ISSUE 59, FOURTH QUARTER 2006 XCELL JOURNAL XILINX, INC.

Issue 59
Fourth Quarter 2006

Xcell journal

THE AUTHORITATIVE JOURNAL FOR PROGRAMMABLE LOGIC USERS

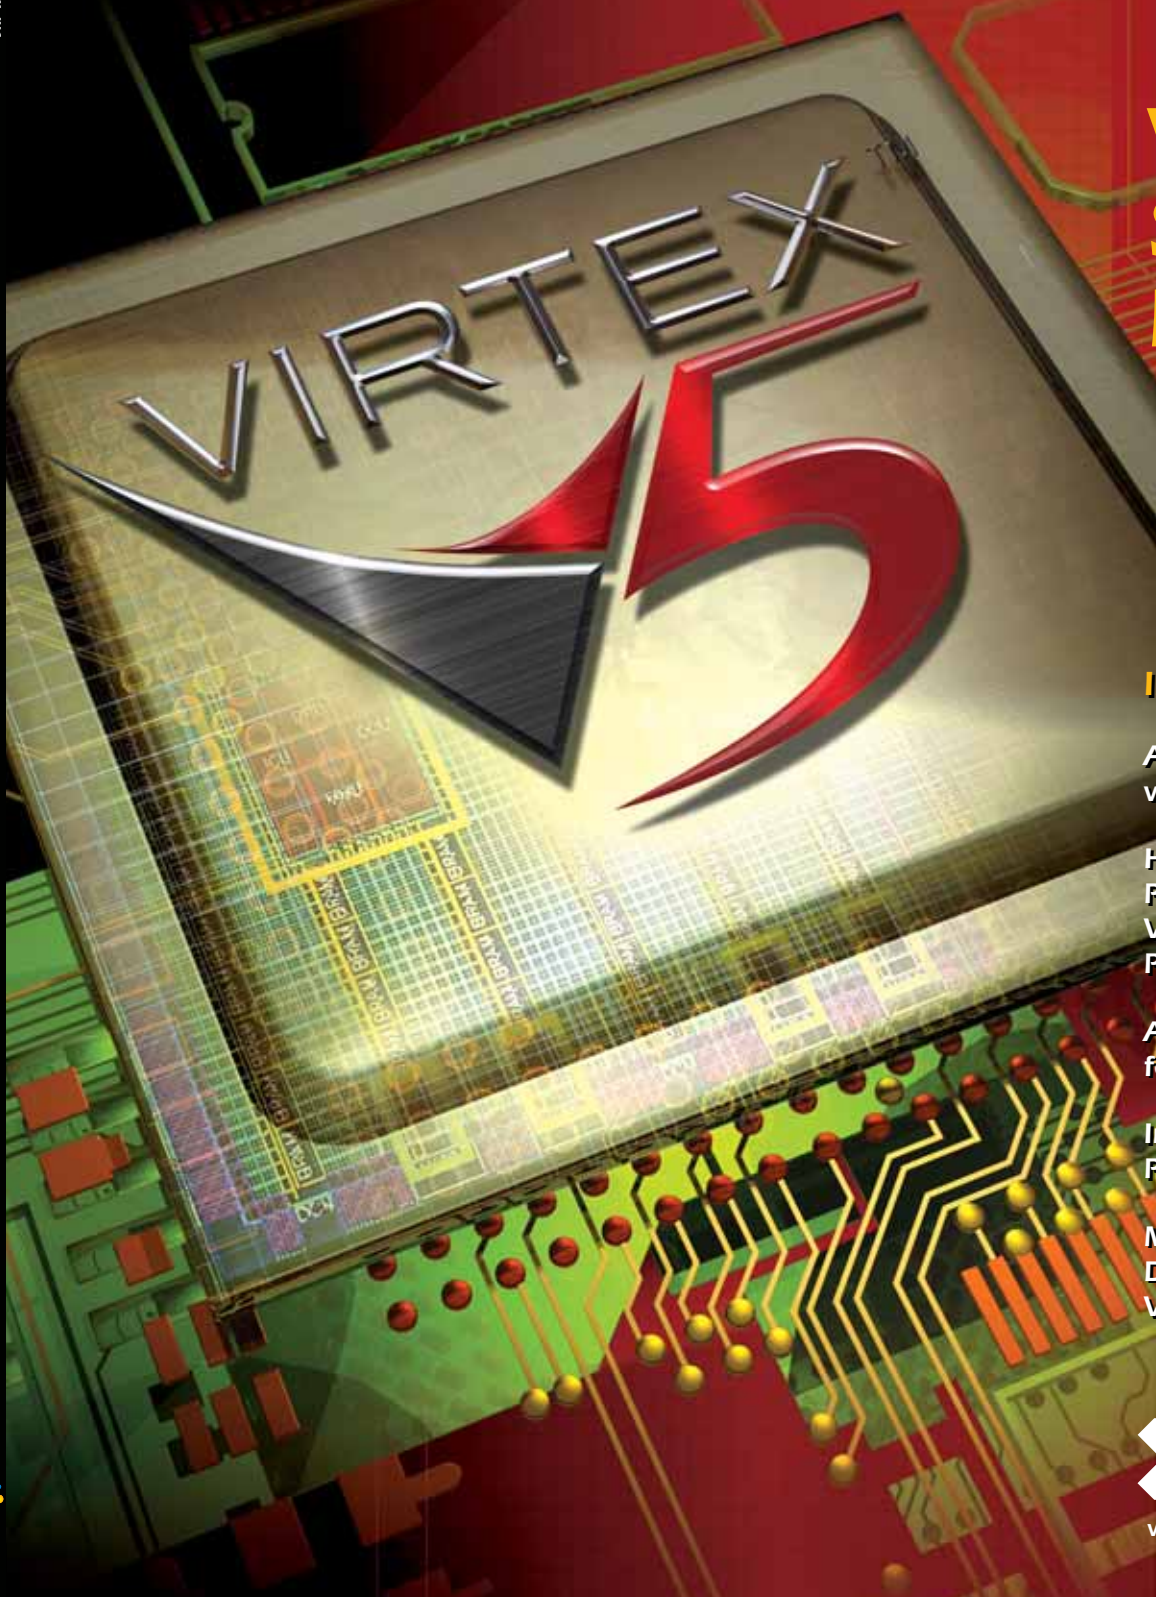

Virtex-5 Special Edition

INSIDE

XILINX®

www.xilinx.com/xcell/

Support Across The Board.™

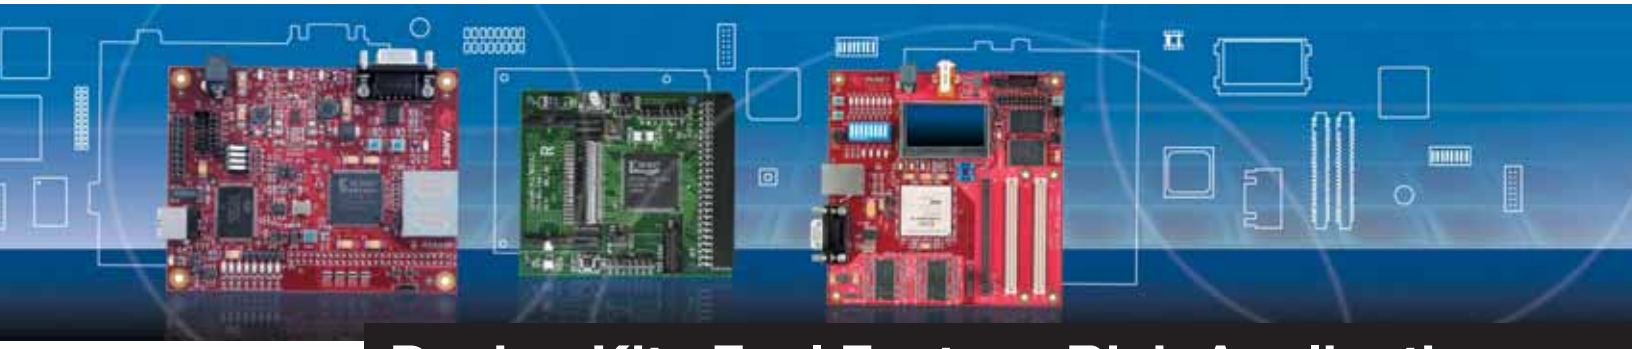

Design Kits Fuel Feature-Rich Applications

Build your own system by mixing and matching:

- Processors
- FPGAs
- Memory
- Networking
- Audio
- Video
- Mass storage
- Bus interface
- High-speed serial interface

Available add-ons:

- Software
- Firmware
- Drivers
- Third-party development tools

Silica designs, manufactures, sells and supports a wide variety of hardware evaluation, development and reference design kits for developers looking to get a quick start on a new project.

With a focus on embedded processing, communications and networking applications, this growing set of modular hardware kits allows users to evaluate, experiment, benchmark, prototype, test and even deploy complete designs for field trial.

By providing a stable hardware platform that enhances system development, design kits from Silica help original equipment manufacturers (OEMs) bring differentiated products to market quickly and in the most cost-efficient way possible.

For a complete listing of available boards, visit **www.silica.com**

XILINX®

SILICA® An Avnet Company

SILICA I The Engineers of Distribution
www.silica.com

Push the Performance of Your Complex Designs with Virtex-5 and Synplify Pro

Synplicity and Xilinx Team Up and Deliver

In May 2006, Xilinx announced Virtex-5 - the industry's first 65nm FPGA. Long before this announcement was made, however, Xilinx knew that an innovative kind of Synthesis tool would be required to enable users to take advantage of the performance and logic density of these revolutionary devices. That's why, once again, Xilinx teamed up with Synplicity.

Engineers at Xilinx and Synplicity have worked closely over the past year to ensure that Synplicity's market-leading synthesis software would provide optimal support for the new Virtex-5 devices. The resulting changes made to the synthesis algorithms allow users to take maximum advantage of these high-capacity Virtex-5 devices.

The enhanced optimizations built into the Synplify Pro software, along with its timing-driven approach to synthesis, allow designers to push the performance of their complex designs while remaining comfortably within their time-to-market goals.

Looking Beyond Today

Future generations of devices will contain even greater density points and capabilities than the current devices, further expanding the reach of advanced FPGA architectures across a wide range of application domains. That's why Synplicity and Xilinx have formed an Ultra-high Capacity Timing Closure Task Force. The purpose of this task force is to enable engineers from both companies to collaborate to define and implement new design flows that maximize the quality of results of design productivity of ultra-high density designs with next-generation 65-nanometer FPGAs.

To learn more about how Synplicity's tools can help you gain maximum performance from your complex designs, contact Synplicity at info@synplicity.com or visit our website at www.synplicity.com.

Synplicity®

Simply Better Results

VIRTEX™ 5

EXPRESSFABRIC

Ultimate Performance…

Achieve highest system speed and better design margin with the world's first 65nm FPGAs.

Virtex™-5 FPGAs feature ExpressFabric™ technology on 65nm triple-oxide process. This new fabric offers the industry's first LUT with six independent inputs for fewer logic levels, *and* advanced diagonal interconnect to enable the shortest, fastest routing. Now you can achieve 30% higher performance, while reducing dynamic power by 35% and area by 45% compared to previous generations.

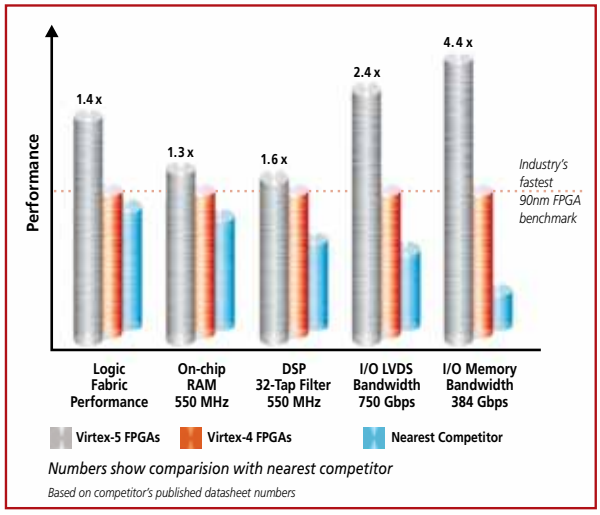

Design systems faster than ever before

Shipping now, Virtex-5 LX is the first of four platforms optimized for logic, DSP, processing, and serial. The LX platform offers 330,000 logic cells and 1,200 user I/Os, plus hardened 550 MHz IP blocks. Build deeper FIFOs with 36 Kbit block RAMs. Achieve 1.25 Gbps on all I/Os without restrictions, and make reliable memory interfacing easier with enhanced ChipSync™ technology. Solve SI challenges and simplify PCB layout with our sparse chevron packaging. And enable greater DSP precision and dynamic range with 550 MHz, 25x18 MACs.

Visit ***www.xilinx.com/virtex5***, view the TechOnline webcast, and give your next design the ultimate in performance.

XILINX®
The Programmable Logic Company℠

www.xilinx.com/virtex5

The Ultimate System Integration Platform

Xcell journal

PUBLISHER Forrest Couch
forrest.couch@xilinx.com
408-879-5270

EDITOR Charmaine Cooper Hussain

ART DIRECTOR Scott Blair

DESIGN/PRODUCTION Teie, Gelwicks & Associates
1-800-493-5551

ADVERTISING SALES Dan Teie
1-800-493-5551

TECHNICAL COORDINATOR Greg Lara

INTERNATIONAL Dickson Seow, Asia Pacific
dickson.seow@xilinx.com

Andrea Barnard, Europe/
Middle East/Africa
andrea.barnard@xilinx.com

Yumi Homura, Japan
yumi.homura@xilinx.com

SUBSCRIPTIONS All Inquiries
www.xcellpublications.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/

Welcome to this special edition of *Xcell Journal*, featuring a broad array of articles on Xilinx® Virtex™-5 FPGAs. In this issue you'll find executive and industry viewpoints; articles on engineering solutions, design challenges, tools, customer successes, and vertical markets; and a technical reference section covering application notes, boards, and IP.

As exciting as this is, I'd also like to let you know about a couple of announcements from *Xcell* Publications.

Xcell Publications Honored with APEX 2006 Award of Excellence

Xcell Publications was recently awarded the APEX 2006 Award of Excellence in two categories – magazine and journal design and layout and custom-published magazines and journals – for two of its flagship *Xcell* Publications, *Xcell Journal* and *I/O Magazine*.

APEX 2006 – the 18th Annual Awards for Publication Excellence – is an international competition that recognizes outstanding publications, including newsletters, magazines, annual reports, brochures, and websites. According to APEX judges, this year's competition was exceptionally intense, with nearly 5,000 entries. Awards were granted based on excellence in graphic design, quality of editorial content, and the success of the entry in conveying the message and achieving overall communications effectiveness.

"We're honored that *Xcell* magazines have been selected for excellence in publishing among such a stellar list of companies by the APEX panel of judges," said Sandeep Vij, vice president of worldwide marketing at Xilinx. "Over the past 18 years, our custom publications have served as a foundational tool, delivering 'how-to' information to a growing base of engineers using Xilinx programmable chips to design a wide variety of electronic systems, ranging from the Mars Rover to high-volume consumer handsets, flat-panel displays and automotive infotainment systems. Being ranked among the industry's best underscores the value and quality of our company's portfolio of custom magazines."

Xilinx joins a prestigious list of award-winning companies from a variety of industries in the APEX competition for custom-published magazines and journals, including Blue Cross Blue Shield, CMP Media/Digital Connect, DaimlerChrysler, IBM Journal of Research and Development, Mac Publishing, National Football League, National Foundation for Advancement in the Arts, Penton Custom Media, and Time Inc. Strategic Communications.

New Digital Editions Available

We now offer digital editions of our magazines. Now you can subscribe for free to the new *Xcell Journal Digital*, requiring no software downloads and visible on any standard Internet browser. This updated publishing technology lets you browse, search, make notes, e-mail authors, and click through to advertisers' websites.

To receive *Xcell Journal Digital*, you have to subscribe. In addition to *Xcell Journal*, we also now offer digital subscriptions of all of our magazines. Please visit our website at *www.xilinx.com/xcell* and click on "Subscriber Services."

I hope you enjoy reading this issue.

Forrest Couch
Publisher

PERFORMANCE

Achieve Higher Performance with Virtex-5 FPGAs
New architectural elements can help you attain higher system-level performance.

PERFORMANCE

HDL Coding and Design Practices for Improving Virtex-5 Utilization, Performance, and Power
These tips and techniques can lead to better Virtex-5 designs.

Introducing the Virtex-5 FPGA Family
The first 65-nm advanced FPGAs raise the bar in performance, power efficiency, capacity, and value.

SERIAL CONNECTIVITY

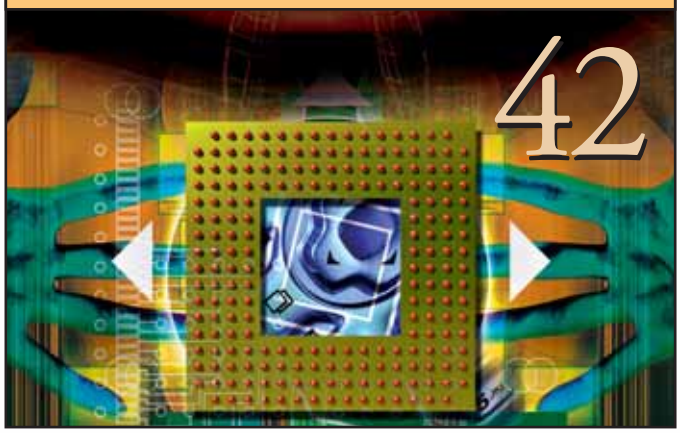

A Multi-Gigabit Transceiver for the Masses
The Virtex-5 GTP transceiver brings versatility, ease of use, power efficiency, and cost-effectiveness to high-volume mainstream applications.

SERIAL CONNECTIVITY

Introducing the Virtex-5 PCI Express Endpoint Block
With PCI Express quickly becoming the standard high-bandwidth interconnect, the Virtex-5 LXT PCIe Endpoint block enables a configurable single-chip solution.

MEMORY INTERFACES

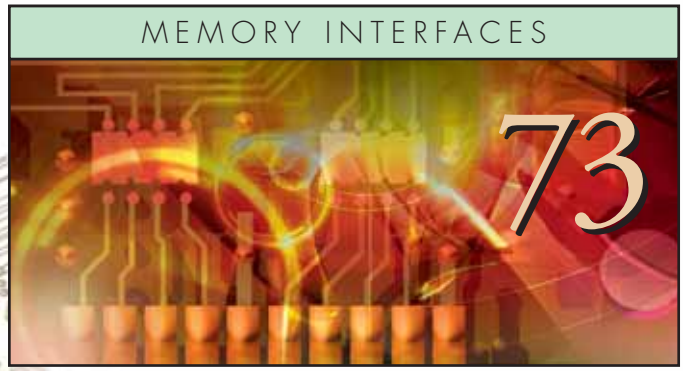

Meeting Memory Interface Design Challenges with Virtex-5 FPGAs
Virtex-5 devices support the latest generation of high-speed memory interfaces.

FOURTH QUARTER 2006, ISSUE 59

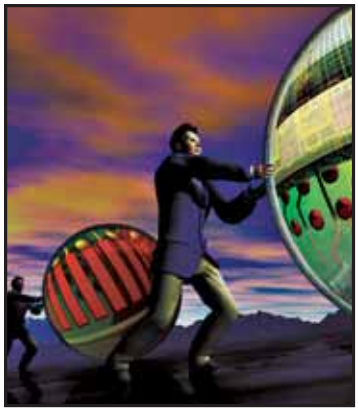

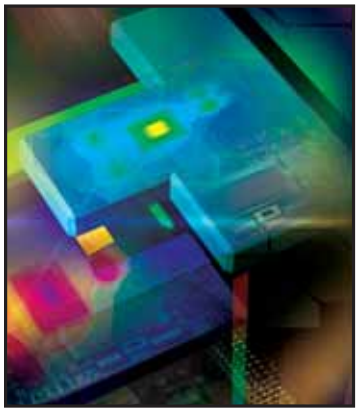

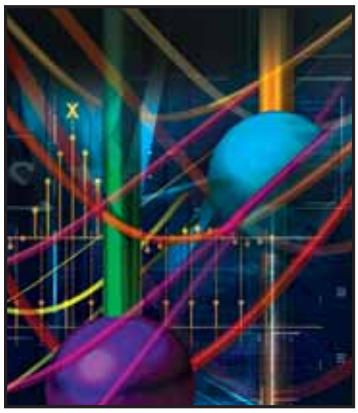

Xcell journal

Introducing the Virtex-5 FPGA Family

The first 65-nm advanced FPGAs raise the bar in performance, power efficiency, capacity, and value.

by Steve Douglass
Vice President
Product Development,
Advanced Product Division
Xilinx, Inc.
stephen.douglass@xilinx.com

Welcome to the Virtex™-5 issue of *Xcell Journal*. The Xilinx® Virtex-5 family is not only the industry's first 65-nm FPGA – it also offers some of the most advanced architecture and highest performance in the world. Continuing our history of developing groundbreaking technology, we listened to leading design engineers in various markets and built on key characteristics that made our Virtex-4 FPGA family a tremendous success:

- Higher performance
- Higher logic density
- Lower power consumption
- More advanced features

The fundamental value propositions of FPGAs include faster time to market, versatility, support for evolving standards, risk mitigation, field upgradability, and lower system costs. Our FPGAs accommodate your demands for continued improvements in performance, capacity, power consumption, and cost.

The Virtex-5 family combines the inherent advantage of state-of-the-art 65-nm process technology with an innovative design that is based on a deeper understanding of the applications our products serve. In this article, I'll provide an overview of the new features in Virtex-5 devices, explain the underlying technology, and offer a glimpse of the design decisions that led to our world-leading FPGA architecture.

Process Technology and Architectural Innovations

Virtex-5 FPGAs are built on 65-nm triple-oxide technology using our Advanced Silicon Modular Block (ASMBL™) architecture and providing additional levels of system integration. This new family offers an advanced platform that meets the growing need for programmable systems with higher performance, higher density, lower power consumption, and lower overall system cost.

It might be easy to deliver on one or two of these items, but our challenge was to deliver all of them at the same time.

We successfully met those challenges through a combination of advanced IC process development and innovative architecture and circuit design. Introduced in the Virtex-4 family, our proven ASMBL chip layout architecture allows us to provide the optimal mix of required device resources (logic, memory, arithmetic, I/O, and IP), thus creating the ideal combination for four new platforms:

- The LX platform, optimized for high-performance logic
- The LXT platform, optimized for high-performance logic with low-power serial I/O
- The SXT platform, optimized for high-performance arithmetic- and memory-intensive DSP with low-power serial I/O
- The FXT platform, optimized for embedded processing and very high-speed serial I/O

Compared to our Virtex-4 family, Virtex-5 devices offer 30 percent higher average speed and 65 percent higher capacity in the largest device. Dynamic power consumption is reduced by 35 percent and chip area is 45 percent smaller, resulting in a lower cost per function.

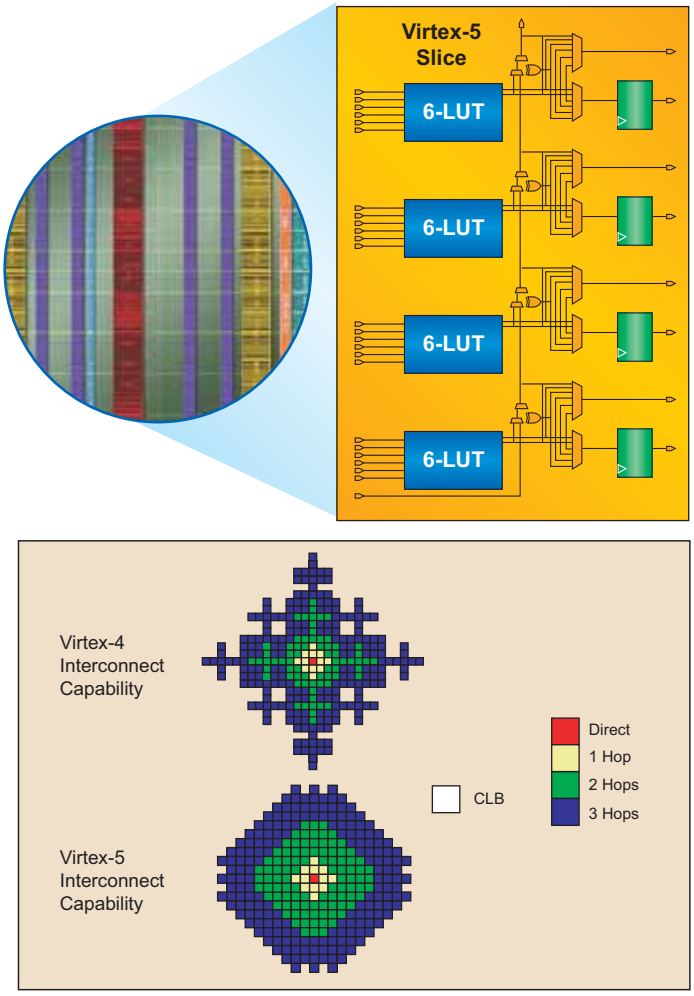

Figure 1 – Virtex-5 ExpressFabric technology

Higher Performance and Density

ExpressFabric™ technology implements logic and local interconnect routing. It incorporates look-up tables (LUTs) with six independent inputs, plus a new diagonal interconnect structure, as illustrated in Figure 1. ExpressFabric technology implements combinatorial logic in fewer LUT levels and uses fewer concatenated connections to neighboring building blocks, as compared to the Virtex-4 architecture. This reduces datapath delays and thus increases design performance.

Advanced 6-LUT Logic Structure

For many years, four-input LUTs were the industry standard. However, at 65 nm, the regular structure of the LUT can be shrunk even more than the remaining circuitry (notably, the interconnect). A six-input LUT (6-LUT) with four times more bits thus increases the CLB area by only 15% – but packs, on average, 40% more logic into each LUT. This higher logic density often reduces the number of cascaded LUTs and can improve the critical path delay, as shown in Figure 2.

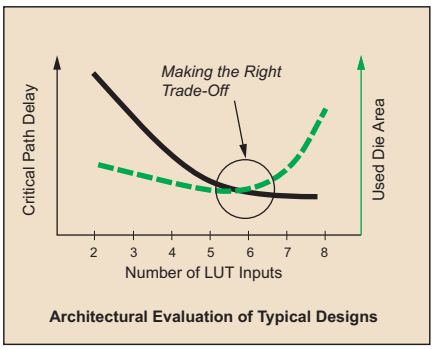

Figure 2 – Optimal performance/area trade-off

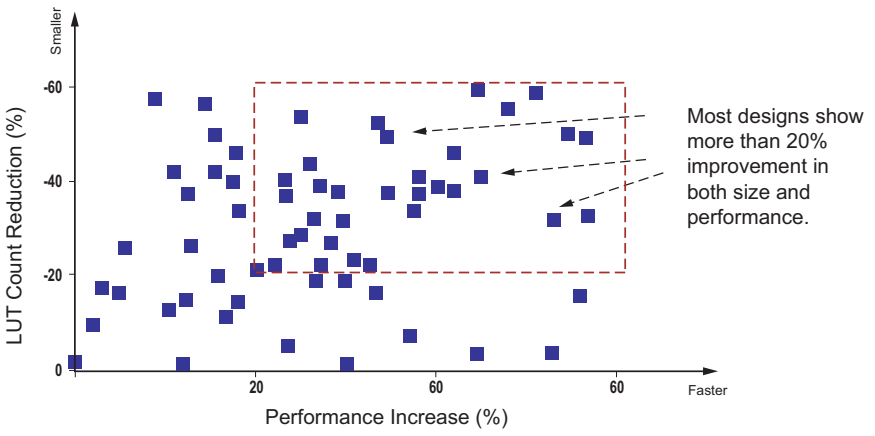

Figure 3 – Virtex-5 FPGA versus Virtex-4 FPGA design suite benchmarks

We took a suite of customer designs and implemented them using ISE™ 8.1i software. For each design, we compared the number of LUTs used with Virtex-4 and Virtex-5 device implementations and correlated this information with the performance increase in megahertz. The scatterplot graph in Figure 3 shows the percentage of performance improvement on the X axis and percentage area savings in terms of LUT count reduction on the Y axis. The new 6-LUT ExpressFabric technology provides a win-win solution in both performance gain and resource savings.

Unlike competing FPGAs, Virtex-5 FPGAs provide real 6-LUTs that you can use as logic or as distributed memories, where a LUT can be a 64-bit distributed RAM (even dual- or quad-ported) or a 32-bit programmable shift register. Each LUT can have two outputs, thus implementing two logic functions of five variables, storing 32 x 2 RAM bits, or acting as a 16 x 2-bit shift register.

New Diagonally Symmetric Interconnect

A new diagonally symmetric interconnect pattern enhances performance by reaching more places in fewer routing hops. A comparison between Virtex-5 and Virtex-4 FPGA interconnect patterns (with each box representing a CLB) is illustrated in Figure 1. The color codes show that with the Virtex-5 FPGA, the pattern is more symmetric, with more CLBs reached in fewer hops. The symmetry thus achieves better results from place and route software tools.

