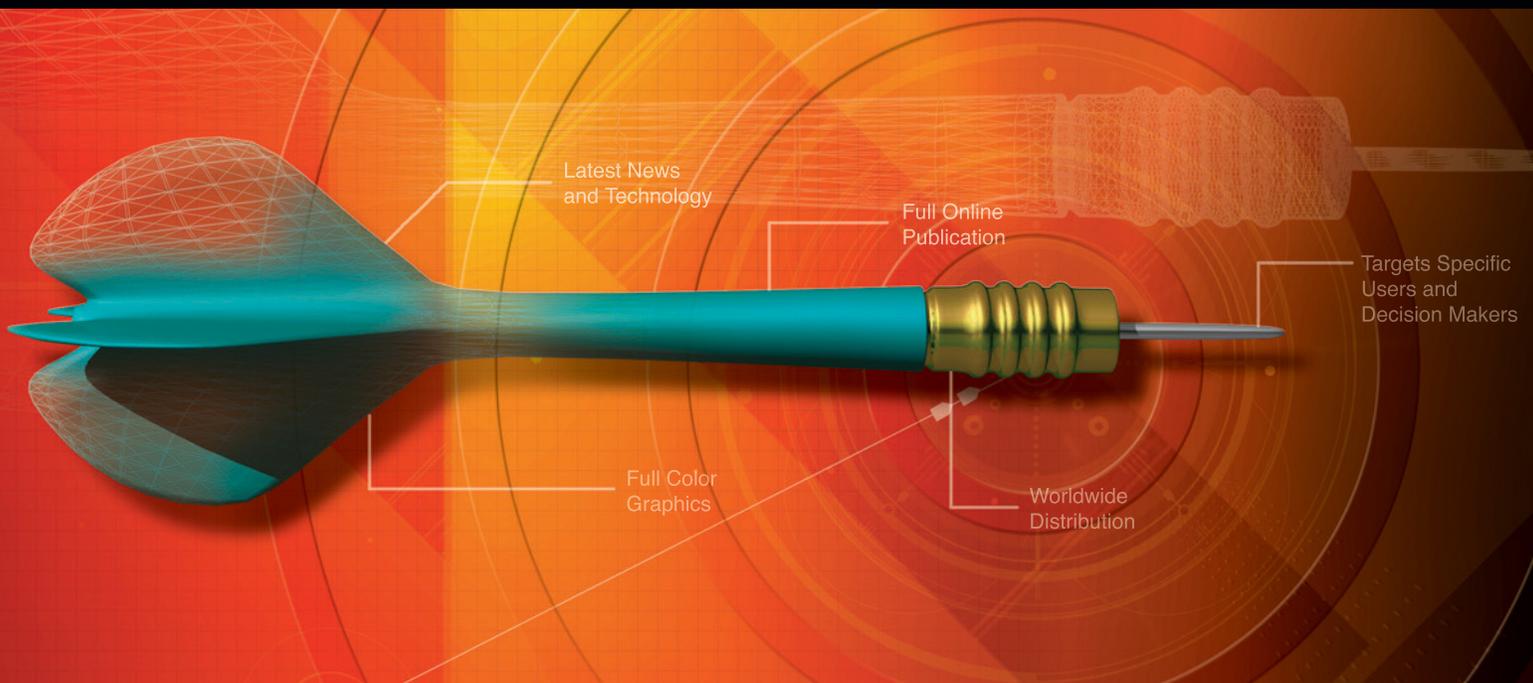
Beyond that, Xilinx Design Services is available for larger projects that require longer periods of engagement on design development and implementation. Please contact your distributor or sales representative for additional information on training, Titanium, QuickStart or XDS.

Clearly, the key to successful design is not just to understand the tasks, challenges and possible risks, but to be able to draw appropriately from a variety of resources. Xilinx continues to invest in providing an entire ecosystem of solutions beyond devices and corresponding tools. One or more of them will no doubt supply the help you need to make your next design a success. 🌈

Barrie Timpe is a Xilinx FAE supporting customers in the Southern Ohio Valley. Prior to joining Xilinx in 2005, he worked as an FAE supporting Xilinx customers with Memec Insight. His FPGA and hardware design experience is in the area of military wired-communication systems and I/O interfacing. Timpe spends his spare time with his wife and four young children. He also actively participates in the Xilinx User Community Forums.

If you need assistance, check in with your local FAE, contact Xilinx tech support at (800) 255-7778 or visit www.xilinx.com/support/clearexpress/websupport.htm.

GET ON TARGET



LET XCELL PUBLICATIONS HELP YOU GET YOUR MESSAGE OUT TO THOUSANDS OF PROGRAMMABLE LOGIC USERS.

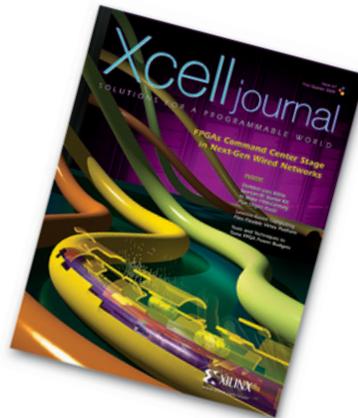
Hit your marketing target by advertising your product or service in the *Xilinx Xcell Journal*, you'll reach thousands of engineers, designers, and engineering managers worldwide!

The *Xilinx Xcell Journal* is an award-winning publication, dedicated specifically to helping programmable logic users – and it works.

We offer affordable advertising rates and a variety of advertisement sizes to meet any budget!

Call today :
(800) 493-5551
or e-mail us at
xcelladsales@aol.com

Join the other leaders in our industry
and advertise in the *Xcell Journal*!



www.xilinx.com/xcell/



Application Notes

If you want to do a bit more reading about how our FPGAs lend themselves to a broad number of applications, we recommend these notes.



XAPP1127: XPS LL Tri-Mode Ethernet MAC Performance with Monta Vista Linux

www.xilinx.com/support/documentation/application_notes/xapp1127.pdf

In this new application note, Brian Hill describes how to use the standard network performance suite Netperf to measure XPS LL Tri-Mode Ethernet MAC (TEMAC) performance with MontaVista Linux 4.0. The note discusses several of the tunable values that can affect Ethernet performance while introducing Netperf as a means of testing and measuring network performance. Netperf 2.4.4 source is included with this application note, along with prebuilt Linux and Cygwin images.

