ISSUE 80, THIRD QUARTER 2012

Xcell journal

SOLUTIONS FOR A PROGRAMMABLE WORLD

Xilinx Rolls World's First Heterogeneous 3D FPGA

Virtex-7 H580T Device Enables 2x100G Transponder-on-Chip for CFP2 Optical Nets

FPGA-Based Instrumentation Withstands the Chill of Deep Space

Partial Dynamic Reconfiguration Fuels FSK Demodulator Design

Demystifying FPGA Mathematics

Ins and Outs of ADCs and DACs

Artix-7 FPGA Brings High-End Value to Low-Cost Market

page 14

XILINX®

www.xilinx.com/xcell/

AVNET electronics marketing

www.zedboard.org

Introducing ZedBoard™ — a Community-Based Development Kit for the Zynq™-7000 EPP

ZedBoard™

www.zedboard.org

Order entry is now available for ZedBoard™, a community-oriented development kit for the Xilinx Zynq™-7000 Extensible Processing Platform (EPP). ZedBoard™ is ideal for hardware and software designers who are interested in exploring and prototyping their application ideas for the new Zynq™-7000 EPP architecture.

The commercial version of ZedBoard™ is available from Avnet for $395.

Zedboard™ Sponsors:

ANALOG DEVICES · CYPRESS PERFORM · MAXIM INNOVATION DELIVERED · samtec · SPANSION · TE connectivity Authorized Distributor

Accelerating Your Success™

XILINX ALL PROGRAMMABLE.

YOUR CHIP FITS HERE

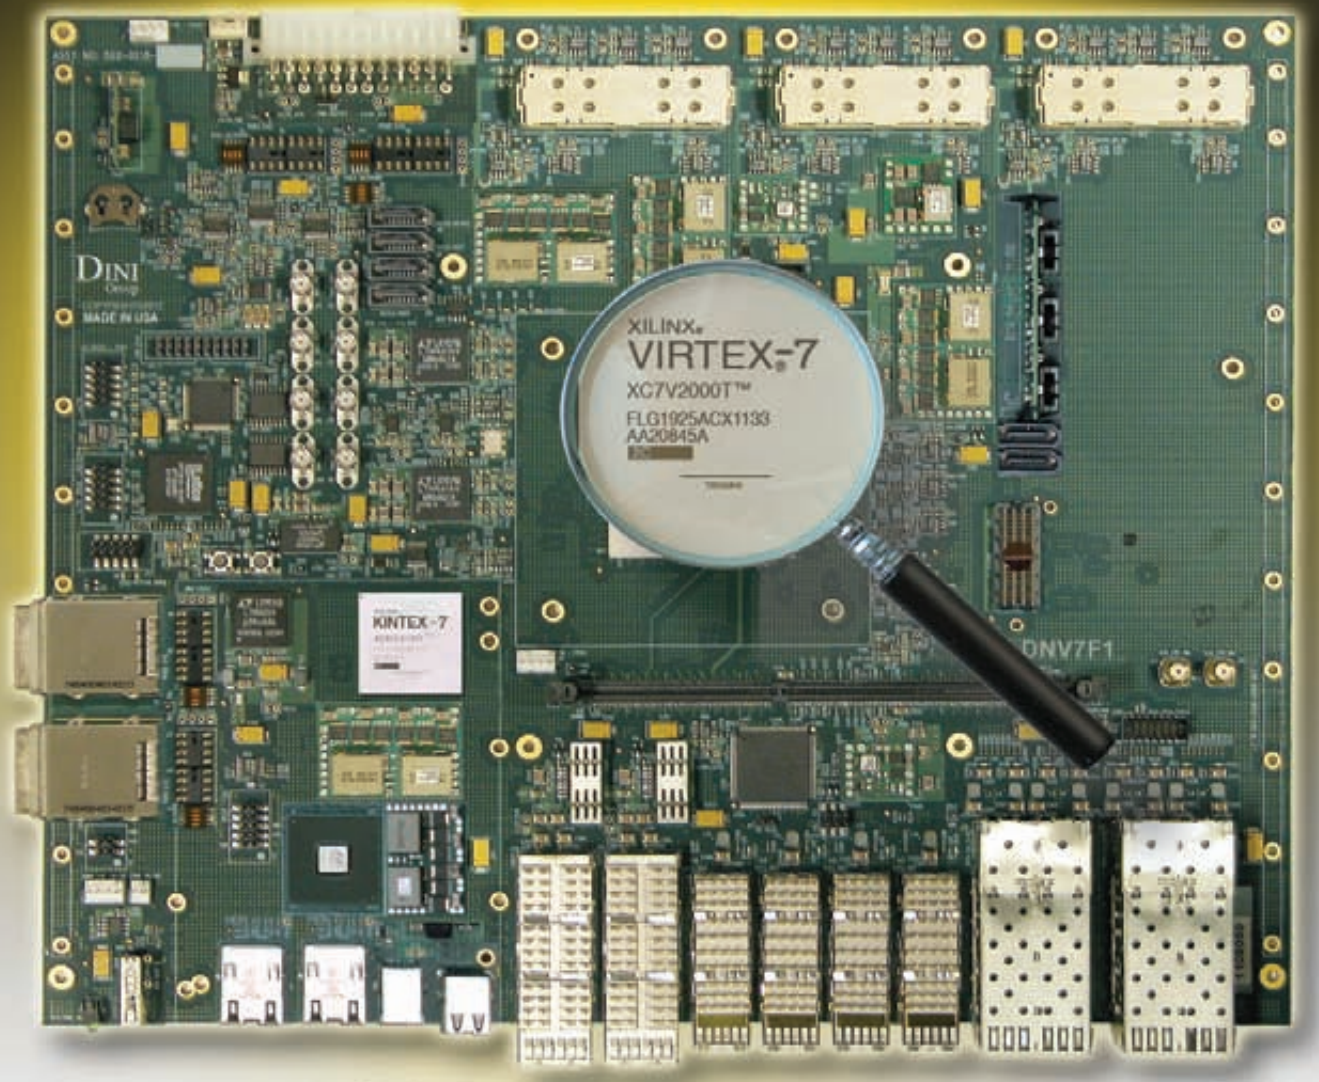

15 Million ASIC Gates for Prototyping Your ASIC

New stacked silicon architecture from Xilinx makes your big design much easier to prototype. Partitioning woes are forgotten, and designs run at near final chip speed. The DINI Group DNV7F1 board puts this new technology in your hands with a board that gets you to market easier, faster and more confident of your design's functionality running at high speed. DINI Group engineers put the features you need most, right on the board:

- 10GbE
- USB 2
- PCIe, Gen 1, 2, and 3
- 240 pin UDIMM for DDR3

There is a Marvel Processor for any custom interfaces you might need and plenty of power and cooling for high speed logic emulation. Software and firmware developers will appreciate the productivity gains that come with this low cost, stand-alone development platform.

Prototyping just got a lot easier, call DINI today and get your chip up to speed.

DINI Group

www.dinigroup.com • 7469 Draper Avenue • La Jolla, CA 92037 • (858) 454-3419 • e-mail: sales@dinigroup.com

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

Miyuki Takegoshi, Japan
miyuki.takegoshi@xilinx.com

REPRINT ORDERS 1-800-493-5551

Xcell PUBLICATIONS
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/

Hardware-Assisted Prototyping: Make vs. Buy?

As long as there are ASICs, ASSPs, microprocessors and other types of digital ICs following the silicon process technology curve, there will be a need to prototype those devices in real hardware before they go to production. But is it easier to make or buy a prototyping system? That's the question panelists attacked in a recent Pavilion Panel session at the 49th Design Automation Conference in San Francisco.

Gabe Moretti, EDA veteran and owner of the popular website *GabeonEDA.com*, moderated the panel, entitled "Hardware-Assisted Prototyping and Verification: Make vs. Buy?" Gabe's panelists were Qualcomm engineering director Albert Camilleri; Austin Lesea of Xilinx Research Labs, and co-author of the book *The FPGA Prototyping Methodology Manual*; and Mike Dini, CEO of the hardware-assisted verification company the Dini Group.

Panelists agreed that with SoCs now reaching well over 100 million gates, it's imperative to emulate the chip you are designing before it goes to manufacturing so as to reduce corner-case bugs and minimize respins. Likewise, as software becomes a greater part of the overall system development, hardware-assisted prototyping systems, such as those offered by the Dini Group and Synopsys' HAPS group, are becoming an imperative.

"I look at the size and complexity of IC designs people are doing today, and I'm not sure I'd want to build a prototype for something this large," said Moretti. "The commercial systems are pretty expensive, but I'm not sure if it makes more sense today to buy one rather than build one. So, should I build one or buy one?" Moretti asked panelists.

"I think a lot of people start out thinking there are some good reasons to build their own, but quickly realize it probably would have been easier to buy one," said Lesea. "In the modern SoC world, about 80 percent of the project is software. The sooner you can get something for the software team to work on, the sooner your project is not stuck. … The time you save in time-to-market pretty much pays for the commercial prototyping system and more."

"I've long argued that unless software guys have real hardware working at some reasonable frequency, they can't really do effective development," added Dini. "Unless they can get a real blue screen of death, if they can't run an interrupt into the weeds and smoke a piece of hardware, generally development slows or stalls … pure software-only simulation can't cover enough corner cases."

Camilleri said he has both built and purchased prototyping systems depending on what a given design required and what commercial offerings were available. "If there is already a commercial offering that will do what you need, then why build one?" Camilleri noted, however, that sometimes a design's performance requirements will force design groups to build a custom system. But he said that in general it's only a good idea if you are building it for a specific project in which everyone on the design team is intimately aware of all the nuances. "If you are building a prototyping system for your division, for example, that can be a huge undertaking," he said. "It's easy to underestimate the time it takes and the quality of EDA software it takes to develop a commercial-class prototyping system."

With a few exceptions in the emulation world, the vast majority of emulation and prototyping systems are FPGA based.

Panelists were especially encouraged by FPGAs vendors' embrace of 3D IC technology to exceed the doubling of capacity for the next generation of FPGAs that one typically expects with each new silicon process technology. "The Virtex®-7 2000T is a game changer," said Dini. Panelists said that the fewer FPGAs there are in an emulation system, the easier it is to partition a design and stabilize it in the prototyping system. An FPGA with massive capacity allows ever-larger designs, with hundreds of millions of ASIC gates, to get to market much sooner.

Mike Santarini
Publisher

www.aldec.com

Your Global Xilinx® Verification Solutions Provider

- VHDL, Verilog, SystemVerilog
- OVM/UVM, VMM
- Code and Functional Coverage
- DSP Co-Simulation (MATLAB® & Simulink®)
- Emulation (Hardware-Assisted Verification)

Headquarters-US
2260 Corporate Circle
Henderson, NV 89074
USA

Phone: +1.702.990.4400
E-mail: sales@aldec.com

Europe
Mercia House
51 The Green, South Bar
Banbury, OX16 9AB
United Kingdom

Phone: +44.1295. 20.1240
Email: sales-eu@aldec.com

Israel
1 Haofe Street
Kadima 60920
Israel

Phone: +972.072.2288316
E-mail: sales-il@aldec.com

Japan
Shinjyuku Estate Bldg. 9F
1-34-15, Shinjyuku
Shinjyuku-ku, Tokyo 160-0022
Japan

Phone: +81.3.5312.1791
Email: sales-jp@aldec.com

China
Suite 2004, BaoAn Building
#800 DongFang Road
PuDong District
Shanghai City, 200122, P.R. China

Phone: +86.21.6875.2030
Email: info@aldec.com.cn

India
#2145, 17th Main
2nd Cross, HAL 2nd Stage
Indiranagar
Bangalore, 560008, India

Phone: +91.80.3255.1030
Email: sales-in@aldec.com

Taiwan
No. 37, Section 2
Liujia 5th Road
Hsinchu County 302
Zhubei City, Taiwan

Phone: +886.3.6587712
E-mail: sales-tw@aldec.com

VIEWPOINTS

Letter From the Publisher
Hardware-Assisted Prototyping: Make vs. Buy?… **4**

Xpert Opinion
FPGAs Head for the Cloud… **20**

20

XCELLENCE BY DESIGN APPLICATION FEATURES

Product Feature
Artix-7 FPGA Brings High-End Value to Low-Cost Market… **14**

Xcellence in Aerospace & Defense
FPGA-Based Instrumentation Withstands the Chill of Deep Space… **24**

Xcellence in Green Technology
Using Spartan Technology to Support Green Energy Development… **28**

Xcellence in Solid-State Storage
Designing a 19-nm Flash PCIe SSD with Kintex-7 FPGAs… **32**

14

28

Cover Story 8

Xilinx Introduces First Heterogeneous 3D FPGA: **Virtex-7 H580T**

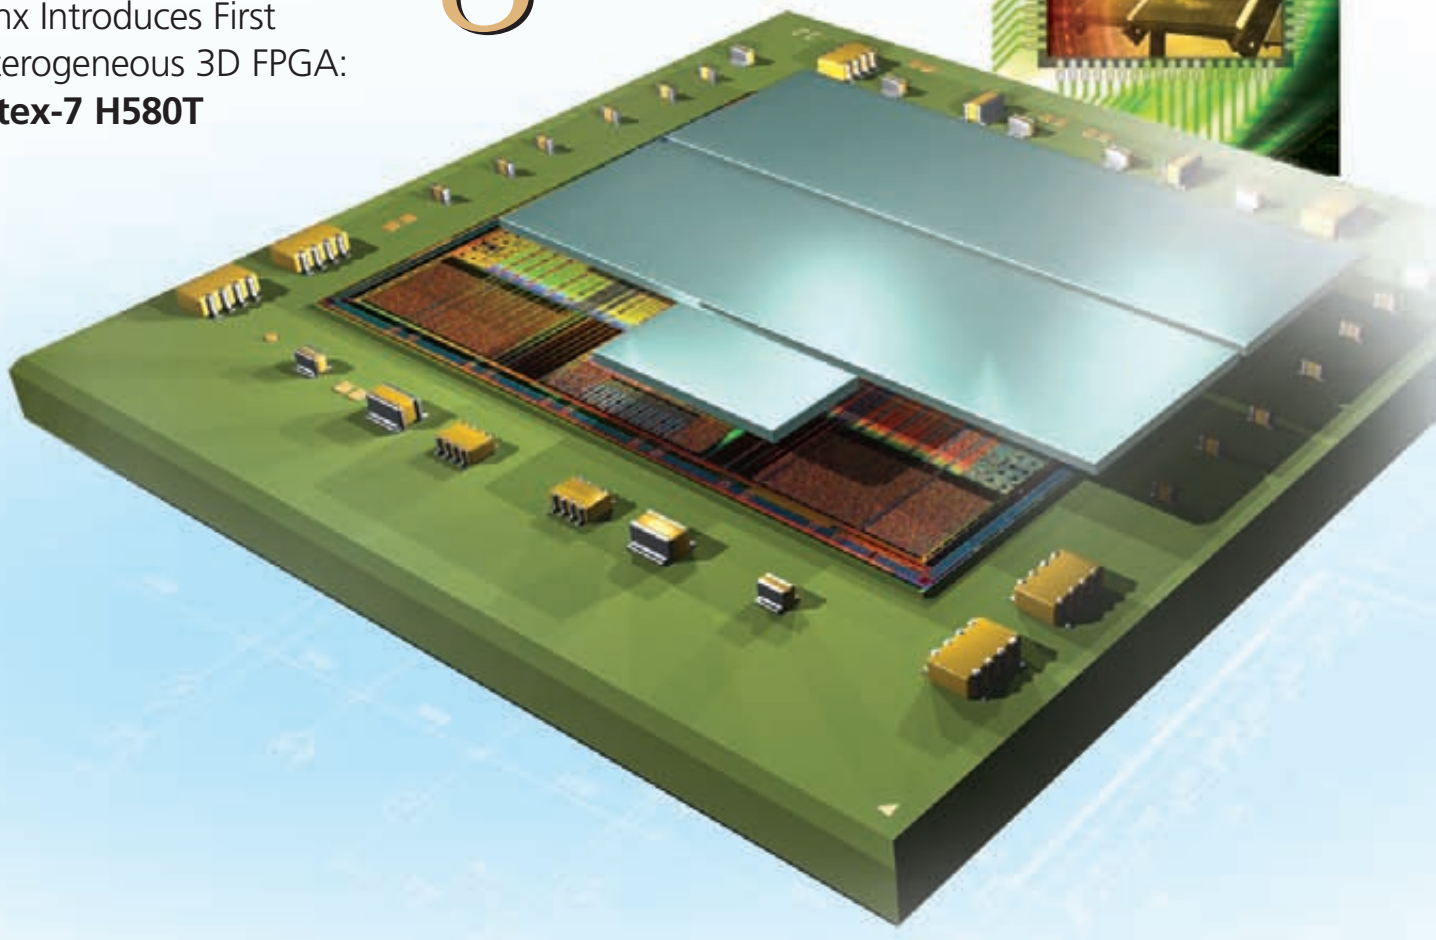

THE XILINX XPERIENCE FEATURES

Xperts Corner

How Partial Dynamic Reconfiguration Helped Build an FSK Demodulator… **38**

Xplanation: FPGA 101

The Basics of FPGA Mathematics… **44**

Xplanation: FPGA 101

The FPGA Engineer's Guide to Using ADCs and DACs… **50**

XTRA READING

Profiles of Xcellence The sky's the limit for SSD enterprise storage startup Skyera… **54**

Tools of Xcellence Xilinx FPGAs improve beamforming system design… **58**

Xtra, Xtra Your questions answered on the new Vivado Design Suite… **64**

Xclamations! Share your wit and wisdom by supplying a caption for our techy cartoon. Three chances to win an Avnet Spartan-6 LX9 MicroBoard!… **66**

APEX® AWARDS FOR PUBLICATION EXCELLENCE

Excellence in Magazine & Journal Writing
2010, 2011

APEX® AWARDS FOR PUBLICATION EXCELLENCE

Excellence in Magazine & Journal Design and Layout
2010, 2011, 2012

Xilinx Introduces First Heterogeneous 3D FPGA: Virtex-7 H580T

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

Built with Xilinx 3D SSI technology, this device enables designers to create a 2x100G OTN transponder-on-a-chip.

Hot on the heels of breaking capacity and transistor-count records with the release of its 28-nanometer Virtex®-7 2000T (the industry's first 28-nm FPGA implemented in 3D stacked-silicon interconnect technology), Xilinx® in May released a device using SSI technology that breaks the record for FPGA bandwidth. The new Virtex-7 H580T device is the world's first heterogeneous 3D FPGA, integrating a dedicated eight-channel 28-Gbps transceiver slice (or die) alongside two transceiver-rich FPGA dice on a single silicon interposer. All told, the new product gives wired communications companies a device with up to forty-eight 13.1-Gbps transceivers as well as the eight 28-Gbps transceivers and 580,480 logic cells—making the Virtex-7 H580T FPGA the only single-chip solution for addressing key 2x100G applications and functions (Figure 1). Product details are spelled out at *http://www.xilinx.com/publications/prod_mktg/Virtex7-Product-Table.pdf*.

"Combined with Xilinx's 100-Gbps gearbox, Ethernet MAC, OTN and Interlaken IP, the Virtex-7 HT devices provide customers with the kind of system integration they need to meet space, power and cost challenges as they transition to 100-Gbps low-power optical modules in the new CFP2 form factor," said Ephrem Wu, senior director of advanced communications at Xilinx. "The 28-Gbps transceivers are independent from the 13.1-Gbps transceivers. Customers can use all available 28-Gbps transceivers without having to give up any 13.1-Gbps transceivers."

The Virtex-7 H580T FPGA is the first of three heterogeneous 3D devices Xilinx will field in its 28-nm family. The Virtex-7 H870T, due in the near future, comprises two eight-channel transceiver dice alongside three FPGA logic dice on a single device, yielding a total of sixteen 28-Gbps transceivers, seventy-two 13.1-Gbps transceivers and 876,160 logic cells on one chip. A third heterogeneous device, the Virtex-7 H290T, places one eight-channel transceiver die alongside one FPGA logic slice on a single device, yielding twenty-four 13.1-Gbps transceivers, eight 28-Gbps transceivers and 284,000 logic cells on one chip.

"Our 3D SSI technology enables Xilinx to jump ahead of the technology curve and offer All Programmable devices that enable the highest levels of integration, system performance, power reduction, BOM cost reduction and productivity," said Wu. "With the Virtex-7 2000T, we used 3D SSI technology to stack four logic slices side-by-side on a silicon interposer to create a device with 6.8 billion transistors and 1,954,560 logic cells—double the capacity of the largest competing 28-nm FPGA and well beyond the expected doubling of transistor counts dictated by Moore's Law. Now, with the Virtex-7 HT devices, we have leveraged our 3D SSI technology to stack 28-Gbps transceiver slices alongside 28-nm FPGA slices on a silicon interposer—all in a single chip."

The SSI technology, Wu said, "enables Xilinx to field a device today that will allow customers to offer distinctly compelling value to their cus-

Figure 1 – Xilinx 3D SSI technology forms the foundation for the Virtex-7 H580T, the world's first heterogeneous FPGA. It marries 28-nm FPGA logic slices and a dedicated 28-Gbps transceiver die on a silicon interposer.

tomers of 100-Gbps optics-enabled equipment and allow the wired communications industry to accelerate the development of next-generation 400G equipment."

THE INSATIABLE BANDWIDTH REQUIREMENT

Alex Goldhammer, senior product line manager at Xilinx for Virtex-7 FPGAs, notes that as more and more systems connect to the Internet and private networks, the demand spirals for more bandwidth to transfer ever-larger files and stream higher-quality video and audio across the globe. To accommodate this demand, service providers want higher-bandwidth wired communications equipment at a lower cost per bit. The wired communications sector in particular is currently building equipment to conform to recently formalized 100-Gbps communications optical transceiver standards—most notably CFP2 optics, OIF CEI-28-VSR and IEEE 802.3ba.

At the heart of the 100-Gbps infrastructure buildout are optical transport network (OTN) transponders and muxponders, and 100G Ethernet cards. Network companies place these OTN cards at the center or core of the optical network—the fastest sections—to ensure the integrity and proper routing of data as it flies across the globe on fiber-optic cables.

Goldhammer said companies already have equipment with first-generation 100-Gbps OTN transponder cards today, each typically composed of a series of one or two ASSPs and one FPGA. These first-generation 100-Gbps OTN cards transmit and receive input from fiber optics via a CFP optical module. (The acronym stands for C Form-factor Pluggable.) An ASSP then takes 10x11.1G OTL 4.10 or CAUI (100-Gbps attachment unit interface) from the CFP and performs 100-Gbps forward error correction (GFEC), OTU-4 framing and 100GE mapping before sending the data to the FPGA via a CAUI. The FPGA is commonly used to translate the protocol to the required form for the backplane to route the data to the next point in the network and, ultimately, its destination.

The CFP optical modules, relatively bulky and relatively expensive, are the main stumbling blocks in these first-generation 100-Gbps OTN transport cards, said Goldhammer. To address this problem, the industry recently created the CFP2 form factor, defining an optical module for 100-Gbps line cards that is half the width (pitch) of CFP with slightly less depth, in the same power envelope. The advent of CFP2 means equipment companies can swap out existing CFP-based line cards for new cards that have two CFP2 channels per unit area, thus doubling the bandwidth of each card slot and potentially doubling the bandwidth of data centers (see Figure 2).

But Goldhammer said that CFP2 comes with new technical challenges. "CFP2 requires 25- to 28-Gbps trans-

Figure 2 – The CFP2 form factor makes it possible to double the bandwidth of 100-Gbps OTN cards. CFP2 cards are half the width of a CFP and use half the power, enabling drastically lower system costs.

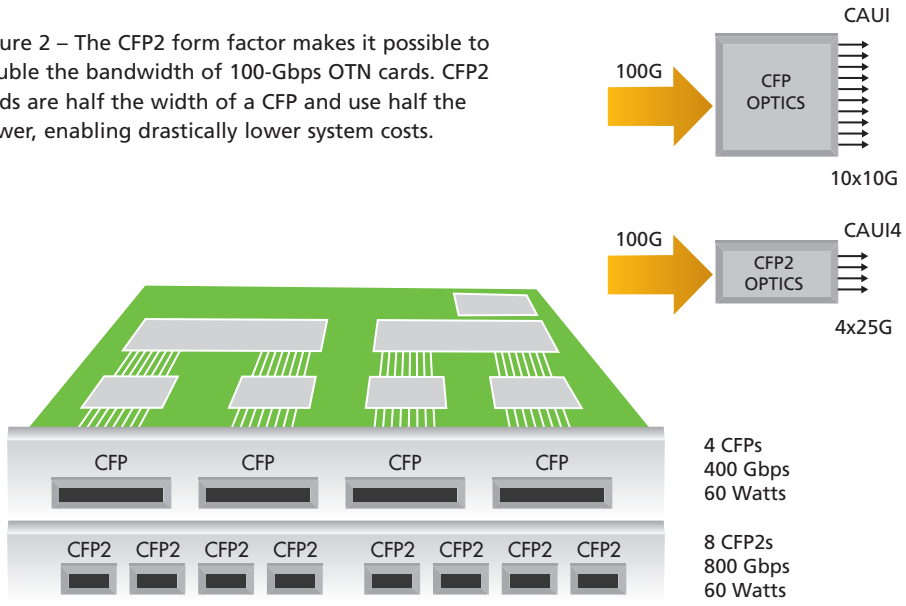

ceivers, PCB channel modeling using IBIS-AMI models and high-speed serial modeling software tools. And each card must maintain the same power budget as the CFP card it is replacing." Although the move from CFP to CFP2 delivers twice the bandwidth per watt, "simply doubling the amount of chips on each card to handle the bandwidth is not going to be viable, especially when it comes to power budgets," he said. "CFP2 requires more sophisticated silicon devices with greater degrees of integration."

One architecture that equipment manufacturers are currently contemplating for CFP2 cards consists of five devices, Goldhammer said: four ASSPs and one FPGA. Each card will have two CFP2 optical modules, which would use a 4x27G OTL 4.4 interface to a gearbox ASSP. The gearbox in its turn demultiplexes the 4x27G OTL 4.4 signal to 10x11.1G OTL 4.10. Then, another ASSP performs 100-Gbps GFEC, OTU-4 framing and 100GE mapping, and transports the data on a CAUI interface to the FPGA. Next, each of the CFP2's two channels sends the data to the one FPGA on the board, which serves as a CAUI-to-Interlaken bridge for the backplane to in turn send data to the next point in the network, and ultimately its destination (Figure 3).

"That configuration will typically require four ASSPs and one FPGA," said Goldhammer. "The biggest problems with that configuration are power, complexity and cost. Simply doubling the ASSPs will blow the power budgets."

While a CFP2 card will enable twice the bandwidth of CFP, each CFP2 module (with two 100-Gbps CFP2 ports) must maintain the power budget allotted to a single-port CFP module so as to stay within the same power budgets allocated for the entire card. Goldhammer said that the operational expense of this equipment is a significant concern for carriers, as there are many of these systems in the carrier's plant and they must maintain strict power caps. "They have to stay in these power budgets, but they want that 2x increase in bandwidth—so much of that burden falls to the semiconductor vendors to also lower power," he said.

With the new Virtex-7 H580T FPGA and Xilinx IP, Goldhammer said, 100-Gbps OTN line card makers can further maximize the value of their CFP2-based OTN cards by using just one Virtex-7 H580T to do the job of what would otherwise take all five chips. "The Virtex-7 H580T FPGA is a groundbreaking device timed perfectly to the CFP2 100-Gbps OTN transponder card market requirements," said Goldhammer.

With the Virtex-7 H580T FPGA and Xilinx IP, companies can implement an architecture for their CFP2-based cards in which two CFP2 channels on a card feed into one Virtex-7 H580T FPGA. The FPGA integrates gearbox, 100-Gbps GFEC, OTU-4 framing, 100GE mapping

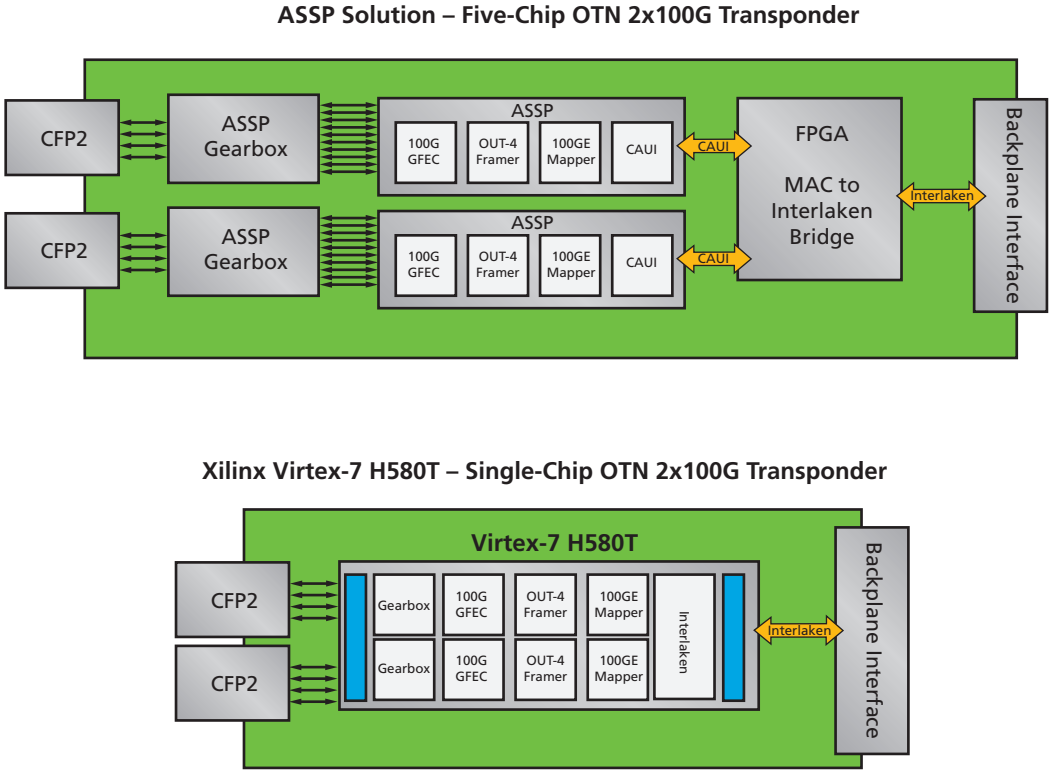

Figure 3 – The Virtex-7 H580T FPGA and Xilinx IP will enable customers to quickly create single-chip, CFP2-based 100-Gbps OTN transponder cards instead of using five-chip cards.