These features are transparent to Virtex-5 FPGA users and are automatically exercised by ISE software, resulting in easier routability and higher overall performance.

Lowest Power Advanced FPGA Solution

The Virtex-5 device family uses our advanced 65-nm, triple-oxide, 11-layer copper CMOS process technology. "Triple oxide" refers to the number of different transistor gate-oxide thicknesses used. The I/O transistors must be 3.3V tolerant and use relatively thick oxide, but the very fast transistors used for logic and other core functions use very thin oxide.

Unfortunately, very thin oxide and very low threshold voltage unavoidably cause high leakage current. There are, however, many transistors in an FPGA that need not be very fast (notably the configuration storage cells). Starting with the Virtex-4 family, Xilinx pioneered a third, intermediate gate thickness for those transistors. This triple-oxide approach allows us to fine-tune the performance and power in the device circuitry; it enables Virtex-5 devices to deliver industry-leading performance while dramatically lowering leakage current and thus static power consumption.

Additionally, the new 6-LUT logic structure combines more logic per LUT, uses fewer local interconnect nodes, and fewer high capacitance nodes between logic functions, reducing the levels of logic and thus the path delay. The new symmetric routing also uses more direct connects between adjacent logic, again lowering routing capacitance.

V_{CCINT}, the core supply voltage, is now 1.0V. All of these factors contribute to a reduction in overall dynamic power consumption. With the success of the Virtex-4 family, we know that many engineers view performance and power consumption as two equally important constraints in their system designs; therefore, we need to offer both high performance and low power.

We completely reengineered the Virtex-5 logic fabric to fully take advantage of the 65-nm triple-oxide CMOS process, resulting in the highest performance fabric ever, with system clock rates in excess of 550 MHz. At the same time, static power is comparable to that of the 90-nm Virtex-4 devices, while dynamic power has been reduced by at least 35%. Just like its predecessor, the Virtex-5 family again provides the lowest power solution of any advanced FPGA family.

Advanced Features for System Integration

In the Virtex-5 family, we have added a phase-locked loop (PLL) to each clock management tile (CMT), which now contains two digital clock managers (DCMs) and one PLL. The CMT thus offers the best of both worlds: the robust versatility and precise incremental phase shift capability of a digital clock manager combined with the jitter reduction from the analog PLL. The largest device in the family has six CMTs capable of generating and manipulating 550-MHz clocks, supporting the performance of Virtex-5 logic and block functions.

Synchronous dual-ported block RAM is an important function. The size of each block RAM has been increased to 36 Kb, but you can also use it as two independent 18-Kb block RAMs. The data bus width is programmable from 1 bit to 36 bits. In simple dual port mode (one port write, one port read) the data bus width can be as high as 72 bits, effectively doubling the data bandwidth. You can turn off unused 18-Kb blocks to save power.

The block RAM has integrated FIFO control logic, simplifying the design of

asynchronous (or synchronous) FIFOs running as fast as 550 MHz without consuming any logic resources.

The 72-bit-wide block RAM now includes 64-bit error checking and correction (ECC) control logic. Like the integrated FIFO support, the integrated ECC improves memory performance and eliminates the cost associated with traditional fabric-based solutions. You can also use the dedicated ECC logic to augment external memory interfaces.

Interfacing to external devices and especially external memory such as DDR, DDR2, QDR II, and RLDRAM II is dramatically enhanced and simplified by our new ChipSync™ technology. A memory development system (ML561) based on our LX50T devices contains fully functional and hardware-proven reference designs for all of today's most popular memory technologies.

In the DSP domain, we are now providing 25 x 18-bit multipliers, mainly for more efficient floating-point designs. These DSP48E slices can be directly cascaded for higher performance in digital filtering or video broadcast applications. Direct cascading also saves power – as much as 40% compared to competing solutions.

Virtex-5 SelectIO™ technology continues to lead the industry. Every pin supports virtually every I/O standard in use today and offers up to 1.25 Gbps LVDS and 800 Mbps single-ended I/O performance.

Beyond the IDELAY option, which offers programmable input delay in steps of 75 ps, the new ODELAY option now offers the same fine granularity at the FPGA output. Either of these functions is individually programmable on every device pin.

The IODELAY function is an important feature to enhance reliable transmit and receive of high-speed source-synchronous data and clocks. The intended application includes compensation for board-level skews, bit alignment in a bus, and alignment between data and clock signals. This enables LVDS I/Os to achieve speeds as fast as 1.25 Gbps per pin pair.

Virtex-5 LXT, SXT, and FXT devices also offer embedded serial transceivers – as many as 24 in the largest LXT device. In designing our fourth-generation RocketIO™ technology of high-speed serial transceivers, we invested significant engineering effort to lower power consumption. At the top speed of 3.2 Gbps, the LXT RocketIO transceiver consumes typically less than 100 mW, making it the lowest power transceiver in any FPGA product (see Figure 4).

Each Virtex-5 LXT RocketIO transceiver is programmable and can implement a myriad of speed and serial standards. Link-layer IP is available for such standards as Ethernet, HD/SDI, Serial RapidIO, FibreChannel, and Aurora. Finally, we anticipated the popularity of PCI Express (PCIe) endpoint applications and integrated the complete PCIe endpoint protocol in hard logic. The Virtex-5 LXT PCIe Endpoint block is fully compliant to PCIe standard specification version 1.1 and can support x1, x2, x4, and x8 lane implementations. The integrated hard IP saves logic resources and improves performance for increasingly popular PCIe applications. For an x4 PCIe lane implementation, the Virtex-5 PCIe subsystem block saves as many as 8,500 LUTs compared to implementation with soft IP.

Virtex-5 devices offer more and smaller I/O banks. The outer I/O banks (as many as eight banks in the largest device) also are arranged to provide a PCB routing advantage that in some cases might save board layers.

To ensure the best simultaneously switching output (SSO) performance and provide the best signal integrity (SI) solution in the FPGA industry, all Virtex-5 devices use Xilinx sparse chevron technology pinout assignments. This ensures that each I/O pin is closely surrounded by power and ground pins, thus minimizing current loop inductance and improving SI.

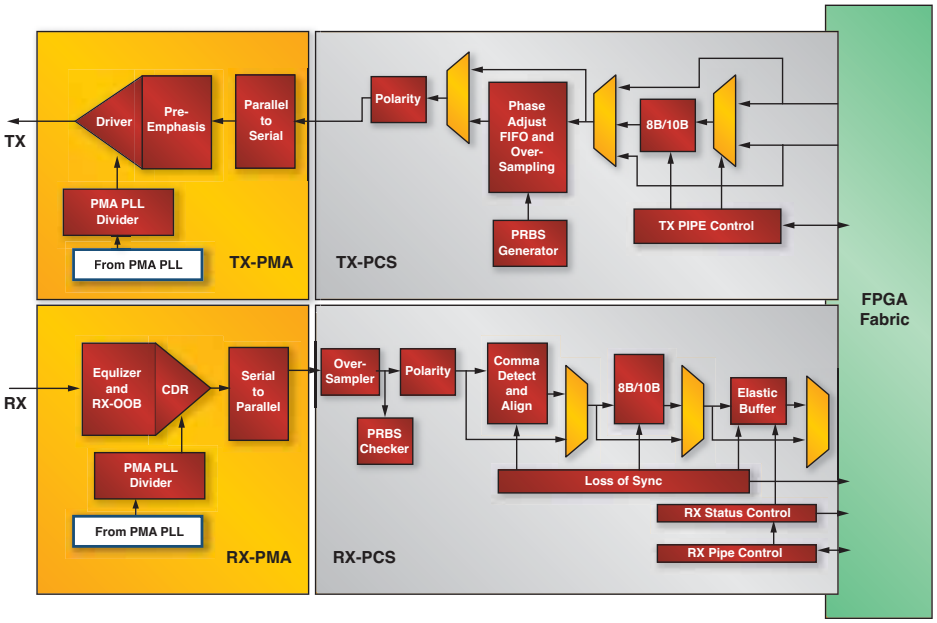

Figure 4 – RocketIO GTP transceiver

Conclusion

I hope that you have enjoyed reading about Virtex-5 devices and the factors that drove their design. At Xilinx, we have truly enjoyed the excitement in the system engineering community about this new architecture. We look forward to seeing your next-generation systems benefit from the Virtex-5 enhanced performance and functionality, taking your complex designs to the next level.

View
from the top

Serial Everywhere — The Triple-Play Challenge

Xilinx is helping to empower the next innovation in the triple-play race.

by Wim Roelandts
CEO and
Chairman of the Board
Xilinx, Inc.

The electronics industry is pressed to its limits as it strives to develop solutions to feed the insatiable appetites of the consumer and enterprise markets for voice, video, and computer data communications on a single network. To the global broadcast and telecommunications industries, the triple-play opportunity is at once a potentially inexhaustible source of revenue and a constant source of frustration. Despite the immeasurable reward for successfully delivering triple-play services to the masses, substantial obstacles continue to impede access.

Perhaps the most central of these obstacles is the inadequacy of legacy infrastructure equipment to support the massive increases in bandwidth. Evolving from voice-only, the legacy infrastructure is a complex web of overlaid networks that represents both a financial and technological burden to service providers. In short, it is neither technologically feasible to deliver triple-play services with existing equipment nor economically practical to replace it with a completely new network. Moreover, legacy customers will not tolerate any interruption of existing services, nor will they pay extra for poor service quality.

Motivated by the promise of substantial rewards to those that enable this massive business food chain, the electronics industry is marshaling every possible resource to find solutions at all levels to the triple-play challenge. It is no surprise that the semiconductor industry endeavors to keep pace with system manufacturers.

Xilinx Serial I/O Solutions: Crossing the Chasm

The evolution of serial I/O solutions in Xilinx® FPGAs is the result of our high-speed serial initiative, which we announced in 2002. The aim of the initiative was (and is) to accelerate the industry's move from parallel to high-speed serial I/O by delivering a new generation of connectivity solutions for system designs that meet bandwidth requirements from 3.125 Gbps to 10 Gbps and beyond.

We began by adding up to twenty-four 3.125 Gbps serial transceivers in our Virtex™-II Pro family, accompanied by IP soft cores for numerous serial connectivity standards, reference designs, hardware development platforms, design software, characterization data, and an in-depth design support program.

The Virtex-4 FX family followed suit in 2005 with a similar complement of broad-range transceivers, this time delivering 622 Mbps to 6.5 Gbps performance, as well as an equally robust set of IP and design support software, hardware, and services.

In each case, one of the key objectives in the introduction strategy of these products – with their attending high-speed serial I/O solution packages – was to reach the early adopters and innovators within the FPGA customer base with a viable alternative to custom ASIC and ASSP serial I/O solutions.

Having successfully proven the viability of FPGA-based serial I/O solutions with these previous product families, there remained a single yet extremely important evolutionary step. To cross the chasm into the mainstream FPGA customer base and truly create equivalency between Xilinx serial I/O solutions and custom solutions required the delivery of fully verified, fully integrated, hard IP-based, turnkey serial I/O solutions.

With our newest 65-nm Virtex-5 LXT platform, we believe that we have indeed crossed the chasm. By offering the industry's first FPGA to deliver hard-coded PCI Express Endpoint and tri-mode Ethernet media access controller (MAC) blocks, Virtex-5 LXT devices are addressing the bandwidth, power, and cost challenges facing equipment vendors working to enable the emerging triple-play services market. The Virtex-5 LXT platform is optimized to enable FPGA designers across a wide range of applications to benefit from serial connectivity by delivering a comprehensive, fully compliant protocol solution with the greatest ease of use.

Virtex-5 Serial Connectivity Solutions

Enabling unconstrained product development for the triple-play market.

by Sandeep Vig
Vice President,
Worldwide Marketing
Xilinx, Inc.
sandeep.vij@xilinx.com

Although "triple play" may be one of the hottest buzzwords and growth drivers in the semiconductor industry, it is insightful to understand the evolution of the technology that was required to realize triple play, the forces behind its explosive growth, challenges that will occur along the way, and the critical role of Xilinx® Virtex™-5 products in the development and deployment of triple-play products and services.

Central to the Virtex-5 platform's value is the recent emergence of two serial I/O standards: Gigabit Ethernet (GbE) and PCI Express (PCIe). In the last three years, these two interfaces have become the defacto connectivity standards for network and computing applications; according to Electronic Trend Publications, GbE and PCIe will account for 80% of all port shipments in 2009.

Disruptive Technology

IP is clearly the preferred protocol in the network market as telecom vendors and service providers transition to an all-IP-based infrastructure supporting Voice over IP, Video over IP, and Data over IP (also known as triple play). Designing carrier-grade to end-user products that support triple play is very challenging, as these products must achieve high levels of performance, manage quality of service (QoS), and be power-efficient and flexible enough to adapt to the seemingly endless evolution of standards and protocols.

In the computing infrastructure market, PCIe has become the predominant host interface for networking, graphics, and backplane connectivity because of its quantum leap in performance, scalability, and pin-count efficiency over the legacy PCI bus. Designing products that span network and compute infrastructures like those in triple-play markets requires system architects and engineers to be well-versed in these new domains, introducing new risks. To this end, Xilinx embarked on a project two years ago to mitigate design risk by introducing a new generation of Platform FPGAs that substantially increase performance, functionality, and device density while reducing cost per gate.

Next-Generation FPGAs

Leveraging our core competence as the premier FPGA vendor and working with our world-class customers and partners, Xilinx developed the Virtex-5 FPGA architecture. With the introduction of the LXT family, Virtex-5 devices now feature integrated multi-GbE and PCIe connectivity technology ideally suited to designs for the triple-play market.

This LXT family is equipped to support high-speed serial connectivity, with features that include:

- Built-in GbE MAC – each Virtex-5 LXT device features four hard-core GbE MACs for multi-port Ethernet connectivity
- Built-in PCIe block – an integrated standards-compliant PCIe Endpoint block supporting one to eight lanes provides as much as 32 Gbps of full-duplex host I/O for extreme performance applications

These features reduce the engineering effort spent on resource utilization, troubleshooting connectivity issues, minimizing power consumption, and optimizing performance, thus giving our customers unconstrained Virtex-5 FPGA resources in designing infrastructure and end-user products for delivering voice, video, and data over IP.

As a programmable platform, the Virtex-5 family positions our customers and partners to enable value-added triple-play technologies such as:

- QoS – customer-specific traffic management solutions enabling tiered services that can change with market conditions
- Digital rights management – enabling hardware-based, adaptive, end-to-end data security for the wide diversity of standards inherent to these markets

Conclusion

In the very dynamic consumer industry where time to market with flexible services is the name of the game, companies are still trying to figure out the right mix of products and services to generate substantial revenue. The Virtex-5 LXT family integrates world-class programmable logic architecture with embedded serial connectivity, providing the performance, density, and connectivity required for delivering voice, video, and data in the emerging triple-play market.

FPGAs for Serial Interconnections

Research by Electronic Trend Publications points to a key role for FPGAs in serial interconnections.

by Steve Berry
President, Electronic Trend Publications
saberry@electronictrendpubs.com
www.electronictrendpubs.com

For most of the last 15 years, networking the world for voice, video, and data has been the key driver of the electronics industry. This worldwide network required that the communications industry connect and converge with the computer processing industry. That convergence has primarily settled on Ethernet for the communications side and PCI for the computer side.

Since its inception, Ethernet has been a serial interface. It has been repeatedly scaled up in bandwidth. Today, 1 Gbps connections are ubiquitous, 10 Gbps connections are becoming more common, and 100 Gbps connections have been proven in the laboratory. Ethernet has vanquished all challengers in the LAN market and is rapidly conquering the MAN and WAN markets.

PCI started out as a parallel interface, and as such ran out of bandwidth when connection requirements exceeded 1 Gbps. Industry groups such as the InfiniBand Trade Association and the RapidIO Trade Association introduced new connections to replace PCI. But PCI is much more than the physical connection between system elements. PCI represents an enormous global investment in software that is not readily replaceable. Only PCI Express has met the challenge of true compatibility with PCI. PCI Express bandwidth will be scaled up repeatedly over the coming years to support the industry's needs.

As a result of the nearly 10-year effort to transition the industry from parallel to serial interconnections, serial interconnections will soon become dominant. Table 1 illustrates the change from parallel to serial. In 2006, serial interconnections will move into the majority. By 2009, serial will represent more than 80 percent of all interconnections.

Although standard semiconductor products will supply the serial interconnection needs of high-volume markets, FPGAs are increasingly important for a wide variety of tasks. There are some key reasons. First, before low-cost standards products are available, FPGAs will provide a mechanism to get to market faster. Second, FPGAs enable system integration with customer algorithms and standards-based serial interfaces. Third, the ability to easily make multi-standard serial connections to FPGAs will dramatically simplify product design.

Thus, the new Xilinx® Virtex™-5 LXT platform – with its built-in PCI Express Endpoint blocks, tri-mode Ethernet MACs, and low-power RocketIO™ transceivers – precisely fits the requirements of today's FPGA market by giving designers a solution that not only saves time, but also reduces power consumption and conserves FPGA logic resources.

RapidIO and Aurora

Although PCI Express and Ethernet will be the overwhelming leaders in the number of serial ports deployed by the industry, a host of other serial interfaces have carved niches for themselves. The Virtex-5 LXT platform also supports nearly all available serial interfaces. Two of these interfaces – RapidIO and Aurora – are emerging as most important to users of FPGAs.

RapidIO is becoming a favorite for high-end, low-volume DSP applications. A number of implementations in this arena use FPGAs (rather than merchant silicon) to implement DSP functions as well as RapidIO interface and switching functions. This should continue to be the case in the future.

Similarly, the Aurora protocol has quietly gained a substantial following in certain high-end embedded markets. Although Xilinx created Aurora, it is an open protocol, free of charge, that designers can implement in any silicon device. Aurora is a scalable, lightweight, link-layer protocol that is used to move data across point-to-point serial links. Aurora enables simple, high-speed connections between fixed points either on a single board or across multiple boards. As many applications in the board-level embedded market use fixed links between various points in the system, there is no need for a complex message-passing protocol.

Conclusion

With its hard-coded PCI Express Endpoint and Ethernet blocks, I anticipate that many will use the Virtex-5 LXT platform to bridge between PCI Express or Ethernet and numerous other interfaces. The Virtex-5 LXT platform is ideally suited for this task.

Serial vs. Parallel Ports	2004	2005	2006	2007	2008	2009
Parallel	75.5%	56.3%	34.8%	25.5%	20.4%	15.9%
Serial	24.5%	43.7%	65.2%	74.5%	79.6%	84.1%

Figure 1 – Serial interfaces are rapidly replacing parallel.

High VELOCITY LEARNING

XpressTrack

Nu Horizons Electronics Corp. is proud to present our newest education and training program - **XpressTrack** - which offers engineers the opportunity to participate in technical seminars conducted around the country by experts focused on the latest technologies from Xilinx. This program provides higher velocity learning to help minimize start-up time to quickly begin your design process utilizing the latest development tools, software and products from both Nu Horizons and Xilinx.

Visit our website and let us know where you reside and what you are interested in learning about and we'll develop a curriculum just for you.

For a complete list of course offerings, or to register for a seminar near you, please visit:

www.nuhorizons.com/xpresstrack

XILINX®

NU HORIZONS ELECTRONICS CORP.

Featured Seminars -

Power PC Embedded System Design

This seminar provides the embedded systems developers with the necessary skills to develop a PPC System on a Programmable Chip system utilizing the Virtex 4 FPGA. Utilizing the Embedded Development Kit (EDK) the embedded systems developers will create a full system based on the Nu Horizons XC4FX12 evaluation board, labs provide hands on experience with the development, verification, debugging, and simulation of an embedded system.

Prerequisites:

- Experience in C programming
- Some HDL modeling experience
- Basic microprocessor experience and understanding of PowerPC™ processor systems
- A basic understanding of FPGA devices and the tools used to program them

Implementing New Features of Virtex-5

This 3-hour seminar will introduce to you the first 65-nm family of Platform FPGAs, the Virtex-5 LX from Xilinx. The Virtex-5 family of FPGAs is the 2nd generation of devices based on ASMBL architecture. Learn how the new features in Virtex-5 can increase logic performance by 30%, reduce area by 45%, and decrease dynamic power by 35% when compared to the 90 nm Virtex-4 family.

Course Outline

- Virtex 4 versus Virtex 5 comparison
- V5's new PLL and Use with DCMs
  - Lab 1 – Introduction to the PLL/Architecture Wizard
- Improved Features in V5
  - Lab 2 – Leveraging Improved Features

DSP Imaging Seminar

Course Outline

- Interpreting images as 2D signals
- Understanding the spectral content of images
- The concept of correlation between target and scene
- Image edge enhancement and its applicability to correlation
- Tradeoffs between correlation calculation methods
- Practical application of correlation to target tracking in video
- Overview of the Xilinx Video Starter Kit (VSK)
- Use of the VSK to perform video target tracking in real time

MicroBlaze Seminar

This 3-hour workshop will introduce you to MicroBlaze™: The Low-Cost and Configurable 32-Bit Soft Processor Solution from Xilinx. It will also introduce Xilinx Embedded Development Kit (EDK). As part of this class you will learn how to build a complete customized MicroBlaze soft processor system including user defined peripherals. You will also be introduced to the "Create and Import Peripheral Wizard" and guide you through process of creating a custom peripheral in the EDK environment and using it in a processor system.

Course Outline

- Overview of MicroBlaze
- Overview of the Embedded Development Kit (EDK)
- Lab 1: Build and Optimize a MicroBlaze Soft Processor System in Minutes
- Lab 2: Custom Hardware Interface Utilizing the MicroBlaze IPIF Interface

Fundamentals of FPGAs

Course Outline

- Basic FPGA Architecture
- Xilinx Tool Flow
  - Lab 1: Xilinx Tool Flow Demo
- Reading Reports
- Architecture Wizard and PACE
  - Lab 2: Architecture Wizard and PACE Demo
- Global Timing Constraints
  - Lab 3: Global Timing Constraints
- Implementation Options
  - Lab 4: Implementation Options
- Synchronous Design Techniques
- Summary

Achieve Higher Performance with Virtex-5 FPGAs

New architectural elements can help you attain higher system-level performance.

by Adrian Cosoroaba
Marketing Manager
Xilinx, Inc.
adrian.cosoroaba@xilinx.com

In FPGA system design, maximizing performance requires a balanced mix of performance-efficient components – logic fabric, on-chip memory, DSP, and I/O bandwidth. In this article, I'll explain how you can benefit from Xilinx® Virtex™-5 FPGA building blocks, particularly the new ExpressFabric™ technology, in your quest for higher system-level performance. I will explore key features of the ExpressFabric architecture with examples that quantify the anticipated performance improvements for logic and arithmetic functions. Benchmarks based on actual customer designs will show that Virtex-5 ExpressFabric technology performs on average 30% better than previous-generation Virtex-4 FPGAs.

With the new logic fabric (in which you can implement functions such as counters, adders, and RAM/ROM storage) and available hard IP blocks, memory, and DSP (optimized to operate at clock rates as fast as 550 MHz), the Virtex-5 FPGA is clearly the platform of choice for high-performance designs.

ExpressFabric Performance

Since the first FPGA was introduced in the mid 1980s, the logic fabric for most FPGAs has been based on the same fundamental four-input look-up table (LUT) architecture. The Virtex-5 family is the first FPGA platform to offer a true six-input LUT (6-LUT) fabric with fully independent (not shared) inputs (Figure 1). Moving to a 6-LUT fabric architecture provides the 65-nm Virtex-5 FPGA family with the most effective trade-off between critical path delay – the determining factor for logic fabric performance – and die size.

With process technology advancements, interconnect timing delay can account for more than 50% of the critical path delay. Xilinx has developed a new interconnect pattern for Virtex-5 FPGAs to enhance performance by reaching more places in fewer hops. The new pattern increases the number of logic connections achievable within two and three hops. Moreover, a more regular routing pattern makes it easier for Xilinx ISE™ software to find the most optimal routes. All of the interconnect features are transparent to FPGA designers, but will translate to higher overall performance and easier design routability. Essentially, the Virtex-5 pattern provides fast, predictable routing based on distance.

The combination of the new 6-LUT structure and special functions like carry chains, dedicated multiplexers, and flip-flops (along with the unique methods by which these elements are connected) creates unsurpassed performance and efficiency for implementing logic and arithmetic functions.

One example that clearly shows the benefits of the ExpressFabric technology is a multiplexer (MUX). Implementing a 4:1 MUX requires two four-input LUTs and a MUXF block in the Virtex-4 architecture. The same 4:1 MUX can now be implemented in a Virtex-5 device with a single LUT. Similarly, an 8:1 MUX requires four LUTs and three MUXF blocks in a Virtex-4 FPGA, while the new Virtex-5 architecture requires only two 6-LUTs. The result is better performance and better logic utilization, as shown in Figure 2.

As in previous Xilinx FPGA families, the Virtex-5 Slice L (logic slice) can implement logic functions, registers, and arithmetic functions using the dedicated carry chain. The slightly more complex Slice M (memory slice) adds the capabilities of implementing distributed RAM and shift registers within the LUT (SRL).

Among the various improvements provided by the ExpressFabric architecture, the new carry chain structure delivers substantially higher performance when used to implement arithmetic operations. Its effect on critical path delay is readily seen for several examples listed in Table 1.

Distributed memory functions such as LUT RAM or ROM also benefit in several ways from the larger LUT structure. The new aspect ratio allows a much denser packing of small memory functions leading to significant performance benefits, as depicted in Table 2.

The performance increases provided by the improved logic fabric with its 6-LUT architecture and interconnect structure are substantial, but this is only the beginning.

Most applications require more on-chip RAM than what LUT-based RAM can provide. With the enhanced Virtex-5 block RAM, you can achieve higher on-chip memory performance.

Block RAM Performance

With the move to 65 nm, the Virtex-5 block RAM inherited a 10% increase in clocking speed to 550 MHz. However, to achieve the desired performance for most applications today, block RAMs need to be more than just faster. They need to be larger.

The Virtex-5 block RAM has doubled in size to 36 Kb. This larger block size (comprising two 18-Kb memories) will support 72-bit data words in simple dual-port mode, thereby doubling block RAM bandwidth. Moreover, the Virtex-5 FPGA provides dedicated connections to enable you to cascade two adjacent 36-Kb block RAMs together in the block RAM column, thereby implementing a 72-Kb memory running at the maximum 550-MHz rate.

The availability of ever-larger FPGAs has accelerated the trend toward integrating more subsystems into a single device, making more common the necessity of interfacing multiple clock domains. Virtex-5 devices accommodate this by providing integrated logic to simplify the implementation of flexible and efficient FIFOs.

Through this combination of enhancements, the Virtex-5 block RAM delivers more on-chip memory, easier to build FIFOs, and higher bandwidth.

DSP Performance

The growing acceptance of FPGAs as a viable solution for high-performance DSP applications is well deserved. Whether as a co-processor or a stand-alone solution for

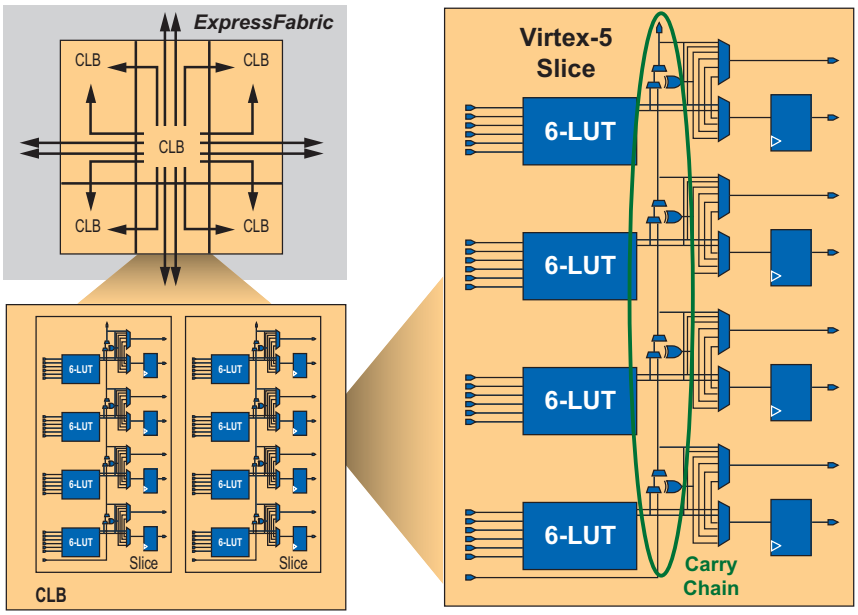

Figure 1 – Virtex-5 configurable logic blocks (CLBs) comprise two slices. Each slice uses four independent 6-LUTs that provide the benefits of fewer logic levels.