XAPP1126: Reference System: Designing an EDK Custom Peripheral with a LocalLink Interface

www.xilinx.com/support/documentation/application_notes/xapp1126.pdf

James Lucero describes the design of an EDK core with a LocalLink interface in this new application note. The note, which includes a reference design, shows you how to use the Create IP Wizard application to generate the PLB v4.6 slave interface and then add the LocalLink interface to the core. To create simple user logic with the LocalLink interface, you can take the payload data transmitted from the Hard DMA (HDMA) and loop it back to the payload on the receive channel. In the note, the XPS LL EXAMPLE core is connected to the Master PLB (MPLB) for the slave registers and the LocalLink interface is connected to the first HDMA block on the PowerPC® 440 processor block. The note shows how you can then use the ChipScope™ Analyzer to verify the core's LocalLink loopback functionality. A standalone software application is included to initiate a single DMA transaction. The receive channel is set up to accept and verify that the Rx and Tx payload data are the same. Lucero uses the Xilinx ML507 Rev. A board for this reference system.

XAPP1103: Simulation of the IEEE 802.16 CTC Encoder and Decoder

www.xilinx.com/support/documentation/application_notes/xapp1103.pdf

In this application note, Michael Francis and Raied Mazahreh describe how to simulate the LogiCORE™ IP IEEE 802.16e CTC Encoder and IEEE 802.16e CTC Decoder together using either Mentor Graphics' ModelSim simulator or hardware-in-the-loop based on Xilinx System Generator software. Simulation using ModelSim gives users a better understanding of the operation of the IP and the design's various control signals, while hardware-in-the-loop provides users with bit-error rate (BER) performance measurements.

The authors describe their use of a noisy-channel model to test the encoder/decoder combination under a variety of additive white Gaussian noise (AWGN) conditions using the AWGN module. They use a combination of Simulink® blocks and MATLAB® scripts to control the hardware, and collect and display the results. This application note does not deal with MATLAB simulation; it's covered in more detail in ug497, "LogiCORE IP 802.16e CTC Decoder v3.0 Bit-Accurate MATLAB Model User Guide" (www.xilinx.com/member/ieee802_16e_ctc_dec_eval/ctc_decoder_v3_0_bitacc_cmodel Ug497.pdf).

XAPP1113: Designing Efficient Digital Up- and Down-Converters for Narrowband Systems

www.xilinx.com/support/documentation/application_notes/xapp1113.pdf

This new application note by Stephen Creaney and Igor Kostarnov demonstrates how you can create efficient digital upconverter and downconverter (DUC/DDC) implementations by leveraging Xilinx's DSP tools and IP portfolio. Digital upconverters and digital downconverters are key components of RF systems in communications, sensing and imaging. While previous

application notes, including XAPP1018 (www.xilinx.com/support/documentation/application_notes/xapp1018.pdf), have provided examples of DUC and DDC implementation in wide-band communications systems, this document concentrates on narrowband systems and the building-block components available to meet their particular requirements.

The note provides you with step-by-step guidance on how to perform simulation of narrowband DUC/DDC systems in MATLAB, how to map functions onto building blocks and IP cores for Xilinx FPGAs with Xilinx's System Generator software and how to verify the implementation against the simulation model. The note includes two examples: a multicarrier GSM system (both DUC and DDC) and a multichannel MRI receiver (DDC only). The examples provide a guide and template for implementation of your own application systems. The authors also outline the advantages and limitations of the approaches and methods they cite, so that you can better tailor these schemes for your particular system design.

XAPP468: Fail-safe MultiBoot Reference Design

www.xilinx.com/support/documentation/application_notes/xapp468.pdf

Jim Wesselkamper in this application note describes a reference design that adds fail-safe mechanisms to the multiboot capabilities of the Extended Spartan®-3A family of FPGAs (Spartan-3A, Spartan-3AN and Spartan-3A DSP platforms). The reference design configures specific FPGA logic via an initial bitstream that determines which application (by alternate bitstreams) to load. The decision is based on the bitstream revision, the number of prior configuration attempts and the integrity of the alternate bitstreams. Users implement the algorithms that test bitstream integrity and which bitstream image to load using a PicoBlaze™ controller. Additional independent modules manage communication with the Internal Configuration Access Port (ICAP) and the Serial Peripheral Interface (SPI) flash device.

XAPP1121: Reference System: Optimizing Performance in PowerPC 440 Processor Systems

www.xilinx.com/support/documentation/application_notes/xapp1121.pdf

In this reference system, James Lucero demonstrates how to improve system performance in the PowerPC 440 processor block on the Virtex®-5 FXT FPGA. The reference system note describes how to connect the XPS Central DMA master interface to the Processor Local Bus (PLB) v4.6 on either PLB Slave 0 (SPLB0) or PLB Slave 1 (SPLB1). You can then modify the parameters for the XPS Central DMA for the DMA engine to allow for parallel reads and writes through the crossbar. For HDMA, the paper discusses setting threshold values for interrupts and changing addresses in main memory for buffer descriptors and transmit/receive buffers. A simple loopback core is connected to the LocalLink interface on one HDMA.

In addition, Xilinx included performance cores for PLB v4.6 master interfaces and HDMA in the system to measure system performance before and after optimizations. The note includes two standalone software applications to demonstrate DMA transactions for XPS Central DMA and HDMA to DDR2. During these DMA transactions, you can measure the core's performance. For XPS Central DMA, the performance is measured between PLB_PaValid and when the XPS Central DMA provides an interrupt (DMA transactions are complete). For HDMA, the Tx and Rx channels are measured between the first frame's start of frame and the last frame's end of frame.

Further, this application note describes how to obtain latency numbers across the crossbar and obtain performance numbers before optimizing the system/software applications. It also details how to optimize the system/software application for performance and obtain performance numbers with the optimized system. The reference system uses the Xilinx ML507 Rev. A board.

XAPP1043: Measuring Treck TCP/IP Performance Using the XPS LocalLink TEMAC in an Embedded Processor System

www.xilinx.com/support/documentation/application_notes/xapp1043.pdf

In this application note, Doug Gibbs illustrates how to measure the network performance of the XPS Local-Link Tri-Mode Ethernet MAC (TEMAC) in an embedded-processor system running the Treck TCP/IP stack. The measurements use a PPC405 processor-based system with the ML405 Evaluation Platform or a MicroBlaze-based system with the ML505 Evaluation Platform. Test software and methods are provided so designers can perform similar tests using different boards. This application note outlines how to acquire the hardware designs and set up the two boards. It details the setup of the Treck TCP/IP software along with how to integrate it with a standalone test application. Test software that can be run on a Windows or Linux PC is also discussed.