Video: A Virtex-7 H580T device demonstrates its ability to deliver the eye and jitter characteristics necessary to reach the performance required to interface to CFP2 optic modules.

and Interlaken bridging in one device (as again shown Figure 3).

"This is a single-chip solution that is not only much lower power than a multiple-chip ASSP or ASIC configuration, but faster, more reliable and of course much less expensive to produce," said Goldhammer. "It eliminates the need for multiple chips and their related power and cooling circuitry. With the Virtex-7 H580T FPGA, we offer customers more value in terms of integration, BOM cost reduction and improved system performance, without exceeding the power cap requirements for CFP2-based OTN transport cards."

What's more, Xilinx also has the right IP to enable communication equipment companies to accelerate their design productivity and get their single-chip 100G optics-based cards to market faster. Through internal development and a series of strategic acquisitions, Xilinx offers the whole package: 100-Gbps gearbox, Ethernet MACs, OTN and Interlaken IP. "We've optimized all these cores for designing into the 28-nm Virtex-7 FPGA logic cell slices on the device," said Goldhammer. "Xilinx manufactured the slices in TSMC's 28-nm high-performance, low-power [HPL] technology, which greatly reduces leakage, for the optimal mix of high performance and low power."

SSI TECHNOLOGY AND 28-GBPS TRANSCEIVERS

One of the biggest challenges of high-speed communications equipment today is ensuring the proper function of transceivers so that they maintain good signal integrity. "Transceivers are analog circuits and that means they can be affected by a number of factors—especially noise," said Goldhammer. "In most mixed-signal devices, transceivers are usually placed in isolation on the edges of devices to shield them from the digital circuitry in the middle of the device. Digital circuitry tends to be noisy, so that's why they usually isolate it from analog."

Over the last decade, to increase bandwidth into the gigabits-per-second range, the industry has turned to high-speed analog transceivers to quickly send and receive fast-traveling signals. Traditionally, the rule of thumb has been that the higher the bandwidth of the transceiver, the harder it is to ensure solid signal integrity.

Goldhammer said that because the Virtex-7 H580T FPGA is a highly integrated one-chip SSI technology solution, CFP2-based line cards built around it will achieve much improved performance. "Moving to 4x25G interfaces greatly reduces the complexity of routing 10x10G interfaces," he said. "Although some are concerned about 25G-to-28G transceivers, with Xilinx SSI technology, Xilinx is able to substantially reduce this complexity. The 28G transceivers, which have sensitive analog circuitry, are physically separated from the digital logic. This architecture ensures good isolation from the digital transceiver-rich dice."

The 28G transceivers are manufactured on a high-speed process technology, Goldhammer said, ensuring they are best in class. "The FPGA slices, by contrast, are manufactured in 28-nm HPL to ensure the lowest total power." The result, he said, is stellar 28-Gbps transceiver performance and signal integrity on the Virtex-7 H580T FPGA. To see these transceivers in action, check out this video on YouTube: *http://www.youtube.com/watch?v=FFZVwSjRC4c&feature=player_profilepage*.

Goldhammer said the physical isolation afforded by the SSI architecture enabled Xilinx to give the Virtex-7 H580T FPGA eight 28-Gbps transceivers—twice the number found in the largest device fielded by the competition.

What's more impressive is that the Virtex-7 H580T FPGA is not even the highest-transceiver device Xilinx will offer in its 28-nm family. The company

will soon roll out the Virtex-7 H870T device, which will have sixteen 28-Gbps transceivers, seventy-two 13.1-Gbps transceivers and 876,160 logic cells. Goldhammer said that if a customer used all the transceiver capabilities of the H580T device, they could conceivably have a design with a serial connectivity totaling 2.78 terabits per second.

"It was impractical and cost-prohibitive to place that many 28-Gbps transceivers on a monolithic FPGA," he said. "Fortunately, SSI technology enabled us to create a scalable FPGA family today that has eight to sixteen 28-Gbps transceivers." ASSP suppliers and other FPGA vendors have at most four 28G transceivers. This seems indicative of the challenges in doing the job monolithically in 40- and 28-nm processes.

The Virtex-7 H870T device is targeted at the next generation in wired communications—the 400G market, Goldhammer said. "The 400G market is a ways off, and if anything, companies are just starting to look at it in their labs and the standards bodies haven't gotten to it yet," he said. "What's beautiful is that we already have a device that's capable of doing it. We can help them speed up development of 400G, speed up the pace of innovation."

In addition to the Virtex-7 H580T and H870T FPGAs, Xilinx will also release as part of the 28-nm family the Virtex-7 H290T. By leveraging Xilinx's 3D SSI technology, the H290T will offer twenty-four 13.1-Gbps transceivers, eight 28-Gbps transceivers and 284,000 logic cells. Goldhammer said the Virtex-7 H290T is particularly well suited for the 2x100G gearbox market.

First silicon of the Virtex-7 H580T FPGAs is shipping to key customers today, with development tool support available in the recently announced Vivado™ Design Suite. Customers interested in using the Virtex-7 H580T device can contact their local Xilinx representative for further pricing and availability details. You can also find new white papers and videos on Xilinx's 28-Gbps Serial Transceiver Technology page: *http://www.xilinx.com/products/technology/transceivers/index.htm.*

Xilinx Artix-7 FPGA Ships High-End Value to Low-Cost Market

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

Xilinx is now shipping the first device in its All Programmable Artix-7 FPGA series, setting new power and performance standards for cost-sensitive applications.

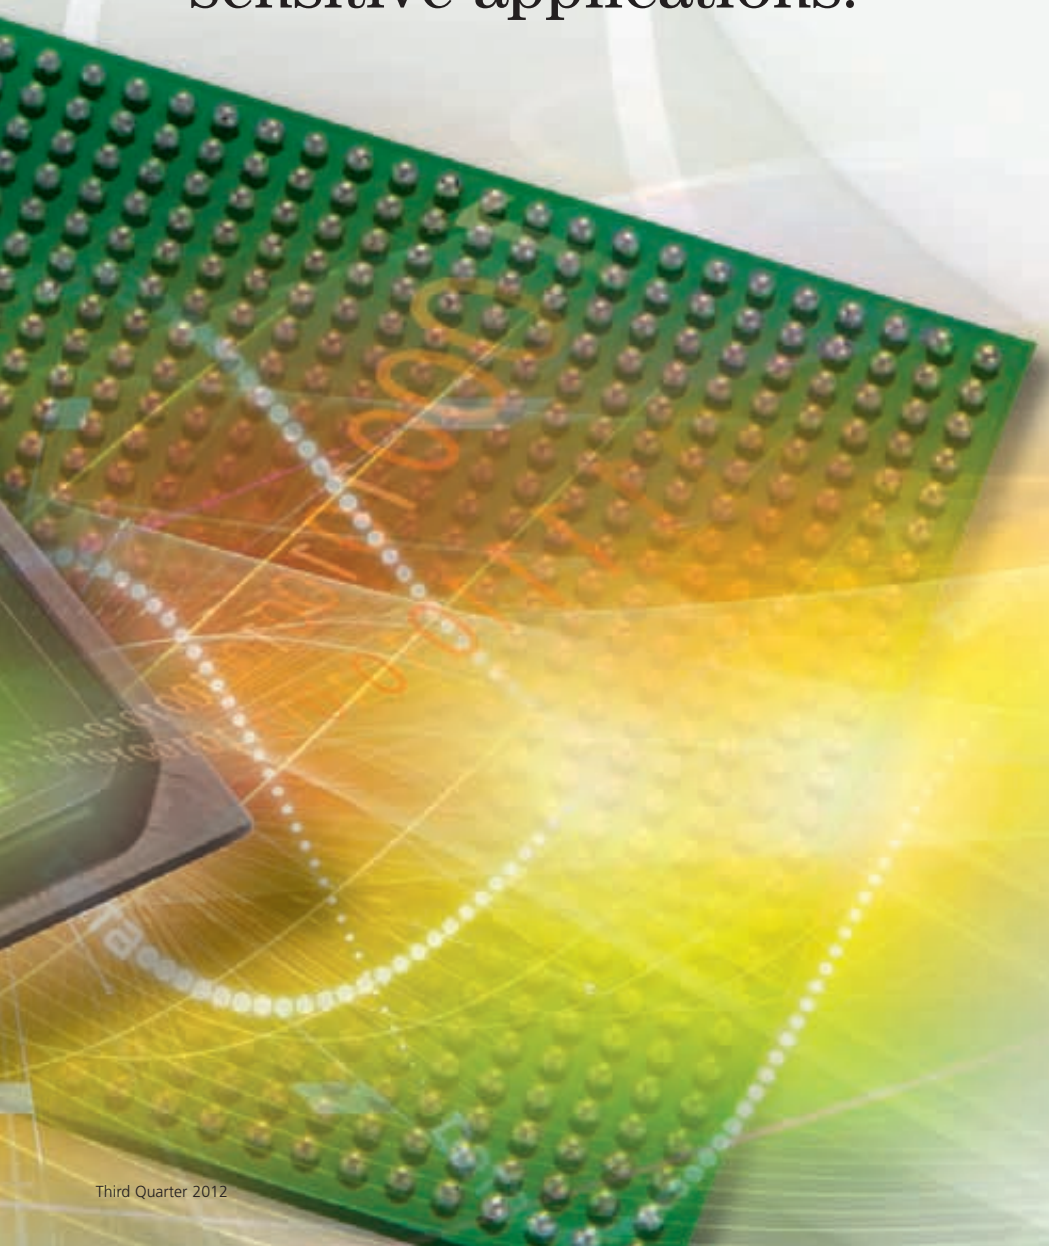

With an eye on helping its customers offer greater value to their customers, Xilinx® in July announced it is shipping its Artix™-7 A100T FPGA, the first of three parts in a feature-rich line of low-cost, low-power, All Programmable devices. The larger Artix-7 A200T and A350T FPGAs will follow in the coming months.

The first Artix-7 device shipment to customers represents another major milestone for Xilinx. It means the company is now shipping FPGAs from all of the families in its 28-nanometer All Programmable device rollout. Xilinx earlier released the world's first 3D IC FPGAs, the Kintex™-7 line, and then broke new ground with the Zynq™-7000 Extensible Processing Platform, which marries an ARM processor and FPGA logic on the same die.

Ehab Mohsen, product marketing manager at Xilinx, predicts the Artix-7 lineup will prove to be a smash hit with customers and will set a new standard for feature-set sophistication, power consumption and ultimately value in what the press has traditionally called the "low end" of the FPGA market. FPGA vendors refer to this sector as the "value-based," "high-volume" or "cost-sensitive" market.

"If you look at the feature set of the Artix-7 family, it's hard to call it 'low-end.' It's certainly the highest-end and highest-value FPGA line in that market to date," said Mohsen. "Where the biggest Spartan®-6 FPGA was 150k logic cells, the Artix-7 family starts at 100k logic cells and runs all the way up to 350k logic cells." Beyond logic cell count, he said, these FPGAs have eight to sixteen 6.6-Gbps transceivers, up to 18,540 kbits of block RAM and as many as 1,040 DSP48E1 slices.

"The Artix-7 family offers twice the performance of the Spartan-6 family and half the power. That's a pretty high-end, 'low-end' FPGA," added Maureen

Smerdon, strategic marketing manager at Xilinx. "In fact, you would have to go to our competitor's 'midrange' line, a more expensive device family, to find a comparable feature set, and even then, the Xilinx Artix-7 family still has advantages."

LEVERAGING HPL AND 7 SERIES' SCALABLE ARCHITECTURE

Reducing power consumption was a top priority for Xilinx's 28-nm generation of devices (see cover story, *Xcell Journal* issue 76). In fact, Xilinx worked very closely with TSMC to formulate TSMC's HPL (high-performance, low-power) 28-nm silicon manufacturing process to a sweet spot for FPGA production. As a result, the entire Xilinx 28-nm line halved total power consumption compared with the previous generation of FPGAs.

"Across all product families, customers had been asking for lower power consumption, but especially so in the cost-sensitive market," said Mohsen. "These devices go into a wide range of applications where lower power is needed for multiple reasons ranging from longer battery life to lower energy costs, better power dissipation, lower BOM (not requiring extra shielding and power circuitry) and smaller end-product form factor."

As such, Mohsen said, the Artix-7 family line takes full advantage of the 50 percent power savings while delivering the needed performance for its target markets. "The 50 percent power reduction provides headroom for additional performance, logic density, I/O bandwidth and signal processing," said Mohsen. That gives designers "the flexibility to either lower power by 50 percent or take advantage of greater performance and capacity at previous power budgets," he said.

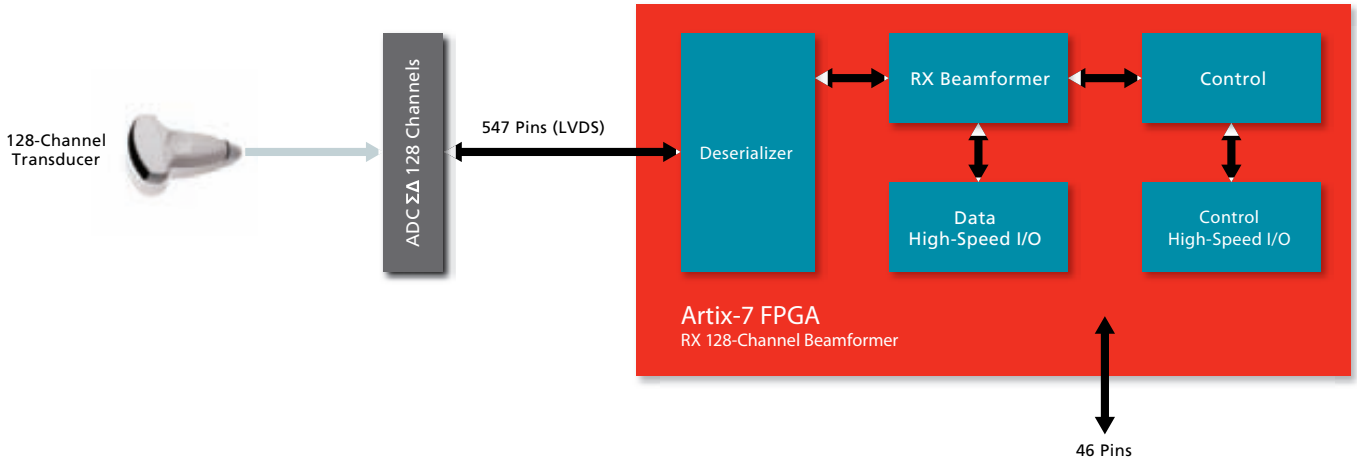

Figure 1 – The Artix-7 FPGA's DSP performance and I/O count can be leveraged for 128-channel portable ultrasound equipment.

Mohsen noted that all of Xilinx's 28-nm All Programmable devices use the same logic architecture. The Artix-7 FPGA's slice architecture is based closely on that of the Xilinx Virtex®-6 and Spartan-6 FPGA families, using the same LUT structure, control logic and outputs. "This scalable architecture provides users with an easy migration path when moving their designs between Spartan-6 and Artix-7 FPGAs," said Mohsen.

MOORE THAN LOGIC CELLS

The Artix-7 is a prime example of how all Xilinx devices are quickly outgrowing the traditional "FPGA" label, which has long tied FPGA advancements merely to a doubling of logic cells every 22 months in keeping with Moore's Law. Even the Artix-7 family—which is in fact Xilinx's smallest 28-nm device—is loaded with programmable system features well beyond logic cells.

Mohsen said that with a block RAM-to-logic ratio of up to 18.5 Mbits within 360k logic cells and 1,040 DSP48E1 slices for the same capacity, the Artix-7 rivals the logic density of competing midrange products while still benefiting from lower power and lower cost. The DSP resources provide up to 1,306 GMACs of DSP performance—three times that of the competition. This signal-processing clout is useful for imaging and communications applications requiring extensive processing capacity.

In addition, the Artix-7 family supports up to 16 configurable 6.6-Gbps transceivers that Xilinx has optimized for low power, giving the Artix-7 the fastest line rates for the cost-sensitive market. These transceivers support pre-emphasis and continuous-time linear equalization (CTLE) to compensate for signal distortion across transmission channels. "With 211 Gbps of

The new device family supports up to 16 configurable 6.6-Gbps transceivers that Xilinx has optimized for low power, giving the Artix-7 FPGAs the fastest line rates for the cost-sensitive market.

total throughput, the Artix-7 is a low-cost alternative for bandwidth-sensitive applications that would otherwise require midrange solutions," said Mohsen.

What's more, Mohsen said that because memory read/write bandwidth can affect overall system performance, the Artix-7 family offers DDR3 data rates of up to 1,066 Mbps, the highest in the industry for FPGAs in its class. The memory solution consists of a flexible controller and physical layer (PHY) for interfacing designs and AMBA® AXI4 slave interfaces to DDR3 and DDR2 SDRAM devices. The controller supports an array of external memories for flexible system design, such as for streamlined access to video and data storage.

As such, the Artix-7 A100T devices are ideally suited for a number of applications that will allow customers to innovate, offer a rich new set of features to their customers and expand their markets. To illustrate, Mohsen described three markets—portable medical equipment, handheld radios and small cellular basestations—that will greatly benefit from the Artix-7 FPGA family's feature set.

PREMIUM VALUE FOR PORTABLE MEDICAL

Mohsen said that companies creating devices for the medical electronics field are eager to expand their product portfolios beyond million-dollar, large-form-factor, hospital-class equipment. They are striving to also offer lower-cost, portable electronic equipment lines to smaller doctors' offices, hospital departments and even individual practitioners.

"Portable ultrasound equipment is a prime example of a market that can greatly benefit from the feature set of the Artix-7 FPGAs," said Mohsen. "Instead of having to wheel a patient into a special room to be tested with a very large ultrasound system, these portable systems are much smaller and can be on a cart or even handheld and brought to the patient. Paramedics can use them in ambulances, and doctors who still perform house calls can use them, too. What's amazing is that with the Artix-7 FPGA line, companies can give their next generation of portable ultrasound equipment many of the advanced features found previously only in high-end systems."

That's not to say that these new classes of equipment will replace those bigger systems, Mohsen added, "because those systems also continue to add incredible new features thanks in part to the rich feature set of our larger Kintex-7 and Virtex-7 FPGA families."

Mohsen said that because the Artix-7 family offers 65 percent lower static

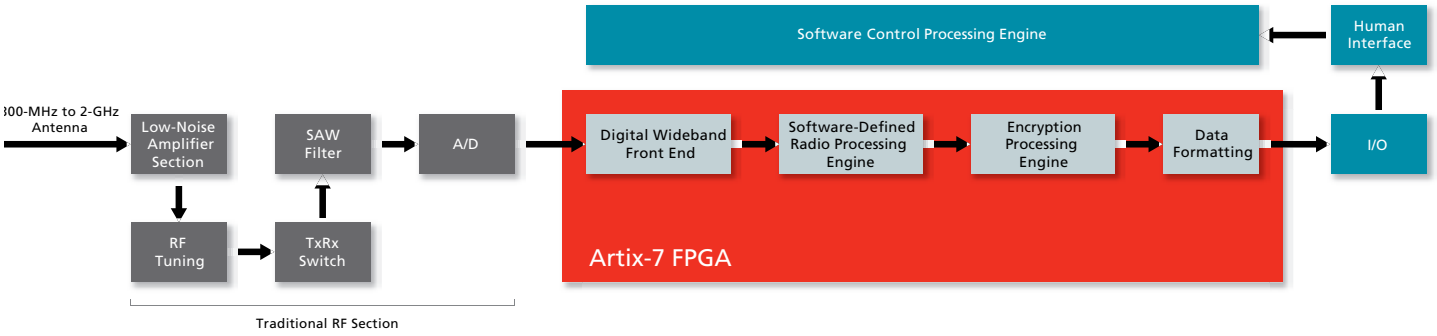

Figure 2 – The system integration and DSP processing capacity in the Artix-7 FPGA is critical for software-radio design.

and 50 percent lower dynamic power consumption than Xilinx's Spartan-6 devices, while delivering up to sixteen 6.6-Gbps transceivers, designers of portable ultrasound equipment can achieve the highest image quality for meeting JESD204B high-speed serial interface standards. At the same time, they can extend battery life and meet safety standards while implementing a 128-channel beamformer at 41 percent less power than alternative FPGA implementations.

Figure 1 shows an example of the All Programmable advantage Artix-7 FPGA affords to the portable ultrasound market.

REDUCING BOM, WEIGHT AND COST FOR MILITARY SDR

Another example of a market that will greatly benefit from the Artix-7 FPGA's rich feature set is military software-defined radio (SDR), Mohsen said. Over the last decade, the U.S. military has been diligently constructing a highly sophisticated worldwide communications network called the Global Information Grid (see cover story, *Xcell Journal* issue 69) that allows U.S. forces and allies to communicate globally and run intelligence and military operations more precisely. While Xilinx's larger Virtex-7 and Kintex-7 FPGAs play an increasing role in the larger communications equipment in the GIG—from networking equipment to aircraft and UAVs—the military is seeking better ways to connect all its assets, even individual soldiers, to the grid more efficiently.

"Many of the portable SDR systems in deployment today suffer from increased power and short battery life," said Mohsen. "They are also too big and heavy as well as expensive. They are fairly complex, too. These systems require extensive DSP processing capabilities to support a variety of radio protocols or waveforms for voice, data and video communications across the globe. Not only is it challenging to support all these waveforms, but they must also be completely secure and able to operate in rugged conditions where radio frequency is difficult. So the military is always looking for better, lighter systems that can run longer and more securely."

All these requirements make the Artix-7 family ideal for SDR systems. Indeed, the new device is particularly well suited for SDR modem management. Mohsen explained that the modem in an SDR system performs baseband signal preprocessing and RF signal improvements, which require enormous amounts of parallel processing and reconfigurability. "FPGAs are a natural fit for this application and most systems today do indeed use FPGAs, but the Artix-7 offers a vast performance improvement," he said. With up to 1,040 DSP slices, the Artix-7 can provide up to 1,306 GMACs of DSP performance—three times the performance of competing FPGAs and far greater than any standalone DSP or GPU.

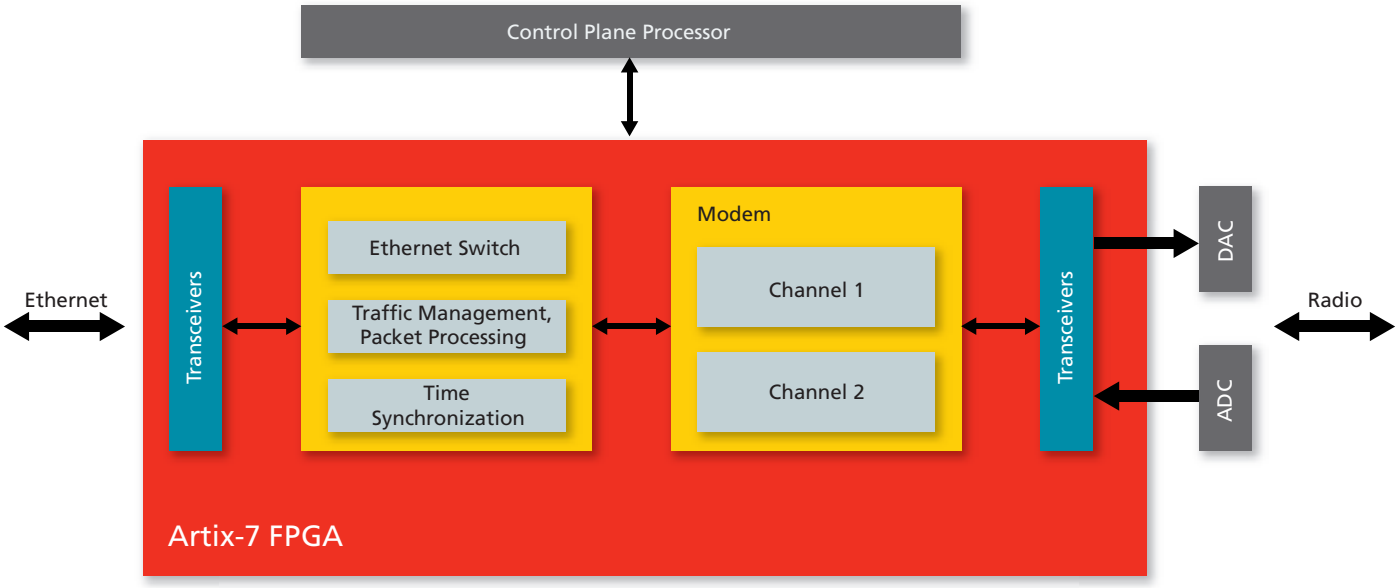

Figure 3 – Designers can integrate multichip functionality for microwave mobile backhaul using the Artix-7 FPGA.

Offering 101,440 logic cells, the Artix-7 is available in a 15 x 15-mm package, which makes it the industry's smallest device at that capacity level. The mix of increased capacity and smaller size allows design teams to create smaller and lighter systems.

Figure 2 shows an example of the All Programmable advantage the Artix-7 FPGA affords in the portable SDR systems market.

WIRELESS BACKHAUL BUILDOUT

Wireless backhaul is another example of an application that will greatly benefit from the Artix-7 family. Mohsen said that the vast majority of the growth in cellular traffic today is occurring in urban and suburban areas. To address this trend, operators plan to boost the capacity of their networks by deploying small cell basestations on lamp posts, traffic lights and even the walls of adjoining buildings. "They need to connect all these small cells in clusters and to the nearest aggregation points, so operators must deploy low-power, low-cost backhaul units whose microwave radio links can span up to tens of miles," said Mohsen.

Whereas a traditional mobile backhaul unit typically supports several Ethernet links, a wireless mobile backhaul forwards the traffic between Ethernet links and radio channels using an internal Ethernet switch. "Both ends of the unit require high-speed transceivers, so that's where the Artix-7 FPGA makes sense as an ideal low-cost alternative to bigger, more expensive devices," said Mohsen. "The Artix-7 family delivers maximal bandwidth with its sixteen 6.6-Gbps transceivers for both Ethernet and RF links using Jedec JESD204B connectivity to data converters."

Artix-7 devices also allow wireless-equipment providers to achieve greater system integration and reduce BOM costs. Half of the backhaul unit contains packet-processing, traffic-management and timing-synchronization functions, Mohsen explained. Meanwhile, the other half of the unit supports modem channels for signal processing. The key requirements for the modem are adequate high-performance DSP processing and high-speed transceivers to interconnect with the data converters for high data throughput.

"The Artix-7 family fits nicely for these functions because it has the right mix of logic density, intellectual-property support and DSP resources," said Mohsen. "The Artix-7 A200T, which will follow the Artix-7 A100T release later this year, has 215,360 logic cells, allowing wireless-equipment companies to create a backhaul solution that integrates all the needed packet-processing, traffic-management, timing and synchronization blocks as well as a single high-speed radio channel into one chip." Likewise, the third member of the family, the Artix-7 A350T device, will allow vendors of wireless-network equipment to integrate two high-speed radio channels on a single chip.

Mohsen also noted that equipment vendors make a concerted effort to ensure the units have a low visual impact and don't appear to clutter the urban and suburban landscapes. This design requirement typically means the units must be very small, which can make it challenging for designers to ensure that each unit effectively dissipates the heat it generates. The Artix-7 family allows equipment vendors to keep power in check while further reducing the overall unit size of their systems.

Figure 3 shows an example of the All Programmable advantage the Artix-7 FPGA affords to the small cell wireless backhaul systems.

The first Artix-7 A100T FPGAs are available today, with production qualification scheduled for the first quarter of 2013. Designers can begin their Artix-7 family designs today using Xilinx design tools. For more information, please visit *www.xilinx.com/artix*.

FPGAs Head for the Cloud

by Michaela Blott
Senior Research Engineer
Xilinx, Inc.
michaela.blott@xilinx.com

Tom English
Research Scientist
Xilinx, Inc.
tom.english@xilinx.com

Emilio Billi
CTO
EB Engineering
e.billi@emiliobilli.com

The last decade has seen the emergence of a new global market for cloud computing. This new paradigm, which delivers computing as a service over the Internet, represents a fundamental shift in the way computers are used. The cloud offers enterprises the means to shift tasks from their local IT infrastructure into remote, optimized computing clusters and thus into the hands of the operator providing the cloud service. For consumers, the cloud delivers storage, video, messaging, social networking, gaming, Web search and many other services coherently across diverse computing devices anywhere in the world.