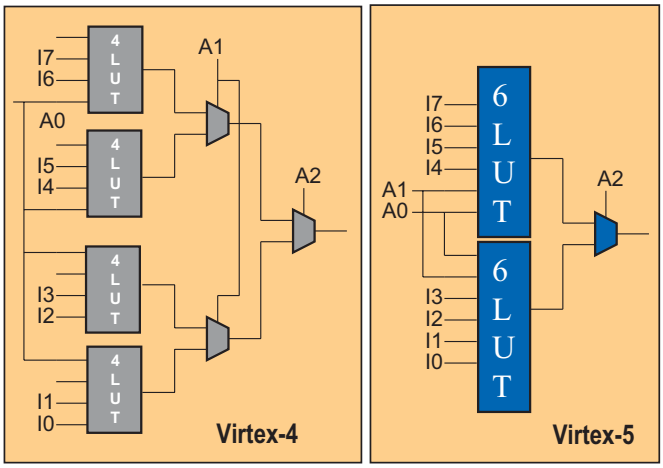

8-to-1 MUX	Virtex-4	Virtex-5	Improvement
Logic Levels	2	1	100%
Path Delay	1.33 ns	1.08 ns	23%

Figure 2 – 8:1 multiplexer implemented with Virtex-5 FPGAs versus Virtex-4 FPGAs

Function	Virtex-4 FPGA Path Delay	Virtex-5 FPGA Path Delay	Improvement
Adder 64-bit	3.5 ns	2.4 ns	46%
Ternary Adder 64-bit	4.3 ns	3.0 ns	40%
Barrel Shifter 32-bit	3.8 ns	2.8 ns	37%
Magnitude Comp. 48-bit	2.4 ns	1.8 ns	34%

Table 1 – Arithmetic functions implemented with Virtex-5 FPGAs versus Virtex-4 FPGAs

Function		Virtex-4	Virtex-5	Improvement
LUT RAM 64 x 1	Logic Levels	2	1	100 %
	Path Delay	1.76 ns	1.26 ns	40 %
LUT ROM 128 x 12	Logic Levels	3	1	200 %
	Path Delay	1.84 ns	1.20 ns	53 %

Table 2 – LUT-based RAM/ROM implementations with Virtex-5 FPGAs versus Virtex-4 FPGAs

more demanding applications, FPGAs continue to provide the best combination of performance, power, and cost.

To keep pace with the seemingly insatiable demand for more DSP performance, Xilinx is leading with Virtex-5 DSP capabilities in terms of both clock rate and precision – the clock rate has increased to 550 MHz and the precision has improved from 18 x 18 bits to 25 x 18 bits.

Xilinx also optimized the Virtex-5 DSP48 slice for adder-chain implementations, a powerful capability that enables the creation of very efficient high-performance filters. Dedicated routing resources on the inputs and outputs of each DSP48 slice permit any number of slices to be chained together within a column. This dedicated routing ensures that every DSP48 slice in the chain will run at full speed without consuming any of the fabric routing or logic resources, as other FPGAs require. Taken together, these improvements reduce by half the number of resources needed to implement common high-precision functions. For example, for a 35 x 25-bit multiply, four DSP48 slices are needed with the Virtex-4 FPGA. With the wider DSP block available in the Virtex-5 FPGA, half as many slices are used to implement this multiply function.

I/O Bandwidth Performance

As performance benchmarks go, the speed with which an FPGA can process data is relevant only in context with the device's I/O bandwidth, which is the speed with which large amounts of data can be moved on and off the device. When using external memory buffers, the interface must be at least two times faster than the data-processing rate because data must be both written out of and read back into the FPGA.

Virtex-5 FPGAs improve on Virtex-4 bandwidth by increasing both the data rate per pin and the number of available I/Os with larger packages. For example, for popular memory interfaces like DDR2 SDRAM, the bandwidth has increased per pin from 534 Mbps to 667 Mbps; the number of data I/Os, when considering SSO requirements, has increased from 432 to 576.

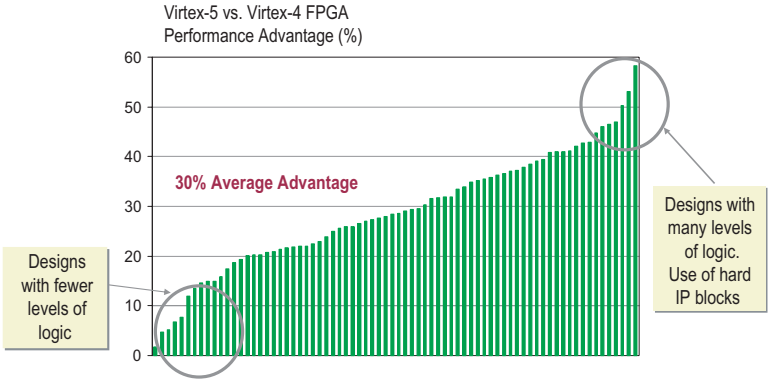

Figure 3 – Comparison based on a suite of 74 customer designs using ISE 8.2i software

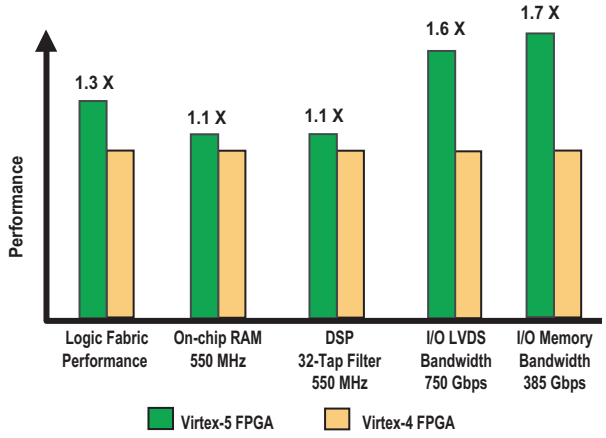

Figure 4 – Virtex-5 FPGA performance improvements

Customer Design Benchmarks

To further evaluate the performance improvements provided by the Virtex-5 FPGA logic fabric, we implemented a set of customer designs using Xilinx ISE software.

These designs were all written in VHDL or Verilog. We implemented some specific design units like memories and FIFOs using direct instantiation of library components or synthesis inference, but many were implemented using EDIF blocks generated by CORE Generator™ software (a part of ISE software).

For these benchmarks, we performed synthesis in a timing-driven fashion with Synplicity's Synplify Pro, using tight, realistic constraints to effectively measure performance. This was done to ensure that all special optimizations and logic replications were employed.

Implementation in ISE software was accomplished with the place and route effort set to high. Clocks were tightened iteratively by 5% increments until the design failed to meet design constraints.

The result was an average performance gain of 30% over designs implemented in Virtex-4 FPGAs, as shown in Figure 3.

Those designs that improved the most have large cones of logic; the critical path implements a large, often complex logic equation. For example, ASIC prototyping designs will typically have very few registers for a large amount of logic in their critical path. These types of designs exhibit a significant improvement with Virtex-5 ExpressFabric technology.

Those designs exhibiting a more moderate improvement either have less levels of logic or provide little opportunity for the use of hard IP blocks or carry-chain structure to improve performance.

Figure 4 summarizes by category the performance improvements of Virtex-5 FPGAs over previous-generation Virtex-4 FPGAs.

Conclusion

With its new ExpressFabric technology and tight coupling to other high-performance hard-IP blocks and I/Os, the Virtex-5 FPGA family represents a significant performance boost compared to previous-generation architectures.

HDL Coding and Design Practices for Improving Virtex-5 Utilization, Performance, and Power

These tips and techniques can lead to better Virtex-5 designs.

by Brian Philofsky
Staff Software Technical Marketing Manager
Xilinx, Inc.
brian.philofsky@xilinx.com

FPGAs have been very flexible in accommodating any HDL coding or design style for digital logic; Xilinx® Virtex™-5 devices are no exception. Although Virtex-5 FPGAs can accommodate many different types of designs written in many different methods, certain recommended constructs and manners can achieve improved optimization in terms of area, performance, and power.

Know Your Target Architecture and Synthesis Tool

Before beginning any project, you should understand the device architecture you are targeting. For Virtex-5 FPGAs, I recommend reading the Virtex-5 Users Guide (*http://direct.xilinx.com/bvdocs/userguides/ug190.pdf*) before starting your first line of code. Once you have a better understanding and vision as to how your code will ultimately result in the base hardware, you can make both large and small design and coding decisions confidently.

For example, if you know of and use Bitslip technology within the ISERDES, you could save time, effort, and resources by capturing input data rather than attempting to describe and build similar circuitry.

In another example, if you know the structure and capability of the DSP48E, you can make better choices as to when and where to place pipeline registers. Dedicated features like the wider multiplier or post adder can also help you achieve better area, performance, and power.

Similarly, knowing the capabilities and current limitations of your synthesis tool can not only help when choosing coding styles to properly infer primitives but can also give you greater insight as to when to instantiate a component or use inference. Review synthesis manuals, application notes, or other relevant materials before starting so that you know the recommended coding styles for the synthesis tool you are using.

You should also update and use the latest versions of synthesis and ISE™ tools before beginning a project. Although initial synthesis support for the Virtex-5 architecture is strong, many improvements in optimization and inference support are still to come with new releases. One easy way to ensure more optimal designs in terms of area, performance, and power is to install the latest version of the software.

Control Signal Polarity

The Virtex-5 architecture can support different control signal polarity (clock enables, resets, or sets). However, to have the most optimal design, I recommend consistent use of active high control signals in your design. The Virtex-5 slice control logic is active high, and when described in this same manner in the code should never require additional LUT resources for a simple signal inversion.

If the signal comes from an external pin and needs an active low polarity, I suggest inverting the signal in the top-level code and using a positive polarity in all processes and sub-modules requiring that signal. This is critical for designs that have several cores, use bottom-up synthesis techniques, have KEEP_HIERARCHY constraints, or employ the use of partitions (Figure 2).

Designs that fall into these categories are more susceptible to the use of additional LUTs per core/netlist/hierarchy/partition for the sole purpose of inverting these control signals, which not only consume extra LUT resources but may also have negative effects on performance and slice

packing. As a general rule, always code sets, resets, and enables with an active high (logic 1 activates) polarity.

Use of Resets

It is common practice to use a global asynchronous reset in the source HDL code to initialize the design; however, in many cases this consumes additional resources. Instead, think synchronous and local. I suggest describing a synchronous set/reset logic to the portions of the design that do need periodical resets. For those portions of the design that do not, you can initialize the signals defined to be registered in the HDL code at the time they are declared (for example, when defining a reg in Verilog or a signal in VHDL). This methodology allows for improved packing density, enhances timing analysis and performance, and can improve area resources.

In terms of FPGA behavior, without a global reset described in the code, a GSR (global set/reset) will occur upon completion of the configuration cycle, initializing all registers to known specified values. This same cycle is also simulated in the gate-level simulation netlist, giving the same known starting point as in the FPGA.

In terms of RTL simulation, having the registers initialized in the code allows for proper RTL or behavioral simulation; this same initialization will be picked up by the synthesis tool and applied to the implemented design. Therefore, for simulation at any stage, a global reset is redundant and unnecessary.

Using a synchronous reset instead of an asynchronous reset also allows for more predictable behavior upon the assertion and release of the reset, because the synchronous signals are automatically analyzed and their behavior is more deterministic when all timing constraints are met. It also allows for the possibility of greater logic optimization and performance because it is not global.

When using synchronous control signals, you can move portions of the logic function to the synchronous set or reset of the flip-flop; this is not possible with asynchronous signals. By only describing a reset where necessary, the synthesis tool can use alternative resource choices like SRLs (shift register LUTs); distributed RAM (LUT-based RAM) memory; or block RAM for the implementation, which would not be otherwise possible nor optimal. The synthesis tool has maximum flexibility to choose the best resource for the described code.

Pipelining

As with previous FPGA generations, properly pipelining your design is necessary to achieve top performance and improved power characteristics. With the introduction of the Virtex-5 architecture, a new logic structure dictates slightly different rules regarding when and how to pipeline.

The Virtex-5 device departs from the traditional four-input LUT in previous FPGA families and has an enhanced six-input LUT (6-LUT), allowing for wider logic functions between pipeline registers while maintaining top performance. You should keep this in mind, as logic functions coded into HDL as optimal code should include six inputs to the logic function between registers to get the most optimal pipelining and LUT resource management.

In cases where it is not practical or possible to have exactly six inputs in a given logic function, the wider input 6-LUT still allows for good performance by

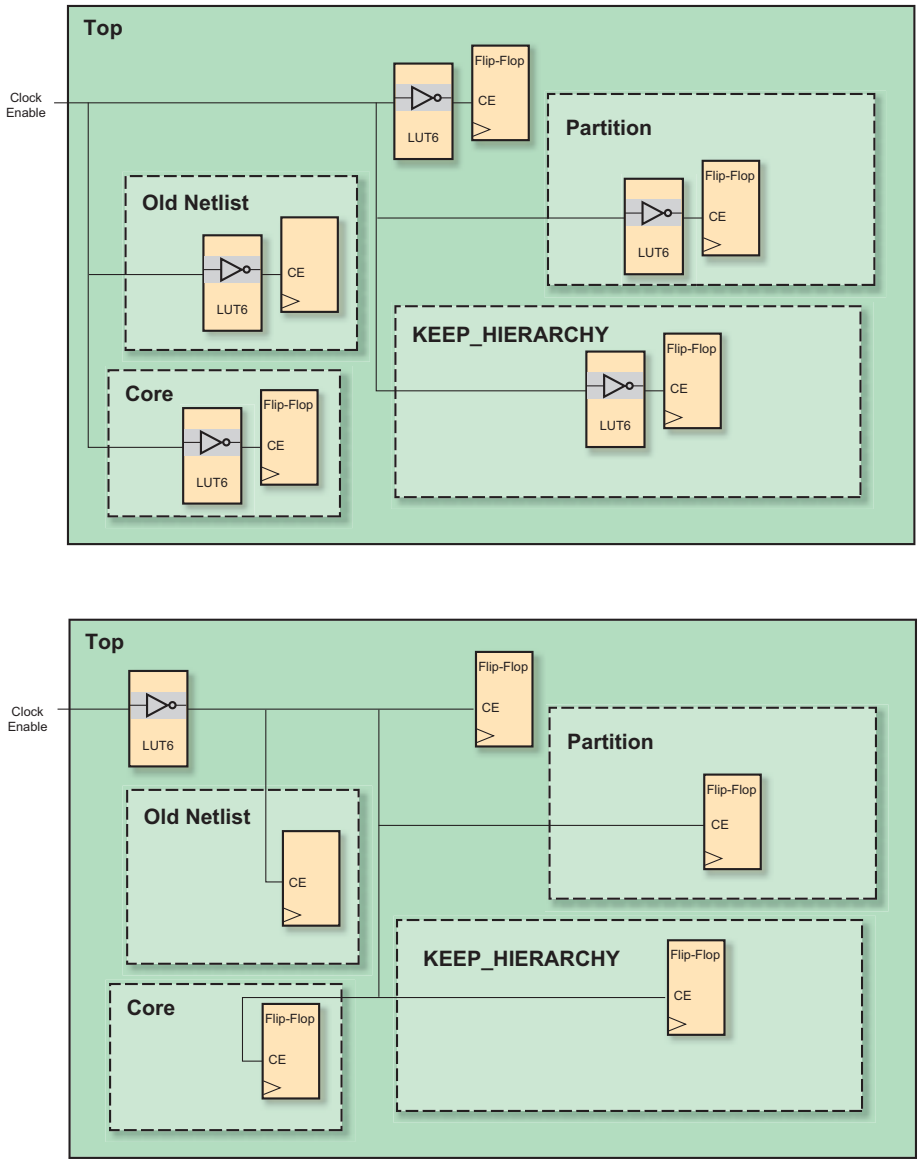

Figure 2 – How clock enable polarity affects LUT utilization in a design

Verilog Coding Example

```
`timescale 1ns / 1ps
//////////////////////////////////////////////////////////////////////////////////
// Company: Xilinx
// Engineer: Brian Philofsky
//
// Create Date:    07:42:58 08/12/2006
// Design Name:    good_design
// Module Name:    good_code
// Project Name:   HDL Coding and Design Practices for Improving Virtex 5
//                 Utilization, Performance and Power
// Target Devices: Virtex 5
// Tool versions: ISE 8.2i
// Description: This is example code employing some good coding practices
//              when targeting a Virtex 5 device.
//
// Revision 0.01 - File Created
//
//////////////////////////////////////////////////////////////////////////////////

module good_code #(
   parameter data_width = 16,
             parity_width = 2)

(  input [data_width-1:0] DATA_IN,
   input DATA_STORE,

   input CLK, RST,
   input READ_DATA,

   output [data_width+parity_width-1:0] DATA_OUT,
   output reg RW_ERROR = 1'b0,
   output DATA_VALID, FULL
);

   // Always initialize registers to known values
   reg [data_width-1:0] data_in_reg = {data_width{1'b0}};
   reg [data_width-1:0] data_in_reg2 = {data_width{1'b0}};
   reg [2:0] data_store_delay = 3'b000;
   reg [2:0] data_valid_delay = 3'b000;
   reg [parity_width-1:0] parity = {parity_width{1'b0}};

   wire read_error, write_error;

   // Use resets only where necessary and make them synchronous
   // Make resets and clock enables active high
   always @(posedge CLK)
      if (RST)
         data_in_reg <= {data_width{1'b0}};
      else if (DATA_STORE)
         data_in_reg <= DATA_IN;

   // Do not use resets where not necessary
   //  In this case an SRL can be used due to the fact no reset is described.
   always @(posedge CLK) begin
         data_store_delay <= {data_store_delay[1:0], DATA_STORE};
         data_in_reg2 <= data_in_reg;
         data_valid_delay <= {data_valid_delay[1:0], READ_DATA};
         RW_ERROR <= read_error | write_error;
         parity[1] <= ^data_in_reg[15:8];
         parity[0] <= ^data_in_reg[7:0];
   end

   //  In general, RAMs should be inferred however in this case, a FIFO is needed
   //   and synthesis can not yet infer the dedicated Virtex 5 FIFO.

   // FIFO18: 16k+2k Parity Synchronous/Asynchronous BlockRAM FIFO
   //         Virtex-5
   // Xilinx HDL Language Template, version 8.2.2i

   FIFO18 #(
      .ALMOST_FULL_OFFSET(12'h080),     // Sets almost full threshold
   .ALMOST_EMPTY_OFFSET(12'h080),    // Sets the almost empty threshold
   .DATA_WIDTH(18),                  // Sets data width to 4, 9 or 18
   .DO_REG(1),                       // Enable output register (0 or 1)
                                     //   Must be 1 if EN_SYN = "FALSE
   .EN_SYN("TRUE"),                  // Specifies FIFO as Asynchronous ("FALSE")
                                     //   or Synchronous ("TRUE")
   .FIRST_WORD_FALL_THROUGH("FALSE") // Sets the FIFO FWFT to "TRUE" or "FALSE
) FIFO18_inst (
   .ALMOSTEMPTY(),                // 1-bit almost empty output flag
   .ALMOSTFULL(),                 // 1-bit almost full output flag
   .DO(DATA_OUT[15:0]),           // 16-bit data output
   .DOP(DATA_OUT[17:16]),         // 2-bit parity data output
   .EMPTY(),                      // 1-bit empty output flag
   .FULL(FULL),                   // 1-bit full output flag
   .RDCOUNT(),                    // 12-bit read count output
   .RDERR(read_error),            // 1-bit read error output
   .WRCOUNT(),                    // 12-bit write count output
   .WRERR(write_error),           // 1-bit write error
   .DI(data_in_reg2),             // 16-bit data input
   .DIP(parity[1:0]),             // 2-bit parity input
   .RDCLK(CLK),                   // 1-bit read clock input
   .RDEN(READ_DATA),              // 1-bit read enable input
   .RST(RST),                     // 1-bit reset input
   .WRCLK(CLK),                   // 1-bit write clock input
   .WREN(data_store_delay[2])     // 1-bit write enable input
);

// End of FIFO18_inst instantiation

endmodule
```

VHDL Coding Example

```
----------------------------------------------------------------------------------
-- Company: Xilinx
-- Engineer: Brian Philofsky
--
-- Create Date:    07:42:58 08/12/2006
-- Design Name:    good_design
-- Module Name:    good_code2
-- Project Name:   HDL Coding Practices for Improving Virtex 5 Utilization,
--                 Performance and Power
-- Target Devices: Virtex 5
-- Tool versions: ISE 8.2i
-- Description: This is an example code employing some good coding practices
--              when targeting a Virtex 5 device.
--
-- Revision 0.01 - File Created
--
----------------------------------------------------------------------------------

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_arith.all;

library UNISIM;
use UNISIM.Vcomponents.all;

entity good_code2 is
generic (
   data_width : integer := 16;
   parity_width : integer := 2
);
port (
   DATA_IN : in std_logic_vector(data_width-1 downto 0);
   DATA_STORE: in std_logic;

   CLK, RST: in std_logic;
   READ_DATA: in std_logic;

   DATA_OUT : out std_logic_vector(data_width+parity_width-1 downto 0);
   RW_ERROR : out std_logic := '0';
   DATA_VALID, FULL : out std_logic
);
end good_code2;

architecture XILINX of good_code2 is

   -- Always initialize registers to known values
   signal data_in_reg: std_logic_vector(data_width-1 downto 0) := (others => '0');
   signal data_in_reg2: std_logic_vector(data_width-1 downto 0) := (others => '0');
   signal data_store_delay: std_logic_vector(2 downto 0) := "000";
   signal data_valid_delay: std_logic_vector(2 downto 0) := "000";
   signal parity: std_logic_vector(parity_width-1 downto 0) := (others => '0');

   signal read_error, write_error: std_logic;

begin

   -- Use resets only where necessary and make them synchronous
   -- Make resets and clock enables active high
   process (CLK)
   begin
      if (CLK'event and CLK='1') then
         if RST='1' then
            data_in_reg <= (others => '0');
         elsif (DATA_STORE='1') then
            data_in_reg <= DATA_IN;
         end if;
      end if;
   end process;

   -- Do not use resets where not necessary
   --  In this case an SRL can be used due to the fact no reset is described.
   process (CLK)
   begin
      if (CLK'event and CLK='1') then
         data_store_delay <= (data_store_delay(1 downto 0) & DATA_STORE);
         data_in_reg2 <= data_in_reg;
         data_valid_delay <= (data_valid_delay(1 downto 0) & READ_DATA);
         RW_ERROR <= read_error OR write_error;
         parity(1) <= (data_in_reg(15) XOR data_in_reg(14) XOR data_in_reg(13) XOR
                       data_in_reg(12) XOR data_in_reg(11) XOR data_in_reg(10) XOR
                       data_in_reg(9) XOR data_in_reg(8));
         parity(0) <= (data_in_reg(7) XOR data_in_reg(6) XOR data_in_reg(5) XOR
                       data_in_reg(4) XOR data_in_reg(3) XOR data_in_reg(2) XOR
                       data_in_reg(1) XOR data_in_reg(0));
      end if;
   end process;

   --  In general, RAMs should be inferred however in this case, a FIFO is needed
   --   and synthesis can not yet infer the dedicated Virtex 5 FIFO.

   -- FIFO18: 16k+2k Parity Synchronous/Asynchronous BlockRAM FIFO BlockRAM Memory
   --         Virtex-5
   -- Xilinx  HDL Language Template version 8.2.2i

   FIFO18_inst : FIFO18
   generic map (
      ALMOST_FULL_OFFSET => X"080",  -- Sets almost full threshold
      ALMOST_EMPTY_OFFSET => X"080", -- Sets the almost empty threshold
      DATA_WIDTH => 18,              -- Sets data width to 4, 9, 18, or 36
      DO_REG => 1,                   -- Enable output register (0 or 1)
                                     -- Must be 1 if the EN_SYN = FALSE
      EN_SYN => TRUE,                -- Specified FIFO as Asynchronous (FALSE) or
                                     -- Synchronous (TRUE)
      FIRST_WORD_FALL_THROUGH => FALSE) -- Sets the FIFO FWFT to TRUE or FALSE
   port map (
      ALMOSTEMPTY => open,   -- 1-bit almost empty output flag
      ALMOSTFULL => open,    -- 1-bit almost full output flag
      DO => DATA_OUT(15 downto 0),   -- 32-bit data output
      DOP => DATA_OUT(17 downto 16), -- 2-bit parity data output
      EMPTY => open,                 -- 1-bit empty output flag
      FULL => FULL,                  -- 1-bit full output flag
      RDCOUNT => open,               -- 12-bit read count output
      RDERR => read_error,           -- 1-bit read error output
      WRCOUNT => open,               -- 12-bit write count output
      WRERR => write_error,          -- 1-bit write error
      DI => data_in_reg2,            -- 16-bit data input
      DIP => parity,                 -- 2-bit parity input
      RDCLK => CLK,                  -- 1-bit read clock input
      RDEN => READ_DATA,             -- 1-bit read enable input
      RST => RST,                    -- 1-bit reset input
      WRCLK => CLK,                  -- 1-bit write clock input
      WREN => data_store_delay(2)    -- 1-bit write enable input
   );

   -- End of FIFO18_inst instantiation

end XILINX;
```

Figure 3 – Sound FPGA coding styles

reducing the number of logic levels, thus requiring fewer pipeline stages to achieve the same as or better performance than previous FPGA architectures.

A good goal is to aim for less than 10 inputs to a given logic function between I/Os, registers, or synchronous blocks (like block RAM or DSP48Es), which generally would represent two logic levels. When you need a significantly higher number of inputs for the design path to meet latency or other requirements, you can attempt to reduce the fan-in to that logic function (when possible) if high performance or low power are your design objectives.

Coding Memories

Among other innovations within the Virtex-5 architecture, Xilinx has enhanced both block RAM and distributed RAM memories with greater capacity and capability. You must make different decisions early in the design process and while coding to get the most from these valuable resources.

General guidelines call for inferring RAMs when possible for easier code changes, faster simulation, and more portable code. However, even when behaviorally describing the RAM, you should keep some important things in mind. The first and most obvious thought is RAM capacity. In terms of block RAMs, the base memory block increased in Virtex-5 devices to 36 Kb of memory storage space. You can configure this block to the wider but shallower 512 x 72 configuration, the deeper single-bit width 32 Kb x 1, or several configurations in between. It is also possible to cascade two 36-Kb RAMs to form a 64-Kb x 1 configuration or break up the 36-Kb RAMs into two separate 18-Kb RAMs capable of 512 x 36 to 16-Kb x 1 configurations.

Distributed RAM have benefited from the larger LUT structure and can now efficiently accommodate 64-bit depths without any area or performance penalties. This is the most optimal size for this type of RAM in the Virtex-5 device, although other sizes can be accommodated. The base RAM sizes are important to remember during memory selection and coding to most efficiently use the limited RAM resources in the device and achieve the best performance.

Both block RAM and distributed RAM memories also have additional capabilities that require different coding and design considerations. For performance, perhaps the most important is the proper use of output registers. For block RAMs, this means enabling the output registers to the block RAM whenever possible. By enabling the output registers, a reduced clock-to-out is realized from the RAM, thus improving timing for the data leaving the RAM. However, an extra clock cycle of latency is added during reads, for which you must account.

Similarly, when using distributed RAM, the output of the RAM can be asynchronous; however, coding it synchronously will allow the use of the register within the slice, providing better timing characteristics and reducing the chance of the RAM being part of the timing bottleneck.

There are more advanced features of the block RAMs, such as FIFO and ECC (error correction circuitry) capabilities. The distributed RAM also has new capabilities such as a quad-port configuration. In some cases, these features cannot be realized by inference within synthesis and instantiation is necessary. If you need such functionality, I suggest instantiating the RAMs either by generating cores within Xilinx CORE Generator™ software or by instantiating the base primitive. Taking advantage of these advanced features can save RAM and logic resources as well as improve area, performance, and power.

Some General Guidelines

A few other general recommendations do not fall into any specific categories but can result in better coding and design choices. First, you should make wise choices in terms of your design hierarchy right from the start. Your choice of hierarchy can have effects on the synthesis and implementation tools' ability to optimize the logic paths.

In general, do not allow timing paths to cross multiple boundaries of hierarchy. This not only limits the tool's ability to optimize logic but may also limit your options for design implementation and design debugging. For instance, you may not be able to use partitions or KEEP_HIERARCHY on certain hierarchies with this practice.

For designs in which some or most of the code was created for an architecture other than Virtex-5 FPGAs, I suggest that you review the code to ensure that it is well suited for implementation into the new architecture. A few minutes of time spent here can save several hours later if you identify and correct suboptimal code.

If your design contains cores or precompiled netlists (EDIF or NGC files) from a previous architecture, you should regenerate those targeting Virtex-5 devices. Unless regenerated, netlists optimized for a previous architecture are more likely than not far less optimal when targeting Virtex-5 architectures.

One last suggestion is to use the HDL language templates within the ISE tools. They not only help with accelerating the generation of VHDL or Verilog code, but also provide assistance in creating more optimal code for FPGAs. They also cut down on the possibility of creating syntax or other simple but common mistakes that can hold up the testing and verifying of HDL code.