XAPP1106 (Updated): Using and Creating Flash Files for the MicroBlaze Development Kit—Spartan-3A DSP 1800A Starter Platform

www.xilinx.com/support/documentation/application_notes/xapp1106.pdf

This application note describes the files for programming the serial flash memory and the StrataFlash memory for the MicroBlaze Development Kit—Spartan-3A DSP 1800A Starter Platform. The reference system executes the HelloWorld software application to run from the serial flash using SPI configuration mode and the BlueCat Linux image from the StrataFlash in BPI configuration mode. The development kit includes all the files you need to run both. This application note describes how to use the provided files, as well as how to create new files and run the reference system successfully from the flash memories. Xilinx recently updated this app note to be in line with the ISE® software 10.1.3 update and added JFFS2 root file system information. 

Mark Moshayedi Drives STEC Enterprise Storage to Greener Pastures

The company and its president and COO seem to have timed the enterprise solid-state disk drive market to perfection.



by Mike Santarini
Publisher, *Xcell Journal*
Xilinx, Inc.
mike.santarini@xilinx.com

Spotting the next big opportunity in electronics and moving quickly to capitalize on it is a skill not many possess. But one who does is Mark Moshayedi, the 46-year-old president and chief operating officer of solid-state disk drive (SSD) OEM STEC Inc. The Santa Ana, Calif., company, which he founded with his two older brothers, has reinvented itself a few times to capitalize on new opportunities. Today, STEC is a public company offering a faster, lower-power replacement for mechanical hard-disk drives at a time when power consumption is the hot topic in electronics and the green movement is on the precipice of becoming the next big economic driver for the U.S. economy. STEC's flagship product, Zeus-IOPS, is built around Xilinx® Virtex® FPGA devices.

"SSDs offer a 97 percent power savings over mechanical hard-disk drives—one SSD can do the work of 30 HDs," said Moshayedi. "Increasingly, customers aren't as concerned with the unit cost as much as the total cost of the product over its life-time." After you buy the equipment, "you must have the space to run it, you have to spend money to power it and then you have to spend more money to cool it," Moshayedi explained. "Today, the purchase price of SSDs is higher than that of mechanical disk drives, but the long-term cost to the user is much less. And SSDs are more reliable."

STEC uses Xilinx Virtex FPGA devices in its drives to monitor bit integrity across 16 channels of flash in the system.

STEC Inc. is a growing company seemingly equipped with the right technology at the right time. But Moshayedi didn't propel it there out of luck. Indeed, he worked hard and overcame many obstacles to get to where he is today, wearing many hats along the way.

In 1982, Moshayedi received his BSEE from the University of California, Irvine, and went to work as a design engineer for Texas Instruments. Two and a half years later, he joined Fujitsu as a field applications engineer and took night classes to get an MBA from Pepperdine. He leveraged that know-how to go into sales, becoming an area sales manager for Sony Semiconductor. In the late 1980s, some of his friends started a rep firm called Westar Rep Co., which became the exclusive representative for Samsung Semiconductor in Southern California. Moshayedi joined them.

"In those days, Samsung was an up-and-coming company and we were mainly selling DRAM for them," said Moshayedi. "A lot of our customers back then were DRAM memory module manufacturers like Kingston, Southland Micro and Viking Components. Those companies did a tremendous business."

Moshayedi said it was around that time when his two brothers, Mike and Manouch, both of whom are civil engineers, sought his advice for starting a new company. "I told them they should look into building and selling memory modules," he recalled. His siblings liked the idea. "I remember we got our start by buying a thousand 1-Gbyte memory modules for \$60 apiece and sold them later the same day for \$64 apiece. They said, 'hey, this isn't a bad gig,'" Moshayedi recalled.

That exercise kicked off Simple Technology Inc. as a business. After 18 months, the fledgling company started to break even, so Moshayedi left Westar Rep to head up engineering, product development and manufacturing. Manouch Moshayedi became the company's CEO, the same job he now holds at STEC, while brother Mike served as its president and has since retired.

In 1992, Mark Moshayedi established an engineering department and the compa-

ny started manufacturing its own memory modules. Then, in 1994, the company began developing CompactFlash cards.

"At the time, the only other supplier was SanDisk, so we started developing our own flash controller and working with a few other controller manufacturers," Moshayedi said. "One in particular was Cirrus Logic, which was doing solid-state controllers. And we ended up in 1995 buying that [controller] group from Cirrus Logic and naming it Lexar Media."

But in 1995-96, the memory market went into free-fall, and soon Simple Technology found that it had to spin Lexar out. "We sold it to a group of investors and it became its own entity," said Moshayedi. The company then signed a deal with Hitachi, to become its exclusive manufacturer of solid-state disk drives in North America. "We would design the products in conjunction with Hitachi and sell them to all the customers," he said. But in 2000, Hitachi opted out of the SSD business, and Simple Technology began developing its own controllers. Simple Technology also made its public offering in September of that year.

By 2002, the company had developed a full line of FPGA- and ASIC-based SSDs and began selling them to the market. In late 2004, Moshayedi said the company saw an opportunity to sell even bigger drives to the enterprise market. Thus began development work on its Zeus line, tailored for enterprise servers.

By 2005, Simple Technology owned 4 percent of the \$250 million SSD market. Also around this time, the company sold off its retail memory module, flash card and portable hard-drive business—along with the name Simple Technology—to Fabrik, to focus exclusively on its SSD OEM business. With the sale came the corporate

name change to STEC Inc. It was a bit of a gamble, but seemingly one that has paid off.

In late 2006, the company began work on its Zeus-IOPS product line for the enterprise market. These devices boast vast performance improvements over mechanical disk drives. STEC announced the Zeus-IOPS first running in a product line with customer EMC in January 2008.

"The product caught on like wildfire and immediately turned on all the other customers in the world to look at using solid-state disk drives," Moshayedi said. "HP, Dell, Sun...several companies in the storage area immediately engaged with us to design-in our product into their product lines."