Robust growth in the data center opens new opportunities for existing and advanced FPGA devices.

At the heart of the cloud computing revolution is the data center, which integrates the compute power, storage and interconnect required to service a global user base. Data centers are experiencing phenomenal growth, which is translating directly into massive investment. According to Synergy Research Group, data center network infrastructure sales grew 22 percent in 2010 alone. Companies such as Google and Facebook are at the forefront of the cloud computing revolution and correctly anticipated the need for colossal data center infrastructure to address a massive global user base.

THE FPGA ADVANTAGE

FPGA technology can bring numerous advantages for computing, storage and networking as data centers strive to become faster, larger, cheaper and greener. Within the networking infrastructure, FPGAs can address the ever-increasing throughput and processing requirements while remaining highly power-efficient. Furthermore, the inherent flexibility of the FPGA is a crucial benefit in this landscape, given the continual arrival of new communications protocols.

On a basic level, FPGAs offer the right physical interfaces and provide the required support and bandwidth for high-speed memory interfaces. They offer sufficient device complexity to implement packet-processing pipelines greater than 100G. Their flexibility allows for the implementation of perfectly optimized custom circuits that operate at maximum efficiency.

Major improvements in high-level synthesis, as for example offered with AutoESL, are helping, to overcome the greatest disadvantage of FPGAs in this space—namely, the low abstraction level of the FPGA programming flow. Finally, a basic FPGA IP portfolio exists that covers the fundamental networking functions. However, more data-center-specific solutions around data center bridging (DCB), VXLAN, virtual switching and other specialized technologies have yet to be developed.

Within servers, FPGAs are an attractive implementation on network interface cards (NICs). Although a plethora of controllers from Intel, Broadcom and others are available for implementing standard adapters for Ethernet and Fibre Channel, FPGAs are ideal when additional processing functions are also integrated on the data path between the network and the CPU. Examples of additional processing functions include encryption, high-frequency trading and TCP offload engines (TOE).

FPGAs are also attractive when either the network interface or the processing function needs to be customized in any way. In these scenarios the FPGA offers high-speed serial transceivers, memory interfaces, PCIe® endpoints and a sufficiently large fabric to accommodate high-throughput data stream processing along with the basic IP blocks. A more sophisticated IP-and-solutions portfolio addressing the specific needs of this market could make FPGAs more competitive in an environment where the end user is accustomed to deploying fully integrated platforms. For example, a more sophisticated TOE IP block capable of thousands of simultaneous sessions (and accompanied by a full Linux driver and TCP/IP stack) would open up a range of new FPGA data center applications.

A special case of such a network adapter is a QuickPath Interconnect (QPI) network adapter. QPI is Intel's proprietary high-bandwidth, low-latency CPU interconnect. Xilinx has developed IP that allows FPGAs to directly attach to the CPU via QPI, significantly reducing latency on the host interface as well as providing higher bandwidth between CPU and network interface. Such a network adapter could be highly attractive within data centers, since latency is rapidly becoming the key performance bottleneck in applications that are already heavily parallelized. A QPI NIC has as much as four times the bidirectional peak bandwidth to the host as a typical PCIe Gen2 server NIC. QPI's higher transfer rates, direct FPGA-to-CPU transfers and shorter headers make it possible to transfer small messages with much

lower latency than in PCIe. As latency becomes the key performance bottleneck in applications that are already heavily parallelized, an ultralow-latency, high-bandwidth QPI NIC becomes a very attractive proposition.

ON THE MOTHERBOARD

We see further opportunities for FPGAs on the motherboard itself. Some common applications in the data center such as in-memory caching are currently implemented on x86-based servers, despite the fact that the x86 is not particularly well-suited to these types of applications. FPGAs could offer a dramatic improvement in performance, power and latency. The trend is to move compute from a distributed number of cores to a more pipelined style of data processing. This approach is beneficial for FPGA architectures. The volume of the silicon fits well with the FPGA opportunities as well. However, the low abstraction level of FPGA programming tools will have to be addressed in order to compete with C compilers on x86-based servers for end users such as Facebook.

On a more speculative note, there is a subclass of servers in data centers that are casually referred to as "wimpy nodes." Xilinx already has many of the key technologies needed to accommodate the new server and SoC architectures emerging for this space, such as ARM processor cores, PCIe interface blocks, memory interfaces and programmable logic. The current Zynq™-7000 Extensible Processing Platform is not yet equipped to compete with ARM-based server SoCs, such as Applied Micro's X-Gene, for this market, but using the technology blocks already available, a future Zynq device could conceivably power a data center server.

Finally, increasing processing demands, especially at the high-performance computing end of the spectrum, can greatly benefit from hybrid computing solutions that combine both FPGAs and CPUs. Existing solutions from Convey and Maxeler showcase the tremendous performance and power benefits of such an approach. For example, Maxeler's implementation of a credit derivative pricing system for a financial customer was up to 37 times faster than software on an Intel E5430 server and reduced energy consumption by more than 97 percent. The QPI technology has the potential to increase these advantages even further, as hardware accelerators can become more tightly coupled with the CPU over a low-latency, high-bandwidth, cache-coherent interface.

DATA STORAGE, WAREHOUSING AND ANALYTICS

Similar to the server and networking scenarios, existing FPGAs can provide competitive implementations within storage, data warehousing and data analytics in three very different ways. First, recent trends integrate flash-based storage closer with the host. A new crop of PCIe SSD controllers allows for direct attachment of flash to PCIe. FPGAs already compete in this space, offering the key functionality and the basic IP building blocks. A further key advantage is the FPGA's flexibility. Currently, the flash-based interface lacks an industry standard, although new standardization efforts such as the Open NAND Flash Interface (ONFi) are on their way.

Then too, FPGAs can assist in acceleration of query processing, handling filtering, decompression and execution of some of the relational operators. This capability could be vital in future storage appliances that need to provide more intelligence in order to cope with throughput bottlenecks. Finally, in so-called "super-storage" applications, FPGAs can play a major role, accelerating file system operations that would otherwise consume considerable CPU cycles. These currently run on separate servers co-located with the storage servers in the storage-area network. FPGA acceleration would allow for a reduction of the control-to-storage server ratio, increasing available storage and performance.

Existing FPGA technology can service these requirements. In particular, the embedded ARM processor in the current Zynq architecture can already tackle OS functionality. Again, a more sophisticated IP-and-solutions portfolio could of course further increase the potential of FPGAs and help accelerate new design developments.

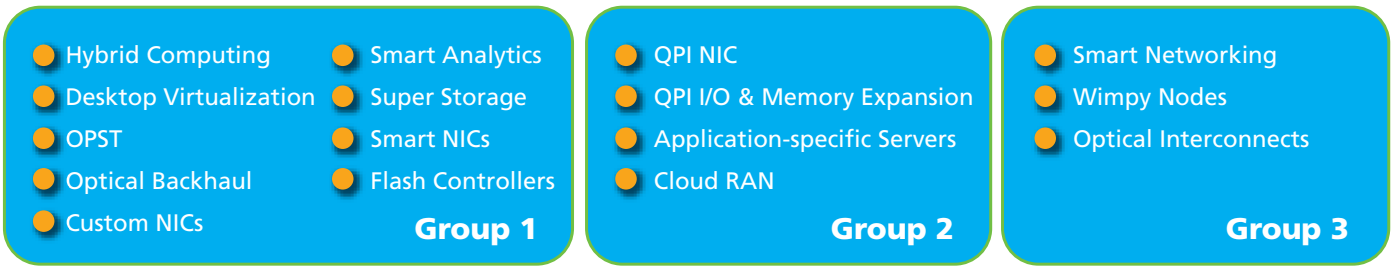

Figure 1 – Data center opportunities in three broad categories await FPGA implementation.

THREE KINDS OF OPPORTUNITIES

As shown in Figure 1, we categorize these varied opportunities in three groups. The first group contains applications that require no additional development effort. Silicon features, IP portfolio, related software and programmability are adequate to address these markets, some of which already employ FPGAs. For example, Intune Networks uses FPGAs to implement optical packet switch and transport (OPST) solutions, claiming up to 300 percent cost reduction in power consumption. Maxeler and Convey Computers offer hybrid computing solutions on an FPGA basis. Smart analytics are based on FPGAs within IBM/Netezza's products. BlueArc demonstrates how super storage can be significantly improved through FPGAs, while FusionIO uses FPGAs for flash controllers. Napatech and Nallatech are among the many vendors offering FPGA-based smart or custom NICs.

The second group addresses opportunities that require some advanced development effort. Most notable in this category are the QPI-related opportunities: QPI NIC and memory and I/O expansion.

The final group includes longer-term opportunities that will involve significant research-and-development effort on silicon features or programming environments. For example, to address wimpy nodes through an FPGA device would take a new generation of Zynq devices that offers more integration of 64-bit ARM processors, as well as wider and faster memory interfaces.

Application-specific servers and C-RAN, or Cloud RAN, lie on the border between groups 2 and 3. Both of them necessitate advanced development efforts to provide necessary infrastructure and platforms, and would benefit from new programming tools with a higher abstraction level. However, the fact that a traditional RTL-based design flow might be adequate to some of the end users moves these opportunities to the second group.

HIGHLY DYNAMIC MARKET

Data centers are a highly dynamic market in which interface standards and protocols are changing rapidly. Particularly within networking equipment, but also for compute functions, this environment offers great opportunities for the deployment of FPGA-based high-speed processing systems. These types of applications are well suited for FPGA implementation with current Xilinx technology, such as the Kintex™ and Virtex® families of devices, leveraging high-speed serial I/O and corresponding IP.

The opportunities for FPGAs in this environment—in particular those within the server space—will be contingent on improvements in FPGA programmability.

The opportunities expand with an additional focus on the needs of the data center market. In particular, memory access (access bandwidth and density), as well as hash and search function support in the form of silicon or IP, are extremely important features within data centers, given that most applications center around large amounts data that must be searched and sorted. Future FPGA devices based on Xilinx's stacked-silicon interconnect (SSI) technology can potentially play a key role here.

Finally, we believe that these opportunities—in particular those within the server space—will be contingent on improvements in FPGA programmability. FPGA programming must be abstracted to a level that is acceptable to programmers of a data center.

FPGA-Based Instrumentation Withstands the Chill of Deep Space

The Xilinx Virtex-5 proves able to endure, operate and survive cryogenic temperatures. No wonder NASA is onboard.

by Alireza Bakhshi
Principal
B&A Engineering Systems Inc.
alib@baengineering.com

Current and future NASA robotic-flight missions to outer planets and asteroids require avionics systems, computers, controllers and data-processing units capable of enduring the extreme low-temperature environments of deep space and lunar and Martian surfaces. With recent technological advances in FPGAs, it has become possible to architect a complete system-on-a-chip (SoC) using a single FPGA. Large FPGAs that are radiation-hardened by design (RHBD) have increased the number of gates per square inch, reduced power consumption per gate and included microprocessors, soft and hard IP, arithmetic modules, sizable onboard memory and analog-to-digital converters.

B&A Engineering (BAENG) conducted studies with the Xilinx® Virtex®-5 mixed-signal RHBD FPGA to address NASA's need for protected, reliable data-acquisition controllers and computer electronics able to operate in cryogenic temperatures. This RHBD FPGA will be the workhorse of future NASA computer and data-handling systems targeted for outer-planet landing, orbiting and sample-retrieval missions.

To conduct the experiment, BAENG designed and built a test board based on a commercial Xilinx XC5VLX30 FPGA and support circuitry (resistors, capacitors and oscillator), as shown in Figure 1. What's remarkable is that we found that the chip works at temperatures well below spec for a commercial part and even well below spec for a space-grade Xilinx FPGA.

The FPGA included circuits using internal phase-locked loops (PLLs), as well as a ring oscillator and a number of basic circuits built from LUTs. During the test, we monitored both circuit functionality and FPGA power consumption.

We disabled all the regulators, switches, resets and configuration-mode pins on the board with the exception of the 100-MHz oscillator. We simulated FPGA and external flash memory voltages and currents, switches, resets and configuration-mode pins and monitored them using external test equipment and power supplies.

Figure 1 – The Xilinx Virtex-5, XC5VLX30 test board

The FPGA was reconfigured at every 10-degree decrement of temperature, from room temperature down to -150°C, from both the JTAG interface (Xilinx IMPACT) and flash (XCF08). We erased and reprogrammed onboard flash during each temperature measurement. Additionally, using Xilinx's ChipSope™ Pro and the internal FPGA system monitor, we monitored Xilinx die temperature, along with 2.5V auxiliary and 1.0V internal voltages. This data provided additional reference

points to the chamber and test equipment monitoring. The internal system monitor can be accessed preconfiguration through the JTAG interface.

TEST RESULTS

Cryogenic testing was done using liquid nitrogen. We started the testing at room temperature of 24°C, having programmed the test chamber to proceed in steps to 10°C, 0°C, -10°C, all the way down to -150°C. Figure 2 charts Xilinx voltage currents vs. the temperature.

In the course of our testing, we made some interesting observations. For starters, we found that the 2.5V/2.5V auxiliary voltage currents remained stable over the test temperature range. Both of these voltages are used for system monitor and JTAG communication. All of the I/Os are tied to 3.3V.

Second, the internal 1.0V current was significantly reduced from 140 mA at +20°C to 81 mA at -150°C. This was no surprise, since a reduction in power is expected at low temperatures.

Finally, we found that the flash memory's 1.8V current remained almost zero down to -50°C and then changed to 10 mA from -50 to -90°C. It dropped to zero from the -100 to -120°C temperature range, and went back to 10 mA from -130 to -150 °C. We are not sure what to make of this finding; it could be due to test measurement errors.

Importantly, both clocks remained stable over the test temperature range (see Figure 3). We used a PLL to both divide and multiply the master onboard 100-MHz oscillator so as to generate the 50- and 150-MHz clocks.

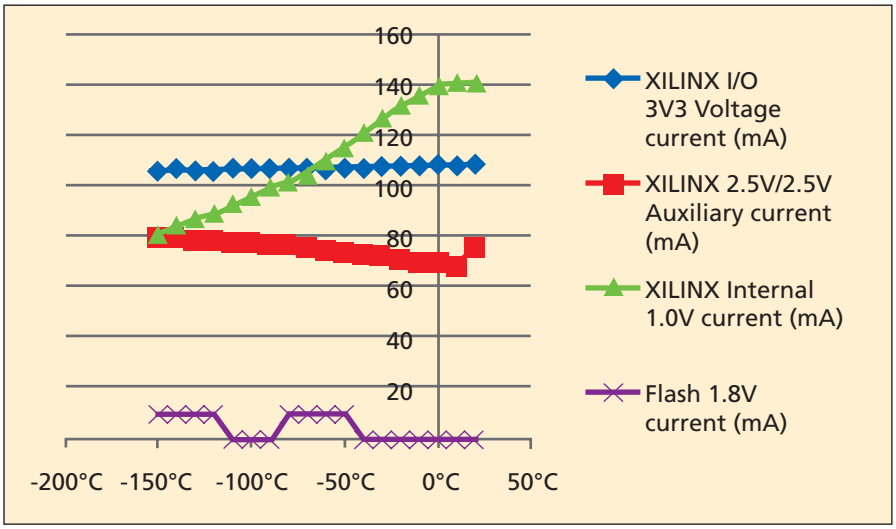

Figure 2 – Xilinx FPGA voltage currents vs. temperature

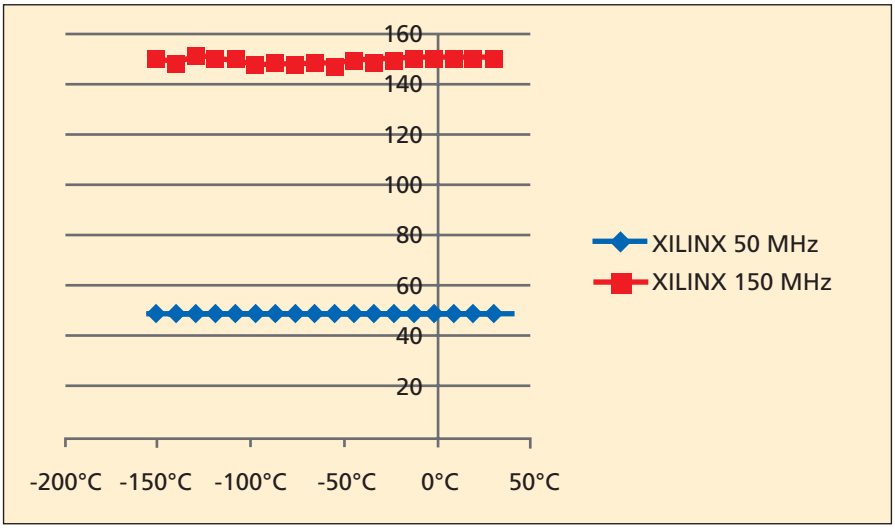

Figure 3 – Xilinx 50/150-MHz clock vs. temperature

XILINX LOGIC

All of the logic circuits remained functional through the test temperature range. However, the same was not true of the flash memory. During testing, we found that the onboard flash memory became unstable at -110°C. Starting at that temperature, it took couple of attempts to program the flash memory from JTAG. Nevertheless we were able to do so after two tries. In addition, the onboard flash memory became nonfunctional at -140°C. We weren't able to program the flash from JTAG thereafter.

The internal 1.0V current increased significantly (384 mA) when we tried to configure the FPGA from nonfunctional flash memory. This result makes sense, since we are not sure what the state of FPGA I/Os would be once the FPGA is configured from nonfunctional flash memory. The 1.0V internal current became normal once the FPGA was configured through the JTAG port. Meanwhile, JTAG communication through IMPACT was functional throughout the test temperature range.

Additionally, ChipScope Pro was functional throughout the test temperature range and with it we were able to monitor the die temperature, 1.0V internal and 2.5V auxiliary voltage. They all tracked very closely with the external LabVIEW and temperature sensor measurements. This is important since it tells us that the Xilinx system monitor including the internal A/D was functional throughout the test.

At the end, we brought the temperature back to ambient, in increments, and left the unit under test to stabilize for 48 hours. In performing our end-to-end test, we were able to configure the Xilinx FPGA from JTAG only once. JTAG communication including the ChipScope Pro stopped thereafter, even though we were able to initialize the JTAG chain through IMPACT. We weren't able to program the flash or configure the FPGA from either JTAG or flash memory. After removing the flash and rewiring the TDI/TDO chain, we were able to configure the FPGA thru the JTAG once again. This is important, since it shows that the FPGA wasn't damaged.

It is important to note that the JTAG chain is serial. The TDO of the IMPACT is connected to TDI of the flash memory, and then the TDO of the flash memory is linked to the TDI of the Xilinx FPGA. This means that the flash memory is sitting between the IMPACT and the Xilinx FPGA.

PROMISING RESULTS

Xilinx FPGA testing using commercial parts showed some promising results as far as reconfiguration at very low temperatures. Reconfiguration through the JTAG interface continued to work down to -150°C. However, due to what appears to be a failure of the flash memory chip at -130°C, we were not able to reconfigure the Virtex-5 chip from flash memory below that temperature.

Internal current steadily declined on the internal 1.0V as expected (and internal power consumption) and ended at 66 percent of where it was at room temperature. Basic circuits, ring oscillator, shift registers and PLL outputs continued to function normally with hardly any detectable changes.

Using Spartan Technology to Support Development of Green Energy

The Xilinx Spartan-3A FPGA augments control algorithm implementation for a multiterminal DRI power inverter.

by Phillip Southard
Senior Design Engineer
PDS Consulting, LLC
phillip.southard@pds-consulting.com

Product development for industrial applications involves extensive research and preparation in an environment of rolling deadlines and ever-evolving product specifications. While time-to-market for this sector may not be as short as it is for consumer electronics, products must ship quickly and with as many essential functions, features and potential hooks for the next generation as possible. Companies vie to be industry leaders in their respective competitive arenas—especially in new markets such as green power, which in their infancy and without defined leaders require pioneers to design, develop and deliver new products. Success depends not only on an inspired, dedicated team of engineers, advanced computing technology and new materials, but also on angel investors or government agencies to provide grants for promising approaches to improved energy generation, distribution, monitoring, metering and consumption.

In the fall of 2011, engineers from Princeton Power Systems (PPS), a New Jersey-based manufacturer of advanced power-conversion products and alternative-energy systems, demonstrated their latest green power product. This demand response inverter (DRI) was the result of a three-year collaboration between PPS, the United States Department of Energy and Sandia National Laboratories' Solar Energy Grid Integration Systems (SEGIS).

The resulting multiterminal DRI (Figure 1) is uniquely flexible to be more reliable, more efficient and more cost-effective than currently available inverters. Equipped with multiple AC and DC terminals, the DRI can route power to the grid, a microgrid, DC energy storage or dynamic loads. Programmable power curves and charge profiles enhance control for generators, loads and batteries, ensuring greater efficiency. And the use of advanced high-capacity long-lifespan switches maximizes reliability.

Princeton Power Systems showcased features of the DRI that improve electrical-grid interconnectivity and efficiency, enhance the performance of renewable energy systems and allow for better integration of electric vehicles and distributed power generation. The DRI was part of the company's "An Island in the Sun" microgrid demonstration (Figure 2), which detailed key advancements in clean technology and manufacturing, including a 200-kilowatt solar array and lithium-ion battery system.

A microgrid can operate independently of a major utility grid to supply reliable, low-carbon-emission energy. PPS' DRI is compatible with AC generators such as diesel or gas, and with photovoltaic (PV) or wind inputs. A small community using a DRI is less dependent on the grid and can reduce its carbon footprint and utility costs. The DRI can also provide grid services, PV with storage and charging for electric vehicles.

XILINX SPARTAN TECHNOLOGY

To meet the demands of industrial product design, companies like Princeton Power Systems leverage flexible development vehicles such as Xilinx's Targeted Design Platforms (TDPs), with their rich ecosystem of design services support. In this case, however, the engineering team faced an initial challenge of determining how to expand the inputs and outputs of the DRI system's digital signal processor, and how to implement control and communication interfaces that functioned in parallel. PDS Consulting

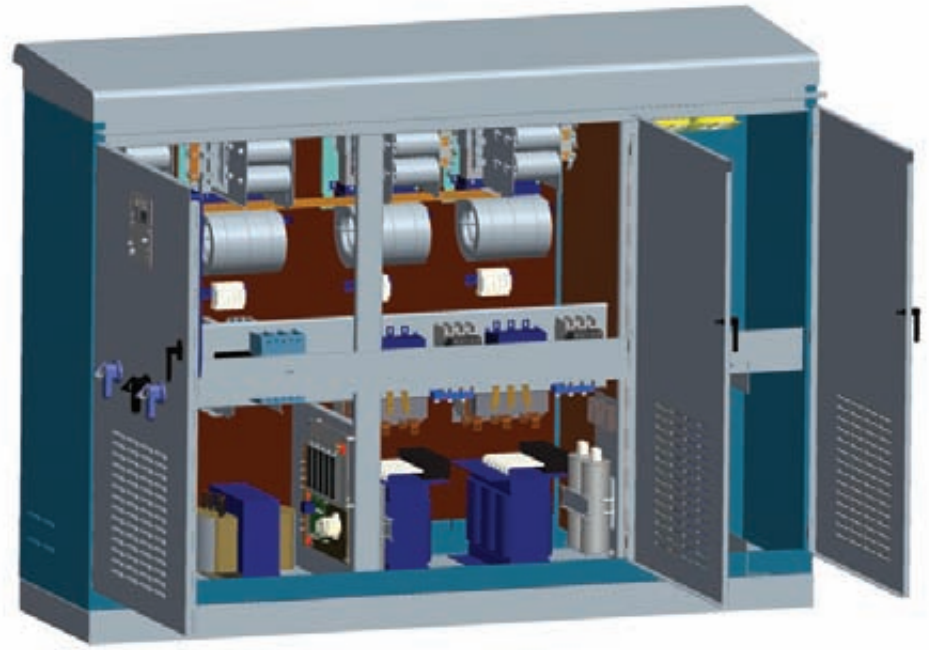

Figure 1 – The flexibility in Princeton Power Systems' demand response inverter comes from FPGAs.

Figure 2 – Princeton Power's flexible, multiterminal DRI is here configured for an electrical microgrid.

offers design services in programmable digital systems for a variety of markets including aerospace and defense, broadcast, industrial, scientific and medical. The firm supported work on the project as a member of Xilinx's Alliance Program.

The PDS Consulting team provided on-site, hands-on system debug and PCB bring-up, as well as off-site RTL and IP design services. We also advised Princeton Power Systems developers on how to implement the system control interface for their green power control algorithm. In the end, engineers chose a Xilinx® Spartan® XC3SD3400A FPGA married to a DSP as a prime system control component (Figure 3).

The Spartan-3A FPGA, with its extensive SelectIO™ capabilities, offered flexibility in implementation, particularly for trigger signals and ADC input channels. Xilinx's Spartan-3A family is a superior alternative to mask-programmed ASICs because these FPGAs permit design upgrades in the field and avoid the high initial cost, lengthy development cycles and inherent inflexibility of conventional ASICs. The integrated technology afforded by the Spartan-3A made the implementation of Princeton Power Systems' patented control algorithm for green power conversion a possibility.

It took more than 300 I/Os to implement the DRI system interface, which enabled access to 8 Mbytes of flash, a 256-Mbit SDRAM and USB/RS-232 at >900 kbits/second. In addition, the team also utilized the generous amount of fast, distributed 32-bit dual-port RAM inherent to the Spartan architecture. The configurable logic block (CLB) lookup tables used as dual-port RAMs enabled the efficient local storage of new energy waveform samples that the ADCs supplied, while the DSP read the previous samples and a PicoBlaze™ embedded processor analyzed new values from the second port concurrently.

THE BENEFITS OF XILINX FPGAS

Princeton Power Systems' algorithms required extensive calculations that can only be accomplished by floating-point DSPs, which traditionally do not have the same features as FPGAs. Some of the features of Xilinx FPGAs that particularly suited the PPS project included multivoltage, multistandard SelectIO I/O pins; configurable logic blocks; block RAM; and memory interfaces that can implement a large number of programmable trigger signals. These signals generate and execute pulse trains that trigger power electronic switches like IGBTs and control a large number of fast ADC channels to read important system measurements on every pulse or custom high-speed serial interfaces.

FPGAs not only allowed Princeton Power Systems to design and implement custom peripherals that matched its specific requirements, but also provided additional computational resources for the processing of input values, which otherwise would have to be done by the DSP. The Spartan-3 FPGA-based design completes several processes: It accomplishes system error checking using the values read from ADCs connected to the DSP. It implements timer-driven activities like reading ADCs precisely when necessary. And it does an averaging of ADC values.

Without the FPGA, some of these functional requirements would have been impossible to implement. Other functionalities would have required more components on the DRI's control board or a significantly more complex software architecture. The PPS team knew it was crucial to avoid the latter,

since the control board acts as the heart of the DRI system.

"While an increasing number of DSPs now offer peripherals that were previously absent, the importance of having an FPGA still remains," said Frank Hoffmann, the R&D manager at Princeton Power Systems. "With each new generation, the amount of computational resources inside the FPGA increases—for example, from a Spartan-3 to a Spartan-6—and it has now become possible to outsource more computational work to the FPGA. And this could mean running our complex control algorithms faster and therefore improving the quality of a generated output like the one in the DRI."

THE BOTTOM LINE

While the technical benefits of using an FPGA are clear (quick prototyping, flexible architecture, advanced support tools like Xilinx's ChipScope™ Integrated Logic Analyzer for quick in-system debug), the decision has also affected Princeton Power Systems' bottom line.

"Using an FPGA has made development much faster, reducing R&D expenses and time-to-market for new and innovative alternative-energy systems," said executive vice president Darren Hammell. "The programming environment was easy to use and enabled us to rapidly develop and test our innovative software. This enabled us to complete the prototype for the demonstration much quicker than otherwise would have been possible." The product is now shipping, and PPS has added two new customers: BMW and SuperPlug have included a DRI in new power system designs.

In fields like green power technology, engineers face new challenges, including determining how to optimize algorithm implementation while retaining necessary functionality. With the right tools, technology and team, enhancements in this field lie just within reach.

For more information on Princeton Power's multiterminal DRI, please visit *http://www.princetonpower.com/prod_demand.shtml*.

You can reach PDS Consulting at *sales@pds-consulting.com*.