Figure 3 shows both Verilog and VHDL code following the guidelines discussed here.

Conclusion

Coding styles are very individual; however, following these suggestions makes it more likely that you will achieve a more optimal result. These guidelines do not represent absolutely everything you need to know to achieve the best Virtex-5 design possible, but I have provided some common strategies that can help in achieving more optimal designs.

Almost any set of valid HDL code likely will result in a functioning design, but following a few simple guidelines can help in terms of improved density, performance, and power, and many times may reduce the amount of time it takes to ultimately complete a design.

For more information, see the Synthesis and Simulation Design Guide at *http://toolbox.xilinx.com/docsan/xilinx82/books/docs/sim/sim.pdf* or White Paper 231, "HDL Coding Practices to Accelerate Design Performance," at *http://direct.xilinx.com/bvdocs/whitepapers/wp231.pdf*.

Getting the Best Results from Virtex-5 FPGAs

Synplicity applies new algorithms and heuristics for optimal Virtex-5 support.

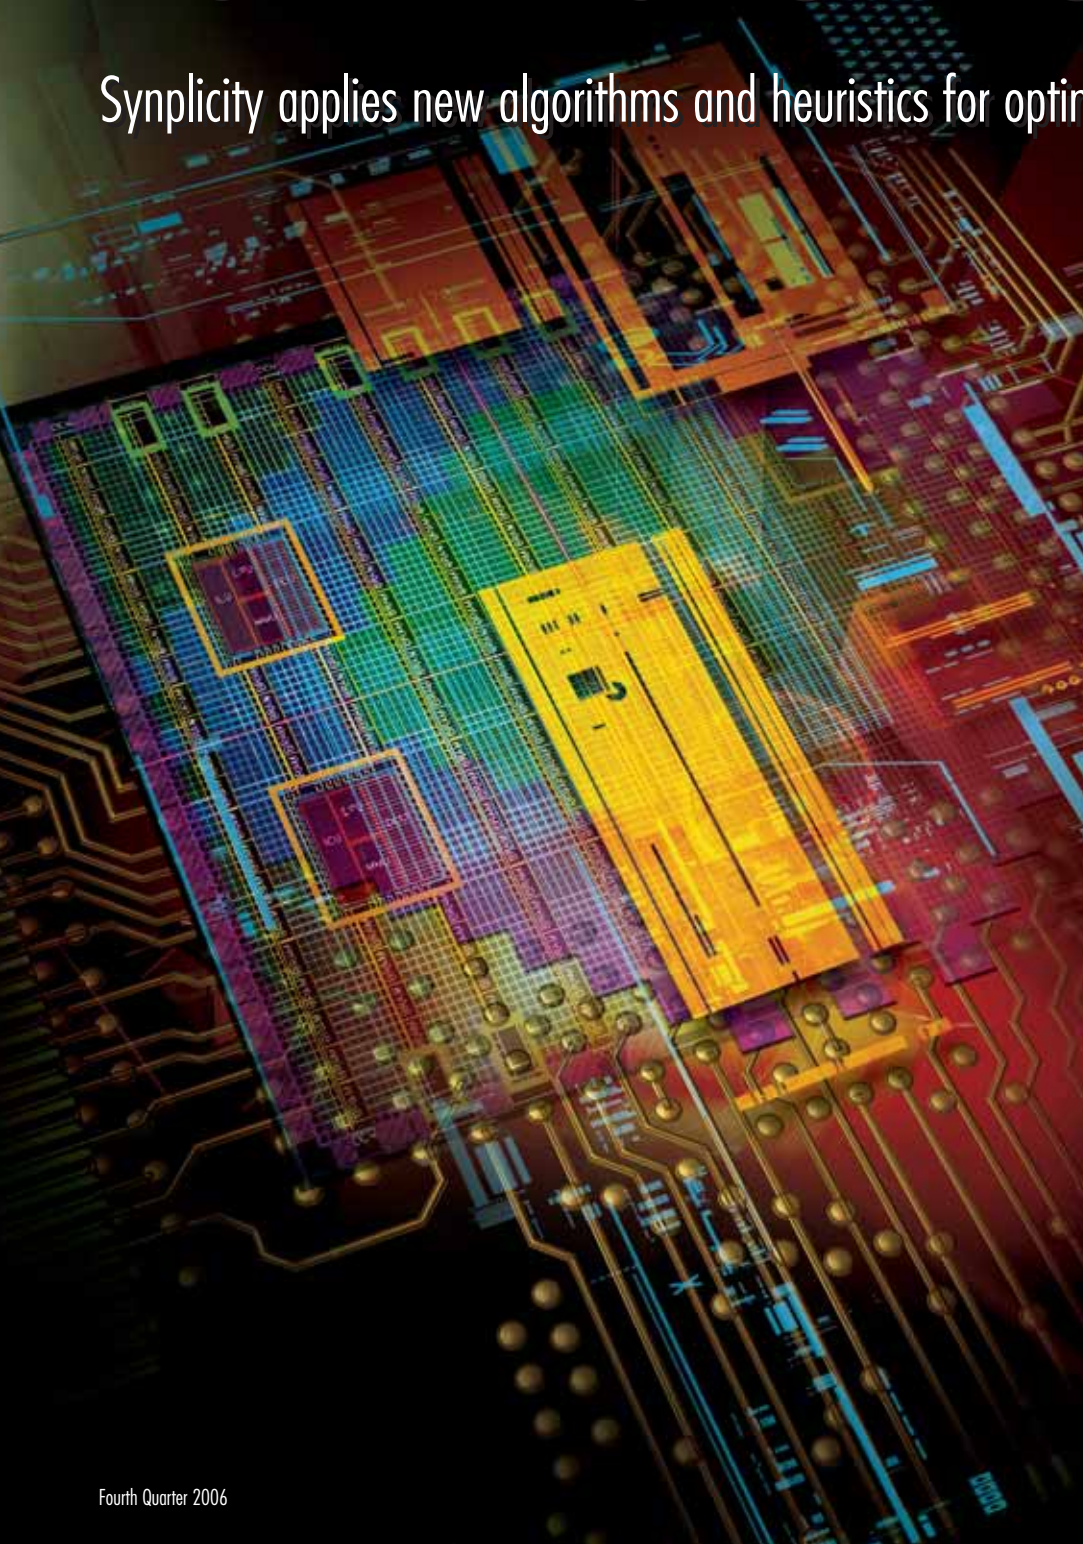

by John Gallagher
Sr. Director Outbound Marketing
Synplicity, Inc.
johng@synplicity.com

The revolutionary capabilities of Xilinx® Virtex™-5 devices can be fully realized only if the EDA technologies to unlock those capabilities are available when the device is. To achieve this, the FPGA architecture and corresponding EDA design tools must be developed simultaneously. The scale of EDA development over previous generations is similar to the difference between Virtex-5 devices and their previous generations.

As the first FPGAs created at the 65-nm technology node, the Virtex-5 family of domain-optimized devices provides as much as 65% more logic cells and 25% more I/O pins than the Virtex-4 family. At the same time, devices in the Virtex-5 family provide 30% higher performance, 35% lower dynamic power dissipation, and consume 45% less silicon real estate when compared to their Virtex-4 counterparts.

To reduce the levels of logic needed to map functions like wide data paths and DSP, the ExpressFabric technology in the Virtex-5 family features LUTs with six independent inputs.

Clearly, the Virtex-5 architecture delivers revolutionary capabilities. To develop a synthesis flow that leverages these new features, Xilinx gave Synplicity early access to the Virtex-5 architecture. By the time the first Virtex-5 devices were introduced, we had been working side-by-side with Xilinx engineers for more than a year to enhance our Synplify Pro synthesis engine. This involved making substantial algorithmic changes to Synplify Pro software to maximize the performance and logic density of Virtex-5-based designs. Because of our partnership, we have the tools and methodologies in place to enable you to rapidly deploy these devices.

In this article, I will describe some of the ways in which we enhanced the Synplify Pro FPGA synthesis engine to take full advantage of the capabilities offered by the Virtex-5 family.

New Algorithms to Synthesize 6-LUTs

Increases in the complexity of the FPGA fabric demanded corresponding increases in the sophistication of the synthesis algorithms. If you applied the same algorithms for a fabric based on 4-input look-up tables (4-LUTs) to a fabric with 6-LUTs, for example, synthesis run times would dramatically increase. To take full advantage of the specialized architectural features in the Virtex-5 family, Synplicity had to either fine-tune or in some cases completely re-craft the underlying synthesis algorithms.

To reduce the levels of logic needed to map functions like wide data paths and DSP, the ExpressFabric™ technology in the Virtex-5 family features LUTs with six independent inputs (Figure 1). This significantly reduces the number of logic levels and LUT area required to implement wide functions. Within the synthesis process, you can use each of these logical elements as a true 6-LUT or as two 5-LUTs that share their five inputs.

However, the combinatorial explosion associated with mapping to 6-LUTs can cause memory utilization and run-time problems if not handled correctly. If you applied the algorithms for a fabric based on 4-LUTs to a fabric with 6-LUTs without significant modification, synthesis run times would be orders of magnitude longer (if they completed at all). In addition, when attempting to find an optimal mapping, traditional algorithms run the risk of becoming trapped in local minima.

Unlike timing-driven engines, most conventional synthesis engines simply attempt to reduce the number of logic levels. This is problematic for LUT architectures in which different input-to-output paths have asymmetric delays. Consider the five shared input pins versus the sixth unique pin in Figure 1; these pins will have very different delays.

The fact that you can use Virtex-5 LUTs in a 2 x 5-input configuration further increases the complexity of the mapping operations. Synthesis tool vendors must conduct a lot of research and development to use structures that share inputs but represent different functions.

Synplicity equipped Synplify Pro software (which features a unique direct-mapping capability) with a variety of sophisticated heuristic algorithms that are tailored to minimize the number of cuts, handle huge capacities, and address these complex mapping and timing scenarios.

Timing Estimation with Diagonally Symmetric Interconnect

Another ExpressFabric feature is a radically new form of diagonally symmetric interconnect that reaches more locations with fewer hops (Figure 2). This diagonally symmetric interconnect pattern was designed to improve speed and predictability. The combination of ExpressFabric 6-LUTs and a diagonally symmetric interconnect pattern results in an average increase of logic performance of 30% over Virtex-4 devices, which equals two speed grades.

The diagonally symmetric interconnect pattern also delivers significantly higher complexity in timing analysis; previous physical synthesis algorithms were based on architectures with "Manhattan routing" or 90-degree routes. To handle the

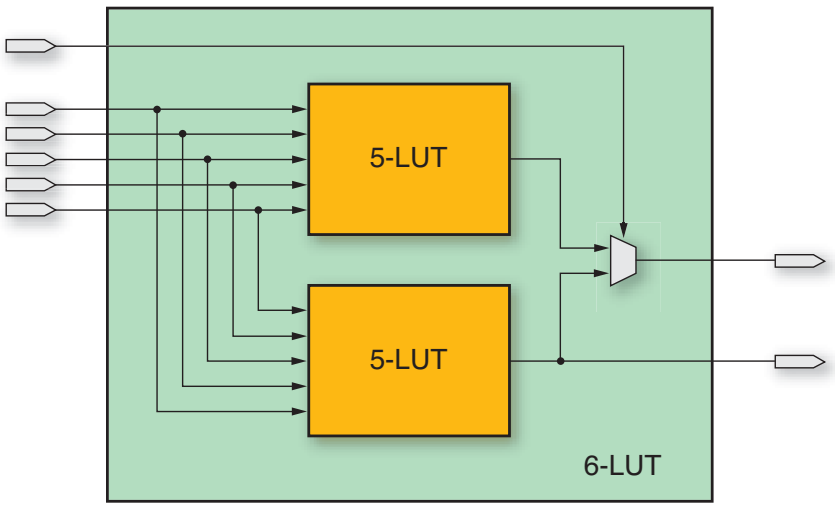

Figure 1 – Virtex-5 six-input LUTs

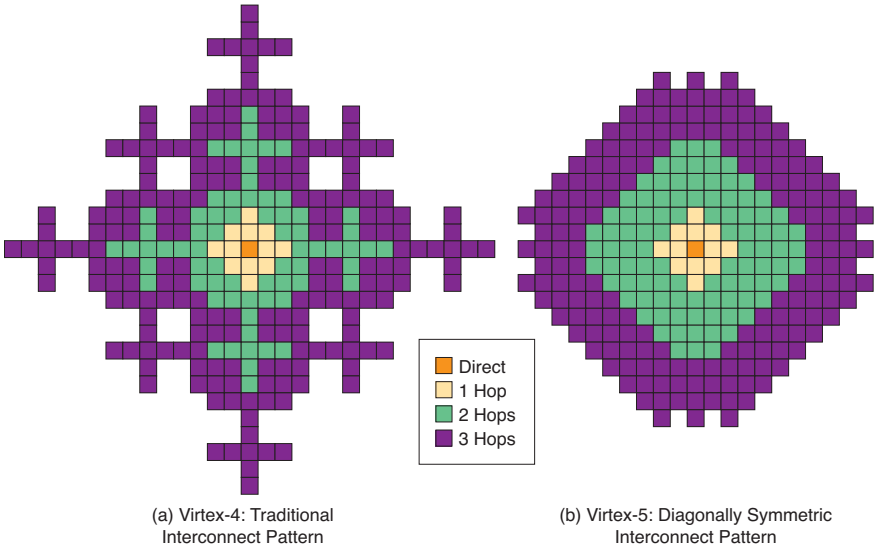

Figure 2 – Diagonally symmetric interconnect routing

combination of both 90-degree and diagonal routes, Synplicity custom-designed new algorithms and delay models. As opposed to simple wire-load models, we engineered the Synplify Pro tool to employ sophisticated netlist-based routing estimation (coupled with known routing values where applicable). In the case of fast carry chains, for example, routing delays are well known and can be directly "plugged in." Similarly, in the case where a cell, driver, load, and specific route are known, an accurate routing delay associated with this path can be plugged in to the routing and timing algorithms.

Synthesizing Fast High-Capacity RAM Blocks

The new block RAM structures (with pipeline) in the Virtex-5 family have increased to 32 Kb in size – twice the size of those found in Virtex-4 components. In addition to offering a simple dual-port mode that can double the RAM's bandwidth, these blocks also contain additional hard IP in the form of FIFO logic and new 64-bit error checking and correction (ECC) logic (Figure 3). Implementing this logic as hard IP frees up other resources and minimizes dynamic power consumption.

As with all hard IP blocks in Virtex-5 devices, these block RAMs have been

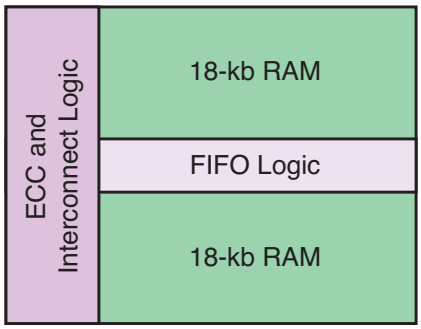

Figure 3 – The Virtex-5 family features as much as 10 Mb of 550-MHz block RAM.

tuned for 550-MHz operation to provide higher on-chip memory bandwidth. The 18-Kb block RAMs are constructed from two physical 9-Kb memories, which are automatically controlled to save power by enabling only one of the 9-Kb sub-blocks for any given read or write operation in most configurations.

For our part, the Synplify Pro synthesis software can perform automatic memory inferencing, including single-port and dual-port implementations, single and multiple clocking schemes, and automatic retiming. Regarding the latter point, Virtex-5 block RAMs are inherently synchronous; however, the design's RTL could describe the memory and registers in such a way as to be technically asynchronous. In such a case, the Synplify Pro tool has the ability to "push" the registers into the RAM.

Similarly, the software will recognize potential conflicts and automatically generate appropriate conflict-resolution logic. In cases such as dual-port RAMs, for example, in which the result of writing two words to the same address may be undefined, the Synplify Pro tool automatically inserts the appropriate logic to resolve the issue such that the memory works in exactly the same way as the RTL will simulate.

Furthermore, Synplify Pro software can automatically analyze the memory described in the design and recognize potential issues in mapping it to preferred memory resources. If you require block RAM, for example, but have used more than what is available on the physical device, the software will automatically move some of the memory into select RAM.

Optimal Use of Faster, Wider DSP Blocks

The Virtex-5 hard DSP slice – called the DSP48E – features a 25 x 18-bit multiplier (versus the 18 x 18-bit multiplier employed in Virtex-4 FPGAs). This increase can lead to fewer cascaded stages, thereby resulting in higher overall performance and utilization (Figure 4).

Tuned for 550-MHz operation, you can configure these high-precision, high-performance, highly flexible slices for DSP, arithmetic, and logic functions and cascade them for adder-chain architectures. The DSP48E slice has 40% lower power consumption compared to equivalent functions in Virtex-4 FPGAs (1.38 mW/100 MHz at a 38% toggle rate).

The sophistication of these DSP slices means that it is unlikely that a data path defined in RTL will exactly match the optimal DSP implementation structure. For example, rather than implementing a function such as "(a + b) + (c + d)" by adding "a" and "b," adding "c" and "d," and then adding the results generated by these operations, it may be more efficient to cascade the DSP slices along the lines of "(((a + b) + c) +d)."

We equipped Synplify Pro software with extremely sophisticated mapping algorithms that perform a lot of data path massaging, creating data path structures that

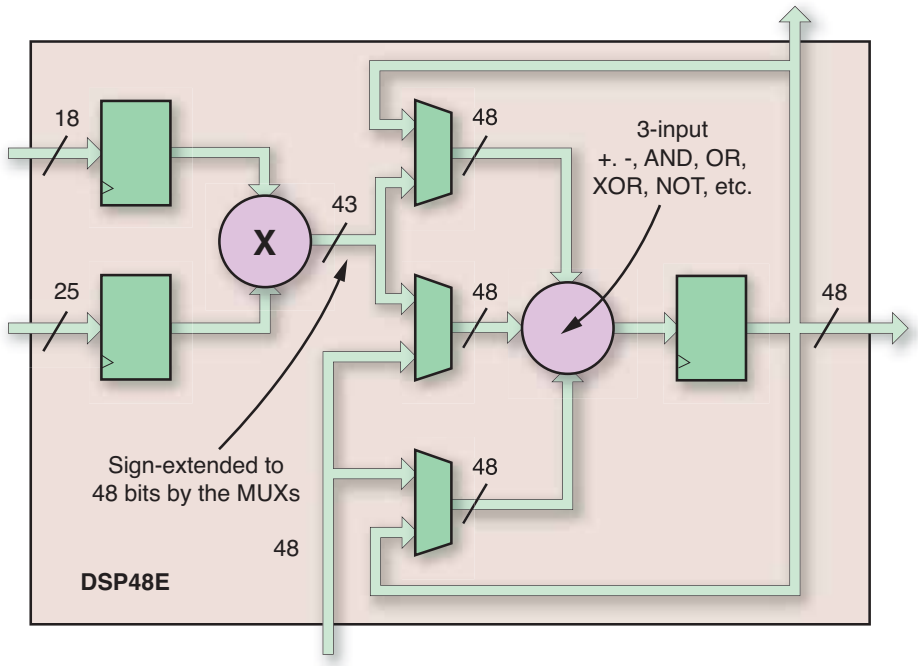

Figure 4 – Virtex-5 DSP slices with 25 x 18 multipliers

map efficiently onto – and take full advantage of – the DSP48E slices featured in Virtex-5 SXT devices.

Note that the DSP48E slice contains a large number of registers (not shown in Figure 4 for simplicity). The Synplify Pro tool can use advanced pipelining and retiming techniques to take full advantage of these embedded register elements. Another consideration is that these internal register elements support only synchronous resets. Thus, if you employ an asynchronous reset in your code, the Synplify Pro tool automatically inserts the appropriate glue logic to restore the required functionality.

Synthesis Considerations with High-Speed I/Os

The number and type of available I/O in a specific device plays a critical role in design implementation, particularly when the application of the FPGA is chip- or system-level verification. With Synplicity tools like Certify ASIC RTL prototyping and the Synplify Pro synthesis solution, we optimized the powerful I/O capabilities in Virtex-5 devices to take into account both signal integration as well as automatic I/O assignments.

Virtex-5 FPGAs offer as many as 1,200 general-purpose input/output (GPIO) pins that you can use to implement industry-standard and custom protocols. The SelectIO™ technology behind these pins provides 1.25 Gbps differential I/O and 800 Mbps single-ended I/O.

Second-generation ChipSync™ source-synchronous technology allows programmable delays to be individually applied to each input and each output. Furthermore, the unique power and ground pin pattern of the Virtex-5 second-generation sparse chevron packaging technology simplifies circuit board layout while minimizing signal integrity and crosstalk effects. Combined, these I/O technologies ensure reliable operation for high-bandwidth interfaces such as DDR2 and QDR II.

We engineered Synplify Pro synthesis software to automatically handle differential signals (and bi-directional I/O signals). For example, if you apply an attribute that identifies the port as being a low voltage differential signal (LVDS) output port, the Synplify Pro tool will automatically insert the appropriate LVDS primitive with one input and two outputs.

Next Steps

As devices move deeper into the submicron domain, the symbiosis between FPGA vendors and EDA vendors increases. Working with engineers at Xilinx for almost a year to enhance our Synplify Pro synthesis engine has yielded tremendous benefits, including having the best synthesis technology available immediately when the device was brought to market, tested against real designs.

Future FPGA platforms will provide even greater densities and capabilities, further expanding the reach of advanced FPGA architectures across a wide range of application domains. These new platforms will require ever-more-sophisticated design flows and synthesis solutions. For this reason, Synplicity and Xilinx formed an Ultra-High-Capacity Timing Closure Task Force. As part of this endeavor, engineering teams from both of our companies will collaborate to define and implement new design flows that maximize design productivity and quality of results for ultra-high-density designs implemented using next-generation FPGAs.

The task force will initially focus on dramatically improving overall quality of results and run times, and ensure the stability of results when small changes are made to designs. Ultimately, the goal of the task force is for designers to realize the benefits of near push-button results for ultra-high-density designs, completing multiple design iterations per day.

Conclusion

The Virtex-5 family from Xilinx has an implementation flow engineered by Xilinx and Synplicity to achieve the best possible results. These new FPGAs boast a wide range of new architectural features, such as 6-LUTs and a diagonally symmetric interconnect fabric.

The matching software solution from Synplicity features new forms of timing estimation for diagonal interconnect, new synthesis algorithms to deal with combinatorial explosion, specialized RAM inferencing and I/O handling, and numerous improvements that enable stable results with minor design changes. Taken together, Virtex-5 devices and Synplify Pro software bring system designers new capabilities that you can design with today.

Triple-Oxide

Ultimate Power Optimization . . .

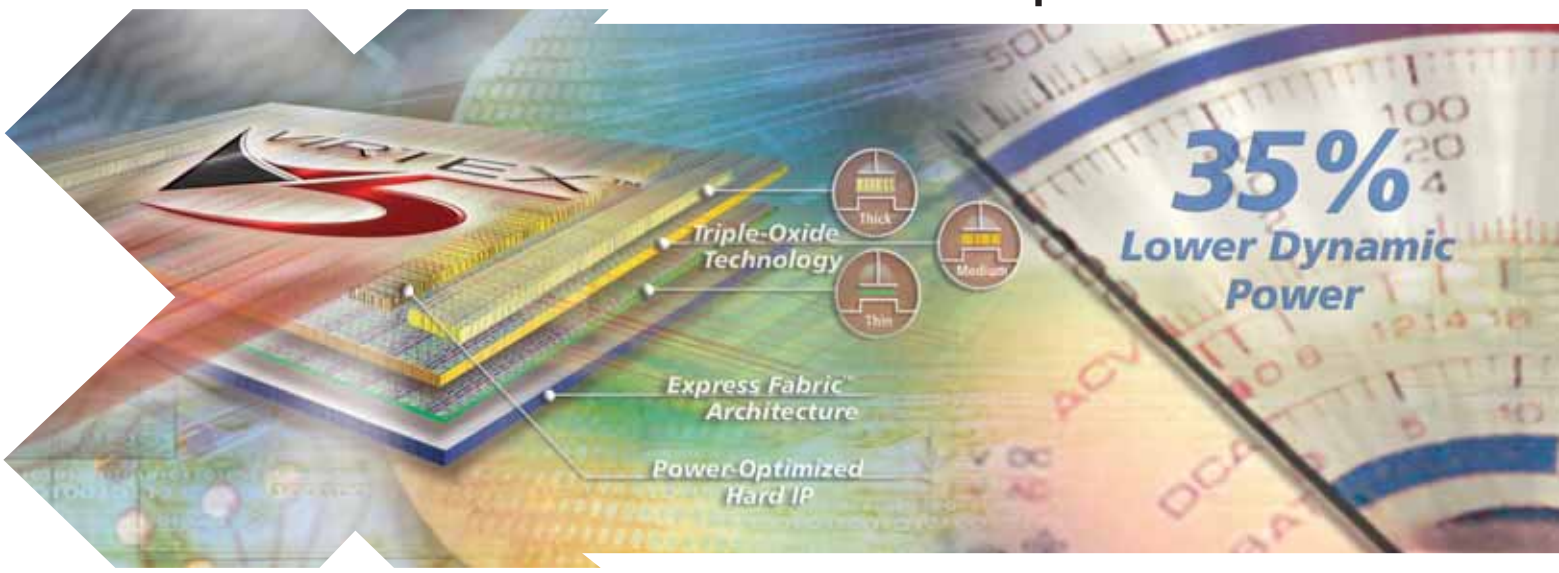

VIRTEX™-5

Reduce power without compromising performance.

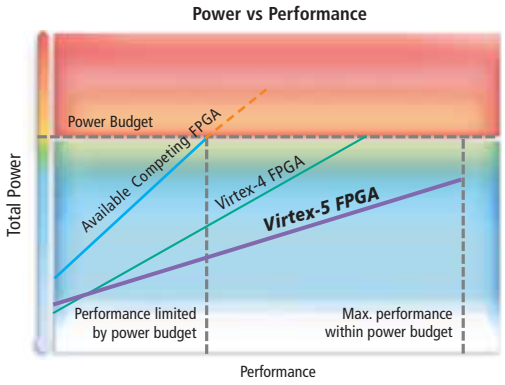

Note: Under worst-case operating conditions (85°C)

Virtex™-5 FPGAs give you unbeatable power savings with the highest performance. The unique combination of 65nm process, second-generation Triple-Oxide technology, ExpressFabric™ architecture, and power-optimized hard IP extends the 1 to 5 Watt power advantage delivered by previous-generation Virtex FPGAs. Achieve higher reliability and a smaller form factor. Save cost on power supplies, heat sinks, and fans. All this, plus the industry's highest performance. No other FPGA vendor comes close.

Meet performance targets within your power budget

Our Triple-Oxide technology optimizes multiple oxide thicknesses to control leakage and keep static power on par with 90nm Virtex FPGAs while maximizing performance. New 65nm ExpressFabric architecture with real 6-input LUTs and diagonally symmetric routing reduces dynamic power by at least 35%. With power-optimized hard IP and automated, block-based power control, you can save even more. With Virtex-5 FPGAs, you can meet your most aggressive performance and power targets. No compromises.

Visit ***www.xilinx.com/virtex5/power,*** view the Virtex-5 power webcast, download the XPower Estimator tool, and read the power analysis white paper.

Virtex-5 LX is the first of four platforms optimized for Logic, DSP, processing, and serial connectivity.

XILINX®
The Programmable Logic Company℠

www.xilinx.com/virtex5/power

The Ultimate System Integration Platform

Maximizing Design Performance for Virtex-5 FPGAs

ISE software gives you the tools to achieve the timing goals of a Virtex-5 design.

by Michelle Fernandez
Software Technical Marketing Engineer
Xilinx, Inc.
michelle.fernandez@xilinx.com

As FPGAs push the performance envelope, maximizing design performance requires knowledge of the device architecture and design software. The 65-nm Xilinx® Virtex™-5 FPGA family delivers the industry's highest performance, with new ExpressFabric™ technology, diagonally symmetric routing, enhanced on-chip memory, DSP slices, and high-speed I/O. To maximize system performance, you should use proper design techniques such as defining timing constraints and selecting options in synthesis and implementation that work best for your design. In this article, I'll describe how to achieve faster timing in the fewest design iterations.

Understanding the Architecture

When evaluating a new FPGA architecture like the Virtex-5 family, it is important to study the user guide and data sheet to understand the hardware features.

The Virtex-5 FPGA family is based on a new ExpressFabric architecture that delivers higher speeds, a new 6-input LUT structure that reduces logic levels, and diagonally symmetric routing that minimizes delays. Each CLB contains two slices that have four 6-input LUTs and four registers configurable in many ways. For maximum slice packing, it is imperative that you understand the slice interconnectivity and any shared resources.