The offering is indeed impressive. "Our Zeus-IOPS product line replaces 30 drives," Moshayedi said. "It not only gets the same performance you get from 30 HDDs from an I/O perspective, but you also cut down latencies to less than a millisecond, from 10 to 15 milliseconds. At the same time, it cuts down power consumption by 97 percent."

One of the main concerns of flash-based SSDs is that after many tens of thousands of writes and reads, the memory bit-error rates increase. STEC uses Xilinx Virtex FPGA devices in its drives to, among other things, monitor bit integrity across 16 channels of flash in the system.

"The Zeus-IOPS lineup uses the Virtex-4 FX60, which has the integrated PowerPC processor," said Moshayedi. "It provides throughput of roughly 300 Mbps for reads and writes. But more importantly, it gives us 50,000 I/Os per second out of a single drive. The Fibre Channel interface is done by the Xilinx chip; then, because it has the PowerPC on board, it does all the ware leveling and all the flash block management—everything is done in that controller."

Moshayedi said that STEC isn't resting on its laurels. The company, which employs 800 people, followed up on the success of its Zeus-IOPS introduction with the release of a 3-Gbyte serial-attached SCSI version in November. It expects to roll out a 6-Gbyte SAS model in 2009. "Those will also be first-to-market products that use Xilinx chips," said Moshayedi.

In Moshayedi's view, the market for enterprise-class SSDs is just now starting to take form.

"I believe this is a disruptive technology that will drive a huge change in the disk drive industry," he said. "It used to be that enterprise-level systems would use one tier of high-speed disk drives to get I/O and then use a second tier of slower drives for capacity. Now we see that second tier completely going away."

STEC is forecasting that by 2013, "nobody is going to be manufacturing 15k RPM drives any longer, and everything is going to be solid state," Moshayedi said. "The enterprise disk drive market today is about a 30 million-units-per-year market. If you assume one-to-one replacement, without any growth in the market—and that one of our drives can do the work of these 30 disk drives—we have an opportunity to sell 1 million new units per year." That prognosis, he added, is "assuming there is no growth in the market: we believe it will indeed grow."

In fact, Moshayedi's estimates seem somewhat conservative considering the fact that many companies are trying to improve their "green" profiles and are con-

stantly looking for ways to reduce operating expenditures. The Zeus-IOPS' 97 percent power reduction means the equipment doesn't generate nearly as much heat, delivering a savings in cooling costs.

Ten years ago, \$50 billion worth of enterprise equipment sold every year, consuming \$1 billion worth of energy. Today, the same amount of enterprise equipment consumes \$5 billion worth of energy. "This growth in energy cost could be addressed by using SSDs in these systems," Moshayedi said.

Time will tell if the market lives up to, or even exceeds, its potential. Certainly, you can be sure that if it does, Moshayedi will have had a hand in its success. If it doesn't, he'll move on to the next hot technology. ●●●

The widest selection

of COTS board-level solutions using Virtex-5 FPGAs

FPE650 - 6U VPX Quad
FPGA Processor Board



AD3000/AD1500 -
3.0GSPS or dual 1.5GSPS
A/D PMC/XMC



PMC-FPGA05 -
Virtex-5 LX110 PMC Board



HPE640 & HPE720
- 6U VPX PowerPC & FPGA
Processor Boards



VPF2 - 6U VXS PowerPC &
FPGA Processor Board



MM-6171 - XMC Buffer
Memory Node



Embedded Computing

Data Recorders & Storage

Bus/Protocol Analyzers

User programmable Xilinx® Virtex® -5 FPGA signal processors and analog, digital and fiber-optic I/O
A new generation of performance

Single or multiple FPGA solutions
Simple solutions for complex tasks

PMC, XMC, VXS and VPX form factors
Flexibility based on open standards

Commercial and Rugged variants
Easily migrate from development to deployed systems

Libraries and Example Code
Easy to use with head-start time-to-market



For more information, please visit
<http://www.vmetro.com/virtex-5> or call (281) 584-0728



Xilinx Global Services. Finish. **FASTER!**



Xilinx Global Services offers comprehensive service solutions that help you design more efficiently and cost-effectively. Seize the competitive advantage today with services designed to lower your costs, reduce speed grades, decrease your time-to-market, and close the gap between consumable and consumed technology.

Xilinx Services Portfolio

Education Services – High quality, comprehensive courses to give you a competitive edge

Titanium Dedicated Engineering – An engineer dedicated to your team on your site or remote, provides the design help you need to get to market faster

QuickStart! – Combines an on-site, customized education course with a dedicated engineer to coach you through your design and increase your design productivity

Xilinx Productivity Advantage Program – A bundled solution offering a one-stop opportunity to purchase all of the tools you need to get it right the first time

To learn more about Xilinx Global Services, we invite you to visit: www.xilinx.com/services

 **XILINX**[®]
www.xilinx.com

Tools of Xcellence

News and the latest products from Xilinx partners

by Mike Santarini
Publisher, Xcell Journal
Xilinx, Inc.
mike.santarini@xilinx.com



Omiino Pioneers 'Virtual ASSP' Business Model Running on Xilinx FPGAs

A startup based in Belfast, Northern Ireland, is on a mission to replace ASSPs in the carrier optical-transport market with FPGA-based designs offered through a pioneering business model. With the "Virtual ASSP," the company provides its full chip designs to customers as completed FPGA netlists ready for implementation in various Xilinx® Spartan®-3 and Virtex®-5 FPGAs as a cost-effective alternative to application-specific standard products.

Founded in 2007 by a team drawn from tier-one equipment and ASSP manufacturers, the 10-person Omiino is currently delivering full Ethernet-over-Sonet/SDH devices. Its 2009 road map calls for the launch of Sonet/Ethernet-over-OTN framers at 10G, 40G and beyond.

"FPGAs are clearly powerful enough to replace ASSPs on boards," said Gary Hamilton, the company's CEO. "There isn't a tier-one or tier-two equipment vendor in the carrier space who isn't doing it or considering doing it, mainly because both they and the ASSP vendors struggle to justify new 65-nanometer ASIC and ASSP designs, given the considerable costs. FPGAs also promise faster design cycles and the ability to field-upgrade, which is invaluable when the standards are evolving so rapidly."

But especially in a tough economy, "the real challenge for the equipment manufacturers in transitioning from ASSP- to FPGA-based designs is to stop R&D budgets from expanding and to lower product costs," Hamilton said. "The lack of off-the-shelf third-party ASSP replacements in FPGA means that equipment vendors are having to grow internal capability and are therefore increasing their R&D budgets." That's why "the resulting FPGA designs are generally bespoke solutions"—highly integrated, but expensive and power-hungry designs, compared with the ASSPs they are replacing, he said. "This is not a scalable business model."