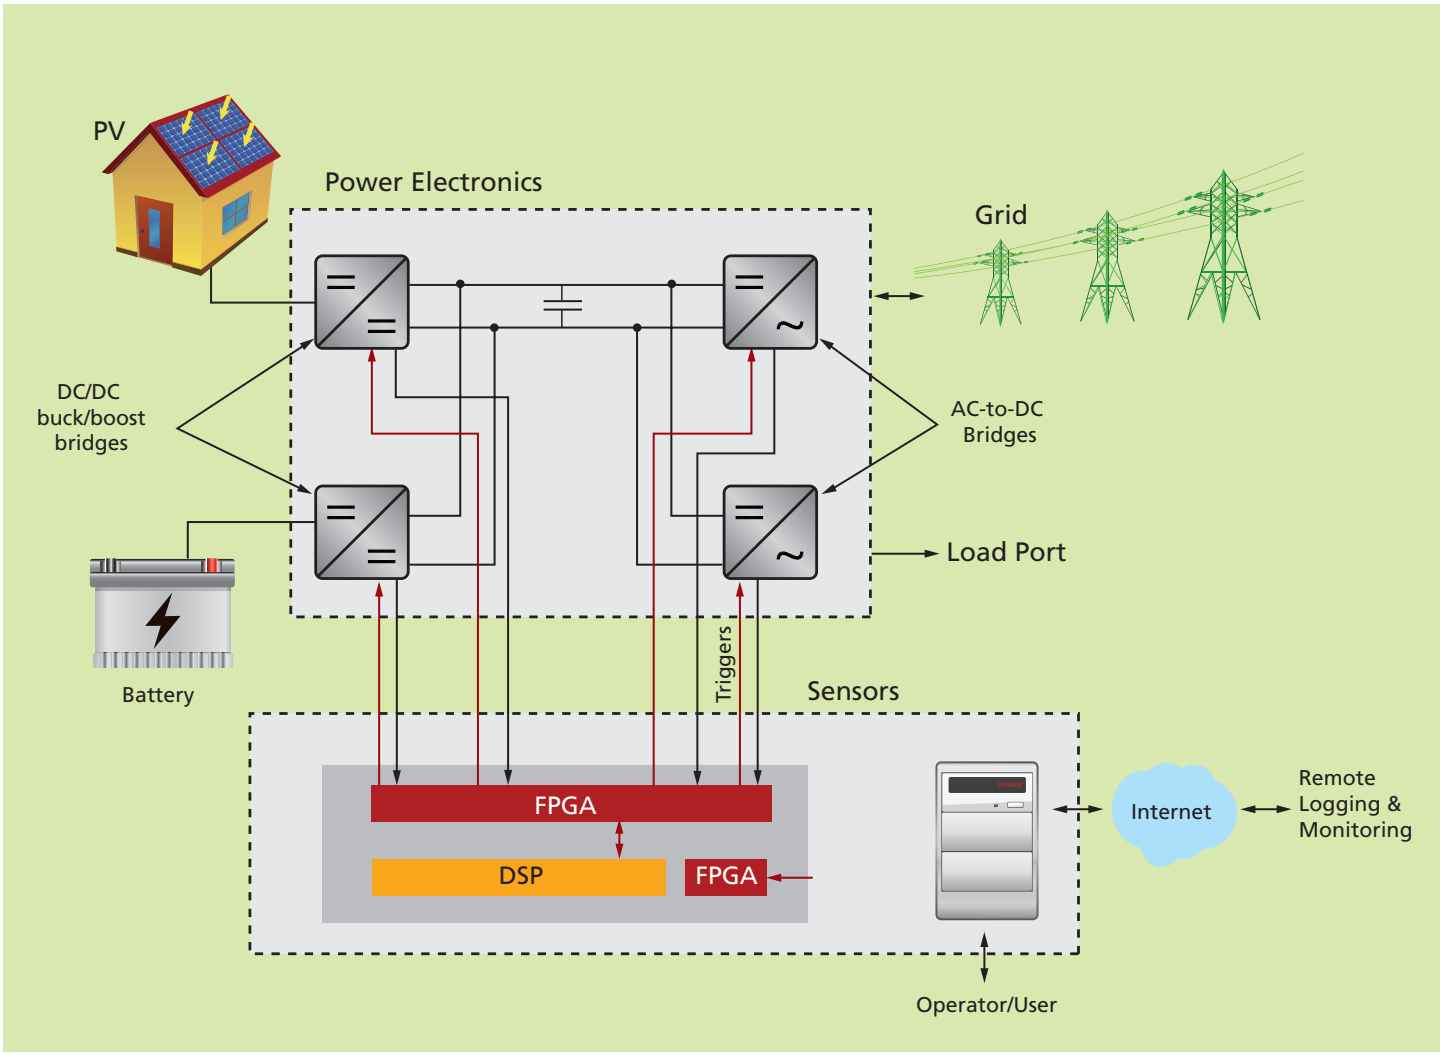

Figure 3 – Engineers chose a Spartan-3A FPGA, with its extensive SelectIO capabilities, as the main system peripheral.

Designing a 19-nm Flash PCIe SSD with Kintex-7 FPGAs

by Yilei Wang
Senior Hardware Engineer
Memblaze China
yilei.wang@memblaze.com

Xiangfeng Lu
CTO
Memblaze China
xiangfeng.lu@memblaze.com

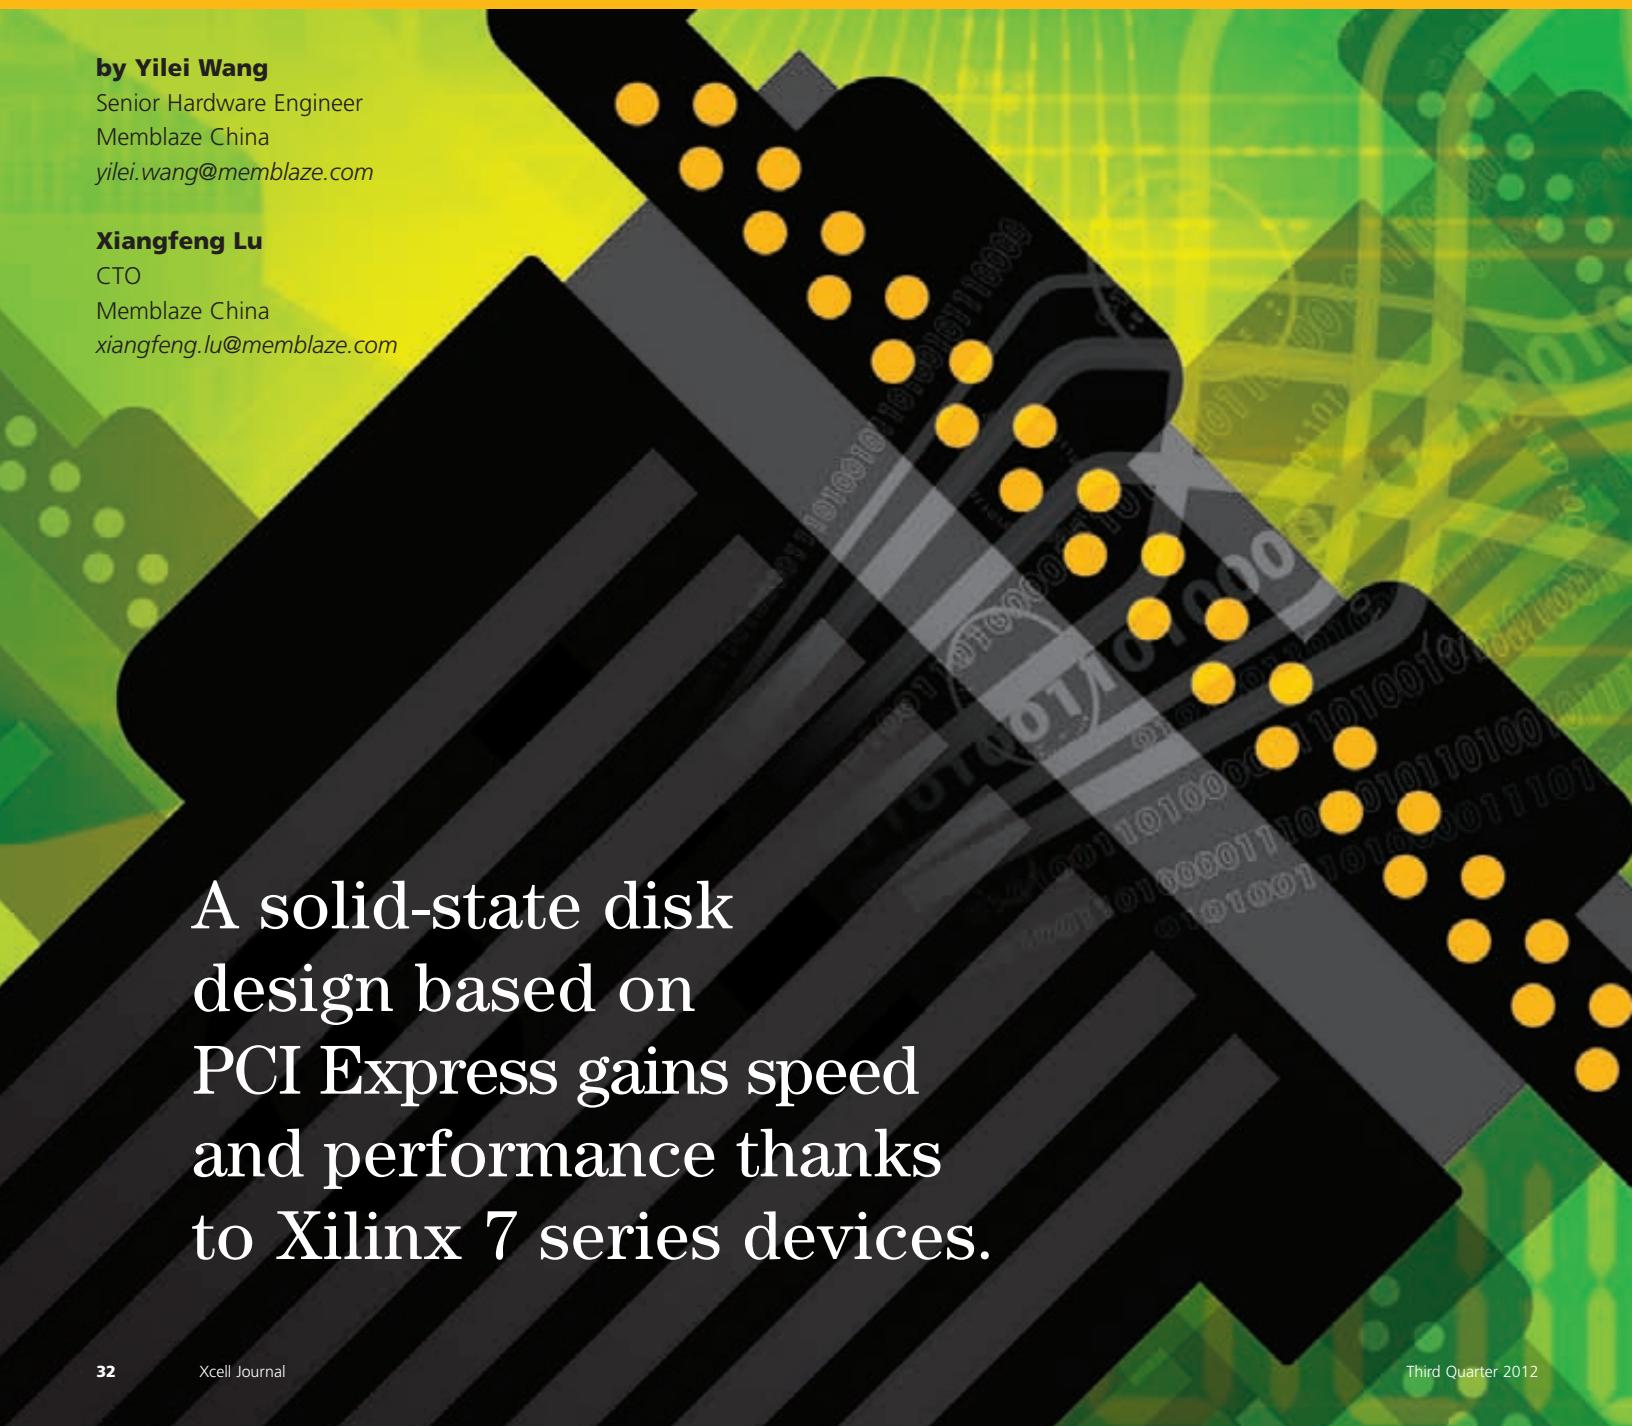

A solid-state disk design based on PCI Express gains speed and performance thanks to Xilinx 7 series devices.

Solid-state disk (SSD) technology based on NAND flash memory provides higher throughput and lower power consumption than traditional mechanical-drive-based storage systems. For that reason, SSD usage has mushroomed over the last decade, moving from handheld devices to laptop and desktop computers and, now, making incursions into the enterprise storage market. The rapid rate of expansion has been further aided by the enterprise storage industry's adoption of SSDs based on the Serial Advanced Technology Attachment (SATA) standard.

However, as SSD manufacturers look toward next-generation systems that achieve new performance and density highs by using flash memory that is implemented in 19-nanometer process technology, SATA hasn't kept up. Even with the latest revision (SATA 3.0), the 6-Gbps physical interface hardly meets the highest throughput of the SSD NAND flash arrays, and thus leaves extra performance on the table.

To break the interface bottleneck, SSDs based on PCI Express® are making a huge impact on the market. PCIe® is an industry-standard local bus with higher performance and scalability than SATA. It is based on multilane high-speed serial links that support one to 16 lanes, each operating at up to 8 Gbps (2.5 Gbps for Gen1, 5 Gbps for Gen2, 8 Gbps for Gen3). The PCIe interface for SSDs supports gigabyte throughput and better margins for the foreseeable future as NAND flash technology evolves.

However, creating a PCIe-based SSD system using 19-nm flash has its share of challenges. The PCIe interface requires more high-speed serial links and more-complex interconnect than SATA. The throughput demands require the PCIe direct memory access (DMA) to operate at a gigabyte bandwidth level. In addition, at the 19-nm process node, flash reliability—or specifically, the metric known as "wear" (the number of times a NAND can read or write before encountering an error)—is a growing issue. At 19 nm, companies must perform wear leveling and error correction faster than ever before.

Xilinx® Kintex™-7 FPGAs establish a new benchmark for FPGA high-end performance at less than half the price of previous-generation FPGAs. The Kintex-7 family is one of four product lines Xilinx built using TSMC's HPL (high-performance, low-power) 28-nm process, designed for maximum power efficiency and delivering a twofold price/performance improvement while consuming 50 percent less power than previous generations. Kintex-7 FPGAs offer high-density logic, high-performance transceivers, memory and DSP, plus Agile Mixed Signal—all to enable higher system-level performance and the next level of integration. These capabilities allow for continued innovation and differentiation in designs at volume price points. As such, Xilinx's Kintex-7 series FPGAs are ideally suited for use as 19-nm flash PCIe SSD controllers.

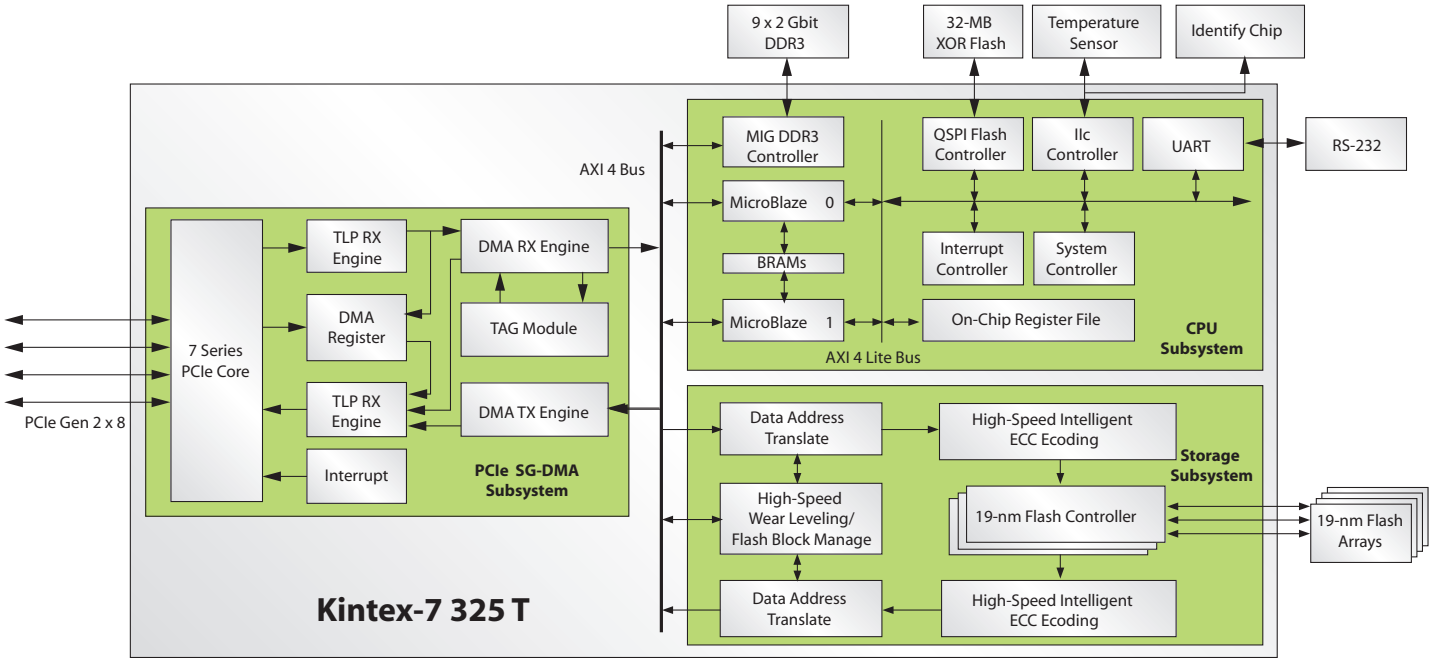

Figure 1 – The Kintex-7 SoC solution for a PCIe 19-ns NAND flash SSD consists of three subsystems: CPU, storage and PCIe SG-DMA.

Figure 1 shows the Memblaze SSD controller architecture, featuring three subsystems interconnected with a high-speed AXI4 bus. The PCIe SG-DMA subsystem, which includes the Kintex FPGA hard core, scatters and gathers data between the host computer and the SSD data buffer (the "SG" stands for scatter and gather). The CPU subsystem manages peripherals and executes SSD access commands, while the storage subsystem manages the SSD sector data processing with a multichannel NAND controller, error-correcting code (ECC) block and wear-leveling block. These three subsystems share a 2-Gbyte DDR3 SDRAM with ECC function. It's easy to generate an ECC DDR3 SDRAM controller with Xilinx Memory Interface Generator (MIG) tools.

In our design, the 7 series PCIe hard core implements the physical-to-TLP layer and allows the design to function as a high-performance PCIe endpoint with minimal latency. The new embedded MircoBlaze® core with ARM® AXI4 interconnect completely removes the bottlenecks of the on-chip bus. The DDR3 hard core provides a 51.2-Gbps ECC solution for the disk cache. Meanwhile, the low-power logic resources make it easy to achieve high performance of wear leveling and intelligent ECC algorithm execution. In addition, abundant high-performance I/O resources provide an easy way to interconnect to the 19-nm NAND flash arrays.

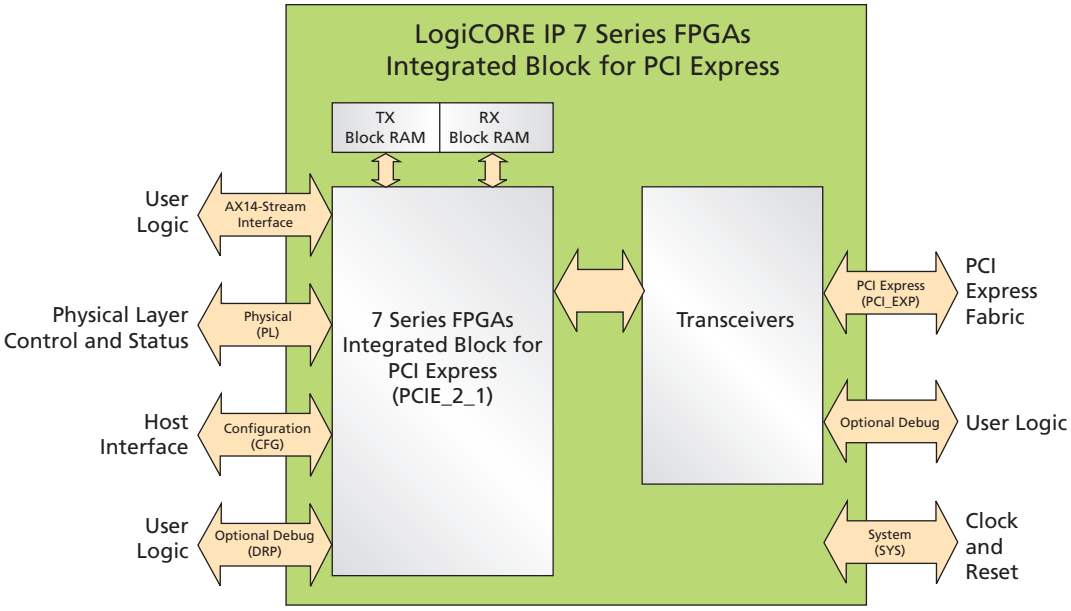

Figure 2 – Top-level functional blocks and interfaces in the PCI Express hard core

PCI EXPRESS SG-DMA

Our design's PCIe interface required a fast DMA controller to implement high-speed communications between the host and the local AXI4 bus. The throughput of the SSD flash arrays can reach up to 2.5 GBps. To simplify the PCIe interface design and get greater margin as flash chips evolve, we chose to use an eight-lane PCIe Gen2/Gen3 architecture.

The PCIe endpoint has many complex protocols to process in the physical, data link and transaction layers. Luckily, designing the PCIe SG-DMA controller in the Xilinx 7 series FPGAs was quick and easy. The PCIe hard core, which Xilinx had implemented in the device's fabric, handled all PCIe operations. This allowed our design team to focus on the functions of the SG-DMA operation itself. The integrated block for the PCIe solution supports one-lane, two-lane, four-lane and eight-lane endpoint configurations at speeds up to 5 GBps (Gen2), compliant with the PCIe Base Specification, rev. 2.1. Table 1 shows the 7 series FPGAs' integrated block for PCIe configurations. The core can be configured as Gen1/Gen2 and for maximum support to x8 lanes, providing up to 40-Gbps bandwidth.

We used CORE Generator™ tools to configure and generate the PCIe endpoint IP, which includes the user guide, source code, simulation code and example design—all of which helped us get up to speed quickly using the core. Figure 2 shows the PCIe hard core's top-level functional blocks and interfaces.

The main function of the SG-DMA core is to process TLP packets from the host and respond. SG-DMA operates as a PCIe master access to the host memory, moving data between the host and local memory. The host

		Artix-7	Kintex-7	Virtex-7 T	Virtex-7 XT	Virtex-7 HT
PCIe	Gen (Integrated block)*	Gen2	Gen2	Gen2	Gen3	Gen3
	Width	x4	x8	x8	x8	x8
	Number of Blocks	1	1	3-4	2-4	1-3
	Serial Date Rate (Gbps)	5	5	8	8	8

*Based on symmetric filter implementation

Table 1 – 7 Series FPGA integrated blocks for PCI Express

sends commands to the DMA controller to control the DMA access. The command code is embedded in the data of a specific host TLP register write. The SG-DMA controller initiates the SG-DMA write request to move data from the local memory to the host memory in response to the host's read commands. Similarly, for host write commands, the SG-DMA controller initiates a DMA read request to move data from the host memory to local memory. Figure 3 illustrates the flow.

AXI4 INTERCONNECT

The AXI interconnect IP connects one or more AXI memory-mapped faster devices to one or more memory-mapped slave devices. The AXI interfaces conform to the AMBA® AXI version 4 specifications from ARM, including the AXI4-Lite control register interface subset. The interconnect IP is intended for memory-mapped transfers only; AXI4-Stream transfers are not applicable. The AXI interconnect IP can be used as a pCORE from Xilinx's Embedded Development Tool Kit (EDK) or as a

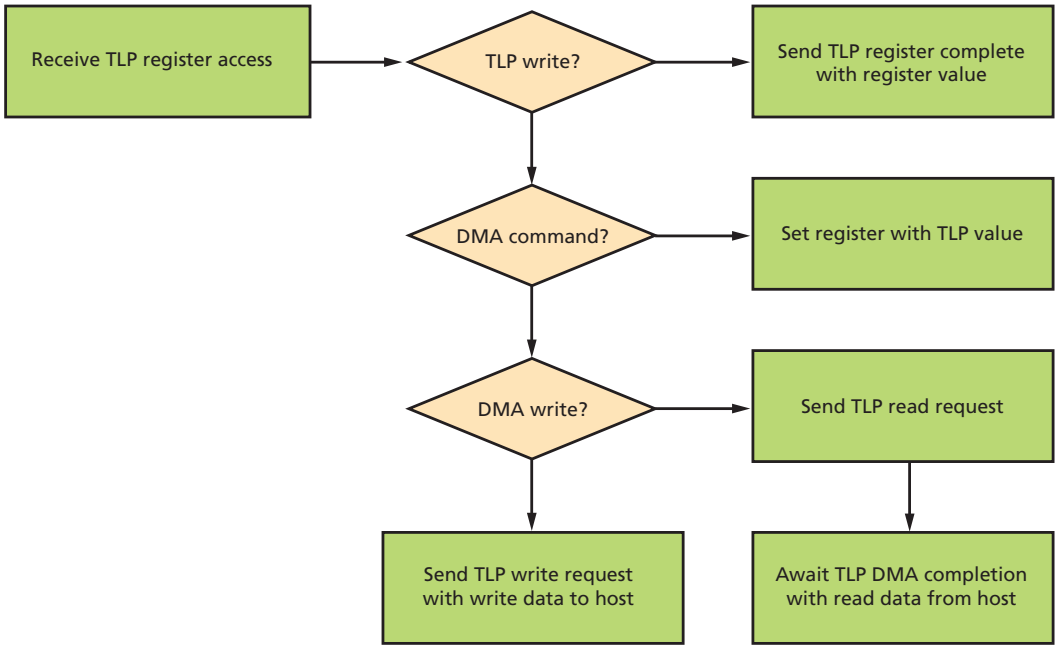

Figure 3 – Operation of the SG-DMA controller

standalone core from Xilinx's CORE Generator IP catalog.

The designer can select from two modes of operation that the Xilinx AXI4 IP supports. The performance-optimized crossbar mode has a shared-address, multiple-data (SAMD) crossbar architecture with parallel pathways for write and read data channels. The area-optimized shared-access mode features shared write data, shared read data and single shared address pathways. Both of these modes support burst lengths up to 256 for incremental (INCR) bursts, and variable data width from 32 up to 1,024 bits. Propagated USER signals are also supported on each channel, if any; an independent USER signal width per channel is optional.

The AXI4 interconnect provides high performance between the PCIe SG-DMA and the DDR3 memory. We found that the AXI4-Lite shared bus is also a perfect solution for the low-speed on-chip interconnect, requiring minimal consumption of logic resources.

WEAR-LEVELING TECHNOLOGY

Wear leveling is a design technique that storage-media companies employ to prolong the service life of various kinds of erasable computer storage types, such as the flash memory used in solid-state drives. There are a few wear-leveling mechanisms used in a flash memory systems, each with varying levels of longevity enhancement.

A flash memory storage system without wear leveling will not last very long if it is writing data to the flash. Without wear leveling, the flash controller must permanently assign the logical addresses from the operating system (OS) to the physical addresses of the flash memory. This means that every write to a previously written block must first be read, erased, modified and rewritten to the same location. This is very time-consuming, and highly written locations will wear out quickly, even when other locations on the flash are completely unused. Once a few blocks reach their end of life, the drive is no longer operable.

The first type of wear leveling is called "dynamic wear leveling." It uses a map to link logical block addresses (LBAs) from the OS to the physical flash memory. Each time the OS writes replacement data, the map is updated to mark the original physical block as invalid data, and a new block is linked to that map entry. Each time a block of data is rewritten to the flash memory, it is written to a new location. However, blocks that never get replacement data sit with no additional wear on the flash memory. The drive may last longer than one with no wear leveling, but some blocks, while still remaining active, will go unused.

Another technique, called "static wear leveling," also uses a map to link the LBA to a physical memory address. Static wear leveling works the same way as dynamic wear leveling except the static blocks that do not change are periodically moved so that other data may access these low-usage cells. This rotational effect enables the SSD to operate until most of the blocks are near their end of life.

Figure 4 shows flash pages with and without wear leveling after a long write/erase operation. The one without wear leveling, with black pages, is broken and can no longer record any data, while the one with wear leveling still functions with all pages.

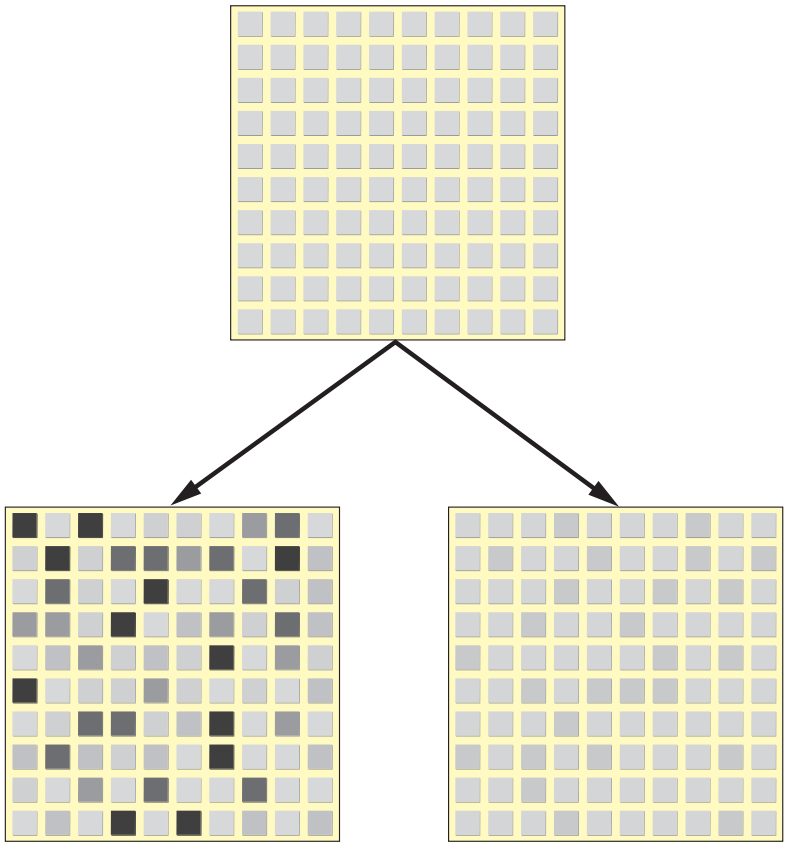

Figure 4 – Flash pages with and without wear leveling

INTELLIGENT ECC ALGORITHM

Another key component of SSD system design is error correction. There are a number of anomalies that can cause bit errors, which in turn can affect data integrity and even the proper operation of the system itself. To deal with these errors, our design team employs complex ECC algorithms that get even more

elaborate when we use new, smaller-geometry flash in these systems.