Virtex-5 FPGAs contain hard IP such as embedded memory (block RAM) and math functions (DSP48E slices) tuned to 550 MHz. Here are some design considerations if any of these hard-IP blocks show up as part of your critical paths:

- Check to see if your design is making the most of the block's features and that the synthesis tool is inferring the features as expected from your RTL.
- When using the embedded block RAM memory or the DSP48E slices, it is important to use their dedicated pipeline registers when possible to reduce setup and clock-to-out timing.
- Another consideration is the mix of block RAMs or DSP48E slices in the design, and the trade-off between using dedicated blocks or implementing the same function in slices to allow for placement flexibility.

The choice of clocking resources can also affect a design's performance. Virtex-5

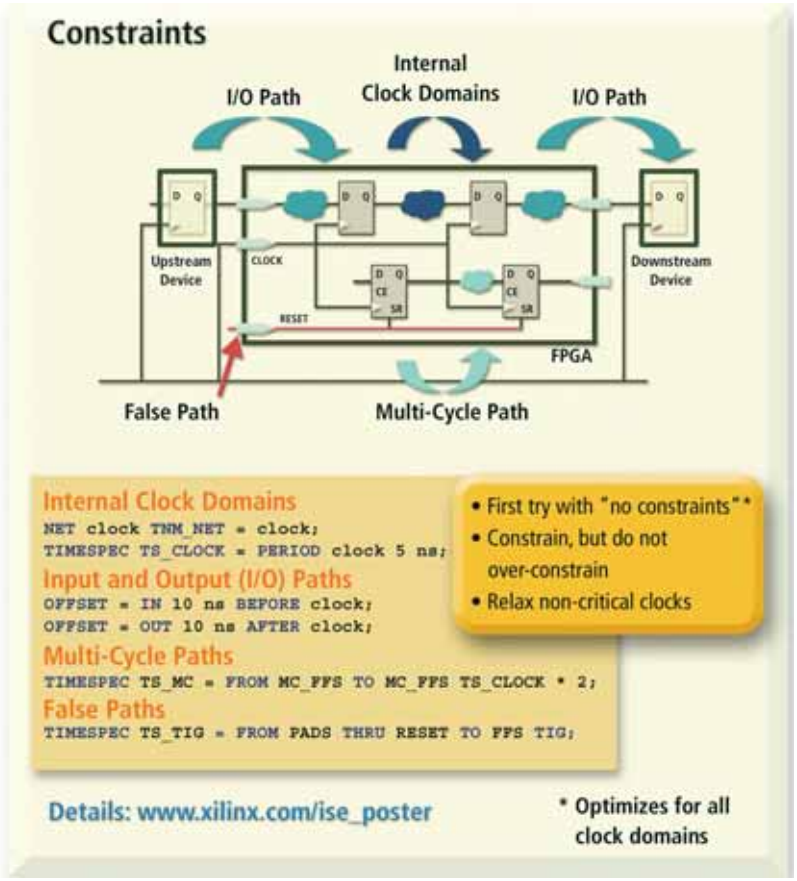

Figure 1 – Proper timing constraints

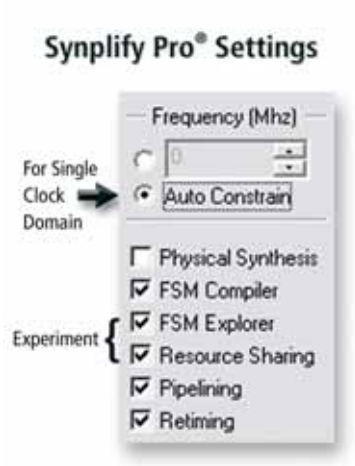

Figure 2 – Recommended Synplify Pro settings

ISE Synthesis "Process Properties" Dialog

	Property Name	Value
Synthesis Options	Optimization Effort	High
	Synthesis Constraint File Create a XCF File (UCF Syntax) `NET clock PERIOD = 5 ns`	☑
	Read Cores	☑
HDL Options	Resource Sharing	Experiment
Xilinx-Specific Options	Register Balancing	YES
	Pack I/O Registers into IOBs	Experiment

Figure 3 – Recommended ISE synthesis (XST) settings

FPGAs have I/O, regional, and global clocking resources. These devices are divided into multiple clock regions, which at most can contain 4 regional clocks and 10 global clocks. During design planning, it is important to analyze how many clock regions you plan to use as well as specific clocks within a clock region. Placing your I/Os so that their interface logic does not require all of the clock resources in a clock region gives ISE™ software greater placement flexibility.

Define Timing Requirements

Synthesis and ISE implementation tools are driven by the performance goals that you specify with timing constraints for internal clock domains, I/O paths, multi-cycle paths, and false paths (see Figure 1). Defining realistic timing constraints will prevent excessive replication and longer run times.

In your synthesis report, check for any replicated registers and confirm that the timing constraints that apply to the original register also cover the replicated registers for implementation. When writing timing constraints, group the maximum number of paths with the same timing requirement before generating a specific constraint to minimize implementation run times and memory usage.

Driving Synthesis

Here are some design considerations for getting optimal results from synthesis tools:

- Use proper coding techniques to ensure that the inference of your RTL by synthesis takes advantage of the architectural features.
- Add any lower level netlists to your synthesis project to better optimize HDL that interfaces to those netlists.
- If critical paths in your implementation are not seen as critical in synthesis, try Synplify Pro's "-route" constraint to force synthesis to focus on that path.
- Explore the synthesis tool settings. (See Figure 2 for Synplicity and Figure 3 for Xilinx Synthesis Technology [XST] suggested tool settings.) There are also a variety of attributes that can affect synthesis optimizations. These attributes are an easy way to affect synthesis without having to re-code (see Table 1).

Certain tool settings, such as retiming in Synplify Pro and register balancing in XST, can impact area. If your design is affected by high fan-out nets and you want the synthesis tool to reduce that fanout, use fan-out attributes specifically on that net versus globally reducing the fan-out limit. Avoid maintaining hierarchy if critical paths cross over the hierarchical boundaries. Before implementation, review the warnings in your synthesis report.

Choosing Implementation Options

Having obtained an acceptable timing estimate from the synthesis tool, you can use the implementation tools to determine the true performance of the design. The ISE default mode is the performance evaluation mode, which enables you to get high-performance results out of your implementation tools without having to specify timing goals.

The next step is to run timing-driven mapping (MAP) and place and route (PAR). Timing-driven MAP performs closed-loop packing and timing-driven placement, while PAR performs the routing of the design. Both MAP and PAR should run with their effort levels set to high to achieve optimal results.

Physical synthesis options in implementation can re-optimize and pack logic based on knowledge of the critical paths of a design, leading to better placement and routing. The physical synthesis options are implemented during the MAP process and include global netlist optimization, localized logic optimization, retiming, register duplication, and equivalent register removal. Details on each of these options can be found in the Xilinx White Paper, "Physical Synthesis and Optimization with ISE 8.1i," available at *www.xilinx.com/bvdocs/whitepapers/wp230.pdf.*

Xplorer Utility

Xplorer is a tool that helps to determine the set of implementation options that result in the best performance for a design. Xplorer has two modes: timing closure and best performance. The timing closure mode evaluates your timing constraints and tries different sets of implementation options to achieve those goals. In best performance mode, you can give the tool a clock domain to focus on; the tool will try to achieve the best frequency for the clock. This is helpful when benchmarking a design's maximum performance.

Evaluating Your Critical Paths

By understanding the characteristics of your critical path, you can make better decisions about what to do for your next design iteration. A datapath comprises both logic and interconnect delay. Individual component delays that make up logic delay are fixed. You can reduce logic delay by reducing the number of logic levels or by redefining the structure of the logic.

In comparison, interconnect delay is much more variable and is dependent on the placement of the logic. Before running your design through PAR, a timing analysis after MAP is recommended. Although this timing report will only have estimates for your routing delays, it can give you an idea of the critical paths the implementation tools are working on. If the critical paths have a high number of logic levels, you may want to work on improving the logic levels versus running it through PAR.

If your design has an excessive amount of logic levels:

1. Try the physical synthesis options in MAP.
2. Go back to synthesis and verify that critical paths reported in implementation match what is reported in synthesis.
3. Review the synthesis inference of your HDL code.

If there are few logic levels but certain datapaths are not meeting timing:

1. Evaluate fan-out on routes with long delay.
2. If the critical path contains hard-IP blocks such as block RAMs or DSP48E slices, verify that the design takes full advantage of the embedded registers. Also understand when to make the trade-off between using these hard blocks or using slice logic.
3. Analyze clock skew.
4. If the logic appears to be placed far apart, floorplanning of critical blocks may be required. Only floorplan where necessary.
5. If area groups were created for a design with a previous version of software or before many design changes, consider removing those area groups.
6. Consider placing hard-IP blocks such as block RAMs for DSP48E slices.

Conclusion

Virtex-5 FPGAs are optimized for high-performance designs, while ISE software has the capabilities you need to quickly achieve design closure, improve productivity, and efficiently verify your designs. Xilinx provides a comprehensive suite of software tools (powered by ISE Fmax technology) that improves design performance. However, the more that you can do upfront with good coding styles, defining timing constraints, and resource planning, the easier it will be for downstream tools to achieve your timing requirements.

	XST	Synplify Pro
Fan-out Control	max_fanout	syn_maxfan
Directs Inference of RAMs to Block RAMs or SelectRAM	ram_style	syn_ramstyle
Directs Usage of DSP48 Slice	use_dsp48	syn_multstyle/syn_dspstyle
Directs Usage of SRL16	shreg_extract	syn_srlstyle
Controls % of Block RAMs Utilized	n/a	syn_allowed_resources
Preservation of Register Instances During Optimizations	Keep	syn_preserve
Preservation of Wires	Keep	syn_keep
Preservation of Black Boxes with Unused Outputs	Keep	syn_noprune

* You can find XST documentation at *http://toolbox.xilinx.com/docsan/xilinx82/books/docs/xst/xst.pdf.* Synplify Pro documentation is located in the tool help documentation.

*Table 1 – Helpful synthesis attributes**

Clock Management in Virtex-5 Devices

Virtex-5 FPGAs give designers fresh choices.

by Ralf Krueger
Sr. Staff Applications Engineer
Xilinx, Inc.
ralf.krueger@xilinx.com

As FPGAs grow in size, quality on-chip clock distribution becomes increasingly important. Clock skew and clock delay impact device performance; managing clock skew and clock delay with conventional clock trees becomes more difficult in large devices.

Traditionally, you would deploy solutions such as a Xilinx® Virtex™-4 digital clock management (DCM) or mixed-signal phase-locked loop (PLL) to achieve clock tree deskew and frequency synthesis, among other functions. Yet each solution has its advantages and disadvantages.

In Virtex-5 devices, for the first time in an FPGA, both digital DCMs and analog PLLs are implemented side by side in a clock management tile (CMT). You can now select the clock management solution best suited for your particular applications.

Each Virtex-5 device has as many as six CMTs. A CMT contains two DCMs and one PLL. You can use either of the two DCMs or the PLL as a stand-alone module, or they can interact with each other. If used as a stand-alone module, the application requirements typically dictate which clock management solution to use. The DCM, for example, supports a fine phase shift, a dynamic phase shift, and a multiply/divide feature that does not depend on any maximum VCO frequency. However, the PLL filters input clock jitter, support a wide range of output frequencies with higher frequencies, and consume less power.

The DCM and PLL are also designed to interact with each other. The PLL can help clean up input or output clocks to the DCM. Dedicated resources within each CMT make the connections and still guarantee a proper deskew of the FPGA clocks. The CMTs are located in the center column of the Virtex-5 architecture. This enables well-matched clock routes to and from every DCM or PLL for enhanced symmetry (see Figure 1).

DCM

Virtex-5 DCMs provide a zero propagation delay buffer, clock division and multiplication capabilities, fixed and dynamic fine phase shift, and multiple phases of the input clock. Along with fully differential global clock trees and low skew between output signals, the application's various clocks are distributed efficiently throughout the device. Each DCM can drive as many as 9 of the 32 global clock routing networks within the device.

The global clock distribution network minimizes skews caused by loading differences. By monitoring a sample of the DCM output clock, the DLL compensates for the delay on the routing network, effectively eliminating the delay from the external input port to the individual clock loads within the device.

In addition to providing zero delay with respect to a user source clock, the DCM provides multiple phases of the source clock. The DLL can also act as a clock doubler or divide the user source clock by as much as 16. The DCM can also act as a clock mirror. By driving the DCM output off-chip and then back in again, the DCM can deskew a board-level clock between multiple devices.

Another submodule provides the ability to phase shift the DCM's output clock in small increments (1/256th of the period). The versatile digital phase shift (DPS) operates in four different modes for maximum flexibility: fixed, variable-positive, variable-center, and direct. The DCM's digital frequency synthesis (DFS) module provides two outputs, CLKFX and CLKFX180, which are derived from the input clock by frequency multiplication and division. You provide valid multiply (M) and divide (D) values, which the DFS implements through a frequency calculator. For example, if you provide an M value of 19 and a D value of 8, they would yield a 2.375 source-clock multiplier.

PLL

The CMT's PLL is a mixed signal block designed to support clock network deskew,

frequency synthesis, and jitter reduction. The PLL block diagram in Figure 2 provides a general overview of the various components.

Input multiplexers (MUXs) are used to select the reference and feedback clocks from the global clock pins, global clock trees, or one of the DCMs. Each clock input has a programmable counter. This pre-scales the reference clock and allows a wide range of frequency synthesis.

The phase frequency detector (PFD) compares both phase and frequency of the input clock and the feedback clock. A signal is generated that is proportional to the phase and frequency error between the two clocks, which is then used to drive the charge pump and loop filter to generate a reference voltage to the VCO. An up or down signal from the PFD determines if the VCO should operate at a higher or lower frequency.

After the PFD determines that the input and feedback clocks are phase- and frequency-aligned, a lock signal is raised, indicating that the PLL output clocks are valid. The VCO continues to compensate for any variations in voltage or temperature. The M counter in the feedback path controls the feedback clock and multiplies the VCO frequency to the desired target frequency. The VCO output clock drives six output counters. Each can be independently programmed to generate a variety of frequencies for the application design.

Additionally, clock switchover, phase shifting, various duty cycles, and bandwidth control are also supported. You can dynamically select one of two input clocks before or during operation. In many cases, alternate phases of a clock are required. The VCO provides eight phase-shifted clocks at 45 degrees each. The higher the VCO frequency, the smaller the phase-shift resolution of the clocks coming out of the 0 counter.

You can individually program each output counter to provide a separately phase-shifted clock. The PLL can also generate non-50/50 discrete duty cycles in each output counter. The resolution and possible output duty cycles depend on the divide value. The higher the output divide value, the higher the resolution setting of the output duty cycle.

Conclusion

Virtex-5 FPGAs give digital designers a choice of either digital or analog clock management. Depending on your particular application, either module – or a combination of both modules – provides you with choices that you never had before.

Together with an abundance of clock tree resources, Virtex-5 devices greatly simplify and improve system-level designs involving high fan-out and high-performance clocks. Virtex-5 devices have powerful frequency synthesis, phase-shifting, and clock deskew capabilities never offered before in an FPGA. Along with comprehensive software support, you can achieve larger, faster, and more complex designs than in any previous-generation FPGA.

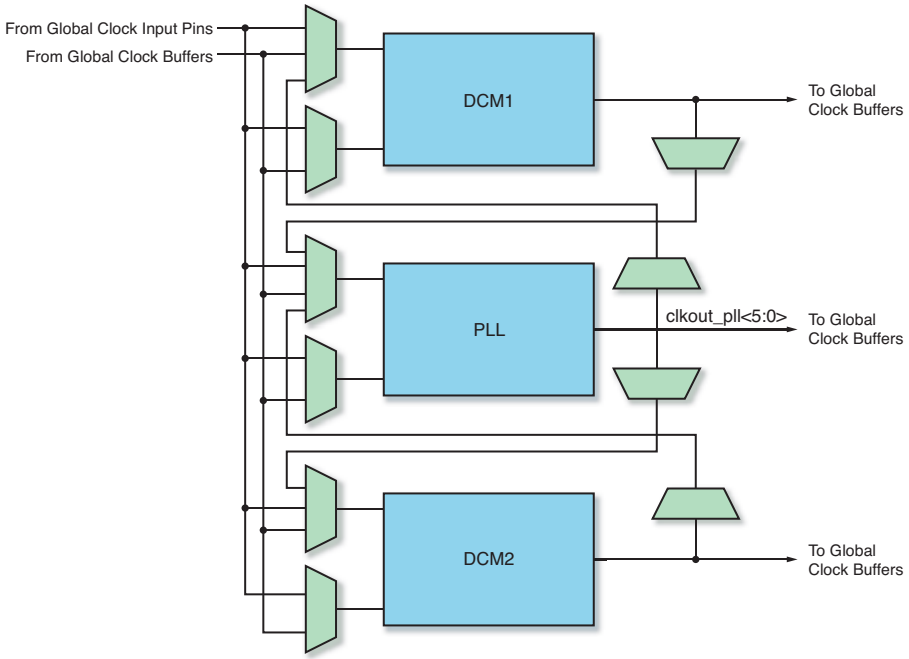

Figure 1 – Virtex-5 CMT block diagram

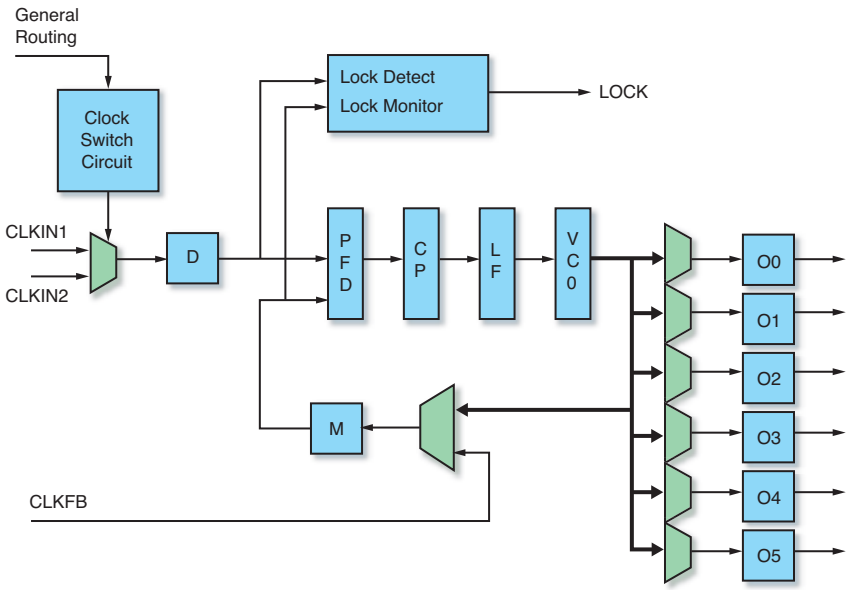

Figure 2 – PLL block diagram

Reduce Power with Virtex-5 FPGAs

The world's first 65-nm FPGAs offer the lowest power without compromising performance.

by Derek Curd
Senior Staff Applications Engineer,
Advanced Products Division
Xilinx, Inc.
derek.curd@xilinx.com

With the introduction of the Virtex™-5 family, Xilinx is once again leading the charge to deliver new technologies and capabilities to FPGA consumers. The move to 65-nm FPGAs promises to deliver benefits traditionally associated with smaller process geometries: lower cost, higher performance, and greater logic capacity. And although these benefits present exciting opportunities for advanced system designers, the 65-nm process node brings with it new challenges.

Power consumption, for instance, becomes increasingly important when selecting an FPGA for your application. More than likely, your next-generation design will require you to integrate more features and higher performance within a similar (or perhaps even smaller) power budget.

In this article, I'll explore the benefits of reduced power consumption. I'll also illustrate the many process and architectural innovations implemented in Virtex-5 devices to offer you the lowest possible power solution without compromising performance.

Benefits of Reducing Power

Implementing a lower power FPGA design offers advantages beyond simply adhering to the device's thermal operating requirements. Although meeting component specifications is obviously critical for performance and reliability, how you achieve this has a significant impact on system cost and complexity.

First, lowering FPGA power consumption allows you to use less expensive power supplies, which have fewer components and consume less PCB area. The implementation cost for a high-performance power system is typically between $.50 and $1 per Watt. Lower power FPGA operation, therefore, contributes directly to lowering overall system cost.

Second, because power consumption is directly related to heat dissipation, lower power operation allows you to use simpler, less expensive thermal management solutions. In many cases, designs will not need heat sinks, or they will need smaller, less expensive heat sinks.

Finally, because lower power operation means fewer components and lower device temperatures, overall system reliability improves. A decrease of 10° C in device operating temperature can translate to a 2x increase in component life, which clearly illustrates the importance of controlling power and temperature for systems with high reliability requirements.

Power: Challenges and Solutions

Total power in an FPGA (or any semiconductor device) is the sum of two components: static power and dynamic power. Static power results primarily from transistor leakage current, the small current that "leaks" from either source-to-drain or through the gate oxide of the transistor even when it is logically "off." Dynamic power is the power consumed during switching events in the core or I/O of the device and is therefore frequency-dependent.

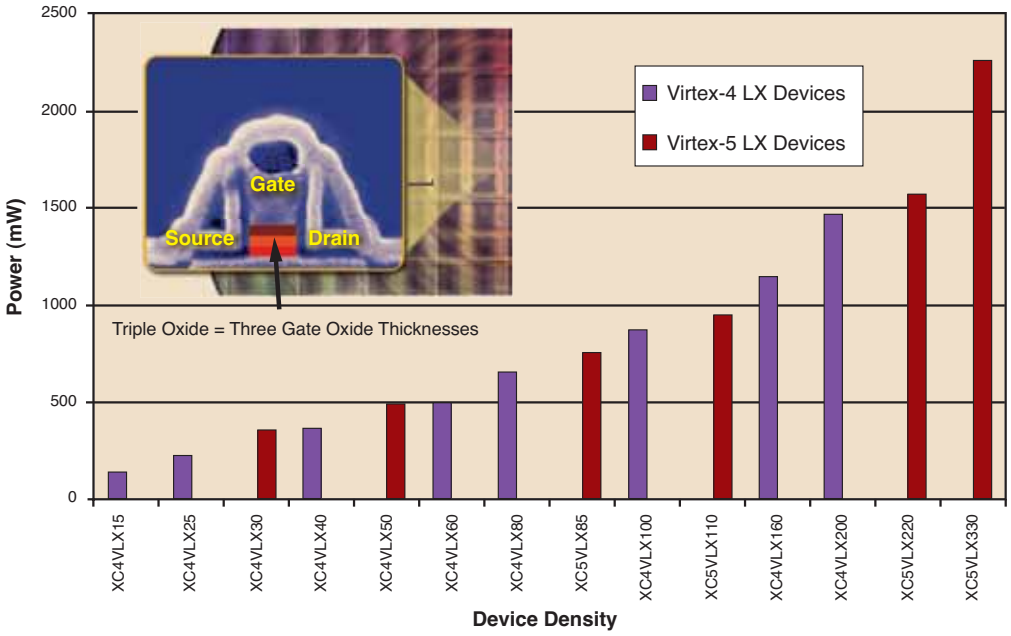

Figure 1 – Static power comparison at 85° C

Static Power

As you shrink transistor size (for example, move from 90-nm to 65-nm devices), leakage current tends to increase. The shorter channel lengths and thinner gate oxides generally used at the new process node make it easier for current to leak, either across the channel region or through the gate oxide of the transistor.

In the 90-nm Virtex-4 family, Xilinx introduced "triple-oxide" process technology, which gave Xilinx® circuit designers a tremendous tool to fight leakage. In older FPGAs, two gate-oxide thicknesses were used: a thin one for the high-performance, lower operating voltage transistors in the FPGA core, and a thicker one for the larger, high-voltage-tolerant transistors in the I/O blocks. Simply put, "triple oxide" refers to the addition of a third, medium-thickness gate oxide (or "midox") transistor that has much lower leakage than the thin-oxide core transistor.

The "midox" transistor is used extensively in the core of the device for non-performance-critical circuits (like configuration memory) or circuits that do not require fast switching times in response to a changing gate voltage (like routing pass gates). The thin-oxide, highest leakage transistors are reserved only for the portions of the speed path that require very fast switching times. The net result is that total device leakage is dramatically reduced, while still offering a substantial performance improvement over previous-generation FPGAs.

The triple-oxide process allowed Virtex-4 devices to reduce static power consumption by an average of more than 70% relative to competing 90-nm FPGAs. The results were so successful that the Virtex-5 family again makes extensive use of this technology to reduce leakage at the 65-nm process node.

Figure 1 illustrates how the triple-oxide process enables 65-nm Virtex-5 devices to achieve comparable static power to similarly sized 90-nm Virtex-4 devices under worst-case (high-temperature) operating conditions, despite industry predictions that 65-nm devices would see a dramatic rise in static power consumption. Thus, the Virtex-5 family retains a substantial static power advantage over competing high-performance FPGAs.

Dynamic Power

Dynamic Power consumption presents other challenges for 65-nm FPGAs. The

equation governing dynamic power is:

dynamic power = CV^2f

where C is the capacitance of the node switching, V is the supply voltage, and f is the switching frequency. The 65-nm process node enables FPGAs that have significantly greater logic capacity and higher performance than older devices. In other words, more nodes are switching at higher frequencies. All else being equal, this tends to increase dynamic power.

However, there is good news with respect to dynamic power at 65 nm. The core FPGA supply voltage (V) and node capacitance (C) generally reduce with each new process node, providing substantial dynamic power savings over previous-generation FPGAs.

In Virtex-5 devices, the core supply voltage (V_{CCINT}) decreases from the 1.2V used in Virtex-4 devices to 1.0V. Node capacitance tends to decrease because of smaller parasitic capacitances (associated with the smaller transistors) and shorter, less capacitive interconnects between logic. Additionally, Virtex-5 devices use a reduced-K dielectric material between metal interconnect layers to minimize routing capacitance.

The estimated reduction in average node capacitance for Virtex-5 devices is 15% compared to Virtex-4 devices. Taken together with the voltage reduction benefit, this translates to a 35-40% reduction – at least – in core dynamic power for Virtex-5 devices.

Although the "process shrink" to 65 nm provides an inherent 35-40% dynamic power reduction, architectural innovations in Virtex-5 devices offer additional power savings for every design. Most of the node capacitance that contributes to dynamic power is attributed to the routing or interconnect between logic functions. The new Virtex-5 architecture fundamentally reduces routing capacitance in two ways:

- Virtex-5 configurable logic blocks (CLBs) are based on a six-input look-up table (6-LUT) logic architecture, as opposed to the 4-LUT architecture used in older devices. This means that more logic is implemented within each LUT, translating to fewer levels of logic and thus a reduced need for higher capacitance routing between logic functions.
- The Virtex-5 routing architecture now includes diagonally symmetric routes, meaning that every CLB now has a direct "one hop" connection to all of its neighbors, including diagonal neighbors. When a connection is required between logic functions, it is now more likely that this connection is a less-capacitive "one hop" connection, whereas previous routing architectures may have required two or more hops for the same connectivity.

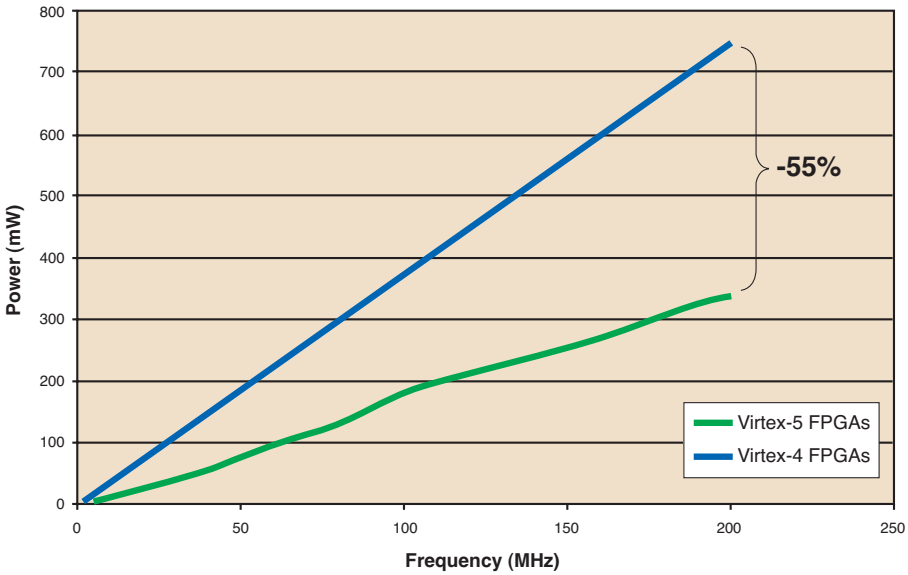

Figure 2 – Dynamic power comparison of counter benchmark design

Together, the 6-LUT architecture and improved routing pattern reduce core dynamic power by lowering average node capacitance beyond the level achieved purely from 65-nm process scaling. Figure 2 shows the core dynamic power measurements from a benchmark design in which a Virtex-5 device and a Virtex-4 device are each filled with 1,024 8-bit counters. These actual silicon measurements illustrate that the combined process and architectural benefits to dynamic power reduction can exceed 50%.