To address these issues, Omiino has developed a suite of "ultracompact" intellectual property (IP) which it will integrate and verify to deliver fully functional, high-performance Virtual ASSPs targeted at very cost-effective Xilinx FPGAs. Hamilton said Omiino designs consume 60 percent to 80 percent fewer FPGA resources than conventional ASSP-targeted IP.

"We can give customers a lower product cost and do it in the familiar ASSP business model, which is why we call our offerings Virtual ASSPs," he said.

Implementing the designs cost-effectively in FPGAs rather than ASSPs or even structured ASICs reduces the delivery time and preserves the reconfigurability, Hamilton said. Retaining reconfigurability is important, especially when a standard isn't quite solidified or changes frequently—a common occurrence in telecom. "We say retain the programmability, allowing customers to go to market early, and then do field upgrades if they are needed," said Hamilton.

Today, the company has three main offerings: the OM1312, OM1411 and OM1412 Virtual ASSPs. It describes each of these devices as highly integrated and ultracompact.

The OM1312 is a 2.5-Gbit/second, high- and low-order Ethernet-over-Sonet/SDH (EOS) packet mapper targeted to the Xilinx Virtex-5 FPGA family. The OM1312 transports 2.5 Gbps of channelized Ethernet over high-order or low-order Sonet/SDH using standards compliant to GFP-F, VCAT and LCAS. Its SPI3 packet interface supports 128 channels.

The OM1411 is a 10-Gbps high-order EOS packet mapper likewise targeted to the Xilinx Virtex-5 FPGA family. It transports 10 Gbps of channelized Ethernet over Sonet/SDH using standards compliant to GFP-F, VCAT and LCAS. The SPI4.2 packet interface supports 192 channels. The Sonet/SDH interface is a 10-Gbps 1+1 protected TFI5 link for high-reliability carrier-grade deployments. The OM1411 handles both real and virtual concatenation (CCAT and VCAT).

The company's third offering, the OM1412, is a 10-Gbps high- and low-order EOS packet mapper targeted to the Xilinx Virtex-5 FPGA. It transports 10 Gbps of channelized Ethernet over high-order or low-order Sonet/SDH using standards compliant with GFP-F, VCAT and LCAS. Up to 25 percent of the 10G bandwidth can be low-order traffic. Its SPI4.2 packet interface supports 192 channels, making it possible to handle more individual customers on a single device, further reducing operating expenditures. The Sonet/SDH interface is 10-Gbps 1+1 protected TFI5.

Like the OM1411, the OM1412 handles both CCAT and VCAT. It also supports the Link Capacity Adjustment Scheme to allow a hitless increase or decrease in group bandwidth as customer need dictates.

Along with its designs, Omiino offers the software drivers as standard, but uniquely, Omiino has developed on-chip debug tools. OmniTest and OmniSpy allow the user to view and capture the internal state information of the device or to set up the device and test it independently of external software.

"The OmniTest and OmniSpy were born from our frustration, as product developers, with time taken to integrate silicon into the product, the dependency on software for board bring-up and the manually intensive and time-consuming root-cause-analysis support models," said Hamilton. "The OmniTest and OmniSpy tools can dramatically accelerate [development] and reduce costs during integration, verification and field support. We have great feedback from all of our customer engagements as they see firsthand the benefits, simplicity and power of the tools."

Despite the fact that Omiino has a number of offerings that are off the shelf and ready to go, the company recognizes that customers have individual requirements, said Hamilton. Omiino will work with them to customize designs for specific applications. In addition, Omiino can supply subsystems of ultracompact IP blocks such as GFP, VCAT and Sonet/SDH framers with speeds from 155 Mbps to 10 Gbps.

For further information, visit www.omiino.com/en/home or contact r.cosgrave@omiino.com.

GateRocket Drive Speeds Verification of Virtex-5 FPGAs 4x to 10x

With a mission to help “ASIC refugees” speed the verification and debug of ever-more-complex system-on-chip (SoC) designs they’re implementing on FPGAs, EDA tool startup GateRocket has released a Xilinx Virtex-5 version of its RocketDrive “native device” verification and debug environment. Using it, engineers can test the behavior of a design running on actual FPGA silicon while still in their HDL verification and debug environment of choice.

because the RocketDrive itself contains an FPGA from the largest platforms in the Virtex-5 family. One model packs a Virtex-5 LX330T, another the Virtex-5 SX240T and a third, the Virtex-5 FX200T device. Orecchio said this native-device approach, among other things, eliminates the nagging question that haunts users of traditional logic emulation systems: is the cause of a given bug really a problem with the design under test or is it due to a flaw in the emulation system running the design?



GateRocket's RocketDrive lets you test the behavior of your design running on actual FPGA silicon while still in your HDL verification and debug environment of choice.

Traditionally, hardware-assisted verification and emulation systems have targeted ASIC designs, but Dave Orecchio, GateRocket's president and CEO, said that today's FPGAs have advanced to the point where designers can now truly implement SoC designs on a single FPGA. “It used to be that for many applications you had no choice but to design an ASIC,” Orecchio said. “Today designers have a choice, and increasingly they are designing with FPGAs.”

Orecchio described RocketDrive as “the first hardware-assisted verification and debug environment specifically targeting these complex FPGA design projects—we’re here to help ASIC refugees.”

GateRocket's offering has a hardware and a software component. The hardware is a drive that fits into any 5.25-inch drive slot in a PC chassis, while the RocketVision software coordinates verification and debug, connecting the drive to Xilinx ISE® and third-party vendor tools of the user's choice.

Customers select which model of drive they need for their designs based on which Virtex-5 device they plan to use. That is

RocketDrive claims to speed verification 4x to 10x over software-based simulation. The larger the design, the greater the speedup, Orecchio said. In addition, he went on, because the drive is device-native, it can help users quickly find the root causes of functional errors in their HDL code. Designers can select which blocks they want to run on the drive and those they want to run on the simulator.

Customers can also use the drive to validate a model of an IP block against the IP actually running on silicon, or to validate that they have selected an optimal pin placement during the early stages of design. RocketDrive also helps users figure out if they have any tool chain faults.