One ECC algorithm we use for 19-nm NAND flash memory is called an "anti-random data error record." The algorithm addresses bit errors caused by temperature changes, noise and reliability of the storage cell. In addition, the storage cell of NAND flash normally has limited lifetimes of erasing/programming. The bit error rate (BER) increases with the accumulation of erasing/programming operations until the limited lifetimes run out. The SSD ECC function requires the algorithm also to detect the BER of each cell and to understand their lifetimes. Designers set a certain BER threshold to indicate that a lifetime has been reached and to identify a replacement block. However, optimizing this threshold is critical. A BER threshold that's too low can cause the system to abandon a reliable cell too early, ultimately reducing the SSD's lifetime. On the other hand, a higher BER threshold runs the risk of losing data as it tries to write to an unreliable cell. Thus, an ECC algorithm must find the proper balance between reliability and lifetime.

The 19-nm NAND flash offers more density in storage but less reliability. That's why our design introduces high-speed, high-level error correction. The ECC part occupies more than 35 percent of the design resources, implementing a parallel computing ability of a maximum 49 bits of error correction in a 1,024-bit sector, at 4-Gbyte read speeds. The new 28-nm Kintex-7 technology increases system-level performance by up to 50 percent and increases capacity twofold while lowering total power consumption by up to 50 percent over the previous-generation FPGAs. Compared with the same ECC block in a Virtex-5 device, our Kintex-7 implementation reduced the area by 5 percent while increasing the performance more than 40 percent and maintaining the same cost.

Xilinx Kintex-7 series FPGAs are ideally suited for 19-nm flash PCIe SSD designs. The hard PCIe core, the performance, capacity and low power make it the perfect chip for this market. With this device, our SSD throughput easily maintains 2 GBps for both writing and reading. The device allows us to bring a great value to our customers and gain a large margin for our 19-nm NAND flash system.

How Partial Dynamic Reconfiguration Helped Make an FSK Demodulator

The ability to reconfigure a portion of a Xilinx FPGA on the fly allowed a European research team to create a more dependable system.

by Fabio Giovagnini
Software Manager
Aurion S.r.l., SESM S.C.A.R.L.
fabio.giovagnini@gmail.com

Antonio Di Marzo
Embedded System Dept. Manager
SESM S.C.A.R.L.
adimarzo@sesm.it

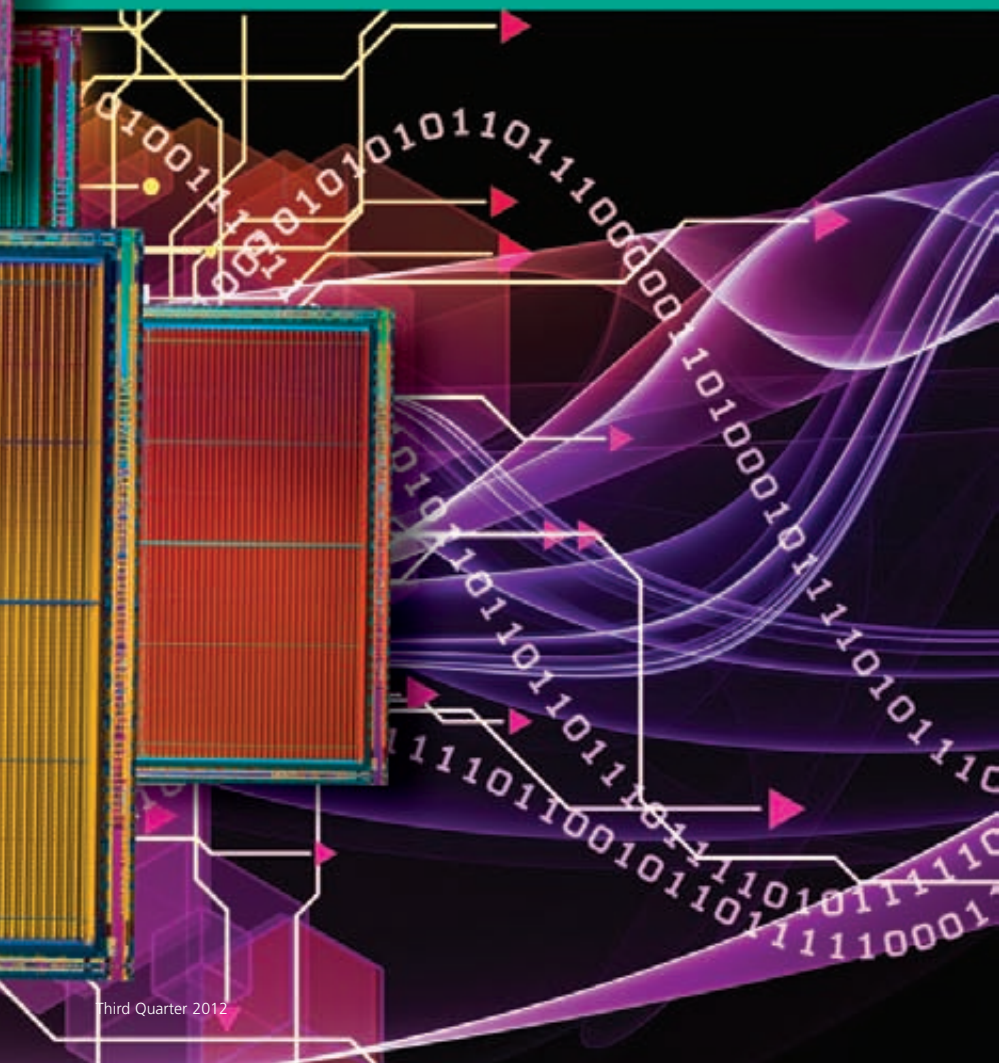

Partial dynamic reconfiguration is a radical new way to configure and reprogram an FPGA. Unlike the standard FPGA reconfiguration process, PDR allows you to change a small part of the device based on the needs of your design, while other parts are still running. There is no need to hold the device in reset while an external controller or internal piece of glue logic reloads a design onto it, the standard reconfiguration methodology. With PDR, critical parts of the design continue operating while a controller, either on or off the FPGA, loads a partial design into a reconfigurable module. The technique pays off in hardware resource optimization and a reduction in power consumption.

The PDR method has emerged as a topic for investigation in the context of the European Union research project pSHIELD. This project aims at pioneering techniques to build security, privacy and dependability (SPD) into embedded systems, rather than tacking them on as "add-on" functionalities. The idea behind pSHIELD is to take a first step toward SPD certification for future embedded systems. The leading concept is to demonstrate the composability of SPD technologies.

In such a context, we have identified PDR as a key technology to implement a secure, dependable and reconfigurable embedded system. Our investigation of this new technology involved implementing a project demonstrator—a reconfigurable frequency shift-keying (FSK) demodulator system—within the Xilinx® PDR design flow.

FREQUENCY SHIFT-KEYING ADAPTABLE DEMODULATOR

The FSK adaptable demodulator is a proof of concept that we developed to demonstrate the pSHIELD SPD paradigm. In fact, it implements a simple system managing a data stream. Figure 1 shows the block diagram of the hardware implementation of the A-FSK demodulator SPD node.

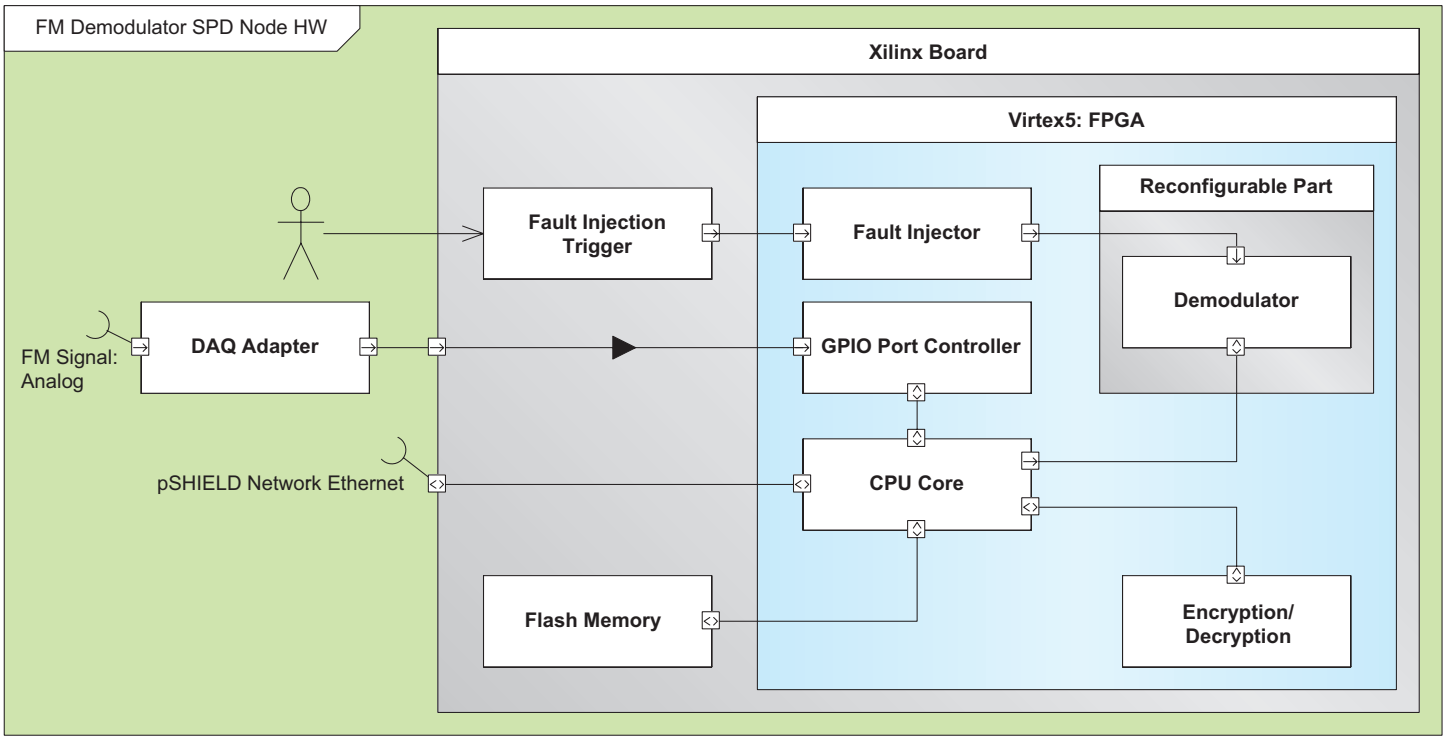

Figure 1 – Hardware implementation of the A-FSK demodulator SPD node

One of the most common forms of digital modulation in the high-frequency radio spectrum, frequency shift keying has important applications in telephone circuits. This technology transmits data by shifting the frequency of a continuous carrier in a binary way. One frequency is designated as the "mark," with frequency f0, and the other as "space," with frequency f1. The mark is associated with the symbol one, the higher frequency, while the space is associated with the symbol zero, the lower frequency. In the example of the FSK signal seen in Table 1, the mark is at frequency 1,031 Hz and the space at 968 Hz.

After the FSK performs the demodulation, the signal and the carrier are multiplied together (or nco and multiplier, I2 and I1 blocks in Figure 2) and then low-pass filtered. This low-pass or loop (I3 block in Figure 2) filter discriminates the symbol mark from the space. The amplitude of the space symbol will be higher than that of the mark.

The output of that loop filter goes to a 16-tap finite impulse response (FIR; I4 block in Figure 2) filter to perform digital low-pass filtering. The FIR filter is essentially an average filter, since its output is equal to the average value of its input over the last n-tap samples, where n is the number of taps used. This configuration needs 16 coefficients, but you can simplify the process by assuming all the coefficients are the same, 1/16. In reality, you can implement a 1/16 multiply by just performing a 4-bit right-shift operation.

The FSK adaptable demodulator is capable of adapting dynamically to a different frequency of the carrier Fc0 and Fc1. In a general communications schema, two modules are present: the modulator and the demodulator. An adaptable demodulator is able to

DIGITAL SIGNAL TO TRANSMIT	
A-FSK Rate	Variable between 100 Hz and 50 Hz
A-FSK MODULATED SIGNAL	
"Space" frequency	968 or 1,937 Hz
"Mark" frequency	1,031 or 2,062 Hz
Amplitude	1 Vpp
Analog-to-digital sampling rate	16 kHz or 32 kHz

Table 1 – Mark and space in an example FSK signal

The first step in the PR design flow is to identify the partially reconfigurable modules in our top-level design. For each of these modules, we must define the interfaces in terms of input and output signals. In our case, we identified a single PRM, named the demodulator.

switch automatically between two different carriers in order to match a carrier switch made by the modulator. The modulator might switch carriers for several reasons, including transmission errors, too much noise or the risk of intrusion on the current carrier.

The FSK adaptive demodulator has an additional built-in block named the carrier controller. This block continually checks the integrity of the transmitted signal by analyzing the consistency of received data. Based on that analysis, the carrier controller drives the reconfiguration condition.

The FSK adaptive demodulator can reconfigure itself in two distinct modes, each able to decode the signal modulated at the given carrier frequency Fc0 and Fc1. The process of configuration takes place in accordance with the partial dynamic reconfiguration methods. Figure 3 shows the general layout of the FSK adaptive demodulator. The carrier controller, which we implemented via software, runs on a PowerPC® 440 as a single task and it performs a data integrity check. In the case of a communication error, the carrier controller will force a reconfiguration event, using the Internal Configuration Access Port (ICAP) software primitives.

We designed our FSK adaptive demodulator using the Xilinx developer board ML507. Equipped with RocketIO™ GTX transceivers, this embedded-system FPGA development board provides a feature-rich, general-purpose evaluation and development platform. It includes onboard memory and industry-standard connectivity interfaces to deliver a versatile development platform for embedded applications.

PDR DESIGN FLOW

A typical static Xilinx ISE® Design Suite flow consists of four main steps:

- Design/edit
- Synthesis
- Implementation
- Device configuration

The partial-reconfiguration design flow is more sophisticated and complex than this. Figure 4 depicts a simplified PDR design flow.

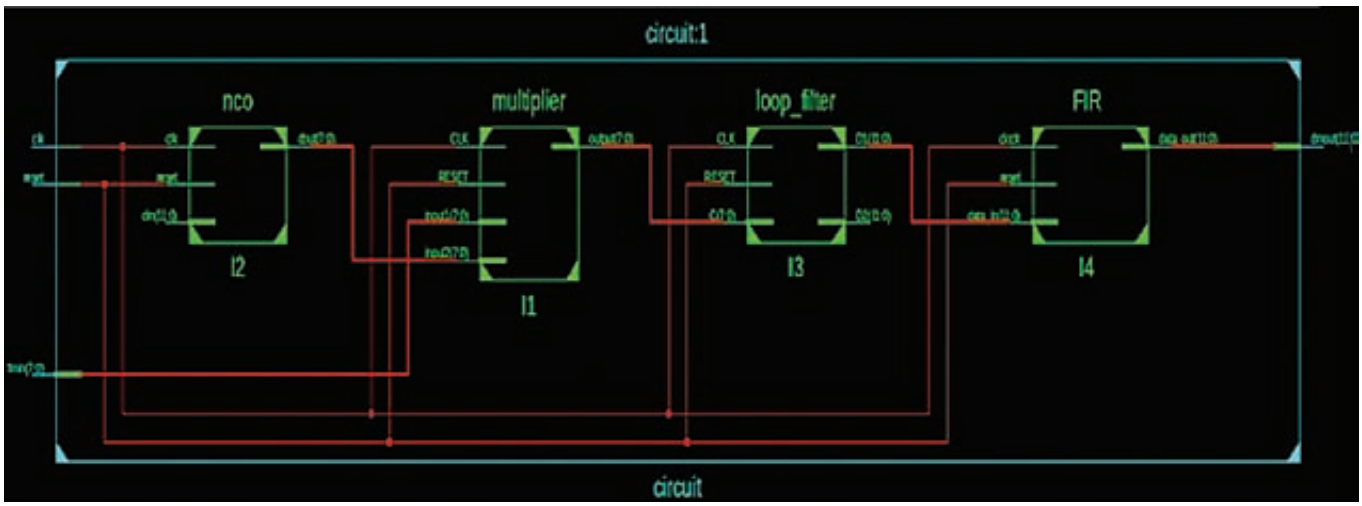

Figure 2 – Block diagram of the FSK demodulator circuit

The first step is to identify the partially reconfigurable modules (PRMs) in our top-level design. For each of these modules, we must define the interfaces in terms of input and output signals. In our case, we have identified a single PRM, named the demodulator. The following coding describes the demodulator interface:

```
ENTITY Demodulator IS
  PORT (
   clk       : IN    std_logic;
--Main Entity clock
   reset     : IN    std_logic;
--High active reset
   fmin      : IN    std_logic_vrc-
tor ( 7 DOWNTO 0);   --Modulated
fsk signal
   dmout     : OUT   std_logic_vrc-
tor (11 DOWNTO 0);   --Pre- demodu-
lated signal
   clko      : OUT   std_logic;
--Synch. FIFO signal
   dbg       : OUT   std_logic;
--debug line
  );
END Demodulator;
```

Given n number of PRMs we need in the current design, the next step is to generate the n PRM netlist files using the XST tools. The output of XST consists of the NGC files. The NGC is a netlist that contains both logical data and constraints. At the end, we have to generate n NGC files. In the event that we will use those NGC files in a partial-reconfiguration project, is necessary to ensure that the IOBUF is disabled.

For our project, we have two NGC files: a 1k and a 2k demodulator. The difference between the two modules is that one has the numerically controlled oscillator (NCO) set at around 1 kHz and the other at around 2 kHz.

Now that we have the two NGC files to use in our system, we have to create a design that will host the PRMs. Using XPS, we build up a system-on-a-chip (SoC), instantiating all the needed modules and controllers via the XPS menu. Additionally, we have to create a black-box IP that will host our PRM modules, previously created by the XST tool. To do that, we can use the option "Create and Import Peripheral Wizard" available in XPS.

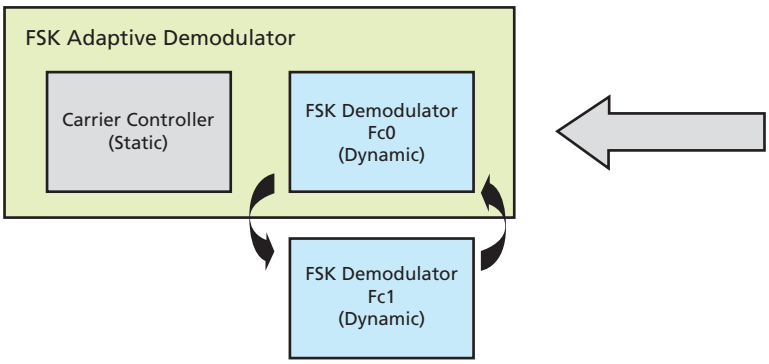

Figure 3 – Overview of the FSK adaptive demodulator design

In our case, we created a black-box module, named FSKDemodulator, using the appropriate options. The wizard process generates two VHDL files, FSKDemodulator.vhd and User_logic.vhd. The FSKDemodulator.vhd is the top level associated with our PRM modules (demodula-

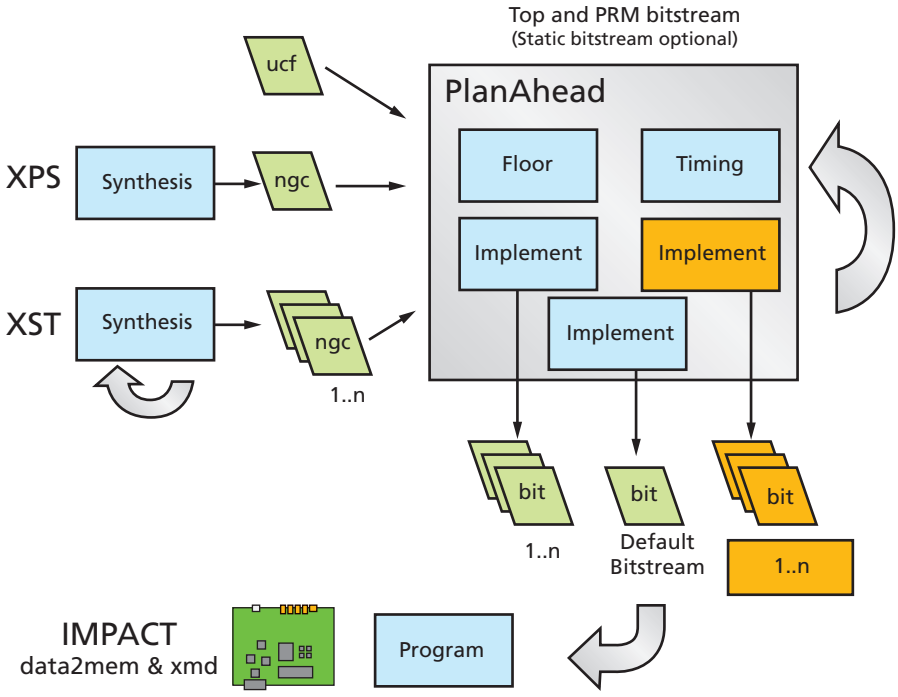

Figure 4 – Simplified view of a PDR design flow using PlanAhead, XST and XPS

tor) at the programmable-system level. This file defines the interface of the PRM module with the programmable-system components such as the Peripheral Local Bus (PLB) v4.6. The User_logic.vhd represents the user logic function, and it includes the instance of the PRM module. Once the programmable-system design has been completed, we can generate the NGC file for this configuration.

Using XPS, we defined the programmable-system as well as the general system architecture. It's important to note that we defined all the PRMs as a black box in XPS.

Using PlanAhead™ we can merge the outputs and NGC files coming from the two processes, XST and XPS, in order to have just a few PRM bitstreams and a default bitstream. A top-level implementation will be defined and built up using the NGC file from XPS and one of the NGC files from XST. The designer then needs to add a partial-reconfiguration region to the design and must specify the NGC file associated with it. The last step is to promote this configuration in order to make this implementation the default system implementation that will load into the system at startup.

The output of this process will be the default bitstream. To build the PRMs' bitstream files, we must reopen PlanAhead and, starting from scratch, check the option "PR project" in order to import all the NGC files. PlanAhead will generate a distinct bitstream for each of those PRMs. In our case, it outputs two PRM bitstreams, one for demodulator1k and the other for demodulator2k.

For debug purposes, we suggest creating n (number of PRMs) different static implementations, one for each PRM. In this case the designer will have n static full implementations, each of which will perform the function of the n^{th} PRM statically attached to the FPGA. In our opinion, this is a good compromise between debug needs and development complexity.

The last step is to download the generated bitstream into our target device. You can program the device using the IMPACT tool as well as the command-line data2mem and the Xilinx Microprocessor Debugger (XMD) tool, in case you need to store the bitstream and System ACE™ file in a CompactFlash. In our case, we opted to program the device by command-line methods, since the ML507 Xilinx board has a CompactFlash on it and the System ACE manages the CompactFlash as a boot device.

TRANSFORMABLE DEVICES

Compared with static reconfiguration, the technique of partial dynamic reconfiguration is extremely efficient in terms of the time it takes for reconfiguration to occur. Although the time is related to the physical dimensions of the PRMs, if these modules can be smaller than the full bitstream by even one gain factor, that implies a time of reconfiguration of tens vs. hundreds of milliseconds. The use of PDR brings FPGA system design to another level, giving the designer an opportunity to drastically reduce the power consumption and the cost of a whole system.

In the context of the EU's pSHIELD research project, where design criteria such as security, privacy and dependability are the main factors, we found the PDR technique extremely useful. The possibility of changing a cryptographic algorithm or communication protocol on the fly while keeping the other functionality alive was a big advantage. With this approach, we think that the FPGA is going to usher in a new era in electronic design. We are envisioning systems that can change their functionality and adapt themselves to a specific scenario or threat. In short, we are thinking of a world made up of transformable devices.

The Basics of FPGA Mathematics

by Adam Taylor
Principal Engineer
EADS Astrium
aptaylor@theiet.org

One of the main advantages of the FPGA is its ability to perform mathematical functions as desired. Here's a refresher on the basic rules and methods involved.

One of the many benefits of an FPGA-based solution is the ability to implement a mathematical algorithm in the best possible manner for the problem at hand. For example, if response time is critical, then we can pipeline the stages of mathematics. But if accuracy of the result is more important, we can use more bits to ensure we achieve the desired precision. Of course, many modern FPGAs also provide the benefit of embedded multipliers and DSP slices, which can be used to obtain the optimal implementation in the target device.

Let's take a look at the rules and techniques that you can use to develop mathematical functions within an FPGA or other programmable device.

REPRESENTATION OF NUMBERS

There are two methods of representing numbers within a design, fixed- or floating-point number systems. Fixed-point representation maintains the decimal point within a fixed position, allowing for straightforward arithmetic operations. The major drawback of the fixed-point system is that to represent larger numbers or to achieve a more accurate result with fractional numbers, you will need to use a larger number of bits. A fixed-point number consists of two parts, integer and fractional.

Floating-point representation allows the decimal point to "float" to different places within the number, depending upon the magnitude. Floating-point numbers, too, are divided into two parts, the exponent and the mantissa. This scheme is very similar to scientific notation, which represents a number as A times 10 to the power of B, where A is the mantissa and B is the exponent. However, the base of the exponent in a floating-point number is base 2, that is, A times 2 to the power of B. The floating-point number is standardized by IEEE/ANSI standard 754-1985. The basic IEEE floating-point number utilizes an 8-bit exponent and a 24-bit mantissa.

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0	•	2^{-1}	2^{-2}	2^{-3}	2^{-4}	2^{-5}	2^{-6}	2^{-7}	2^{-8}

Due to the complexity of floating-point numbers, we as designers tend wherever possible to use fixed-point representations. The above fixed-point number is capable of representing an unsigned number between 0.0 and 255.9906375 or a signed number between –128.9906375 and 127.9906375 using two's complement representation. Within a design you have the choice—typically constrained by the algorithm you are implementing—to use either unsigned or signed numbers. Unsigned numbers are capable of representing a range of 0 to $2^n - 1$, and always represent positive numbers. By contrast, the range of a signed number depends upon the encoding scheme used: sign and magnitude, one's complement or two's complement.

The sign-and-magnitude scheme utilizes the left-most bit to represent the sign of the number (0 = positive, 1 = negative). The remainder of the bits represent the magnitude. Therefore, in this system, positive and negative numbers have the same magnitude but the sign bit differs. As a result, it is possible to have both a positive and a negative zero within the sign-and-magnitude system.

One's complement uses the same unsigned representation for positive numbers as sign and magnitude. However, for negative numbers it uses the inversion (one's complement) of the positive number.

Two's complement is the most widely used encoding scheme for representing signed numbers. Here, as in the other two schemes, positive numbers are represented in the same manner as unsigned numbers, while negative numbers are represented as the binary number you add to a positive number of the same magnitude to get zero. You calculate a negative two's complement number by first taking the one's complement (inversion) of the positive number and then adding one to it. The two's complement number system allows you to subtract one number from another by performing an addition of the two numbers. The range a two's complement number can represent is given by

$$-(2^n\text{-}1) \text{ to } +(2^n\text{-}1 - 1)$$

One way to convert a number to its two's complement format is to work right to left, leaving the number the same until you encounter the first "1." After this point, each bit is inverted.

FIXED-POINT MATHEMATICS

The normal way of representing the split between integer and fractional bits within a fixed-point number is x,y where x represents the number of integer bits and y the number of fractional bits. For example, 8,8 represents 8 integer bits

and 8 fractional bits, while 16,0 represents 16 integer and 0 fractional. In many cases you will determine the correct number of integer and fractional bits required at design time, normally following conversion from a floating-point algorithm. Thanks to the flexibility of FPGAs, we can represent a fixed-point number of any bit length; the number of integer bits required depends upon the maximum integer value the number is required to store, while the number of fractional bits will depend upon the accuracy of the final result. To determine the number of integer bits required, use the following equation:

$$\text{Integer Bits Required} = Ceil\left(\frac{LOG^{10}\,\text{Integer\_Maximum}}{LOG^{10}2}\right)$$

For example, the number of integer bits required to represent a value between 0.0 and 423.0 is given by

$$9 = Ceil\left(\frac{LOG^{10}\,423}{LOG^{10}2}\right)$$

That means you would need 9 integer bits, allowing a range of 0 to 511 to be represented. Representing the number using 16 bits would allow for 7 fractional bits. The accuracy this representation would be capable of providing is given by

$$Accuracy = \left(\frac{Actual\_Value - FPGA\_Value}{2^{\text{Fractional Bits}}}\right) = 100$$

You can increase the accuracy of a fixed-point number by using more bits to store the fractional number. When designing, there are times when you may wish to store only fractional numbers (0,16), depending upon the size of the number you wish to scale up. Scaling up by 2^{16} may yield a number that still does not provide an accurate enough result. In this case you can multiply up by the power of 2, such that the number can be represented within a 16-bit number. You can then remove this scaling at a further stage within the implementation. For example, to represent the number $1.45309806319x10^{-4}$ in a 16-bit number, the first step is to multiply it by 2^{16}.

$$65536 \bullet 1.45309806319x10^{-4} = 9.523023$$

Storing the integer of the result (9) will result in the number being stored as $1.37329101563x10^{-4}$ (9 / 65536). This difference between the number required to be stored and the stored number is substantial and could lead to an unacceptable error in the calculated result. You can obtain a more accurate result by scaling the number up by a factor of 2. The result will be between 32768 and 65535, therefore still allowing storage in a 16-bit number. Using the earlier example of storing $1.45309806319x10^{-4}$, multiplying by a factor of 2^{28} will yield a number that can be stored in 16 bits and will be highly accurate of the desired number.

$$268435456 \bullet 1.45309806319x10^{-4} = 39006.3041205$$

The integer of the result will give you a stored number of $1.45308673382x10^{-4}$, which will provide for a much more accurate calculation, assuming you can address the scaling factor of 228 at a later stage within the calculation. For example, multiplying the scaled number with a 16-bit number scaled 4,12 will produce a result of 4,40 (28 + 12). The result, however, will be stored in a 32-bit result.