Hard IP Blocks

Virtex-5 devices contain more hard IP blocks (circuitry dedicated to commonly used functions) than any other FPGA in the industry. FPGA designs that utilize these blocks see additional power savings in comparison to implementing these functions in general-purpose FPGA logic.

Unlike the FPGA fabric, these dedicated blocks contain only the transistors necessary to implement the required function. And there are no programmable interconnects, so routing capacitance is as small as possible. Fewer transistors and lower node capacitance benefit both static and dynamic power consumption. The net result is that these dedicated blocks can perform the same function in as little as one-tenth the power of an equivalent implementation using the general-purpose FPGA fabric.

In addition to adding new types of dedicated blocks, many blocks that existed in Virtex-4 devices have been redesigned in Virtex-5 devices to add features, improve performance, and reduce power. For example, the 18-Kb block RAM memories in the Virtex-4 family have been sized up to 36-Kb block RAMs in Virtex-5 devices; each of these block RAMs can be broken into two independent 18-Kb memories for backward compatibility to Virtex-4 designs.

Interestingly, from a power perspective, each of the 18-Kb sub-blocks is constructed from two 9-Kb physical memory arrays. For the majority of block RAM configurations, any given read or write request to the block RAM only needs to access one of the 9-Kb

physical memories at a time. The other 9-Kb memory can therefore be effectively "powered down" while it is not being accessed. This reduces power consumption by nearly an additional 50% beyond those reductions resulting from the 65-nm process migration. This "ping-pong" accessing of the 9-Kb blocks is inherent to the new block RAM architecture, meaning that no user or software control is required to take advantage of this capability. It occurs dynamically and automatically, providing dramatic power reductions for all designs that use block RAM without compromising block performance.

The dedicated DSP elements in Virtex-5 devices have also received a significant design overhaul to incorporate more functionality at higher performance and lower power consumption. In a slice-versus-slice comparison, the new Virtex-5 DSP slice has roughly 40% lower dynamic power consumption relative to the Virtex-4 DSP slice. This is mostly attributable to the voltage and capacitance scaling factors of the 65-nm process discussed earlier.

However, because the new Virtex-5 DSP slice has greater functionality and wider interfaces, many DSP operations experience even greater dynamic power reduction by taking advantage of these additional capabilities. In many cases, you can achieve dynamic power reductions as high as 75% when utilizing the full capability of the new DSP slice. If you are not designing a DSP application, keep in mind that you can use the DSP slices for many standard logic functions (counter, adder, barrel shifter) at a substantial power savings compared to implementing the same function in standard FPGA logic.

As a final example of redesigned dedicated blocks, the LXT platform of the Virtex-5 family includes integrated multi-gigabit serial transceivers, running at rates as fast as 3.125 Gbps. These "SERDES" blocks are implemented with an emphasis on reducing power consumption. Each full-duplex transceiver in a Virtex-5 LXT device consumes less than 100 milliwatts of total power at 3.125 Gbps, representing roughly a 75% reduction relative to Virtex-4 serial transceivers.

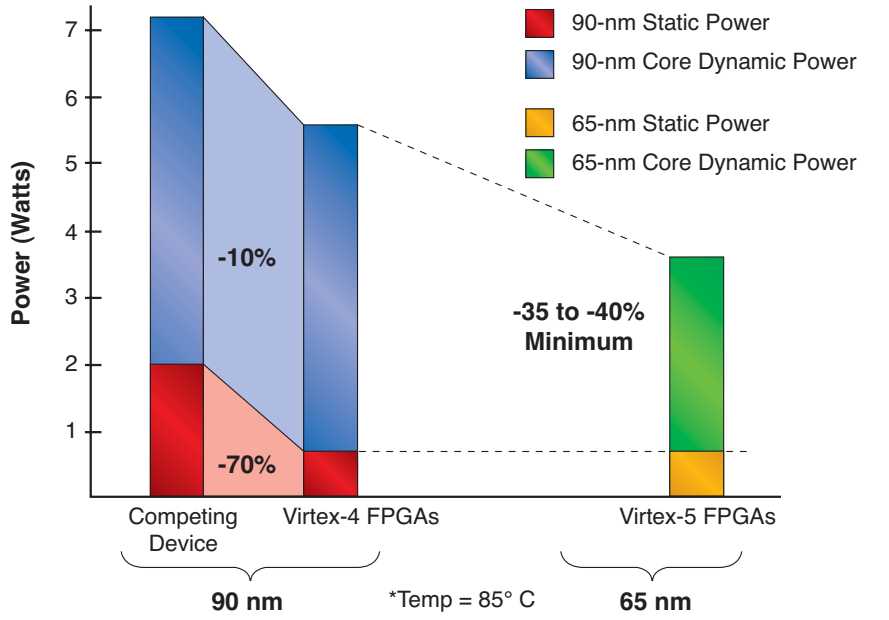

Figure 3 – Power comparison between available FPGAs for a typical design

Conclusion

Xilinx has a long history of innovation dating back to the invention of the first FPGA more than 20 years ago. So it is no surprise that Xilinx was the first FPGA company to make reducing power a top priority in deep sub-micron technologies. As with the Virtex-4 family, Virtex-5 devices employ a number of process and architectural innovations aimed at offering the lowest possible power consumption, while still enabling performance increases of 30% or more.

As Figure 3 illustrates, with static power levels comparable to Virtex-4 devices, the Virtex-5 family provides a clear advantage relative to competing FPGAs. As the only available 65-nm FPGA, Virtex-5 devices also offer a minimum of 35-40% core dynamic power reduction over other high-performance FPGAs on the market. Architectural innovations such as the new 6-LUT and diagonally symmetric routing are likely to enable actual core dynamic power savings up to 50% or more. And taking advantage of the unprecedented level of dedicated blocks lowers power consumption even further.

To find out more about how you can harness the low power of Virtex-5 devices, visit *www.xilinx.com/power*.

Xilinx Power Estimator (XPE)

Introduced in January 2006, the Xilinx® Power Estimator (XPE) spreadsheet-based power tool supports the Virtex™-4 and now Virtex-5 and Spartan™-3E FPGA families. XPE was designed to replace the Web Power tool as the premier pre-design power estimation tool for all new Xilinx FPGA families. The key advantages of XPE over previous power estimation tools are an improved user interface, improved accuracy, and better presentation of important data.

XPE's summary page displays a complete summary of power usage, first by resource type and then by voltage supply. You can use the navigation buttons on the summary page to access more detailed information. XPE automatically displays several graphs to complete the power usage picture.

Since the initial release, Xilinx has introduced newer versions of XPE that include many additional features and improvements in accuracy. These versions, plus those supporting Virtex-5 and Spartan-3E devices, are available at *www.xilinx.com/power*.

– Kevin Bixler
Power Tools Product Marketing Engineer
Xilinx, Inc.

Compact Power Supplies for Xilinx-Based Systems

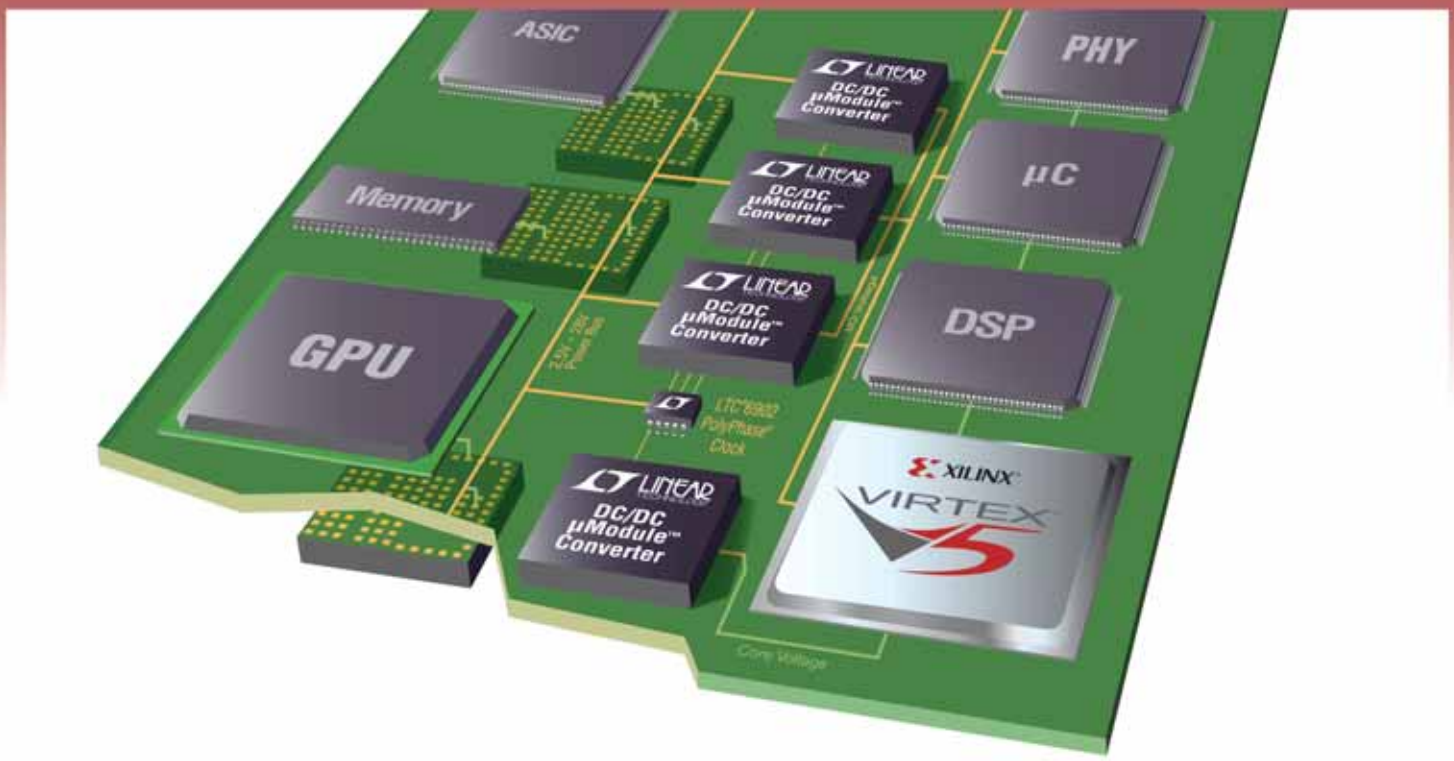

Tiny µModule™ Power Supplies Fit on Both Sides of PC Board

Our µModule DC/DC point-of-load power supply family is complete with built-in inductor, MOSFETs, bypass capacitors and compensation circuitry. At only 2.8mm height, these tiny, lightweight (1.7g) point-of-load regulators fit the tightest spaces on top and bottom of your board. Small size and impressive low thermal impedance allow high power conversion from a wide range of voltages. Our µModule DC/DC converters simplify the design of your Xilinx-based system and are backed by rigorous testing and high product reliability.

µModule DC/DC Converters for Core, I/O, Clock & System Power

V_{IN}: 4.5V-28V; V_{OUT}: 0.6V-5V						LGA Package (15°C/W)	
Part No.	I_{OUT} (A)	Current Share	PLL	Track, Margin	Remote Sense	Height (mm)	Area (mm)
LTM®4602	6	Combine two for 12A to 24A or 4x LTM4601 for ≤48A				2.8	15x15
LTM4603	6		✔	✔	✔		
LTM4603-1	6		✔	✔			
LTM4600	10						
LTM4601	12		✔	✔	✔		
LTM4601-1	12		✔	✔			
V_{IN}: 2.25V-5.5V; V_{OUT}: 0.8V-3.3V							
LTM4604*	4	2x for 8A	✔			2.3	9x15

*Future Product

Info & Online Store

Nu Horizons Electronics Corp.
Tel: 1-888-747-NUHO
www.nuhorizons.com/linear

FREE
LTM4600
µModule Board
to qualified customers.

NU HORIZONS ELECTRONICS CORP.

LINEAR TECHNOLOGY

Applying Compact Thermal Models

Xilinx gives you one more tool in the FPGA thermal management toolbox.

by Abu Eghan
Principal Engineer
Xilinx, Inc.
abu.eghan@xilinx.com

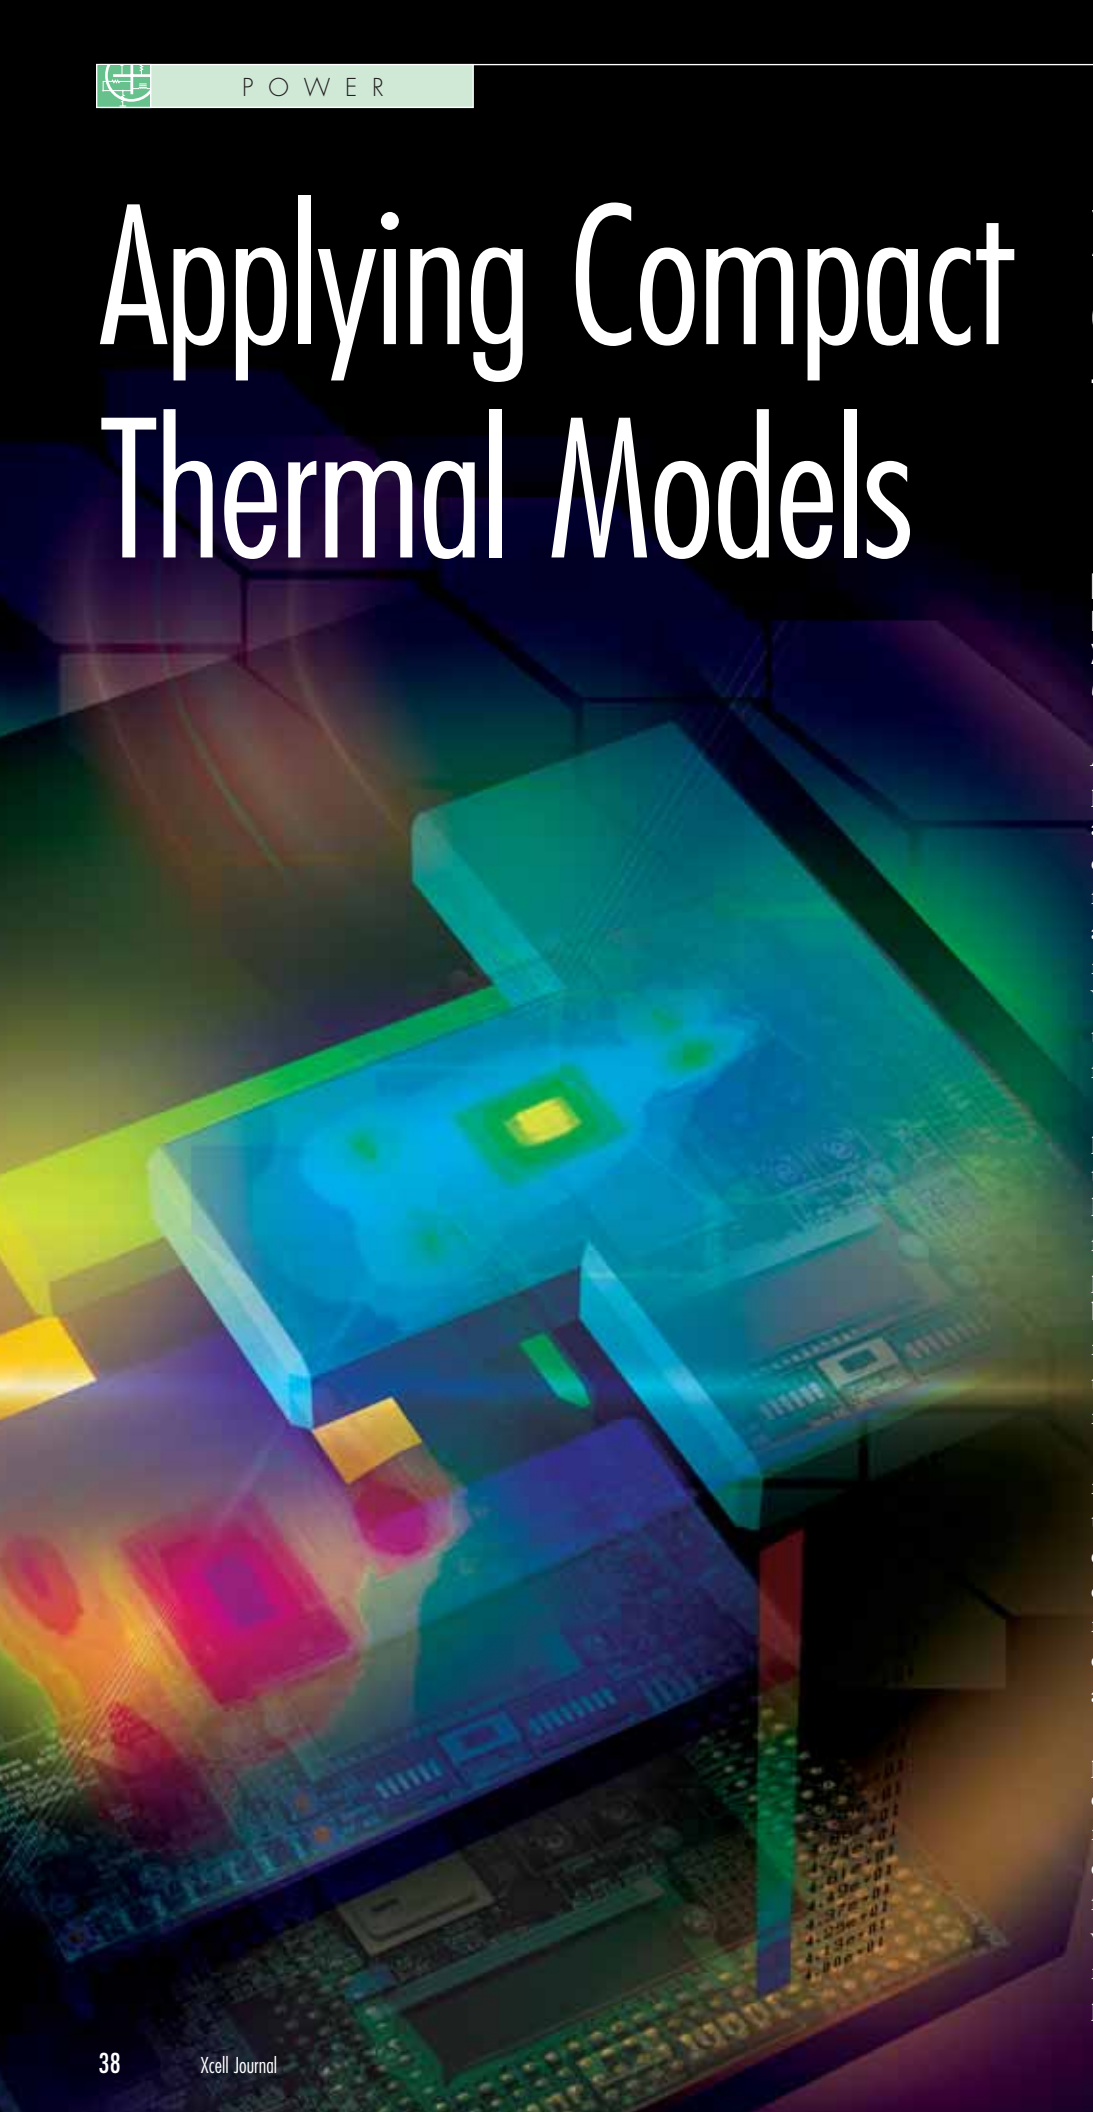

Although Xilinx has made substantial progress at the silicon level to reduce static and dynamic power in FPGAs, each successive family takes advantage of reduced feature sizes to increase transistor density and performance – thus resulting in thermal concerns for the top end of the family. You should not underestimate the importance of power consumption mitigation for these devices.

Designers cannot afford inaccurate temperature predictions when power consumption is high and thermal budget margins are low. Flip-chip packages used in high-performance FPGAs have multiple heat-flow paths and are thermally efficient. Using the basic "one-resistor" figure of merit thermal resistance – Theta-ja (Θja) – in estimating temperature does not do justice to the thermal efficiency of the packages.

Thus, a need exists for an alternate and more accurate approach to obtain Tj predictions on these components in an end-user's environment. This is where the boundary condition-independent compact thermal model (BCI-CTM) becomes useful. You can conveniently use these models to make faster and more accurate Tj predictions.

In this article, I'll discuss better ways to predict temperature for these faster and denser FPGA components in a system environment. I'll also introduce the availability of and support for compact thermal models for Virtex™-4 and Virtex-5 devices as one way to help system designers and component selectors estimate temperatures in the pre-design and implementation phases.

Motivation for Better Predictive Models

In a specific system implementation, the actual component Tj may be different from the arithmetic predictions using the published Θja. The prediction depends on the environment and the prevailing conditions in the system. The following equation governs the relationship:

$$\Theta ja = \frac{Tj - Ta}{P}$$

Or, stated in Tj prediction form:

$$Tj = Ta + P * \Theta ja$$

where

Θ*ja* is the thermal resistance between the device junction and ambient

Tj = junction temperature of the device

Ta = ambient temperature

P = package power dissipation

Although you can easily determine Tj, Ta, and P, representing the thermal resistance in an application is not easy, particularly for packages with multiple thermal paths. The single parameter Θja is strongly influenced by the application environment and therefore does not represent a suitable thermal resistance.

Theta-ja – The Misunderstood Model

Theta-ja has become the base thermal parameter most engineers gravitate toward when estimating component Tj with known Ta. But for a more demanding, higher wattage component on a large multilayer system board – particularly with other components around it – this approach often leads to an erroneous prediction of Tj.

In a design with loose margins in the thermal budget, the simple prediction using published Θja data may not be an issue. Indeed, it will likely lead to a system running at a lower than predicted Tj, because most common board types are more efficient than the largest standardized thermal board. Increasingly, with higher wattage components where margins are tight, "conservative" data may be the difference between selection and rejection of the component in a specific program.

The key point here is that Θja was not meant to be used in these types of predictions. JEDEC, the semiconductor engineering standardization body of the Electronic Industries Alliance, explains in EIA/JESD51-2 that "the intent of Theta-ja measurements is solely for a thermal performance comparison of one package to another in a standardized environment. This methodology is not meant to and will not predict the performance of a package in an application-specific environment."

A typical implementation of a one-cubic-foot Θja still-air standardized environment is depicted in Figure 1. This is clearly not a typical system environment. Ideally, you should use these numbers to compare package efficiency, reserving any serious Tj prediction for other tools using models that are more relevant.

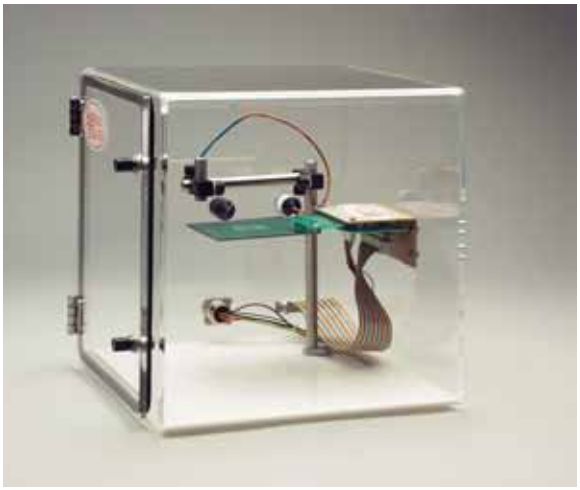

Figure 1 – The Analysis Tech implementation of Theta-ja standardized environment

To illustrate the pitfalls and potential discrepancies in Tj predictions, let's look at a Virtex-5 flip-chip component – XC5VLX50T- FF1136. The published Θja is 10.8° C per Watt. Although the Tj prediction expression will suggest a 43.2° C above ambient for 4W dissipation, actual detailed simulation shows a much lower number – and thus suggests a lower effective Θja – of close to 5° C per Watt.

Table 1 shows the corresponding Tj for the same component dissipating 4W on various FR4 board sizes and layer counts. This illustrates the power of the environment or boundary conditions on the effective Θja, and the type of Tj prediction discrepancy that can result.

Note that while in general the effective Θja tends to be lower on larger board environments, it can also trend higher and under-predict Tj on small cards in confined places like PDAs or cell phones. The same rationale is at play – Θja is not boundary condition-independent. A component with Θja = 22° C per Watt on a JEDEC board can easily exhibit a 30° C per Watt-effective Θja on a 30 mm x 30 mm card.

Some application engineers have suggested that because most high-performance devices use denser and larger PC boards, component suppliers should provide Θja using a larger "JEDEC/network board" – a board that may be closer to network application boards. This seems like a good argument and should be advocated at the next JEDEC forum. However, regardless of the board used for data gathering, the prediction will be wrong for some applications. Additional JEDEC boards and standard-

Xilinx 35 x 35mm FF1136-5VLX50T*		Board Size		
		4" x 4" Board	10" x 10" Board	20" x 20" Board
Layer Count of Mounted Board**	4	68.2° C	64.3° C	-
	8	63.0° C	50.9° C	48.3° C
	12	60.4° C	47.0° C	45.7° C
	16	59.1° C	46.6° C	44.9° C
	24	-	45.3° C	44.0° C

* Single component considered at 25° C ambient
**All layers have 1oz Cu with 80% coverage except outer layers that have 2 oz with 20% coverage.

Table 1 – Tj matrix for FF1136-XC5VLX50T on various boards

ized enclosures will only lead to more flavors of Θja, further confusing the issue. There ought to be a better way.

What Should an Engineer Do?

Engineers should view Θja with caution when predicting Tj in specific environments. Xilinx will continue to publish Θja and other thermal resistance data because those are the prevailing standards. They have their uses and should be deployed with their limitations in mind.

To address these limitations and to make more accurate Tj predictions in a system environment, a more refined model of the package is needed. Recognizing this need, Xilinx now supports compact thermal model data for high-performance FPGA devices.

What is a Compact Thermal Model?

A compact thermal model is a behavioral model that seeks to accurately predict the temperature of the package at selected nodes: junction, case, top, bottom, and balls, for example. It cannot predict the temperature at any other part of the package that is not predefined. It can be viewed as a reduced node abstraction of the response of a component to various boundary conditions. It is also more computationally efficient than the corresponding detailed model. These models are supplied for use in compatible computational fluid dynamics (CFD) tools for thermal simulations in place of detail models.

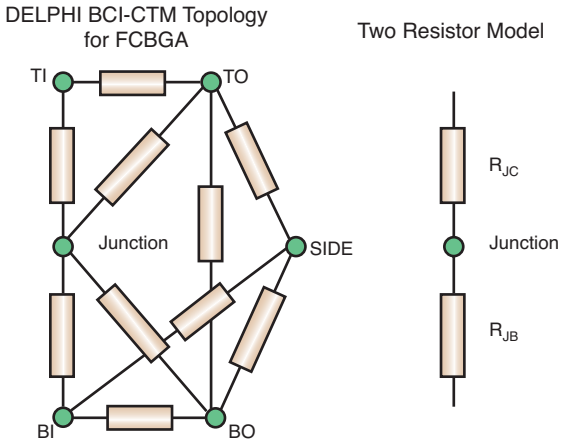

Figure 2 – CTM topologies

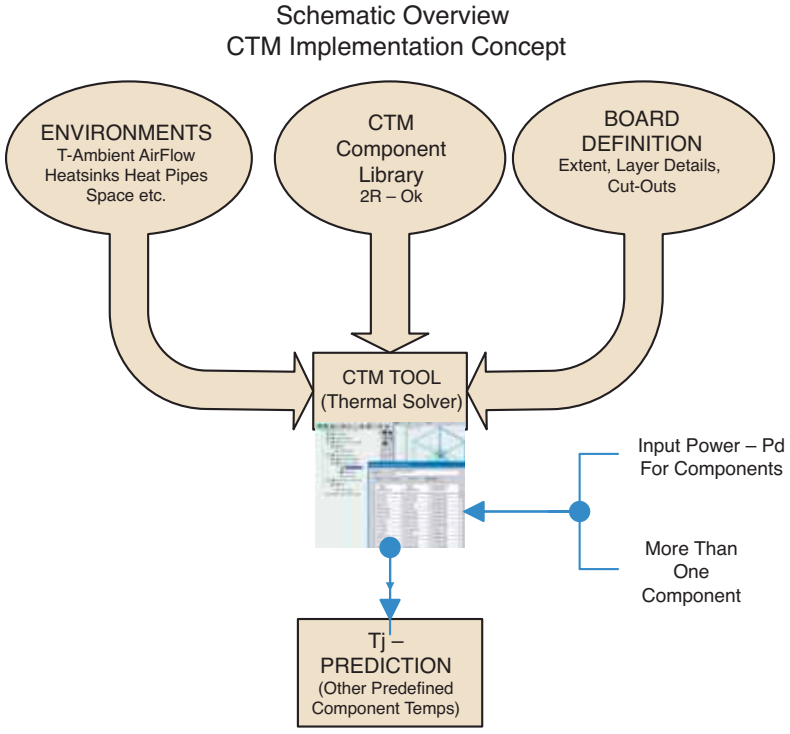

Figure 3 – CTM application schematics

Xilinx offers two model types for FPGA products:

1. Two-resistor (2-R) compact models comprising the familiar Theta-jc and Theta-jb for the package. There is no geometrical information. Although 2-R models are useful and give better predictions than traditional Θja estimations, they are not as accurate as Delphi models.