“You can use the RocketDrive at every phase of the design, from RTL creation all the way through onboard test,” said Orecchio.

The RocketDrive sells for \$45,000, while the RocketVision software is priced at \$9,500. You can learn more about GateRocket and even sign up for a free Webinar at www.gaterocket.com.

GE Fanuc Aims Versatile Virtex-5 FPGA Board at A&D COTS Market

Building on the success of its board offerings built around Virtex-4 FPGAs, GE Fanuc has rolled out a Virtex-5 FPGA-based mezzanine card targeting DSP-centric applications developers in the aerospace-and-defense COTS markets.

The company designed its Intelligent Platforms XMCV5 FPGA mezzanine board in a way that allows customers to use the Virtex-5 device that best suits their application needs, said Ramon Mitchell, principal engineer of the MNG Group at GE Fanuc.

The board supports Virtex-5 devices based around the 1136 package. As such, Mitchell said, it is possible to fit members of the Xilinx LXT, SXT and FXT families into the XMCV5 system.

Mitchell noted that GE Fanuc also offers processor-based boards and switches, which essentially allows A&D COTS customers to mix and match boards to create systems that best fit their requirements.

“During the development of an FPGA-based system, the system may require a logic-based FPGA, then a DSP-based one at some other point in the system and maybe one that will benefit from the processors and MGTs [multigigabit transceivers] in the FXT,” said Mitchell. “With these almost bite-size FPGA boards, you can slot these in throughout your system. Flexibility is the key, because this is a commercial off-the-shelf product and it isn’t designed to a specific requirement. Every pin in the FPGA is connected to something; we didn’t want to miss an opportunity.”

In addition to a Virtex-5 FPGA, the XMCV5 FPGA mezzanine board includes two banks of QDR SRAM, DDR2 SDRAM and SPI flash. For back I/O it has VITA 42.3 PCI Express Gen2 (to 5 GHz) or a VITA 42.2 Serial RapidIO build option, along with RocketI/O™ to 6.5 GHz via GTX ports, LVDS, GPIO, Gigabit Ethernet and serial ports. Front I/O includes Gigabit Ethernet, serial, LVDS and GPIO.

The company has also created its own design care package called VWRAP. “It is based on work we’ve done to develop all the required interface IP inside the FPGA to make our customers’ lives easier when they build projects,” said Mitchell. Among other things, VWRAP includes DDR and DDR2 controllers, LVDS drivers

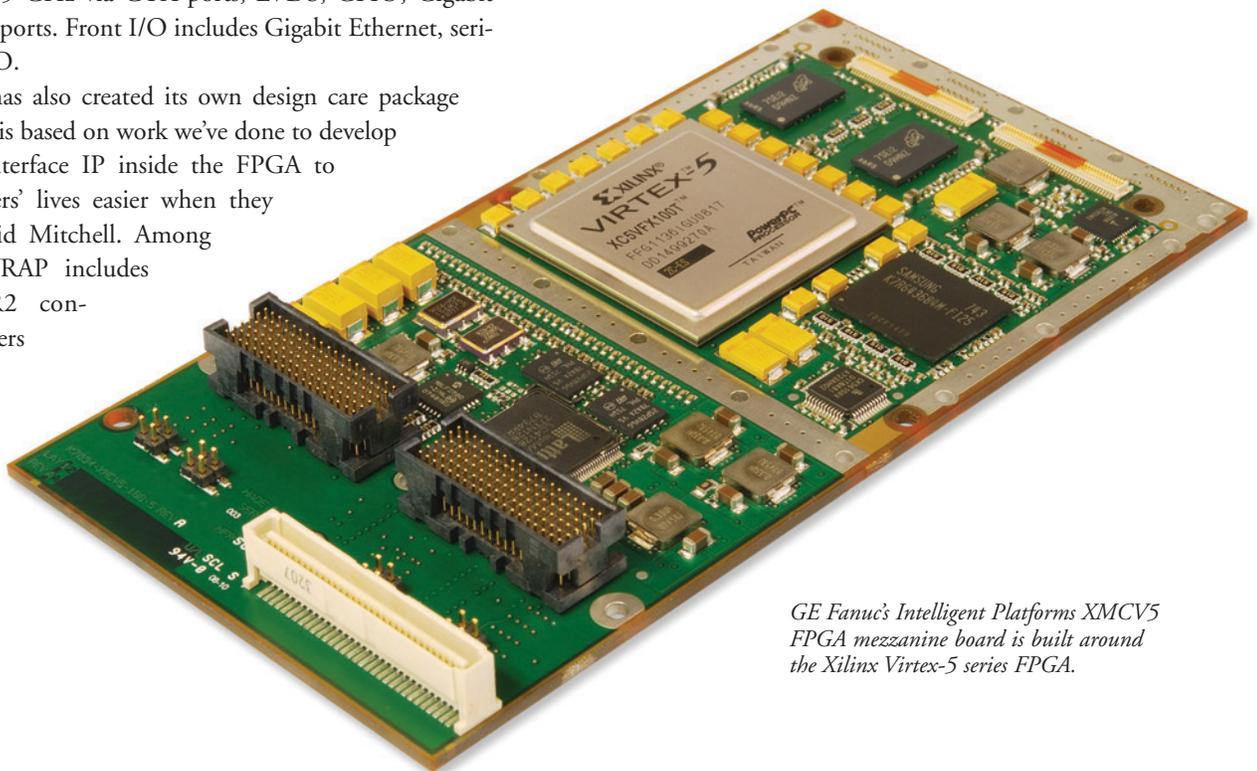
and preconfigured PCI Express end points, along with an embedded-processor example system. Mitchell noted that the company developed all of these designs with the Xilinx ISE tool suite, and sticks to Xilinx’s recommendations whenever possible so that when customers integrate the cores into their designs, there are no compatibility surprises.

GE Fanuc has also extended its internally developed tool, Axis, to now support its FPGA systems. Axis allows customers at a high level of abstraction to assign algorithms to processors running on the company’s processor-based cards, said Mitchell. “We’ve now added the hooks in to allow this to work with our FPGA nodes as well,” he said. “So someone that wants to do a system that uses both processors and FPGAs can use Axis to partition the design and hook up the various interfaces.”

GE Fanuc also provides a BIST tool called VBIT. “It’s basically a preconfigured image which will test all the interfaces on the board,” said Mitchell. “Customers run that at power-up to check the board is working properly and then program their own image in once they know it is good. We do all this to make life easier for our customers.”