FIXED-POINT RULES

To add, subtract or divide, the decimal points of both numbers must be aligned. That is, you can only add to, subtract from or divide an x,8 number by a number that is also in an x,8 representation. To perform arithmetic operations on numbers of a different x,y format, you must first ensure the decimal points are aligned. To align a number to a different format, you have two choices: either multiply the number with more integer bits by 2^X or divide the number with the fewest integer bits by 2^X. Division, however, will reduce your accuracy and may lead to a result that is outside the allowable tolerance. Since all numbers are stored in base-two scaling, you can easily scale the number up or down in an FPGA by shifting one place to the left or right for each power of 2 required to balance the two decimal points. To add together two numbers that are scaled 8,8 and 9,7, you can either scale up the 9,7 number by a factor or 2^1 or scale the 8,8 format down to a 9,7 format, if the loss of a least-significant bit is acceptable.

For example, say you want to add 234.58 and 312.732, which are stored in an 8,8 and a 9,7 format respectively. The first step is to determine the actual 16-bit numbers that will be added together.

$$234.58 \bullet 2^8 = 60052.48$$

$$312.732 \bullet 2^7 = 40029.69$$

The two numbers to be added are 60052 and 40029. However, before you can add them you must align the decimal points. To align the decimal points by scaling up the number with a largest number of integer bits, you must scale up the 9,7-format number by a factor of 2^1.

$$40029 \bullet 2^1 = 80058$$

You can then calculate the result by performing an addition of

$$80058 + 60052 = 140110$$

This represents 547.3046875 in a 10,8 format (140110 / 2^8).

When multiplying two numbers together, you do not need to align the decimal points, as the multiplication will provide a result that is X1 + X2, Y1 + Y2 wide. Multiplying two numbers that are formatted 14,2 and 10,6 will produce a result that is formatted as 24 integer bits and 8 fractional bits.

You can multiply by a fractional number instead of using division within an equation through multiplying by the recip-

rocal of the divisor. This approach can reduce the complexity of your design significantly. For example, to divide the number 312.732, represented in 9,7 (40029) format, by 15, the first stage is to calculate the reciprocal of the divisor.

$$\frac{1}{15} = 0.0666'$$

This reciprocal must then be scaled up, to be represented within a 16-bit number.

$$65536 \bullet 0.06666 = 4369$$

This step will produce a result that is formatted 9,23 when the two numbers are multiplied together.

$$4369 \bullet 40029 = 174886701$$

The result of this multiplication is thus

$$\frac{174886701}{8388608} = 20.8481193781$$

While the expected result is 20.8488, if the result is not accurate enough, then you can scale up the reciprocal by a larger factor to produce a more accurate result. Therefore, never divide by a number when you can multiply by the reciprocal.

ISSUES OF OVERFLOW

When implementing algorithms, the result must not be larger than what is capable of being stored within the result register. Otherwise a condition known as overflow occurs. When that happens, the stored result will be incorrect and the most significant bits are lost. A very simple example of overflow would be if you added two 16-bit numbers, each with a value of 65535, and the result was stored within a 16-bit register.

$$65535 + 65535 = 131070$$

The above calculation would result in the 16-bit result register containing a value of 65534, which is incorrect. The simplest way to prevent overflow is to determine the maximum value that will result from the mathematical operation and use this equation to determine the size of the result register required.

$$\text{Integer Bits Required} = Ceil\left(\frac{LOG^{10}\,\text{Integer\_Maximum}}{LOG^{10}2}\right)$$

If you were developing an averager to calculate the average of up to fifty 16-bit inputs, the size of the required result register could be calculated.

$$50 \bullet 65535 = 3276750$$

Using this same equation, this would require a 22-bit result register to prevent overflow occurring. You must also take care, when working with signed numbers, to ensure there is no overflow when using negative numbers. Using the averager example again, taking 10 averages of a signed 16-bit number returns a 16-bit result.

$$10 \bullet -32768 = -327680$$

Since it is easier to multiply the result by a scaled reciprocal of the divisor, you can multiply this number by 1/10 • 65536 = 6554 to determine the average.

$$-32768 \bullet 6554 = -2147614720$$

This number when divided by 216 equals -32770, which cannot be represented correctly within a 16-bit output. The module design must therefore take the overflow into account and detect it to ensure you don't output an incorrect result.

REAL-WORLD IMPLEMENTATION

Let's say that you are designing a module to implement a transfer function that is used to convert atmospheric pressure, measured in millibars, into altitude, measured in meters.

$$-0.0088x2 + 1.7673x + 131.29$$

The input value will range between 0 and 10 millibars, with a resolution of 0.1 millibar. The output of the module is required to be accurate to +/-0.01 meters. As the module specification does not determine the input scaling, you can figure it out by the following equation.

$$4 = Ceil\left(\frac{LOG^{10}\,10}{LOG^{10}2}\right)$$

Therefore, to ensure maximum accuracy you should format the input data as 4 integer and 12 fractional bits. The next step in the development of the module is to use a spreadsheet to calculate the expected result of the transfer function across the entire input range using the unscaled values. If the input range is too large to reasonably achieve this, then calculate an acceptable number of points. For this example, you can use 100 entries to determine the expected result across the entire input range.

Input (millibar)	Output (meters)
0	131.2900
0.1	131.4666
0.2	131.6431
0.3	131.8194
0.4	131.9955
0.5	132.1715
0.6	132.3472

Once you have calculated the initial unscaled expected values, the next step is to determine the correct scaling factors for the constants and calculate the expected outputs using the scaled values. To ensure maximum accuracy, each of the constants used within the equation will be scaled by a different factor.

The scaling factor for the first polynomial constant (A) is given by

$$8 = Ceil\left(\frac{LOG^{10}\,133.29}{LOG^{10}2}\right)$$

The second polynomial constant (B) scaling factor is given by

$$1 = Ceil\left(\frac{LOG^{10}\,1.7673}{LOG^{10}2}\right)$$

The final polynomial constant (C) can be scaled up by a factor of 2^{16}, as it is completely fractional.

Polynomial Constant	Unscaled	Scaled
A	133.29	33610
B	1.77	57910
C	-0.01	-577

These scaling factors allow you to calculate the scaled spreadsheet, as shown in Table 1. The results of each stage of the calculation will produce a result that will require more than 16 bits.

The calculation of the Cx^2 will produce a result that is 32 bits long formatted 4,12 + 4,12 = 8,24. This is then multiplied by the constant C, producing a result that will be 48 bits long formatted 8,24 + 0,16 = 8,40. For the accuracy required in this example, 40 bits of fractional representation is excessive. Therefore, the result will be divided by 2^{32} to produce a result with a bit length of 16 bits formatted 8,8. The same reduction to 16 bits is carried out upon the calculation of Bx to produce a result formatted 5,11.

The result is the addition of columns Cx^2, Bx and A. However, to obtain the correct result you must first align the radix points, either by shifting up A and Cx^2 to align the numbers in an x,11 format, or shifting down the calculated Bx to a format of 8,8, aligning the radix points with the calculated values of A and Cx^2.

In this example, we shifted down the calculated value by 2^3 to align the radix points in an 8,8 format. This approach simplified the number of shifts required, thus reducing the logic needed to implement the example. Note that if you cannot achieve the required accuracy by shifting down to align the radix points, then you must align the radix points by shifting up the calculated values of A and Cx^2. In this example, the calculated result is scaled up by a power of 2^8. You can then scale down the result and compare it against the result obtained with unscaled values. The difference between the calculated result and the expected result is then the accuracy, using the spreadsheet commands of MAX() and MIN(), for the maximum and minimum error of the calculated result that can be obtained across the entire range of spreadsheet entries.

Once the calculated spreadsheet confirms that you can achieve the required accuracy, you can write and simulate the RTL code. If desired, you could design the testbench such that the input values are the same as those used in the spreadsheet. This allows you to compare the simulation outputs against the spreadsheet-calculated results to ensure the correct RTL implementation.

Input Scaled	C	B	A	Result	Result Scaled	Expected Result	Difference
0	0	0	33610	33610	131.289	131.2900	0.0009
409	-6	361	33610	33655	131.465	131.4666	0.0018
819	-24	723	33610	33700	131.641	131.6431	0.0025
1228	-52	1085	33610	33745	131.816	131.8194	0.0030
1638	-93	1447	33610	33790	131.992	131.9955	0.0033
2048	-145	1809	33610	33835	132.168	132.1715	0.0035
2457	-208	2171	33610	33880	132.344	132.3472	0.0035
2867	-283	2533	33610	33925	132.520	132.5228	0.0033

Table 1 – Real results against the fixed-point mathematics

RTL IMPLEMENTATION

The RTL example uses signed parallel mathematics to calculate the result within four clock cycles. Because of the signed parallel multiplication, you must take care to correctly handle the extra sign bits generated by the multiplications.

```
ENTITY transfer_function IS PORT(
sys_clk : IN std_logic;
reset : IN std_logic;
data : IN std_logic_vector(15 DOWNTO 0);
new_data : IN std_logic;
result : OUT std_logic_vector(15 DOWNTO 0);
new_res : OUT std_logic);
END ENTITY transfer_function;

ARCHITECTURE rtl OF transfer_function IS
-- this module performs the following
transfer function -0.0088x2 + 1.7673x +
131.29
-- input data is scaled 8,8, while the
output data will be scaled 8,8.
-- this module utilizes signed parallel
mathematics

TYPE control_state IS (idle, multiply,
add, result_op);
CONSTANT c : signed(16 DOWNTO 0) :=
to_signed(-577,17);
CONSTANT b : signed(16 DOWNTO 0) :=
to_signed(57910,17);
CONSTANT a : signed(16 DOWNTO 0) :=
to_signed(33610,17);
SIGNAL current_state : control_state;
SIGNAL buf_data : std_logic; --used to
detect rising edge upon the new_data
SIGNAL squared : signed(33 DOWNTO 0); --
register holds input squared.
SIGNAL cx2 : signed(50 DOWNTO 0); --regis-
ter used to hold Cx2
SIGNAL bx : signed(33 DOWNTO 0); -- regis-
ter used to hold bx
SIGNAL res_int : signed(16 DOWNTO 0); --
register holding the temporary result

BEGIN
fsm : PROCESS(reset, sys_clk)
BEGIN
IF reset = '1' THEN
 buf_data <= '0';
 squared <= (OTHERS => '0');
 cx2 <= (OTHERS => '0');
 bx <= (OTHERS => '0');
 result <= (OTHERS => '0');
 res_int <= (OTHERS => '0');
 new_res <= '0';
 current_state <= idle;
ELSIF rising_edge(sys_clk) THEN
buf_data <= new_data;
CASE current_state IS
WHEN idle =>
 new_res <= '0';
 IF (new_data = '1') AND (buf_data = '0')
 THEN --detect rising edge new data
   squared <= signed( '0'& data) *
   signed('0'& data);
   current_state <= multiply;
 ELSE
   squared <= (OTHERS =>'0');
   current_state <= idle;
END IF;
WHEN multiply =>
 new_res <= '0';
 cx2 <= (squared * c);
 bx <= (signed('0'& data)* b);
 current_state <= add;
WHEN add =>
 new_res <= '0';
 res_int <= a + cx2(48 DOWNTO 32) +
   ("000"& bx(32 DOWNTO 19));
 current_state <= result_op;
WHEN result_op =>
 result <= std_logic_vector(res_int
 (res_int'high -1 DOWNTO 0));
 new_res <= '0';
 current_state <= idle;
END CASE;
END IF;
END PROCESS;
END ARCHITECTURE rtl;
```

The architecture of FPGAs makes them ideal for implementing mathematical functions, although the implementation of your algorithm may take a little more initial thought and modeling in system-level tools such as MATLAB® or Excel. You can quickly implement mathematical algorithms once you have mastered some of the basics of FPGA mathematics.

The FPGA Engineer's Guide to Using ADCs and DACs

by Adam Taylor
Principal Engineer
EADS Astrium
aptaylor@theiet.org

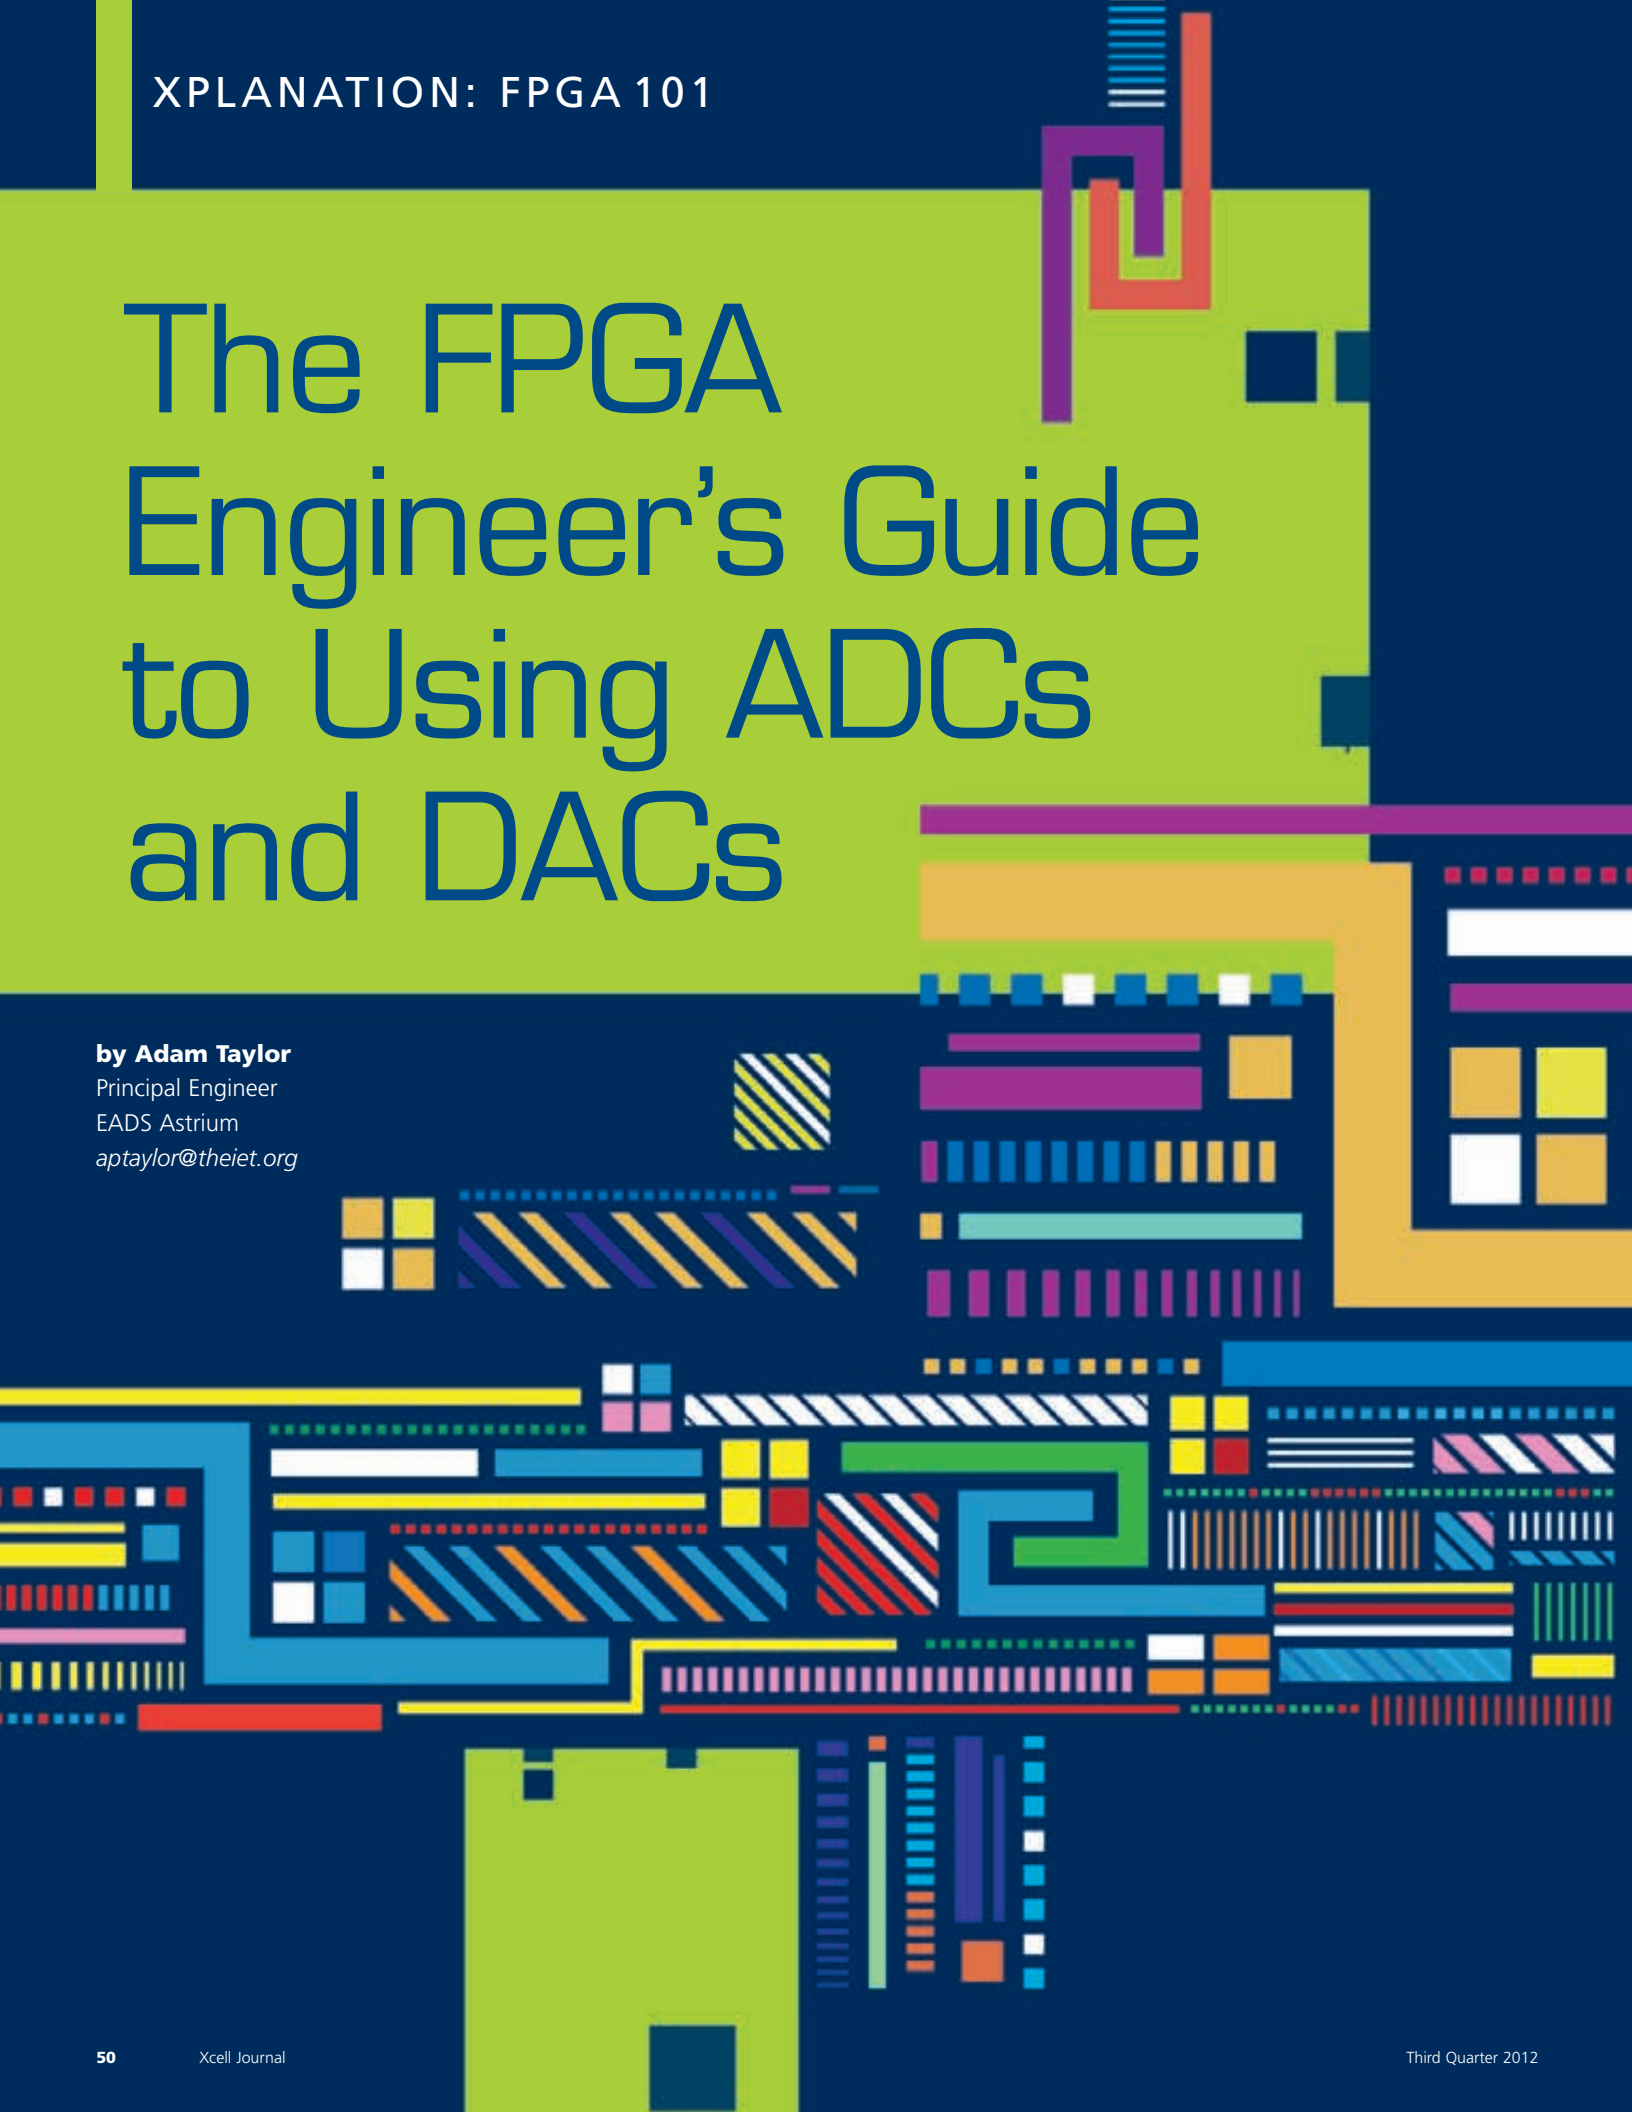

Interfacing the signal-processing FPGA with the real world requires the use of an analog-to-digital or digital-to-analog converter.

Once it's performed the task it was designed to do, an FPGA-based system next has to interface with the real world, and as every engineer knows, the real world tends to function around analog as opposed to digital signals. That means conversion is going to be required to and from the digital domain from the analog realm. Just as you face a plethora of choices in selecting the correct FPGA for the job at hand, so too will you find an abundance of riches when choosing the correct ADC or DAC for a system.

The first thing to establish is the sampling rate you will need to convert the signal. This parameter will drive not only the converter selection but will also impact your FPGA choice as well, to ensure the device can address the processing speed and logic footprint required. The sampling rate of the converter needs to be at least twice that of the signal being sampled. Therefore, if you need to sample a signal at 50 MHz, your sampling rate must be at least 100 MHz. Otherwise the converted signal will be aliased back upon itself and will not be correctly represented. This aliasing is not always a bad thing; in fact, if the converter bandwidth is wide enough, you can employ the aliasing to fold signals back into the usable bandwidth.

ADC AND DAC KEY PARAMETERS

Analog-to-digital converters are constructed by many different techniques. Some of the most common are flash, ramp and successive approximation.

- Known for their speed, **flash converters** use a series of scaled analog comparators to compare the input voltage against a reference voltage; ADCs use the outputs of these comparators to determine digital code.
- **Ramp converters** utilize a free-running counter connected to a digital-to-analog converter, comparing the output of the DAC against the input voltage. When the two are equal, the count is held.
- **Successive-approximation converters** are an adaptation of ramp converters and also utilize a DAC and a comparator against the analog input. However instead of counting up, the SAR converter determines whether the analog representation of the count is above or below the input signal, allowing a trial-and-error-based approach to determining the digital code.

Digital-to-analog converters also come in several implementations, some of the most common being binary-weighted, R-2R ladder and pulse-width modulation.

- **Binary-weighted** is one of the fastest DAC architectures. These devices sum the result of individual conversions for each logic bit. For example, a resistor-based DAC will switch resistors on or out depending upon the current code.
- **R-2R ladder converters** use a structure of cascaded resistors of value R-2R. Due to the ease with which precision resistors can be produced and matched, these DACs are more accurate than the binary-weighted types.
- **Pulse-width modulation,** the simplest type of DAC architecture, passes the PWM waveform through a simple low-pass analog filter. These devices are commonly used in motor control but also form the basis for delta-sigma converters.

Many manufacturers of specialist devices have developed their own internal conversion architectures to provide the best possible performance in specific areas depending upon the intended use. Each of these varieties has pros and cons relating to the speed of conversion, accuracy and resolution. As when selecting an FPGA, you will look at the number of I/Os, I/O standards supported, clock management, logic resources and memory, along with parameters specific to the device type: the maximum sampling rate, signal-to-noise ratio (SNR), spurious-free dynamic range (SFDR) and effective number of bits (ENOB).

The sampling frequency is pretty simple; it's the maximum rate at which an ADC can digitize its input. SNR represents the ratio of the signal to the noise level, which is assumed to be uncorrelated with the input signal. You can determine the SNR theoretically using the equation

```
SNR = 6.02N + 1.76 dB
```

where N is the resolution. This equation is valid for a full-scale sine wave.

You can determine the actual SNR during system test by taking a fast Fourier transform (FFT) of the output and measuring between the value of the input signal and the noise floor.

The SFDR, meanwhile, is the ratio between the input signal and the next

Nyquist Zone	Range Lower	Range Upper	Aliasing
First	DC	0.5 FS	None
Second	0.5 FS	FS	Folds
Third	FS	1.5 FS	Direct
Fourth	1.5 FS	2 FS	Folds

Table 1 – Nyquist zones and aliasing

highest peak, usually a harmonic of the fundamental. Normally the SFDR is given in terms of dBc and will degrade as the input signal power drops.

From these measurements of the converter, you can calculate the effective number of bits using the equation

```
ENOB = (SNR – 1.77 / 6.02)
```

When performing this testing, take care to ensure that the FFT you're using is correctly sized, so as to ensure that you haven't inadvertently miscalculated the noise floor. An incorrect FFT size will throw off your calculations. The FFT noise floor is given by

```
FFT noise floor
= 6.02N + 1.76 dB + 10
LOG10(FFT size / 2)
```

You should perform these steps using a single-tone test—normally, a simple sine wave—to reduce the complexity of the output spectrum. To guarantee you will get the best results, be sure to take coherent samples of the output. Coherent sampling takes place when there is an integer number of cycles within the data window. For this to be correct, then

```
FS / Fin = Ncycles / FFT
```

THE FREQUENCY SPECTRUM

On another front, you must also be aware of the Nyquist criteria when implementing your system, to ensure the signal is correctly converted or quantized. This means that you must sample at least twice the maximum frequency of the signal of interest to ensure correct conversion. When signals are sampled outside of this criterion aliasing will occur, potentially resulting in unwanted performance if it is not correctly understood.

It is also for this reason that ADCs require anti-aliasing filters to prevent signals or noise being aliased back into the quantized signal. However, aliasing can be very useful to the engineer, especially if the ADC has a wide input bandwidth. With careful consideration, aliasing can allow you to directly convert signals without the need for down-converters. For this reason, the frequency spectrum is divided up into a number of zones.

Using the information presented in Table 1, it is possible to alias signals from one Nyquist band to another if the converter has a wide enough bandwidth.