2. A Delphi compact model comprising several thermal resistors that connect a junction node (representing the die) to several surface nodes. Thermal links are also allowed between the surface nodes. Figure 2 shows the topology for a flip-chip BGA Delphi compact model. The matrix of resistors has been optimized through a Delphi optimization algorithm so that they can be used in various environments without compromising prediction accuracy.

Table 2 depicts a typical Delphi half-matrix model for flip chip. The resistance data is usually saved along with the node definitions and package extents to complete the model.

JEDEC has proposed a neutral file format in XML for CTM distribution. Xilinx plans to support the format when CFD tools adopt and support it. In the interim, Xilinx is offering the CTM files in two CFD tool formats, Flotherm and Icepak, selected from a pre-introductory survey of Xilinx customers. These tools cover the majority of those end-users who answered the survey. If you do not use one of these tools, you can request ASCII data for manual or script-based entry into your tool.

Application Examples

Figure 3 shows a typical flow for a CTM application. Normally, the component CTM data is stored in a library; as the user, you will bring in the CTM data as a library item. You then specify the board attributes and boundary conditions of your assembly, adding other items like component power and heat contributions from other components for the Tj prediction.

Delphi Compact Model (C/Watt)	Top Inner (TI)	Bottom Inner (BI)	Top Outer (TO)	Bottom Outer (BO)	Side
Junction	0.22	1.25	1.05	14.97	–
Top Inner (TI)		16.14	4.45	–	–
Bottom Inner (BI)			–	9.47	11.1
Top Outer (TO)				14.55	3.18
Bottom Outer (BO)					4.72

Table 2 – Delphi CTM resistors for FF676-XC5VLX50

Let's examine the benefits of CTM prediction in the cases below.

Case Study #1:
Battleboard Temperature Estimations
The "battleboard" is a 24-layer 20 x 16-inch board that Xilinx uses to assess signal integrity issues on Virtex components. In this case, lab measurements showed that a Virtex-4 component case temperature was well below what you would expect from the Θja prediction. A discrepancy of about 20° C was apparent. The Icepak CFD model using CTM inputs with radiation "on" showed a more realistic Tj – very comparable to those measured in the lab. Table 3 summarizes the observations (note that reported temperature is T-case).

Case Study #2:
Small Board with Multiple Components
A 3.75 x 2-inch board can illustrate the small board size effect and the influence of adjacent components on the Tj prediction of the component of interest. I have used the BCI-CTM approach with the aggressor components in active and inactive states to show the impact on Tj prediction.

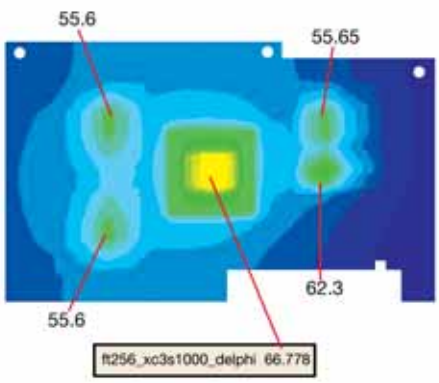

Figure 4 – Contours of static temperature

The high-density interconnect (HDI) board used in this case is smaller than the JEDEC standard 2S2P board. The Xilinx XC3S1000-FT256 component deployed on this board has a Θja of 19.7° C per Watt. With a 20° C ambient, and without any board input, you can predict a Tj of 39.7° C (20 + Pd * Θja [100LFM]).

The BCI-CTM model with 100 linear-feet-per-minute (LFM) airflow shows a Tj of 55° C with Ta = 20° C for the single component on the board. The predicted Tj is already worse than the JEDEC prediction. This case is in line with what you would see with smaller components that use smaller, thinner boards in consumer products – PDAs, MP3 players, or GPS systems.

Figure 4 shows the board temperature contour with four chip-size package (CSP) components doing 0.25W each. The 8-mm-square CSP components used 2-R data (published jb and jc) in the model. Figure 4 shows the component temperatures. The XC3S1000-FT256 component yielded a Tj of 65° C – a further 10° C rise over the single component case. Both single and multi-component runs took less than four minutes on a conduction-based tool using the Delphi CTM for the Xilinx component. These predictions are clearly different from what the basic Θja parameter predicted.

Hand Calc	Reported	ICEPAK CFD solution < 4 min	
Tj – JA = 10.6° C/watt (1-R model Ta = 25° C)	Battleboard Tc Ta = 25 - 27° C	Tj; 2-Resistor Model Ta = 25° C	Tj: Delphi Model as Published (Ta = 25° C)
		JB = 2.6 and JC = 0.19° C/watt	
67.4° C	46.4° C	48.8° C	43.9° C

Table 3 – "Battleboard" temperature predictions

Conclusion

Relying on the basic Θja metric to predict junction temperature for high-performance devices in a system is inadequate and can cause errors that could lead you to preclude a perfectly good component in your system. To address this shortcoming, Xilinx provides compact thermal models to assist in predicting Tj – in stand-alone calculations as well as system deployments.

You can use these models in CFD tools to make Tj predictions that take your environment and board conditions into consideration. Although you can accomplish the same predictions with a detailed package model, note that these CTMs offer reduced node benefits of faster solutions that are also computationally efficient.

You can download Virtex-4 and Virtex-5 CTM data at *www.xilinx.com/xlnx/xil_sw_updates_home.jsp*. Future high-performance FPGA products will have the Delphi models available for download as part of thermal collateral. Xilinx will also support legacy products such as Virtex-II and Virtex-II Pro devices, older Spartan™ FPGAs, and CPLD products on a by-request basis.

A Multi-Gigabit Transceiver for the Masses

The Virtex-5 GTP transceiver brings versatility, ease of use, power efficiency, and cost-effectiveness to high-volume mainstream applications.

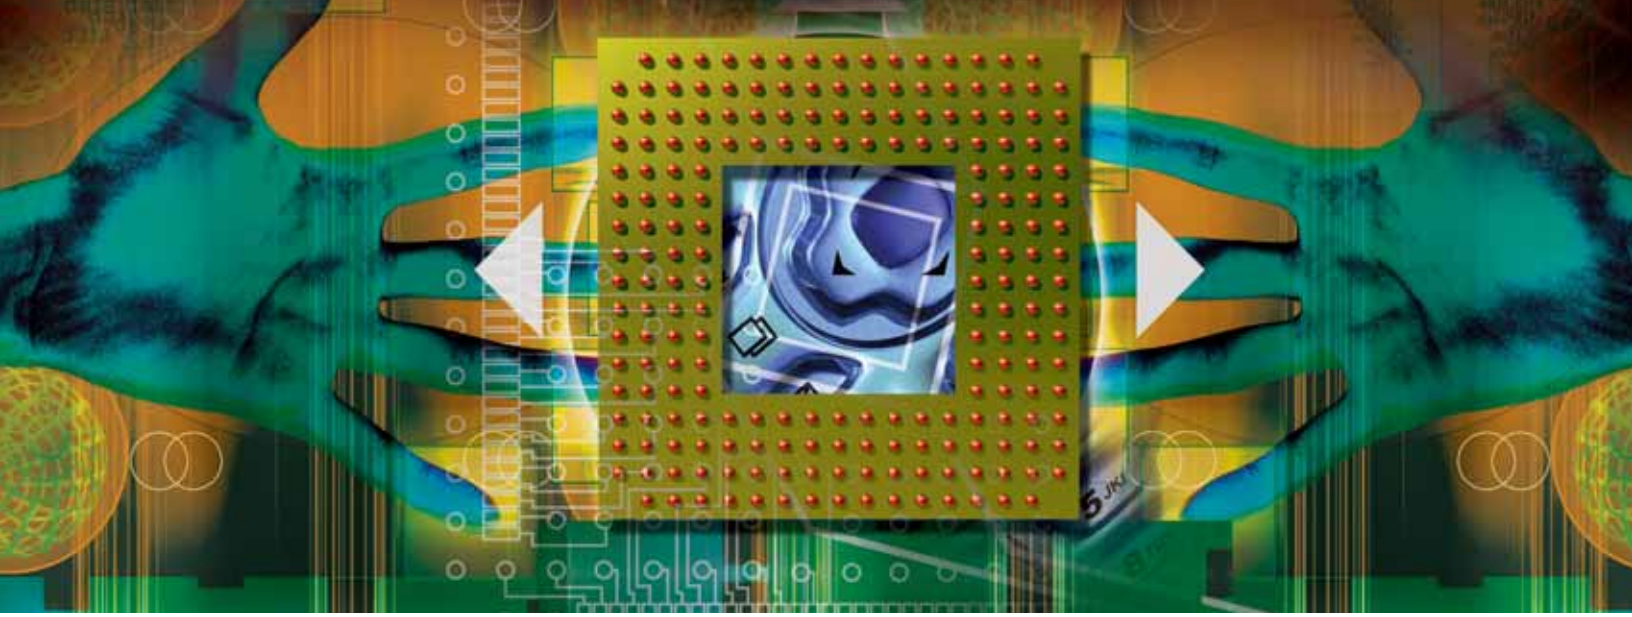

by Gang Sun
Senior Product Marketing Manager, High-Speed Serial I/O
Xilinx, Inc.
gang.sun@xilinx.com

The incessant demand for ever-increasing bandwidth has led designers away from parallel buses and low-speed transceivers toward serial transceiver-based interfaces. High-speed signals solve many design challenges; they offer new levels of bandwidth and lower overall system cost and power consumption.

These successes have led engineers to believe that the industry can continue to lower overall cost and power simply by increasing transceiver speed indefinitely. However, going beyond 3 Gbps can in some cases lead to fundamentally different engineering challenges that make it harder to lower overall system cost and power consumption. The explanation is simple; maintaining signal integrity becomes increasingly difficult at ultra-high speeds, and the extra overhead can sometimes outweigh the benefits associated with increased data rates.

Transceivers in Transition

Figure 1 shows the frequency loss and crosstalk associated with a legacy backplane channel. At 1.6 GHz, the loss is reasonably manageable, making transceiver implementation at or below 3.2 Gbps relatively cost-effective and power-efficient.

However, at 3 GHz, the loss becomes significant. Consequently, the implementation of a 6 Gbps backplane transceiver requires different feature sets. You will likely need advanced techniques such as decision feedback equalization (DFE) to maintain signal integrity, and these advanced capabilities require a different set of optimized features.

This explains why a 3 Gbps transceiver typically consumes less than 100 mW per channel, whereas a DFE-enabled 6 Gbps transceiver consumes at least twice as much power. For applications requiring these advanced features, this extra power consumption is a worthwhile trade-off. But it becomes advantageous to offer both a low-power 3.2 Gbps transceiver and a high-performance transceiver for cutting-edge applications – in essence offering the best tool for the job.

At 5 GHz, the signal-to-noise ratio (SNR) becomes negative. In that case, you would have to redesign the entire backplane with more expensive materials and more sophisticated manufacturing technologies to enable 10 Gbps transmission. Consequently, achieving a 10 Gbps serial transmission over a backplane incurs a higher cost in terms of die area and power consumption.

The preceding example clearly shows that transceivers running at or below 3.2 Gbps are at a sweet spot; they are more cost-effective and power-efficient than both parallel interfaces and ultra-high-speed transceivers (running at 6 Gbps and 10 Gbps) for a large

majority of interconnect applications. This phenomenon has led to two diverging trends in the transceiver market:

1. Bandwidth-hungry applications (such as a backplane interconnect for terabit routers) have needs for 6 Gbps and 10 Gbps transceivers. These applications continue to push the performance envelope while trading off cost and power.
2. High-volume applications are well served by transceivers running at or below 3.2 Gbps.

The Virtex-5 RocketIO GTP Transceiver

Xilinx clearly recognizes the different requirements of the high-performance market segment, noting that the high-volume market segments have in some cases conflicting requirements. The vast majority of serial protocols run at or below 3.2 Gbps; examples include PCI Express Generation 1, Gigabit Ethernet, XAUI, SATA I and II, Serial RapidIO, CPRI, OBSI, and HD-SDI. Many emerging protocols such as JEDEC's data converter interface and VESA's DisplayPort also run at these relatively slow data rates. In reality, these established and emerging protocols represent more than 90% of current transceiver applications. Therefore, transceivers running at or below 3.2 Gbps are "transceivers for the masses."

Xilinx has taken a truly innovative step and developed two different transceivers for its Virtex™-5 FPGA family. The first transceiver, the Virtex-5 RocketIO™ GTP transceiver, is designed for high-volume applications and covers data rates from 100 Mbps to 3.2 Gbps. Targeting the majority of system designers, the GTP transceiver is versatile, easy to use, power-efficient, and cost-effective.

The GTP transceiver is versatile because it is designed to support not only 8B/10B-based protocols such as the PCI Express Wrapper but also scrambling-based protocols such as SONET. (Table 1 is a complete list of applications supported by GTP transceivers.) Consequently, the spectrum of applications that can be supported by the GTP transceiver is limitless. In addition, validation and characterization of the GTP transceiver occurs in application-specific settings to ensure standards compliance. The combination of these design and characterization approaches ensures the universal appeal of the GTP transceiver.

The GTP transceiver is easy to use because it enjoys the support of the best FPGA CAD tools. The Xilinx® Virtex-5 RocketIO GTP transceiver wizard offers an intuitive GUI interface that allows you to select the GTP, clocking option, FPGA fabric interface, protocol stack, and encoding/decoding mechanism. After you have completed your selections, the tool generates a GTP wrapper with the necessary features.

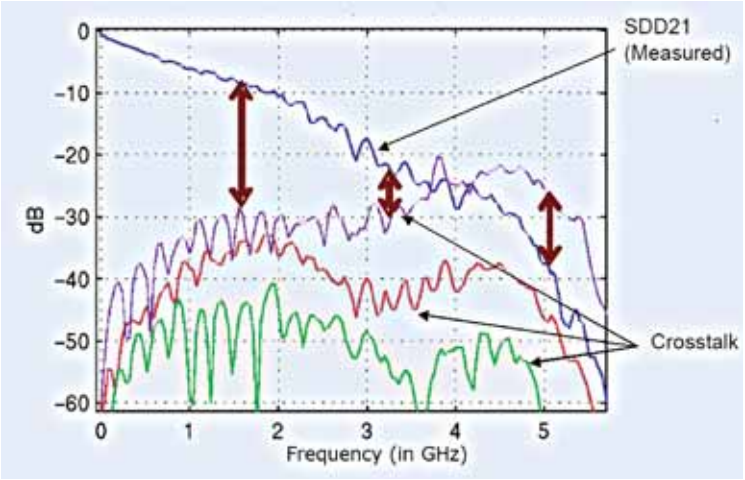

Figure 1 – Channel S-parameter and crosstalk

Figure 2 – ChipScope IBERT console

The Xilinx ChipScope™ Analyzer offers self-testing capabilities for the GTP transceiver by leveraging the integrated bit-error-rate tester (IBERT) feature built into the transceiver. The ChipScope IBERT console is shown in Figure 2.

Among the advanced features offered by the ChipScope Analyzer are channel performance measurement capabilities, automated eye scan capability for finding the best Tx and Rx settings, and transceiver and link status reporting. This comprehensive set of tool offerings greatly simplifies design and manufacturing efforts based on the GTP transceiver and is a key enabler for a large variety of applications.

As PCB boards become increasingly crowded, transceiver power consumption becomes a critical issue. Therefore, power efficiency was one of the top design objectives for the GTP transceiver. Average power consumption per GTP transceiver is substantially below 100 mW. In some cases, per-transceiver power consumption is as low as 60 mW. The universal appeal of low-power requirements further enhances the competitiveness of the GTP transceiver for power-sensitive applications.

As high-volume applications start to use embedded transceivers, cost has also become an important consideration. Consequently, Xilinx offers certain solutions in hard logic rather than in look-up tables (LUTs). For example, a hard-coded PCI Express protocol stack includes a physical layer based on the GTP transceiver, a link layer, and a transaction layer. This approach significantly lowers overall solution costs, and the increased cost-effectiveness makes GTP transceiver-based solutions even more attractive to high-volume/high-margin applications.

Conclusion

Transceivers at 3.2 Gbps or below are at a sweet spot; the vast majority of transceiver-based applications fall into this data-rate range. With its versatility, ease of use, power efficiency, and cost-effectiveness, the Virtex-5 GTP transceiver from Xilinx is ideally positioned in this market, a true multi-gigabit transceiver for the masses.

Market	Standard	Speed (Bits per Second)	Key Features
Telecom	OC-3/SDH STM-1	155 Mbps	• FIFOs can be Bypassed for Synchronous Operation
	OC-12/SDH STM-4	622 Mbps	
	OC-48/SDH STM-16	2.488 Gbps	
	OBSAI (Issue 1.0)	768 Mbps 1.536 Gbps 3.072 Gbps	
	CPRI (Version 2.0)	614 Mbps 1.228 Gbps 2.457 Gbps	
	SFI-5	2.448 -3.125 Gbps	• Synchronous Clocking (Bypass FIFOs)
Datacom	1G Ethernet (802.3z D5.0)	1.25 Gbps	
	XAUI (802.3ae D5.0)	3.125 Gbps	• Loss of Signal (LOS)
	10G Base CX-4	3.125 Gbps (x4)	
Computing / Communication	PCI Express Specification (Rev 1.1)	2.5 Gbps	• Tx Receive Detect • Loss of Signal (LOS)/Idle state detect • Low Power States and OOB Beacon • Ground Referenced Termination
	Serial Rapid IO	3.125 Gbps	• Supports All Data Rates from 1.25-3.125G
	InfiniBand	2.5 Gbps	
Storage	Fibre Channel (Rev4.0)	1.0625 Gbps 2.125 Gbps	• Rate Negotiation, Allows Tx and Rx to Operate at Different Speeds
	SATA (Rev1.0a)	1.5 Gbps 3.0 Gbps	• Rate Negotiation for Gen1/Gen2 • LOS and Out-of-Band Signaling Beacon
	SAS (Rev5)	1.5 Gbps 3.0 Gbps	
Video	SDI DVB-ASI	143 Mbps 176 Mbps 270 Mbps	• Internal AC Coupling Caps can be Bypassed for Video Standards • 2.97G is the New HD-SDI Standard in Development

Table 1 – Applications supported by GTP

Introducing the Virtex-5 PCI Express Endpoint Block

With PCI Express quickly becoming the standard high-bandwidth interconnect, the Virtex-5 LXT PCIe Endpoint block enables a configurable single-chip solution.

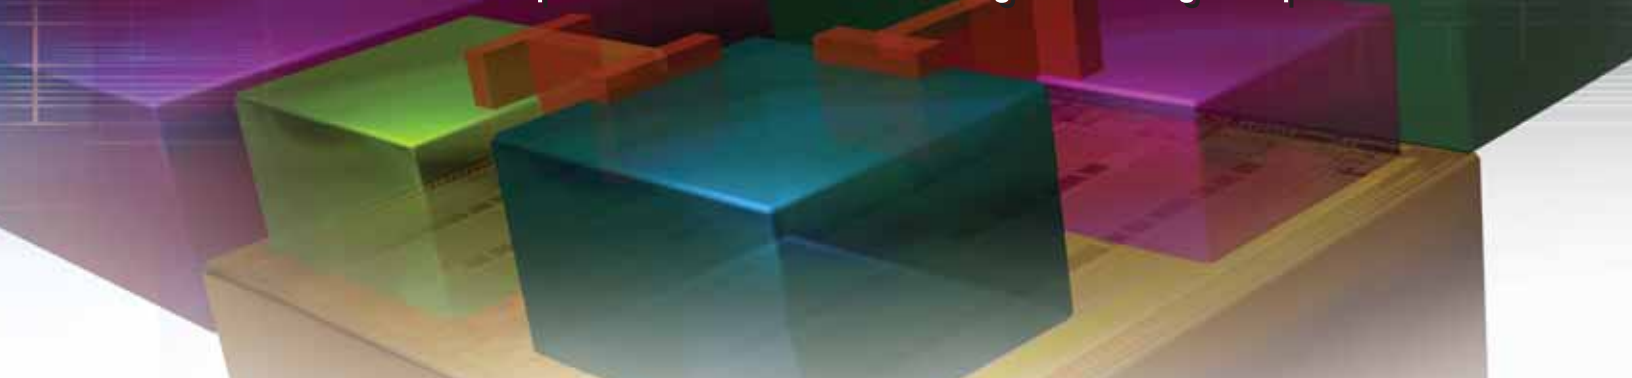

by Doug Kern
Staff System Design Engineer
Xilinx, Inc.
doug.kern@xilinx.com

Currently dominating the desktop PC motherboard and graphics markets, the PCI Express (PCIe) interconnect is poised to supplant PCI and PCI-X as the dominant high-bandwidth interconnect for the server, enterprise, mobile, workstation, networking, communications, industrial control, and medical equipment markets.

With more than 58 form factors, including Express Card, Advanced TCA, Compact PCI Express, Com Express, and a cable spec, the PCIe protocol is becoming ubiquitous. The PCI Special Interest Group (PCI-SIG) maintains the PCIe specification (along with the PCI and PCI-X specifications) and holds compliance workshops.

The PCIe subsystem is a point-to-point interface that replaces and overcomes the limitations of bus-based PCI and PCI-X standards. PCIe Generation 1 (Gen1) offers 2.5 Gbps speed with low-voltage differential signaling (LVDS), embedded 8B/10B encoding, dual-simplex signaling, and message-based serial protocol.

With plans in place to increase bandwidth to 5 Gbps in Generation 2 and 10 Gbps in Generation 3, the PCIe bus is expected to be the dominant high-bandwidth interconnect for several years to come. (For more information on the PCIe specification or compliance information, visit *www.pcisig.com*.)

With scalable lane widths from x1 to x32 lanes and advanced features such as traffic classes, virtual channels, hot-plug, and power management, the Xilinx PCIe block provides support for a wide range of applications, from a simple upgrade from PCI to an x1 PCIe endpoint device to advanced high-bandwidth x8 PCIe communications endpoint devices.

Figure 1 shows the topology of a PCIe system. The CPU is connected to a root device and is responsible for configuring and enumerating all plug-and-play PCI Express endpoint devices in a system. Because the PCIe system is point-to-point, switch devices are necessary to grow the number of devices or endpoints in a system. A switch has one upward facing port and numerous downward facing ports. These downward facing ports connect to the working devices or endpoints of a system.

Although only one root exists in a system, there are one or more endpoint devices. For example, a standard PC motherboard provides three to seven expansion PCIe slots. With the integrated PCI Express Endpoint block, Xilinx® Virtex™-5 LXT FPGAs allow you to rapidly develop and deploy high value-added PCIe endpoint devices. The numerous value-added endpoint designs are the target applications for the FPGA-based configurable Virtex-5 LXT PCI Express Endpoint block.

The Virtex-5 LXT PCIe Endpoint Block

The Virtex-5 LXT PCIe Endpoint block (see Figure 2) implements the physical layer (PHY), data link layer (DLL), transaction layer (TL), and configuration layers of a PCIe endpoint device. The implementation of a small reset circuit and clock generation blocks require you to use the FPGA fabric.

The PCI Express Endpoint block capabilities include:

- Compliance with the PCI Express base specification, revision 1.1
- Choice of PCI Express Endpoint block or legacy PCI Express Endpoint block implementation
- x8, x4, x2, or x1 lane width
- Easy-to-use user interface similar to the familiar Xilinx LocalLink interface
- Integration of RocketIO™ GTP transceivers
- Spread-spectrum clocking support
- Low power operation
- Power management support
- Ability to use on-chip block RAMs for buffering
- Fully buffered transmit and receive
- Management interface to access PCIe configuration space and internal configuration
- Support for full range of maximum payload size (128 to 4,096 bytes)
- Capable of as many as two virtual channels (VCs)
- VC arbitration: round robin, weighted round robin, or strict priority
- 6 x 32-bit or 3 x 64-bit base-address registers (BARs) or a combination of 32-bit and 64-bit BARs
- BARs configurable for memory or I/O
- Memory BAR checking/filtering
- Non-memory transaction layer packet (TLP) ID checking/filtering
- Implements one PCI Express function
- Signals to the programmable fabric for statistics and monitoring
- Full documentation and reference example design

Virtex-5 GTP transceivers interface to the serial differential electrical signals of the PCIe specification. The PCIe block completes the physical logic that provides lane deskew. The DLL is responsible for data integrity and implements a user-configurable-sized retry buffer to retransmit packets that are received incorrectly without re-requests from the applications software. The TL provides Tx and Rx buffers and orders the packets to be transmitted. With capability for eight traffic classes and two virtual channels, great flexibility for packet arbitration is available.

High-Level Intregration

The Virtex-5 PCI Express Endpoint block allows you to implement a single endpoint device with one FPGA while leaving almost all of the FPGA programmable fabric avail-

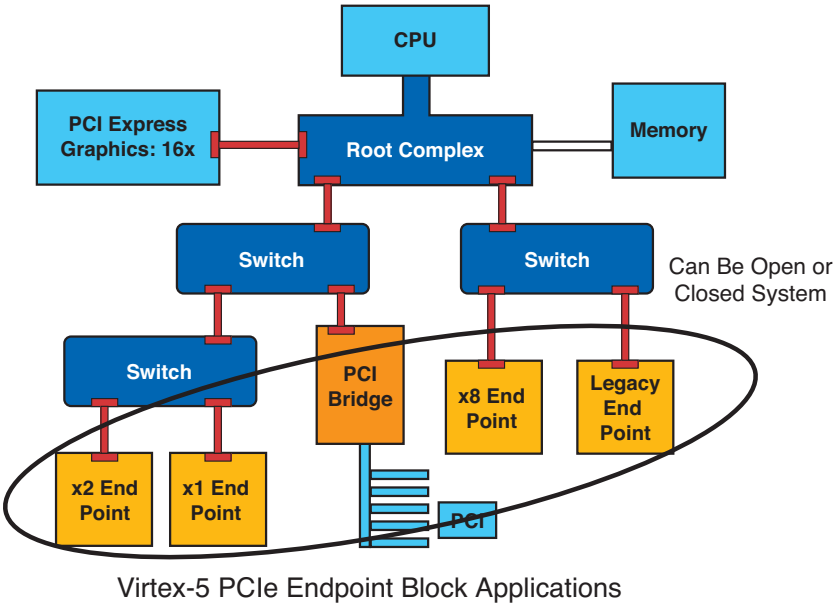

Figure 1 – PCI Express topology

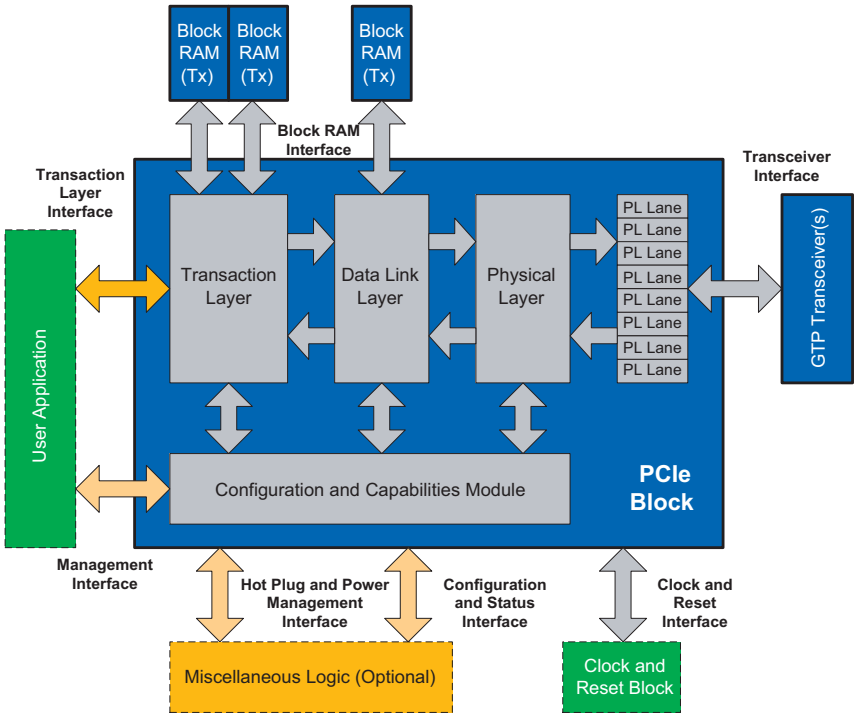

Figure 2 – Xilinx Virtex-5 LXT PCI Express Endpoint block

able for value-added endpoint application design functionality. The combination of the PCIe block, GTP transceivers, and the block RAM incorporate the majority of the logic required for a low-power, high-bandwidth, configurable PCIe endpoint port. The GTP transceivers support Gen1 2.5 Gbps serial rates while being electrically compliant to the PCIe specification. Some of the new transceiver features include power management support such as beacon and electrical idle detect and the spread-spectrum reference clock required in PC system motherboards.