Because the system is targeting harsh environments, GE Fanuc has rigorously tested the board running at -45°C to 90°C and has done extensive signal integrity testing on the system as well. “It’s extremely robust,” said Mitchell.

For more information, visit www.gefanucembedded.com/products/2260.



GE Fanuc's Intelligent Platforms XMCV5 FPGA mezzanine board is built around the Xilinx Virtex-5 series FPGA.

Xilinx's Support Network: Our Success Is Your Success

A stellar support staff and long-term partnerships with customers make all the difference.

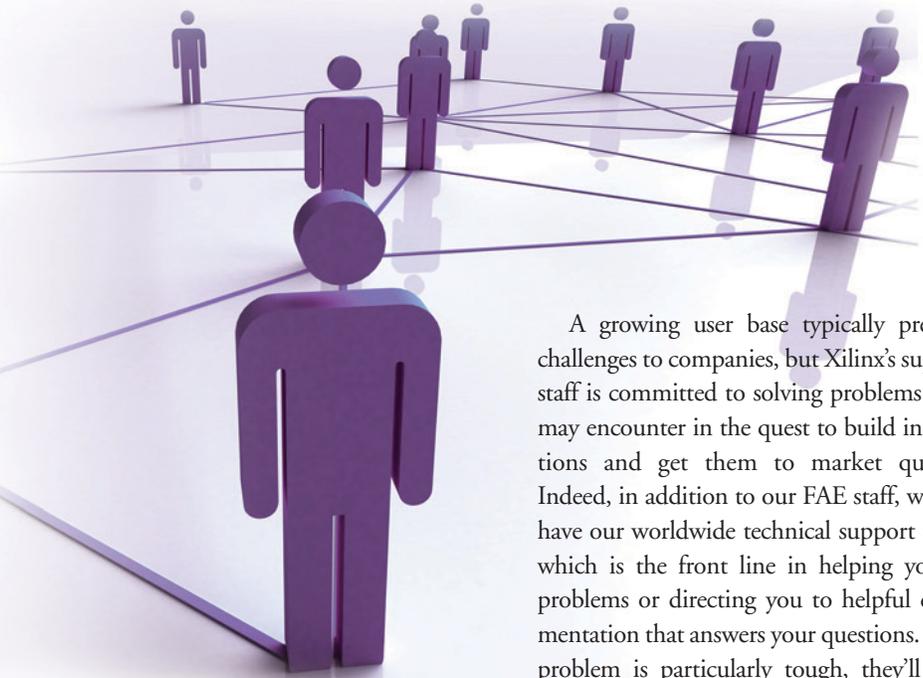


by Bruce Kleinman
Vice President, Technical Sales
Xilinx, Inc.
bruce.kleinman@xilinx.com

Over my many years here at Xilinx, I've come to understand the true value that comes from open communication and real partnership with our customers. Throughout our history, Xilinx has made a sizable investment in establishing an extensive customer support network and staffing it with not just great application engineers and technical support staff, but people who share a passion for solving problems.

Indeed, I have deep respect for the many hardworking people in our technical sales and support network, and awe for their abilities. One would think that solving problems from a variety of customers worldwide, who vary greatly in experience levels and are targeting a large swath of applications, would be a bit overwhelming. But I'm proud to say that Xilinx's staff (more than 400 people in technical sales and support) is full of unique individuals who each get charged when a new challenge arises. Many of our FAEs have done the job for 10 to 15 years and all remain passionate about helping users solve today's problems.

In addition, they play a vital role in delivering constructive feedback to our tool, IP and FPGA silicon development groups about what features next-generation FPGA platforms should include. Sometimes that feedback is not eagerly received; but more times than not, it keeps our tool, IP and silicon developers keenly in tune with reality so they can quickly fix problems that arise and integrate those lessons learned into next-generation products.



Of course, it would be great if all design techniques, tools, methodologies and silicon features were crystal-clear to all customers on the day the features were released. For some customers that may be the case, but as Xilinx remains committed to fielding the most advanced FPGA platforms available, design is certainly becoming more complex and FPGAs are adapting themselves to an ever-growing number of applications. Today, we at Xilinx have technical sales and support personnel who specialize not just in FPGAs, but in targeting those FPGAs in application spaces such as wired and wireless communications, aerospace and defense, automotive, industrial, scientific and medical.

What's more, over the last few years we've seen Xilinx's user base expand beyond the traditional hardware engineer. Each year we see a growing number of embedded-software engineers, DSP algorithm developers and system architects using our devices.