Tap	Coefficient
1	-6.22102953898351E-003
2	9.56204928971727E-003
3	-1.64864415228791E-002
4	3.45071042895427E-002
5	-0.107027889432584
6	1.166276
7	-0.107027889432584
8	3.45071042895427E-002
9	-1.64864415228791E-002
10	9.56204928971727E-003
11	-6.22102953898351E-003

Table 2 – The first 11 coefficients for a DAC compensation FIR filter

COMMUNICATION CHOICES

As with all external devices, ADCs and DACs come with several interfacing options, either parallel or serial. Typically, higher-speed devices will implement a parallel interface while slower ones will utilize a serial interface. However, your choice of application may drive you down a particular route. It is easier, for instance, to detect a stuck-at bit in a serial interface than it is in a parallel one. Really high-speed interfaces may provide multiple output buses (I and Q) or use double-data-rate (DDR) outputs; some devices might even offer both options. Having multiple buses or DDR outputs allows you to maintain the data rate while reducing the frequency of operation required for the interface. For example, an interface sampling at 600 MHz will produce an output at 300 MHz (half the sampling frequency).

It is much easier to recover if the clock frequency is 75 MHz (FS/4) and there are two data buses that provide samples from the device using DDR. This kind of ADC allows the relaxation of the input timings you must achieve. Many high-speed converters utilize LVDS signaling on their I/O, as the lower voltage swing and low current reduce the coupling that can occur with other signaling standards such as LVCMOS. This coupling problem can affect the converter's mixed-signal performance.

DAC FILTERING

Most DACs will hold the analog output until the next sampling period. The result is an interesting effect upon the output frequency domain. You will

notice that both images exist across the output spectrum, and the output signal in all of the Nyquist zones exhibits a roll-off due to the sinc effect being nearly 4 dB (3.92 dB) lower at 0.5 FS, as shown in Figure 1. The solution to both of these issues is to utilize filters.

You can implement the first filter digitally within the FPGA, before the DAC. This simple digital filter can correct for the roll-off. But if you use an anti-image filter, you will need to position it after the DAC output, since the images are a consequence of the reconstruction.

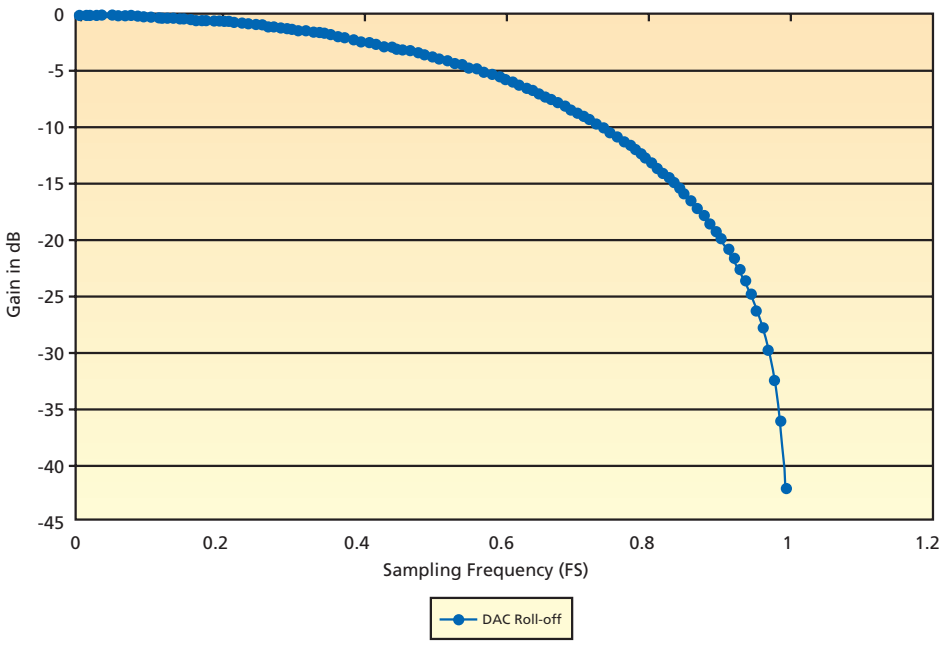

Figure 1 – DAC roll-off between 0 and FS

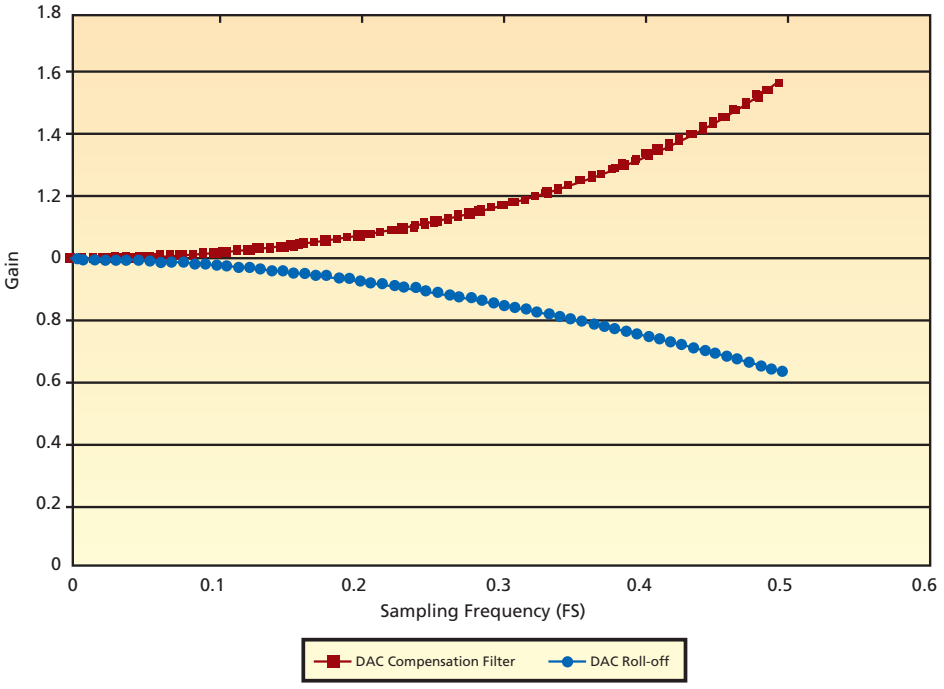

Figure 2 – DAC roll-off and compensation filter to FS/2

The sinc correction filter can be implemented very easily as a FIR filter. The simplest method to undertake in developing this filter is to plot the sinc roll-off using the equation

$$A = \frac{\sin\frac{(\pi * Fin)}{FS}}{\frac{\pi * Fin}{FS}}$$

Create the correction factor, which is the reciprocal of those calculated for the roll-off, and then take an inverse Fourier transform to obtain the coefficients you will need to design the filter. Typically, you will be able to implement this filter with a few taps. Table 2 shows the first 11 coefficients for the filter, while Figure 2 shows the correction against the roll-off.

IN-SYSTEM TEST

Many of these systems will require the converter to achieve specific performance characteristics for the end application, for example CDMA or GSM. Doing the testing necessary to achieve this performance can require a significant investment in test systems (arbitrary waveform generators, logic analyzers, pattern generators, spectrum analyzers and the like). However, the reprogrammable flexibility of the FPGA allows you to insert specific test programs into the device to either capture and analyze the output of an ADC or provide the stimulus for a DAC, reducing the need for additional extra test equipment.

CONVERSION 101

Because FPGAs are commonly called upon to interface with ADCs and DACs, it is crucial for any FPGA engineer to acquire least a basic understanding of the parameters of importance in these devices. This is especially true when you are planning to have the FPGA test the performance of the converter using the reprogrammable flexibility of the FPGA as part of the design-proving and commissioning process.

The Sky's the Limit for SSD Enterprise Storage Startup Skyera

Serial entrepreneur and noted chip architect Radoslav Danilak's latest innovation leverages Spartan-6 FPGAs.

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

With a rich history of creating bleeding-edge IC designs for microprocessors, graphics processors and ASICs, Radoslav Danilak and his team at startup Skyera are now set to introduce a product that promises to ignite the market for solid-state storage systems in the enterprise storage market. Creating innovative architectures isn't new for Danilak, but in this design the processing heart of the Skyera solid-state drive (SSD) storage system isn't implemented in an SoC or standalone processor—it's implemented in a Xilinx® Spartan®-6 FPGA.

'Not only did we have to increase performance and reliability 10 times over competing designs, but we had to figure out how to fit the design into an FPGA that is 10 times smaller than what an ASIC would give us. Having the right balance of capacity, performance and cost is essential.'

Over the last 15 years, the data storage sector has gone through a remarkable renaissance thanks to the advent of flash memory-based solid-state storage systems. For well over a decade, the slowest link in most compute-intensive systems has been a hard drive's read-and-write speed. However, in the early 2000s, memory vendors started using NAND flash memory as a faster, lower-power alternative to mechanical disks. NAND products found their first big uses in mobile phones and handset devices such as Apple's iPod. Soon afterward, higher-capacity configurations found their way into laptop computers, mainly for power-savings reasons, and into desktop computers for performance.

The price of SSD systems is typically much higher than that of traditional mechanical hard drives. However, over the last five years, as the capacity of NAND flash has ballooned, the price per byte of NAND has declined markedly, making it inevitable that SSD will sooner rather than later push mechanical hard drives into obsolescence even in the enterprise storage market.

For the last company Danilak founded, SandForce, Danilak's team created a custom memory controller SoC that sits at the heart of many of today's first-generation enterprise SSD systems. With his new company, Skyera, he's now on the verge of introducing an enterprise system that the company claims will offer a 10x improvement in both performance and capacity at a price point in parity with slower and lower-capacity enterprise storage systems. Skyera will introduce the first commercial version of its product in the coming months, and so at the time of this interview Danilak was not at liberty to disclose its full feature set. However, he said that a key attribute of the product, beyond massive capacity and performance at a competitive price point, will be its ability to extend the life of the flash memory in the system.

The Achilles' heel of NAND flash is that as process geometries continue to shrink, the physical size of the memory cells in the NAND is smaller, and so is the charge they hold. Repeated programming and erasing degrade the ability of each cell to maintain a charge reliably. Danilak said that where a 43-nanometer single-level-cell (SLC) NAND could perform 100,000 writes before encountering an uncorrectable error, a 15-nm multilevel-cell (MLC) NAND will perform only 1,000 writes before encountering an uncorrectable error.

This phenomenon is often referred to as NAND "wear." To counter wear as it gets worse with every process shrink, vendors of SSD systems must develop ever-more-sophisticated NAND controllers to perform wear leveling along with a number of proprietary techniques to get the most out of their systems. Danilak said that vendors have traditionally developed these controllers in ASICs, with every generation a bit more complex than the previous one to ensure better endurance, reliability and performance.

So what's especially noteworthy about Skyera's new product is that to get this revolutionary system to market quickly, Skyera is implementing its unique controller functions on a Xilinx Spartan-6. "The challenge was, not only did we have to increase performance and reliability 10x over competing designs, but we had to figure out how to fit the design into an FPGA that is 10x smaller than what an ASIC would give us," said Danilak. "Having the right balance of capacity, performance and cost is essential for capturing market share in the SSD enterprise market today. We could have implemented the controller SoC for our system in an ASIC design, but that would have taken us 18 months to do and millions of dollars in development costs. If there was any kind of a problem in design or manufacturing, it would cost

us millions more to pinpoint and then fix the problem."

To get to market ASAP, Danilak's group instead chose to implement the production design in a Spartan-6 FPGA. What's more remarkable is that his team was able to get several critical blocks—for error correction and encryption—running at 250 MHz. "The Xilinx data sheet said that 250 MHz was the theoretical limit for the clock tree and so that was our target," said Danilak. "At the same time, we needed to shrink the area of the design by 20 times to fit in an FPGA. We had to throw out conventional thinking and rethink the architecture, but we were able to do it."

Danilak said that he helped come up with the unique architecture for the design, but that his skilled team under Rod Mullendore, chief architect at Skyera and previously the chief architect at flash controller company SandForce, gets the credit for implementing the architecture to squeeze the performance out of the design. Mullendore said the team paid careful attention to the coding of the RTL in Verilog and performed a little, but not excessive, manual layout to achieve the desired performance. "We were very careful in how we structured the RTL, and we did a certain amount of placement to minimize long lines and keep high-speed blocks grouped close together," said Mullendore. "We also limited the levels of logic. With an FPGA we can use an iterative process where we program the design to find long paths and bottlenecks, and fix them."

Danilak has a long and storied career as an IC architect that started when he was a PhD student and a teacher at the Technical University of Kosice in the Slovak Republic, where he studied math and developed database architectures. From there he went on to develop subsequent processor-centric architectures at his first startup, DanSoft (advanced VLIW). Then came stints at Gizmo Technologies (64-bit, x86 architecture), Toshiba (memory and PlayStation II group), Nishan Systems (network processors), NVIDIA (nForce4 chip sets and Fermi HPC) and his prior startup SandForce (memory controllers).

"I went from a PhD in software database systems and mathematics to microprocessors, to chip sets, to graphics processors—and then to flash controllers and now, even an FPGA architecture for a company still in stealth mode … always something new," said Danilak. "At every step in my career, I've learned something new about technology and about the business of technology."

Indeed, one of those "business of technology" lessons involves finding the best scenario for funding in which the company is either bootstrapped or the number of backers is minimized so the startup can focus its efforts on fielding innovative technology rather than on placating investors.

"Funding can be very complex and if you need funding for a custom chip, you need a lot of money," said Danilak. "Implementing this on an FPGA meant we didn't get hit with a huge NRE and we could show to backers the entire system working, rather than have them tapping their watches waiting for the silicon to come back. That's of course in addition to getting the product to customers faster. "

Danilak said that the plan of record is to implement this and Skyera's next-generation system on Xilinx FPGAs and use the FPGAs in production, not just prototyping. "I've used FPGAs for prototyping in the past but today their performance, capacity, power consumption and cost mean we don't have to go to an ASIC or ASSP," said Danilak. "That's not to say we may not do an ASIC down the road, but it's great that we don't have to do one now or in the near future."

For more information on Skyera, visit the company's site at *http://skyera.com/*.

XILINX KINTEX 7

inrevium

A Revolutionary Innovation Platform

Kintex™-7 FPGA ACDC

(Acquisition, Contribution, Distribution and Consumption)

1.0 Baseboard

The Kintex-7 FPGA ACDC (Acquisition, Contribution, Distribution and Consumption) 1.0 Baseboard. The processing performance and flexibility provided by 28nm Kintex-7 FPGAs allows manufacturers to build stunningly immersive technology into new consumer TV screens that come with features such as multiple windows/picture-in-picture, 3D graphics for games and ultra realistic viewing that goes far beyond HDTV.

- **FPGA**：**Kintex-7 FPGA XC7K325T-FFG900**
- **Memory**：**DDR3 2Gbit x 4** (address shared)
- **FMC connector**
  - HPC (High Pin Count) x 2
  - LPC (Low Pin Count) x 2
- **SERDES**：**16ch x 12.5Gbps** (GTX)

TOKYO ELECTRON DEVICE LIMITED

Headquarter : Yokohama , Japan
E-mail:psd-sales@teldevice.co.jp
http://solutions.inrevium.com/

US Offices : Fremont , CA
E-mail:psd-nasales@teldevice.co.jp

China Offices : Shanghai
E-mail:sales@teldevice.cn

XILINX ALLIANCE PROGRAM PREMIER MEMBER

DIGILENT BEYOND THEORY

Buy Now

http://www.digilentinc.com/INREVIUM

Xilinx FPGAs Improve Beamforming System Design

Virtex-6 FPGA-based board with native and custom IP is foundation for eight-channel demo system.

by Rodger H. Hosking
Vice President and Co-founder
Pentek, Inc.
rodger@pentek.com

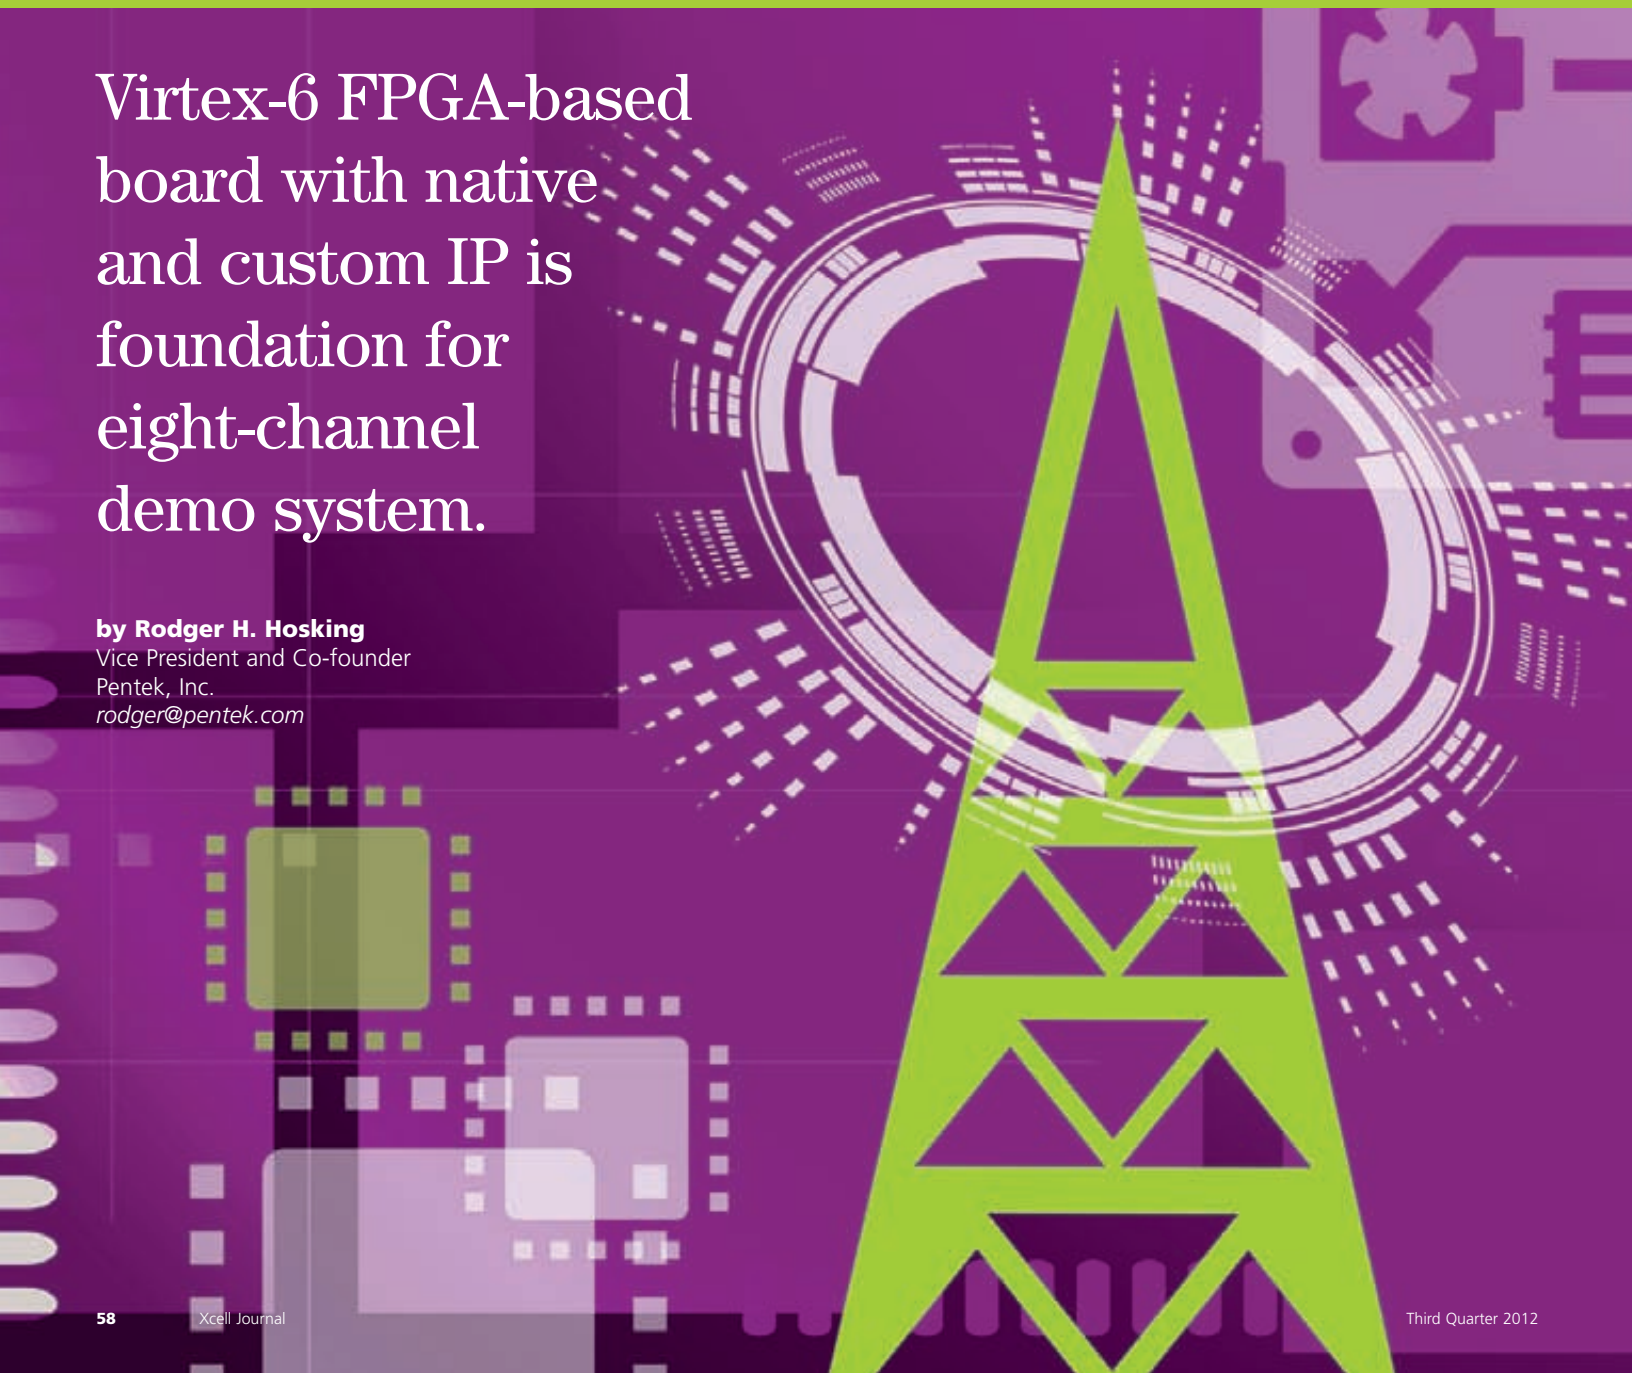

Beamforming is a signal-processing technique that utilizes an array of sensors to achieve directionality, increase the strength of transmitted signals and improve the quality of received signals. Communications, radar, countermeasures, weapons systems, oil and mineral exploration, medical imaging and direction finding make extensive use of beamforming.

In direction finding, we steer the beamformed antenna to locate the arrival angle of a signal source. We can use two or more arrays to triangulate the exact location of the source, which is essential for many signal intelligence and counterterrorism efforts. The accuracy of this technique depends on the exact settings of gain and phase among the beamforming channels. We used a Pentek product built around a Xilinx® Virtex®-6 FPGA with native and custom IP to achieve these fine adjustments that improve system performance and accuracy.

PRINCIPLES OF BEAMFORMING

We typically use beamforming with an array of sensors or antennas to improve receptivity in a specific direction, for example, from a single cell phone as shown in Figure 1. The signal from this source arrives at each antenna based on the distance between the source and the antenna, so the antenna signals have relative phase and amplitude offsets.

The beamforming process adjusts the gain and phase of each antenna signal to compensate for the different delays in the signal paths. The adjustments align the signals from each antenna with the signals arriving from one particular direction. When the signals are summed together, the non-aligned signals arriving from other directions cancel each other, while the signals from the beamformed direction add constructively for greatly improved signal-to-noise ratio. In this way, by electronically adjusting the gain and phase in each path, we effectively steer the antenna toward the direction of the signal source.

EIGHT-CHANNEL SYSTEM

In this system, we arranged eight antennas in a linear array, as shown in the overall block diagram of Figure 2. The antenna frequency here is 2.5 GHz, so each antenna signal needs to be amplified, filtered and then downconverted to an intermediate frequency (IF) so that an A/D converter can digitize it at a reasonable sampling rate. It is mandatory to use synchronous sampling across all eight channels in order to preserve a fixed phase relationship for beamforming.

We then downconvert samples from each A/D to the baseband's complex I+Q signals in a digital downconverter (DDC), which also includes channel-specific phase and gain adjustments for the beamforming "weights." Finally, we add together all eight baseband signals in summation blocks to produce the beamformed

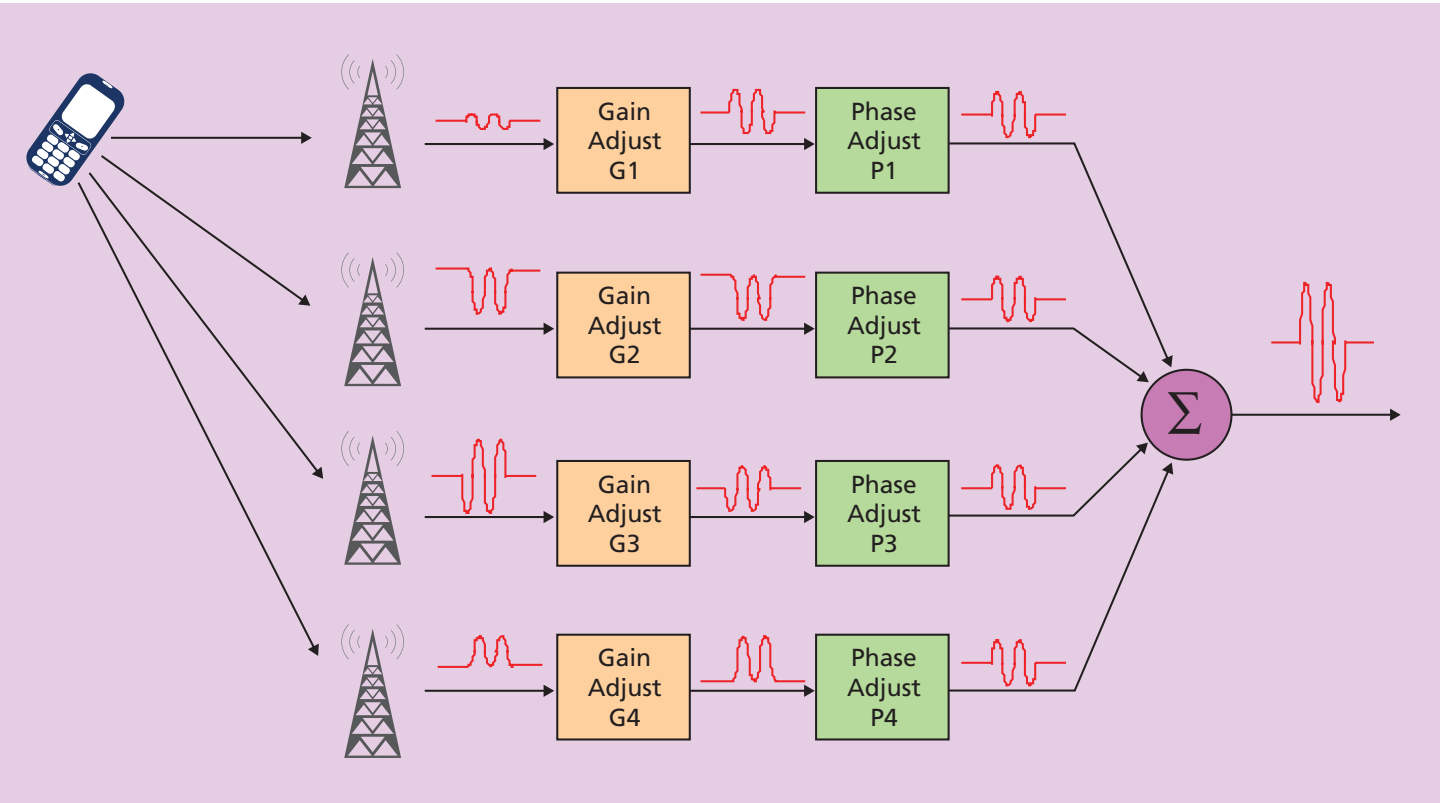

Figure 1 – Typical cell phone beamforming system

The FPGA has access to all of the board's data and control paths, enabling factory-installed functions such as data multiplexing, channel selection, data packing, gating, triggering and memory control. Each function exists as an IP block.

sum signal. A CPU analyzes this sum signal and makes adjustments to the phase and gain coefficients to track or adapt to new targets.

PENTEK MODEL 53661 BEAMFORMING BOARD

The Pentek Model 53661 software radio board is a 3U OpenVPX Cobalt board shown in the simplified block diagram of Figure 3. It features four 200-MHz, 16-bit A/D converters; a timing, clock and synchronization section; and a Xilinx Virtex-6 FPGA.

The FPGA has access to all of the board's data and control paths, enabling factory-installed functions such as data multiplexing, channel selection, data packing, gating, triggering and memory control. The Cobalt architecture organizes the FPGA as a container for data-processing applications where each function exists as an intellectual property (IP) module.

We can use a variety of different FPGAs to match the specific requirements of the processing task. Supported FPGAs include the LX240T, LX365T, SX315T and SX475T. The SXT parts feature up to 2,016 DSP48E slices and are ideal for modulation/demodulation,encoding/decoding, encryption/decryption and channelization of the signals between transmission and reception.