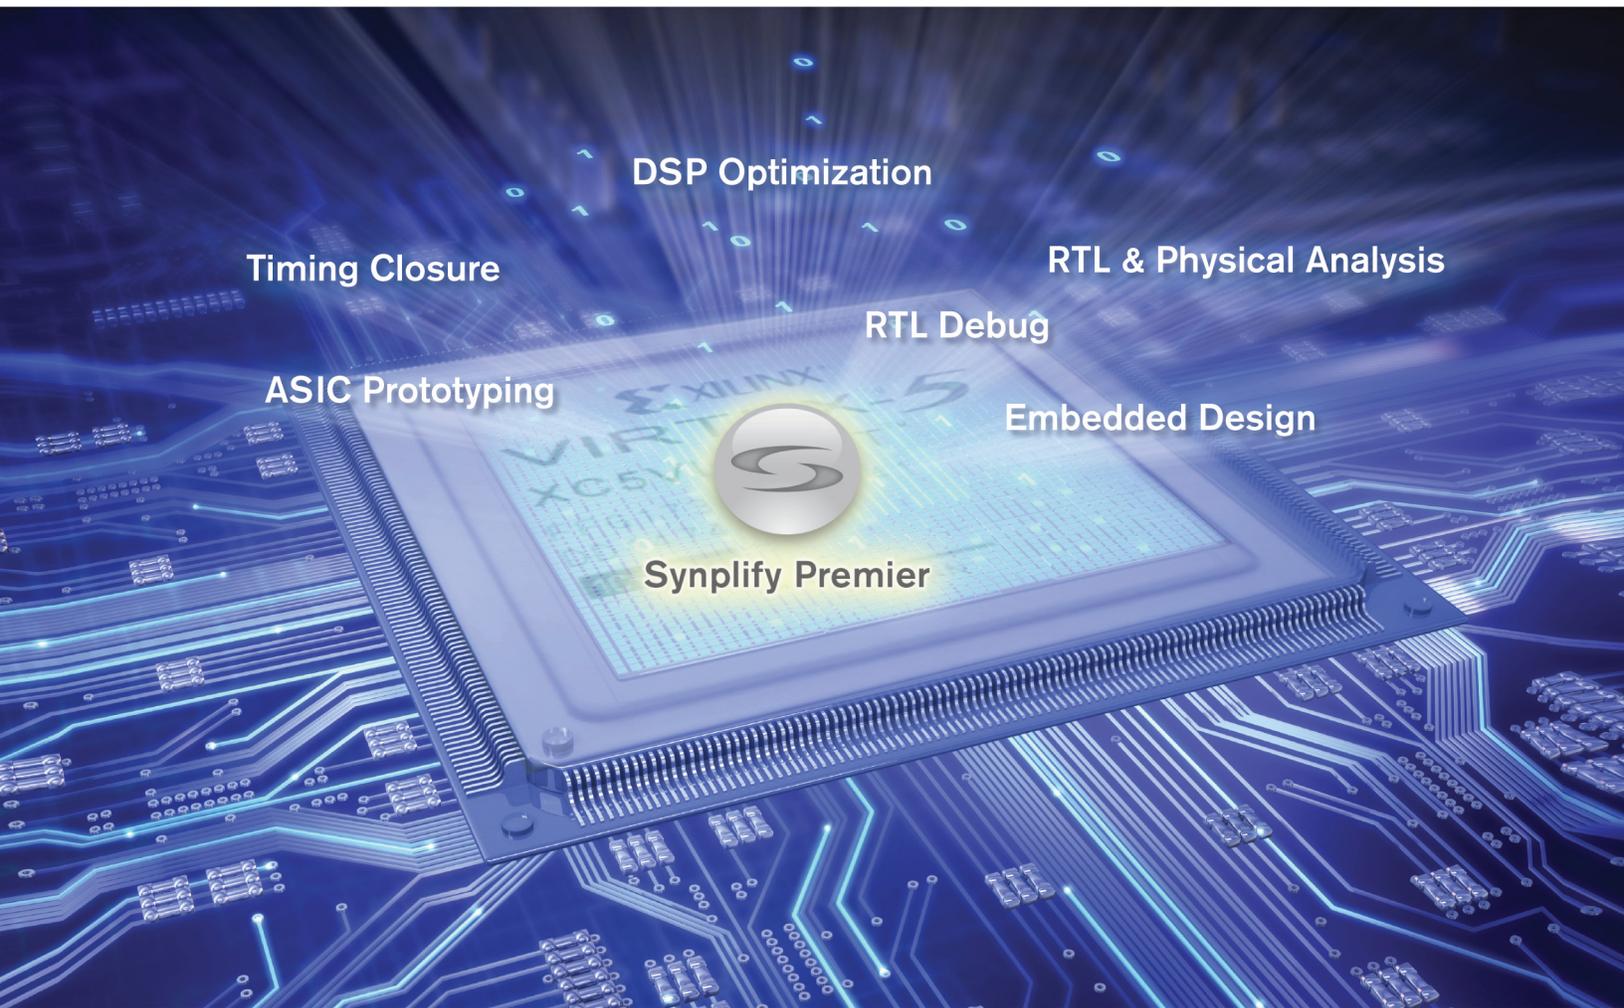
A growing user base typically presents challenges to companies, but Xilinx's support staff is committed to solving problems users may encounter in the quest to build innovations and get them to market quickly. Indeed, in addition to our FAE staff, we also have our worldwide technical support team, which is the front line in helping you fix problems or directing you to helpful documentation that answers your questions. If the problem is particularly tough, they'll have the best experts assist you.

Some customers may be fresh out of college, others may be ASIC refugees who are trying to ramp up on FPGA technology. Still others may be embedded-software designers or DSP programmers who want to create super-advanced embedded systems or use the ultra-advanced DSP capabilities of an FPGA. For those folks, we offer training courses throughout the year and have an extensive network of training partners.

And for companies that need additional specific expertise to get a project done, we have a technical services group packed with experienced engineers who boast broad applications knowledge.

It's no small task to design with an advanced FPGA, but it shouldn't be a formidable task either. Xilinx is keenly aware that our success is directly linked to yours. The company, and our extensive support network, is committed to your success. 

Accelerate FPGA Design



Synplify Premier, the Ultimate in FPGA Implementation

The Synplify® Premier software from Synopsys® is the preferred synthesis and debug environment for complex FPGA designs. It provides a comprehensive suite of tools and technologies for advanced FPGA designers as well as ASIC prototypers targeting a single FPGA. The Synplify Premier solution addresses the biggest FPGA design challenges including timing-closure, logic verification, IP support, ASIC compatibility, DSP implementation, and RTL debug while providing tight integration with Xilinx® back-end tools. Synplify Premier supports DesignWare® components and performs detailed logic placement which is passed on to the Xilinx router for final implementation. With final placement knowledge during synthesis, design iterations are significantly reduced resulting in shorter development schedules.

To learn more about how the Synplify Premier software
can help you achieve your design goals, visit
http://www.synplcity.com/synplifypremier/xcell_premier.html

Copyright © 2008 Synopsys, Inc. All rights reserved. Synopsys, Synplcity, the Synplcity logo, DesignWare, and Synplify are registered trademarks of Synopsys, Inc. All other names mentioned herein are trademarks or registered trademarks of their respective companies. 1108.CE.WO.08-16815


Synplcity
Simply Better Results

SYNOPSYS
Predictable Success

Design Simply

Design Completely

Design Today



The Virtex[®]-5 Family: The Ultimate System Integration Platform

- Increase Your System Performance
- Lower Your System Cost
- Design with Ease

The Virtex-5 family delivers unparalleled system integration capabilities for driving your most mission-critical, high-performance applications. With a choice of five platforms optimized for logic, serial connectivity, DSP and embedded processing with hardened PowerPC[®] 440 processor blocks, the Virtex-5 family delivers an unprecedented combination of flexibility and performance—backed by world class application support.

Only the Virtex-5 family offers you a complete suite of design solutions built on proven 65nm technology in devices shipping today.

Get started on your Virtex-5 design. Visit www.xilinx.com/ise for a free 60 day evaluation of any ISE[®] Design Suite 10.1 product.



www.xilinx.com/virtex5