Factory-installed in the FPGA are four DDC IP cores, each capable of accepting A/D samples from any of the four A/Ds. Each DDC has a decimation range of 2k to 64k and can deliver downconverted baseband bandwidths from 2.5 kHz to 80 MHz. Each DDC has programmable gain and phase shift controls accessible to the processor across the VPX backplane. In this system we will be assigning one A/D to each DDC.

A power meter at the output of each DDC calculates the downconverted signal power. Each power meter is equipped with a threshold detector that generates a system interrupt if the output power exceeds the upper threshold or falls below the lower threshold. These features simplify gain calibration and signal-monitoring tasks that the system processor would otherwise have to do in software.

The 53661 FPGA also includes a native Aurora summation block that

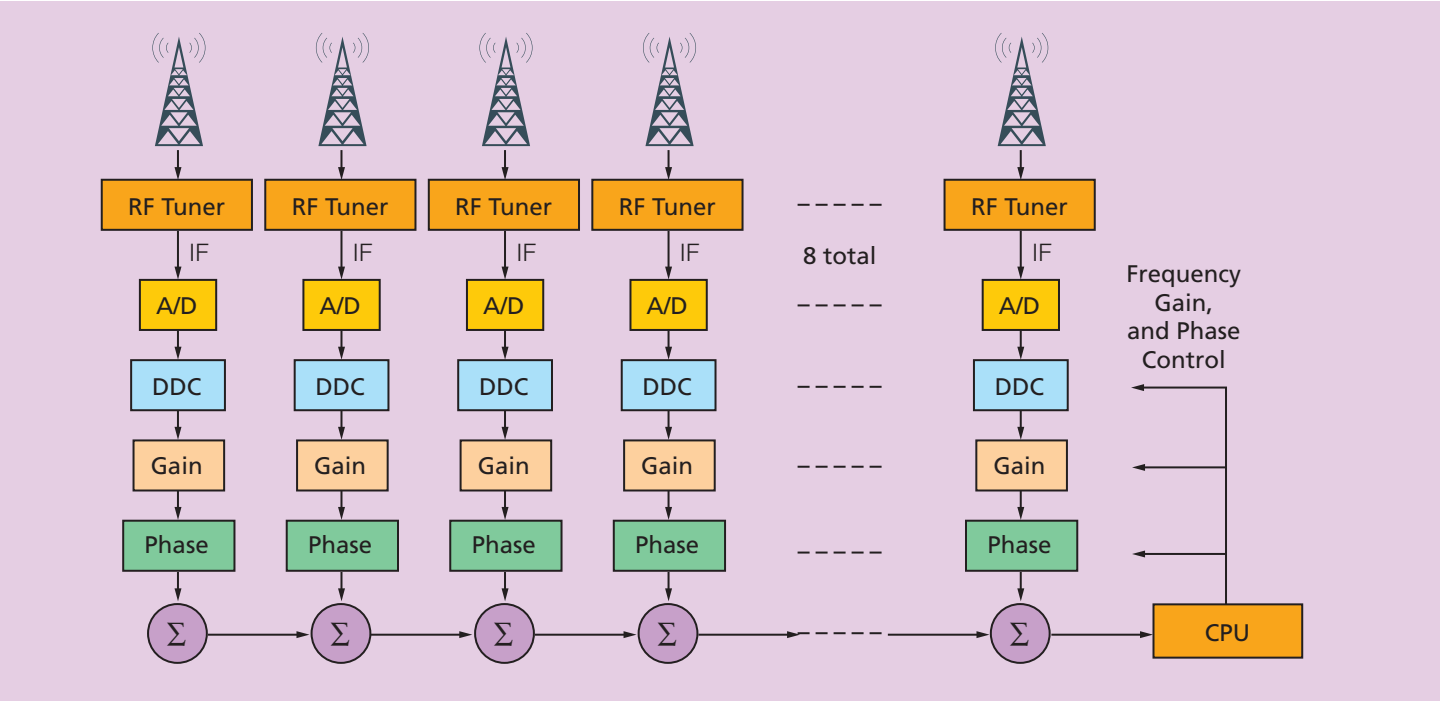

Figure 2 – Eight-channel beamforming system block diagram

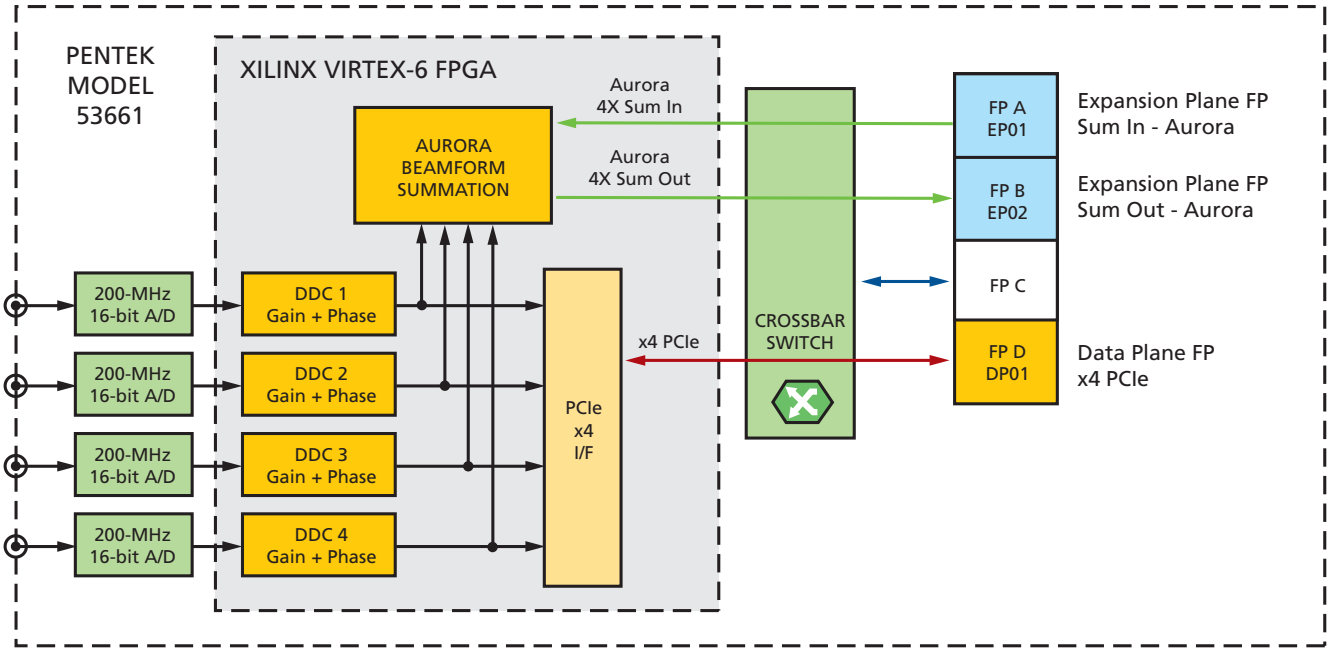

Figure 3 – Pentek Cobalt 53661 OpenVPX beamforming software radio board with a Xilinx Virtex-6 FPGA

adds the four DDC outputs together to perform the channel combining for beamforming. Aurora is a lightweight link-layer gigabit serial protocol for Xilinx FPGAs. In this board, the Aurora interface accepts a propagated sum on one input port with four serial links (4X), and delivers the new propagated sum on a 4X output port, including the contributions from the four onboard channels. Operating at a clock bit rate of 3.125 Gbits/second, each 4X link can transfer data at 1.25 Gbytes/s.

A native PCIe® x4 interface IP operating at a 2.5-Gbit/s serial clock rate provides a 1-Gbyte/s link to the control processor for programming the DDC and the beamforming parameters. This PCIe link also supports delivery of the four DDC outputs as well as the beamforming summation output.

A programmable gigabit serial crossbar switch connects the two 4X Aurora summation links and the x4 PCIe link to the VPX P1 backplane connector. The flexibility of this crossbar switch allows the 53661 to operate in a variety of OpenVPX backplane topologies and slot profiles. In this system, we map the Aurora links onto the OpenVPX expansion plane; we likewise map the PCIe interface onto the OpenVPX data plane, which also assumes the role of control plane.

EIGHT-CHANNEL 3U OPENVPX BEAMFORMING SYSTEM

The complete eight-channel OpenVPX beamforming system is shown in Figure 4. Two Model 53661 boards are installed in slots 1 and 2 of an OpenVPX backplane, along with a CPU board in slot 3. Eight dipole antennas designed for receiving 2.5-GHz signals feed RF tuners containing low-noise amplifiers, local oscillators and mixers. The RF tuners translate the 2.5-GHz antenna frequency signal down to an IF of 50 MHz.

The 200-MHz 16-bit A/D converters digitize the IF signals and perform further frequency downconversion to baseband, with a DDC decimation of 128. This provides I+Q complex output samples with a bandwidth of about 1.25 MHz. Phase and gain coefficients for each channel are applied to steer the array for directionality.

The CPU board in VPX slot 3 sends commands and coefficients across the backplane over two x4 PCIe links, or OpenVPX "fat pipes."

We process the first four signal channels in the upper left portion of the 53661 board in VPX slot 1, where the four-channel beamformed sum is propagated through the 4X Aurora sum-out link across the backplane to the 4X Aurora sum-in port of the second 53661 in slot 2. We then add the four-channel local summation from the second 53661 to the propagated sum from the first board to form the complete eight-channel sum. This final sum is sent across the x4 PCIe link to the CPU card in slot 3.

Assignment of the three OpenVPX 4X links (OpenVPX fat pipes) on the Model 53661 boards is simplified through the use of the crossbar switch shown in the previous block diagram. This allows the 53661 to operate with a wide variety of different backplanes. Because OpenVPX does not restrict the use of serial protocols across the backplane links, the system supports mixed-protocol architectures like the one shown.

BEAMFORMING DEMO SYSTEM

Pentek engineers have set up an eight-channel beamforming demo system equipped with a control panel that runs under Windows on the CPU board. An automatic signal scanner detects the strongest signal frequency arriving from a test transmitter. This frequency is centered around the 50-MHz IF frequency of the RF downconverter. Once the frequency is identified, the eight DDCs are set accordingly to bring that signal down to 0 Hz for summation. The control panel software also allows specific hardware settings for all of the parameters for the eight channels including gain, phase and sync delay.

An additional display shows the beamformed pattern of the array. This display is formed by adjusting the phase shift of each of the eight channels to provide maximum sensitivity across arrival angles from –90° to +90° perpendicular to the plane of the array.

The theoretical seven-lobe pattern of an ideal eight-element array for a signal arriving at a 0° angle (directly in front of the array) is displayed for comparison with an actual plot. Below the lobe pattern is a polar plot showing a single vector pointing to the computed angle of arrival. This is derived from identifying the lobe with the maximum response.

An actual plot of a real-life transmitter is also shown for a source directly in front of the display. In this case, the perfect lobe pattern is affected by physical objects, reflections, cable length variations and minor differences in the antennas. Nevertheless, the directional information is computed quite well. As the signal source moves left and right in front of the array, the peak lobe moves with it, changing the computed angle of arrival.

This demo system is available online at Pentek. If you'd like to view a live demonstration, please let us know of your interest by visiting *http://pentek.com/go/xcellbf*.

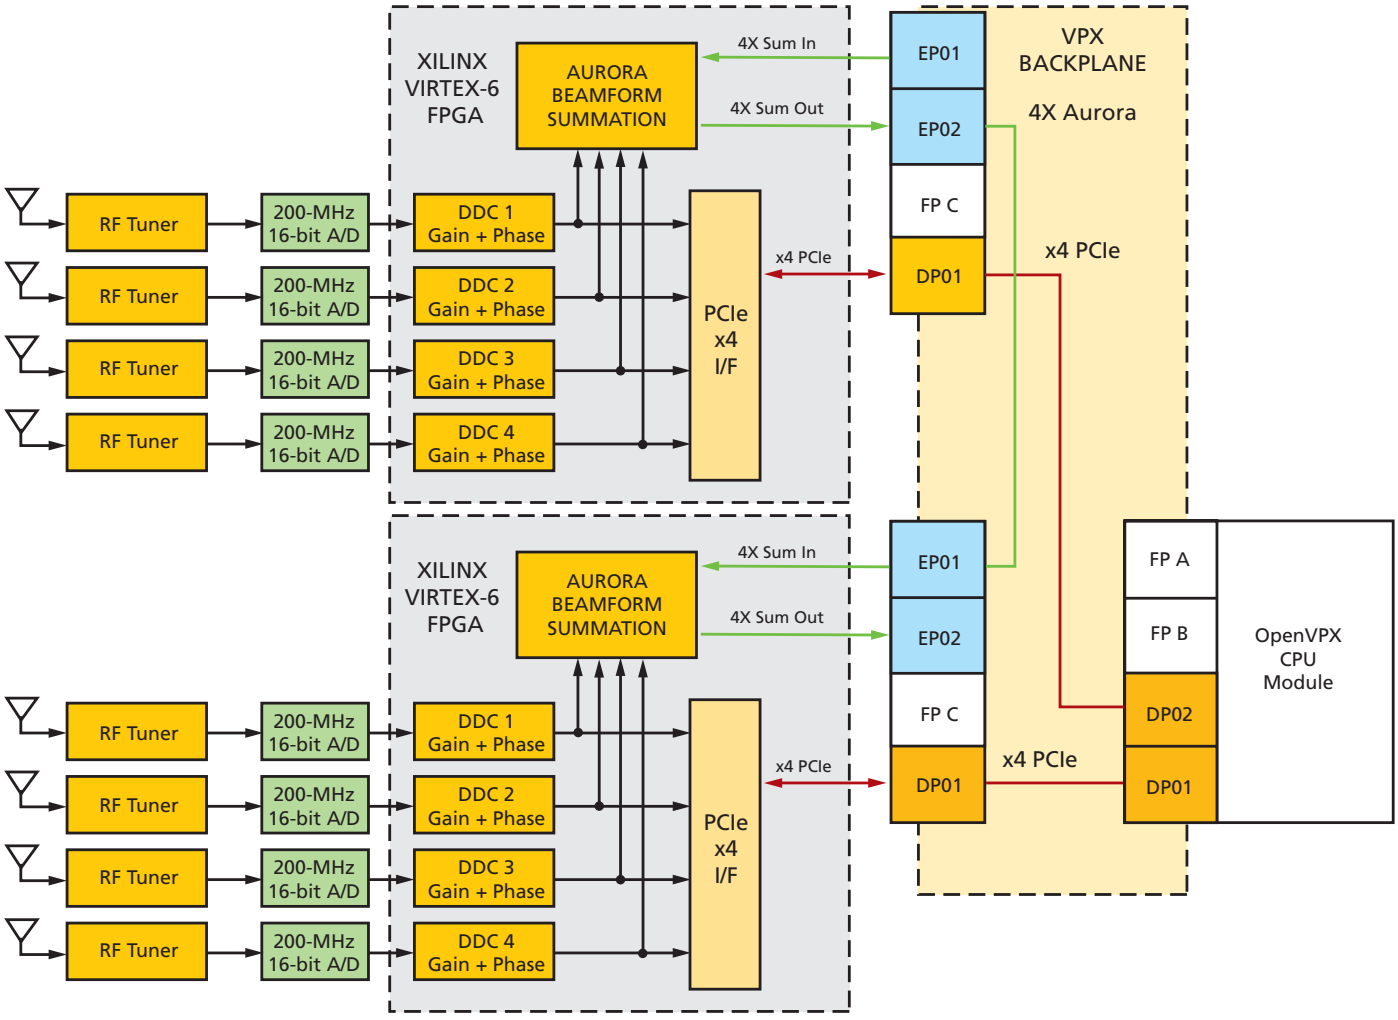

Figure 4 – Eight-channel OpenVPX demo beamforming system utilizes two Pentek Cobalt 53661 beamforming boards.

Sundance Multiprocessor Technology

embedded signal processing solutions

SUNDANCE

FlexTiles Development Platform

Self-Adaptive Heterogeneous MultiCore Solutions based on DSPs and FPGAs

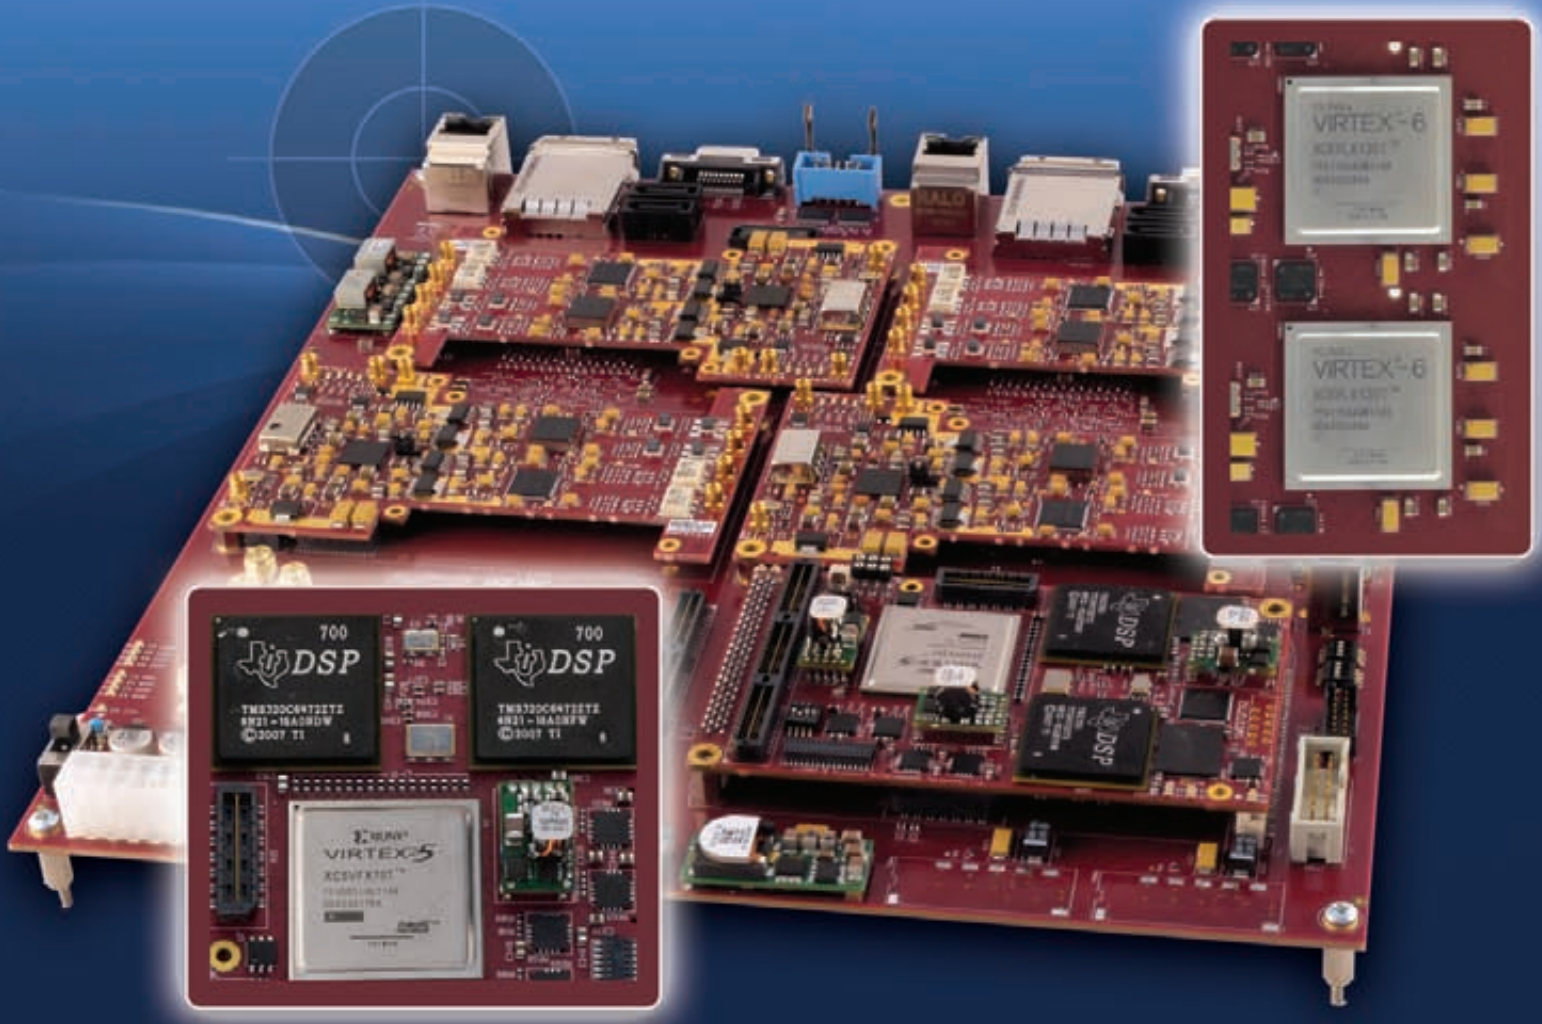

www.FlexTiles.biz

Test & Measurement instruments • Communications, MIMO & Satellite receivers • RADAR ECM systems
Space, Land, Sea, Air equipments • Research, Science & Technology • Industrial control & video solutions

Xilinx Tool & IP Updates

The Vivado™ Design Suite 2012.2 is now available, at no additional cost, to all Xilinx® ISE® Design Suite customers that are currently in warranty. The Vivado Design Suite provides a highly integrated design environment with a completely new generation of system-to-IC-level features, including high-level synthesis, analytical place-and-route and an advanced timing engine. These tools enable developers to increase design integration and implementation productivity.

WHAT IS THE VIVADO DESIGN SUITE?

It's all about improving designer productivity. This entirely new tool suite was architected to increase your overall productivity in designing, integrating and implementing with the 28-nanometer family of Xilinx All Programmable devices. With 28-nm manufacturing, Xilinx devices are now much larger and come with a variety of new technologies including stacked-silicon interconnect, high-speed I/O interfaces operating at up to 28 Gbps, hardened microprocessors and peripherals, and Agile Mixed Signal. With these larger and more complex devices, developers are faced with multidimensional design challenges that can prevent them from hitting market windows and increasing productivity.

The Vivado Design Suite is a complete replacement for the existing Xilinx ISE Design Suite of tools. It replaces all of the ISE Design Suite point tools such as Project Navigator, XST, implementation, CORE Generator™, Timing Constraints Editor, ISim, ChipScope™, Xilinx Power Analyzer (XPA), FPGA Editor, PlanAhead™ and SmartXplorer, among others. All of those capabilities are now built directly into the Vivado Integrated Design Environment (IDE), leveraging a shared scalable data model.

With the Vivado Design Suite, developers are able to accelerate design creation with high-level synthesis and implementation by using place-and-route to analytically optimize for multiple and concurrent design metrics, such as timing, congestion, total wire length, utilization and power. Built on Vivado's shared scalable data model, the entire design process can be executed in memory without the need to write or translate any intermediate file formats, accelerating runtimes, debug and implementation while reducing memory requirements.

Vivado provides users with upfront metrics that allow for design and tool-setting modifications earlier in the design process, when they have less overall impact on the schedule. This capability reduces design iterations and accelerates productivity.

Users can manage the entire design process in a pushbutton manner by using the Flow Navigator feature in the Vivado IDE, or control it manually by using Tcl scripting.

SHOULD I CONTINUE TO USE THE ISE DESIGN SUITE OR MOVE TO VIVADO?

The ISE Design Suite is an industry-proven solution for all generations of Xilinx's All Programmable devices. The Xilinx ISE Design Suite continues to bring innovations to a broad base of developers, and extends the familiar design flow for 7 series and Xilinx Zynq™-7000 Extensible Processing Platform (EPP) projects. ISE 14.2, which brings new innovations and contains updated device support, is available for immediate download.

The Vivado Design Suite 2012.2, Xilinx's next-generation design environment, supports 7 series devices including Virtex®-7, Kintex™-7 and Artix™-7 FPGAs. It offers enhanced tool performance, especially on large or congested designs.

Xilinx recommends that customers starting a "new" design on a Kintex K410 or larger device contact their local FAE to determine if Vivado is right for the design. Xilinx does not recommend transitioning during the middle of a current ISE Design Suite project, as design constraints and scripts are not compatible between the environments.

For more information, please read the ISE 14.2 and Vivado 2012.2 release notes.

WHAT ARE THE LICENSING TERMS FOR VIVADO?

There is no additional cost for the Vivado Design Suite during the remainder of 2012. A single download at the Xilinx download center contains both ISE Design Suite 14.2 and Vivado 2012.2. All current, in-warranty, seats of ISE Design Suite are entitled to a copy of the Vivado Design Suite beginning with the 2012.2 release.

For customers who generated an ISE Design Suite license for versions 13 or 14 after Feb. 2, 2012, your current license will also work for Vivado. Customers who are still in warranty but who generated licenses prior to February 2 will need to regenerate their licenses in order to use Vivado. For license generation, please visit *www.xilinx.com/getlicense*.

IS VIVADO DESIGN SUITE TRAINING AVAILABLE?

Vivado is new and takes full advantage of industry standards such as powerful interactive Tcl scripting, Synopsys Design Constraints, SystemVerilog and more. To reduce your learning curve, Xilinx has rolled out 10 new instructor-led classes that include information on how to use the Vivado tools. We also encourage you to view the Vivado Quick Take videos found at *www.xilinx.com/design-tools*.

ARE THERE DIFFERENT EDITIONS OF THE VIVADO DESIGN SUITE?

The Vivado Design Suite is available in either the Design or System edition (see Table 1). In-warranty ISE Design Suite Logic and Embedded Edition customers will receive the new Vivado Design Edition. ISE Design Suite DSP and System Edition customers will receive the new Vivado System Edition. Vivado is not yet available for WebPACK™ users. Vivado WebPACK is currently planned for later this year.

For more information about Xilinx design tools for the next generation of All Programmable devices, please visit *www.xilinx.com/design-tools*.

VIVADO DESIGN SUITE EDITIONS

PILLARS OF PRODUCTIVITY	FEATURES	WEBPACK (DEVICE LIMITED)	DESIGN EDITION	SYSTEM EDITION
IP Integration and Implementation	Integrated Design Environment	●	●	●
	Software Development Kit (SDK)	●	●	●
Verification and Debug	Vivado Simulator	Limited	●	●
	Vivado Logic Analyzer		●	●
	Vivado Serial I/O Analyzer		●	●
Design Exploration and IP Generation	Vivado High-Level Synthesis			●
	System Generator for DSP			●

Table 1 – Vivado Design Suite editions; WebPACK support will roll out later in the year.

Xpress Yourself in Our Caption Contest

DANIEL GUIDERA

We hope this issue's challenge won't find you at a loss for words. If you have a yen to Xercise your funny bone, we invite you to step up and submit an engineering- or technology-related caption for this cartoon showing a technical presentation by a mime. The image might inspire a caption like "The demand to present his system diagram before the board of directors left Edgar speechless."

Send your entries to *xcell@xilinx.com*. Include your name, job title, company affiliation and location, and indicate that you have read the contest rules at *www.xilinx.com/xcellcontest*. After due deliberation, we will print the submissions we like the best in the next issue of *Xcell Journal*. The winner and two runners-up will each receive an Avnet Spartan®-6 LX9 MicroBoard, an entry-level development environment for evaluating the Xilinx® Spartan®-6 family of FPGAs (*http://www.xilinx.com/products/boards-and-kits/AES-S6MB-LX9.htm*).

The deadline for submitting entries is 5:00 pm Pacific Time (PT) on Oct. 1, 2012. So, pull on your beret and get writing!

GREG ZAKER, senior hardware design engineer at ViaSat Inc., won a shiny new Spartan-6 evaluation board with this caption for the scene of the hypnotic image on the computer monitor from Issue 79 of *Xcell Journal*:

"No, I said let's look at the Virtex code, not vortex."

Congratulations as well to our two runners-up. Like our winner, they will each receive an Avnet Spartan-6 LX9 MicroBoard.

"... and this is Bob.
He's taking our online fast-track management course ..."

– Richard M. Myers, VP, senior staff development engineer, AISD Inc.

"Attempting an innovative computer penetration technique, John became a victim of his own hack."

– Tom Drake, electrical engineer, Harris

No Room for Error

Synopsys FPGA Design Tools are an Essential Element for Today's High Reliability Designs

There is no room for error in today's communications infrastructure systems, medical and industrial applications, automotive, military/aerospace designs. These applications require highly reliable operation and high availability in the face of radiation-induced errors that could result in a satellite failing to broadcast, a telecom router shutting down or an automobile failing to respond to a command.

To learn more about Synopsys FPGA Design tools visit www.synopsys.com/fpga

SYNOPSYS®
Predictable Success


```
pename T> class TCLE

E*,tcltypes::tclobj,
base<NATIVE*,tcltype

ase<NATIVE*,tcltypes

cltypes::tclobj> tcl
nfo<NATIVE*,tcltypes

cl_Obj* tobj)
```

Productivity. Multiplied.

Vivado™ Design Suite. For the Next Decade of All-Programmable Devices.

You innovate.

We attack the major design bottlenecks.

Vivado tools not only speed the design of programmable logic and I/O, but accelerate programmable systems integration and implementation into devices incorporating 3D stacked silicon interconnect technology, ARM® processing systems, Analog Mixed Signal (AMS), and a significant percentage of intellectual-property (IP) cores.

Your design productivity just got multiplied.

www.xilinx.com/designtools

XILINX®
ALL PROGRAMMABLE

PN 2521