The block RAM provides a scalable, user-configurable retry memory along with Tx and Rx FIFO for any packet size supporting one or two virtual channels. With complex configuration options of the PCIe block, GTP transceivers, and block RAMs, along with clock and reset logic, software automation provides quick and accurate configuration and interconnect of these functions.

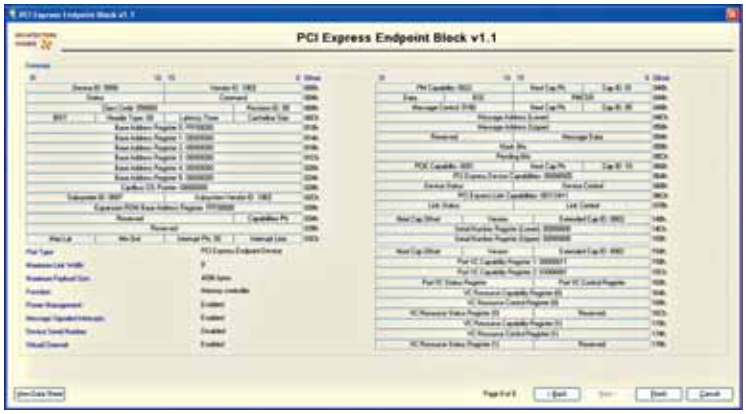

Figure 3 – Xilinx CORE Generator GUI

Figure 4 – ML523, ML555, and ML505 Virtex-5 LXT PCIe hardware reference boards

Using the PCI Express Endpoint Block

The Xilinx PCIe LogiCORE™ solution delivers wrappers through the CORE Generator™ software GUI flow of the ISE™ tool, which makes it easy to use the PCIe block and still provides full flexibility of the configurable features and capabilities of the high-bandwidth PCI Endpoint block. The configuration capabilities of the PCIe block are abstracted into several self-checking and instant-feedback menus that walk you through the configuration of key design elements (Figure 3).

The LogiCORE wrappers connect the PCIe block with the GTP transceivers and block RAMs. They also create and connect the clock and reset logic blocks to the PCIe block. You can customize the clock and reset RTL blocks to address advanced system requirements. In addition to connecting the various hardware resources to the PCIe endpoint block, the advanced wrapper provides value-added features such as the user-friendly interface similar to Xilinx LocalLink, the memory BAR, and non-memory TLP ID checking and filtering.

Ready to Provide High Bandwidth

The launch of the Virtex-5 PCIe block in the Virtex-5 LXT device not only provides the block and PCIe wrapper support, but also includes extensive system design aids. Memory endpoint and programmed input/output (PIO) reference designs are included in the deliverables. These designs serve as a training aid, quickly bringing up a simple user application to test in hardware.

In addition, the successful PCI-SIG compliance achieved with this design greatly speeds up the path to compliance for the users of the PCIe block in Virtex-5 devices. Xilinx provides a suite of hardware reference boards such as the ML523 characterization board, ML505 embedded design reference board, and the ML555 x8 high-data-bandwidth PCIe board to enable you to build and test PCIe systems (Figure 4).

Compliance testing requires a complete design with a demonstrable function, hardware board, software device driver, and application software to demonstrate interoperability with PC motherboard systems. The memory endpoint and PIO reference designs provide all of the above. Sample device drivers for Windows XP, Windows Server 2003, Windows Vista, or Linux are available by request. As the reference design function emulates a memory aperture, a simple test is provided to demonstrate the design operation.

This Virtex-5 LXT device and Xilinx reference hardware boards are listed on the PCI-SIG integrators list. For PCI-SIG-complaint Xilinx solutions, visit *www.pcisig.com/developers/compliance_program/integrators_list/pcie.*

Conclusion

The Virtex-5 LXT PCI Express Endpoint block, combined with the GTP transceivers and block RAMs, provide an extremely high level of integration for you to efficiently and quickly build high-performance, fully compliant PCIe systems in a single device. Xilinx developed and delivered the Virtex-5 PCI Express solution with a focus on end-user requirements for ease of use, legacy design migration, powerful yet flexible features and capabilities, system-level compliance, and cost.

For more information, visit *www.xilinx.com/virtex5.*

30% faster than last year's model...

Here are 6 of the new, faster, bigger, Virtex-5 FPGAs on a 12 Million ASIC Gate Board that offers unmatched performance to ASIC Prototypers, IP Designers, and FPGA Developers. The V5 65nm process, with 6 input LUT and advanced interconnect, enables 30% faster clock speeds in your application. The Dini DN9000k10PCI captures this performance on an easy to use board with these handy features:

- 33/66 MHz PCI bus or stand-alone operation
- 6 - DDR2 SODIMM Sockets
- 7 - Global Clock Networks
- 3 - 400pin FCI-MEG Connectors for daughter cards
- Easy configuration via CompactFlash, USB or PCI

All necessary operating software, including reference designs and Synplicity Certify™ models to simplify partitioning, is supplied with the board. If your need is speed — visit The Dini Group web site www.dinigroup.com for complete details on the fastest FPGA board ever.

The DiNI Group

1010 Pearl Street, Suite 6
La Jolla, CA 92037
(858) 454-3419
sales@dinigroup.com

PCI Express Markets, Trends, and Applications

The Virtex-5 LXT device's built-in PCI Express solution enables significant power and area savings.

by Navneet Rao
Technical Marketing Manager, Horizontal Platform Solutions
Xilinx, Inc.
navneet.rao@xilinx.com

End users are adopting multimedia-enabled devices rapidly; you need look no further than iPod video or YouTube-like blog sites. As users consume this type of rich data, the need for efficient storage and higher connectivity speeds becomes critical.

Today, the megahertz debate has been replaced with the gigabit-per-second debate, as the focus shifts from processing speed to high-speed interconnect. A host of serial standards have come into play. The key market requirements governing these standards are:

- Scalable performance
- An extensible feature set to adapt to various use models (chip-to-chip, backplanes, cable)
- Interconnects suitable for multiple market segments and applications
- Implementation of cost-effective solutions in mainstream high-volume technology

One of the key serial standards to emerge is PCI Express (PCIe), a third-generation I/O interconnect introduced in 2002 to provide a scalable path from PCI and PCI-X (see Table 1). PCIe has become the standard interconnect of the PC industry and is rapidly gaining momentum in other applications as well (Figure 1). It promises scalability, an extensible feature set, multiple market suitability, and cost-effectiveness.

Key highlights of PCIe include:

- A high-speed serial standard offering bidirectional communication at 2.5 Gbps line rates per lane
- Layered packet-based architecture, enabling modular design
- Bandwidth enhancement (as much as 80 GB) through easier scalability – 1, 2, 4, 8, 16, and 32 lanes
- Advanced features like reliability, power management, and hot plug
- Support for next-generation three-dimensional/multimedia traffic through virtual channels, traffic classes, and quality of service (QoS)
- Ease of use through new form factors and innovative designs, enabling applications targeted to multiple market segments
- Software preservation by supporting legacy PCI architecture and infrastructure

Tremendous acceptance, design wins, and strong customer feedback have propelled our understanding of the inherent benefits of PCI Express in our customers' applications. To create solutions for solving tomorrow's problems today and keep pace with rapidly changing times, Xilinx has introduced a hard PCI Express Endpoint block in its Virtex™-5 LXT devices (Figure 2).

The salient features of the PCIe Endpoint block are:

- Full-featured and compliant to PCIe base specification v1.1
  - Highly configurable PCIe end-point solution
- Passed compliance/interoperability tests at PCI plug-fest (*www.pcisig.com/developers/compliance_program/integrators_list/pcie*)
- Supports 1-, 2-, 4-, or 8-lane implementations
- Meets all key requirements
  - Electrical signaling
  - Protocol (CRC, automatic retry)
  - QoS
  - Hot-pluggable

PCI Specification	Bus Width	Transfer Rate	Lane Width	Line Rate	Max Data Bandwidth
PCI 1.0	32 bits	33 MHz			133 Mbps (half-duplex)
PCI 2.x	64 bits	33-66 MHz			266-533 Mbps (half-duplex)
PCI-X 1.x	64 bits	133 MHz			Up to 1 Gbps (half-duplex)
PCI-X 2.0	64 bits	266-533 MHz			Up to 4 Gbps (half-duplex)
PCI Express 1.x			1 lane	2.5 GHz	Up to 500 Mbps
			2 lane	2.5 GHz	Up to 1 Gbps
			4 lane	2.5 GHz	Up to 2 Gbps
			8 lane	2.5 GHz	Up to 4 Gbps
			16 lane	2.5 GHz	Up to 8 Gbps
			32 lane	2.5 GHz	Up to 16 Gbps
PCI Express 2.0*			1-32 lanes	5 GHz	Up to 32 Gbps

* The PCI Express 2.0 specification is "still under construction."

Table 1 – PCI/PCI-X/PCIe specification and bandwidths

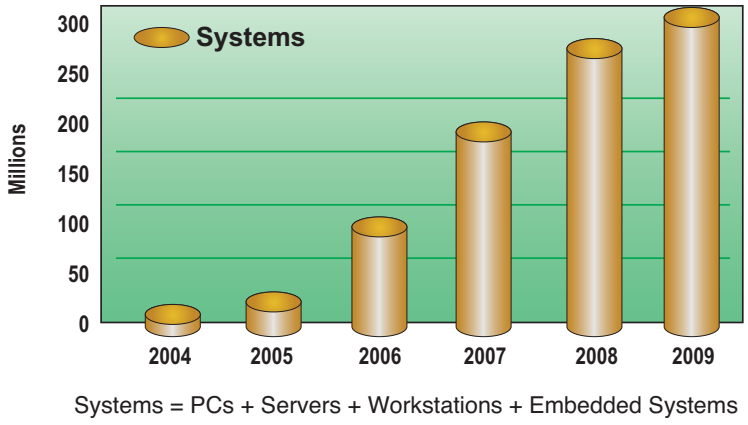

Figure 1 – PCI Express momentum

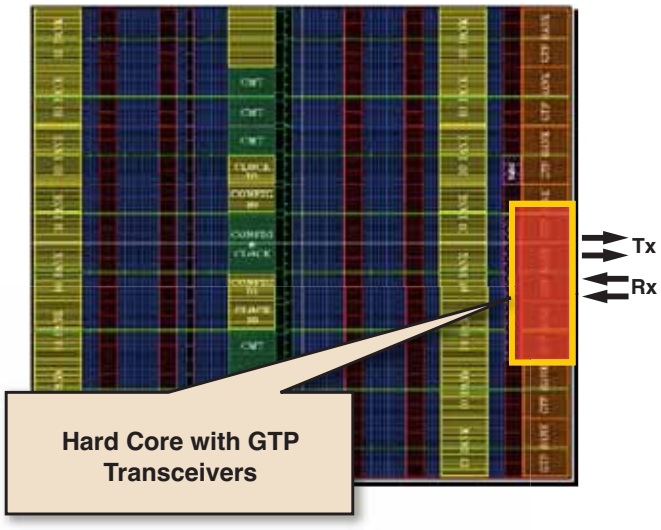

Figure 2 – PCIe Endpoint block in the Virtex-5 LXT FPGA

- Uses Xilinx® RocketIO™ GTP transceiver blocks
  - PCI Express electrical support
  - 100-MHz direct reference clock
- Saves resources
  - Integrated in all Virtex-5 LXT devices
  - Adjacent to GTP transceivers
- Ease of design
  - Shortens design cycles
  - Simplified, intuitive design flow
- Low cost and low power
- Packet buffering with configurable block RAM
  - Rx buffer
  - Tx buffer
  - Retry buffer
- Simple transaction layer interface for easy integration
- Signals available to fabric for statistics and monitoring
  - credit status, max payload size, error signals
- As many as two virtual channels for QoS
  - Round robin, weighted round robin, or strict priority

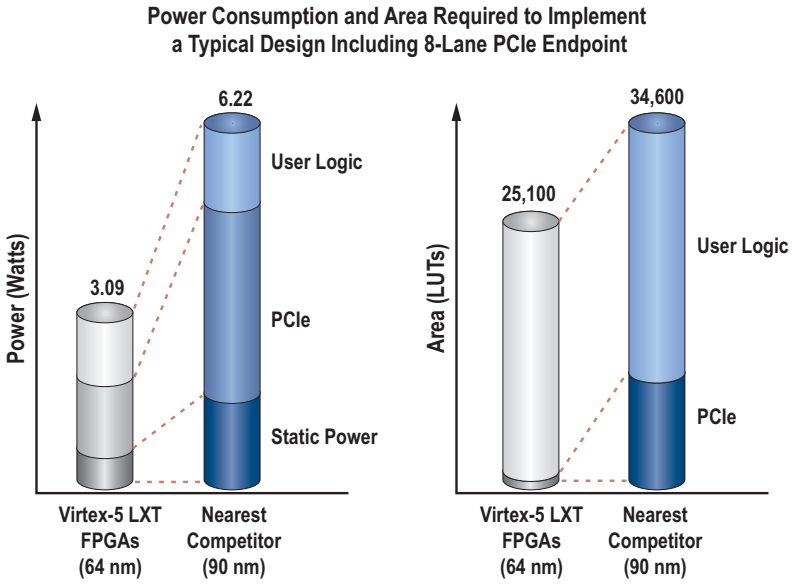

XC5VLX30T versus 2SGX60D. Target Frequency = 200 MHz. Worst-case process 25K LUTs, 17K flip-flops, 1 Mb on-chip RAM, 64 DSP blocks, and 128 2.5V I/Os. Based on Xilinx tool v8.2 and competitor tool v6.01.

Figure 3 – Power and area savings of a high-performance Virtex-5 LXT PCIe solution

Designing with the Virtex-5 LXT PCIe Block

PCI Express has gained considerable momentum, with broad acceptance in the PC industry. Engineers designing with Virtex-5 LXT FPGA-based PCIe endpoints can also lead the proliferation of PCI Express in new markets by leveraging these advantages:

- Faster time to market. Existing ASSPs do not support PCIe today; FPGAs enable bridging between proprietary parallel interfaces and PCIe. In addition, evolving add-ons to the PCI Express standard discourage ASIC/ASSP starts until a broad market base has been created. A case in point is the recent "Geneseo" architecture announcement by Intel and IBM at the Intel Developers Forum in September 2006. Xilinx also endorses this initiative, which extends the PCIe architecture to enable emerging application accelerators.
- Low power and reduced area. Applications that need higher performance but are designed to smaller form factors can use the Virtex-5 LXT solution (Figure 3). The PCIe Endpoint block allows you to be able to choose a smaller device and still achieve significant power and cost savings.
- Bridging legacy and disparate standards to PCIe. The movement of legacy applications to new form factors optimized for PCI Express requires bridging functions between legacy standards and PCI Express. The new Virtex-5 LXT platform offers the customization and logic resources to enable this transition as well as bridging to other serial standards.
- Scalable solution. The PCI Express protocol is here to stay, but the protocol itself and use models are in rapidly

Xilinx PCI History

Xilinx has been at the forefront of the PCI/PCI-X/PCIe technology. Significant achievements include:

1996 – Industry's first PCI core for FPGAs

1999 – Industry's first 64-bit, 66-MHz PCI solution

2000 – Industry's first 64-bit, 133-MHz PCI-X solution

2003 – Industry's first PCIe solution

2005 – Industry's first PCIe PIPE solution – Xilinx + NXP Semiconductors

2006 – Industry's first FPGA Express Card solution – Xilinx + NXP Semiconductors

2006 – Industry's first FPGA with built-in PCI Express Endpoint block

evolving phases. Designing with Virtex-5 LXT PCIe Endpoint blocks enables you to scale from 1- to 4- to 8-lane link-widths in the same Virtex-5 family. This allows you to future-proof the system and the equipment. In addition, because PCIe is inherently compatible with legacy PCI and PCI-X architectures, scaling and designing Virtex-5 LXT FPGA-based PCIe solutions will preserve software investments and extend infrastructure life.

- Form factors supported. Virtex-5 LXT RocketIO GTP transceivers offer significant power advantages over competing FPGA/ASSP solutions. This enables designers to consider Virtex-5 FPGAs in new markets. You can use the inherent advantages of the 65-nm FPGA to support multiple form factors by using scalable logic density for different solutions. For example, a desktop solution in an add-in card form factor can be scaled to a lower power Express Card form factor using similar/identical FPGA resources. Conversely, a desktop add-in card form factor PCIe solution in a Virtex-5 LXT FPGA can be easily scaled up to support the transition to high-performance form factor solutions like ATCA, uTCA, and server I/O module.

Applications	Form Factors	Link Width (Typical)	Data Bandwidth	Prominent Feature Required
Enterprise	HBAs, Server I/O Module	x1 x4 x8	250 Mbps 1 Gbps 2 Gbps	High Reliability Scalability Error Recovery Reduced Board Space Reduced Power Budgets
Desktop	Add-In Card	x1 x4 x8 x16	250 Mbps 1 Gbps 2 Gbps 4 Gbps	Legacy Support to Existing (PCI) Software Reduced Board Space Reduced Power Budgets Ecosystem Existence High Availability
Mobile	Express Card, Mini-Card	x1	250 Mbps	Reduced Power Budgets High Reliability Power Management Capability
Communications	HBAs, ATCA, Server I/O Module	x4 x8	1 Gbps 2 Gbps	High Availability High Performance Interoperability Reliability
Embedded Platforms	Integrated Endpoints, Custom Cards, Mini-Card	x1	250 Mbps	Low Cost Reliability High Availability Ease of Use and Integration

Table 2 – Virtex-5 LXT PCIe Endpoint applications

PCI ExpressFabric Topology

The PCI ExpressFabric™ topology, referred to as a hierarchy, comprises a root complex (RC), multiple endpoints (I/O devices), a switch, and a PCI Express/PCI bridge, all interconnected through PCI Express links.

An RC denotes the root of an I/O hierarchy that connects the CPU/memory subsystem to the I/O. A root complex may support one or more PCI Express ports, for example, Intel chipset(s).

A switch is defined as a logical assembly of multiple virtual PCI-to-PCI bridge devices, which forward transactions using PCI bridge mechanisms, namely address-based routing such as the IDT PCI Express switch.

Endpoint refers to a type of device that can be the requester or completer of a PCI Express transaction, either on its own behalf or on behalf of a distinct non-PCI Express device, for example, a PCI Express-attached graphics controller.

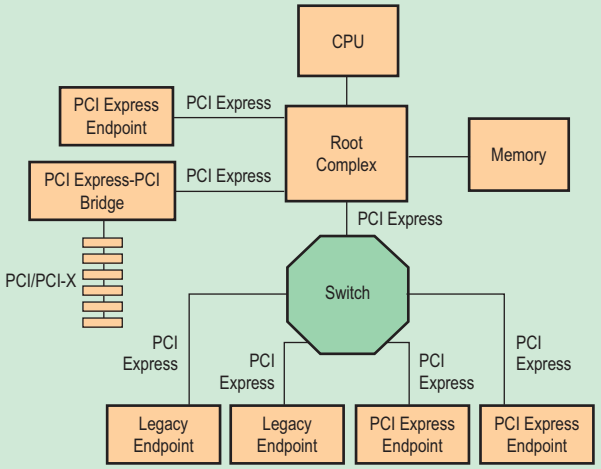

Figure 4 – PCI Express topology

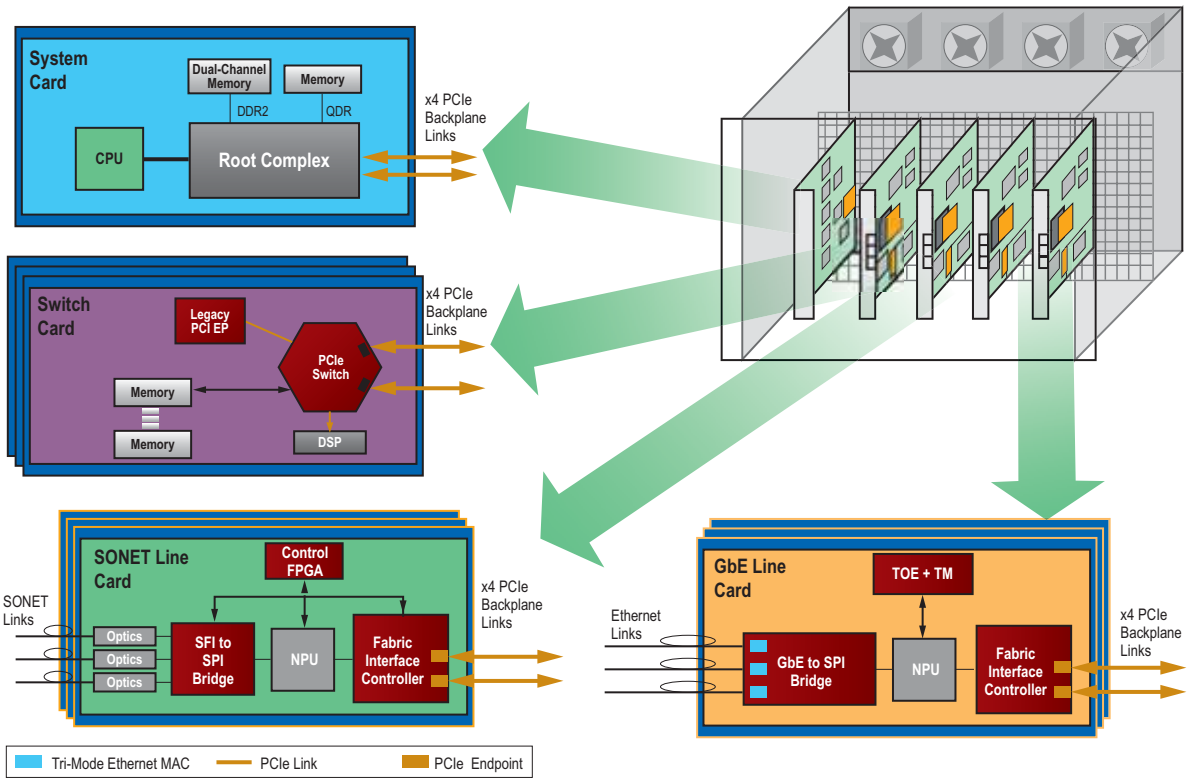

Figure 5 – PCIe in a communication system

Virtex-5 LXT FPGAs with built-in PCIe Endpoint blocks can easily be designed in all form factor applications, as shown in Table 2.

Figures 5 and 6 outline applications using Virtex-5 LXT PCIe Endpoint block capabilities to aggregate multiple-source traffic and bridge protocols to PCI Express.

Conclusion

The Virtex-5 LXT platform, with built-in PCIe Endpoint blocks and RocketIO GTP transceivers, offers great value by providing a full-featured, fully compliant PCIe solution. Say goodbye to IP licensing and hello to lower power and fewer utilized logic resources. You can achieve significant cost savings by targeting smaller FPGA devices with 50% of the power of soft-IP alternatives. Built-in hard blocks deliver guaranteed functionality and ease of use by reducing design time.

Thus, the Virtex-5 LXT platform offers unique built-in PCIe capabilities in a fast, low-power, 65-nm FPGA, launching a new era of efficient PCIe system development.

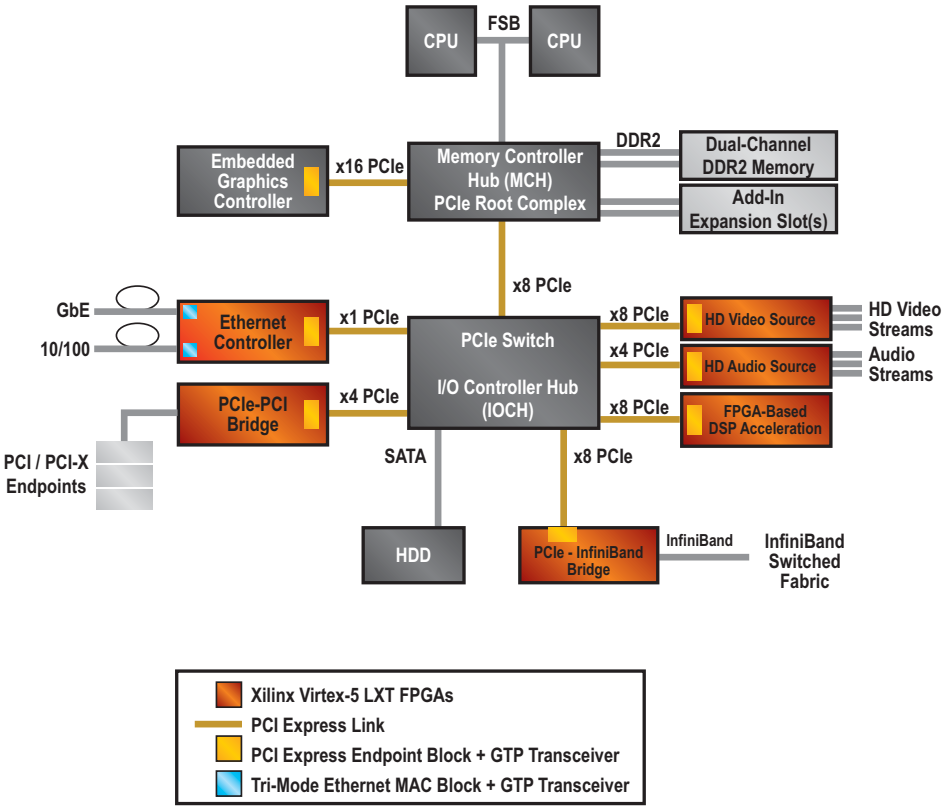

Figure 6 – PCIe in a high-end desktop/server system

Designing with Virtex-5 Embedded Tri-Mode Ethernet MACs

You can implement flexible Ethernet systems using the Virtex-5 10/100/1000 Ethernet MAC.

by Nick McKay
Senior Design Engineer
Xilinx, Inc.
nicholas.mckay@xilinx.com

Soma Potluri
Senior Design Manager
Xilinx, Inc.
soma.potluri@xilinx.com

Stuart Nisbet
Senior Design Manager
Xilinx, Inc.
stuart.nisbet@xilinx.com

Ethernet is the dominant wired connectivity standard. The Xilinx® Virtex™-5 Ethernet media access controller (Ethernet MAC) block provides dedicated Ethernet functionality, which together with Virtex-5 RocketIO™ GTP transceivers and SelectIO™ technology enables you to connect to a wide variety of network devices. The Ethernet MAC block is integrated into the FPGA as a hard block in Virtex-5 devices.

The Ethernet MAC is available in the Xilinx design environment as a library primitive, named TEMAC. The primitive contains a pair of 10/100/1000 Mbps Ethernet MACs. Each Virtex-5 LXT device contains four Ethernet MAC blocks; thus, a Virtex-5 LXT design can incorporate two TEMAC primitives. Using standard Xilinx products, you can create a range of customized packet processing and network end-point products. Xilinx has also provided an overclocking mode to enable backplane connectivity at speeds as fast as 2,000 Mbps.

Xilinx developed the Virtex-5 Ethernet MAC from the Virtex-4 FX Ethernet MAC, making improvements in the areas of global clock usage, serial interface flexibility, and software control complexity.

In this article, we'll review the feature set of Ethernet MAC blocks in Virtex-5 devices. We'll also describe the differences between Virtex-5 and Virtex-4 FX Ethernet MACs, illustrate some potential applications, and describe how to use standard Xilinx tools to integrate an Ethernet MAC into your design.

Supported Interfaces

The Virtex-5 Ethernet MAC is fully compliant to the IEEE802.3 specification. Figure 1 shows a block diagram of the Ethernet MAC.

Physical Interfaces

You can independently configure the physical interface of each Ethernet MAC to operate as one of five different Ethernet interfaces.

The Media Independent Interface (MII), Gigabit Media Independent Interface (GMII), and Reduced GMII (RGMII) are parallel interfaces. These are typically connected to an external physical layer (PHY) chip to provide BASE-T functionality at 10/100/1000 Mbps. Half-duplex operation is supported at 10/100 Mbps; full-duplex operation is supported at all speeds.

Serial GMII (SGMII) and 1000 BASE-X are serial interfaces that use the physical coding sublayer (PCS) and physical medium attachment (PMA) sections of the Ethernet MAC. These interface to the Virtex-5 RocketIO GTP serial transceivers. SGMII, as with the parallel interfaces, provides 10/100/1000 Mbps full-duplex BASE-T functionality. The serial interface significantly reduces the number of pins required to connect to the external PHY chip.

When the Ethernet MAC is configured in 1000 BASE-X mode, the PCS/PMA block, along with the RocketIO transceiver, provides all of the functionality required to connect directly to a gigabit interface converter (GBIC) or small form-factor pluggable (SFP) optical transceiver. This removes the need for an external PHY chip for 1000 BASE-X network applications.

Control Interfaces

The host interface provides access to the configuration registers of the Ethernet MAC block. Examples of configuration options include jumbo frame enable, pause and unicast address settings, and frame check sequence generation.

The host interface is accessible through either a generic host bus or a device control register (DCR) bus (when connecting to a processor). In addition, each Ethernet MAC has an optional management data

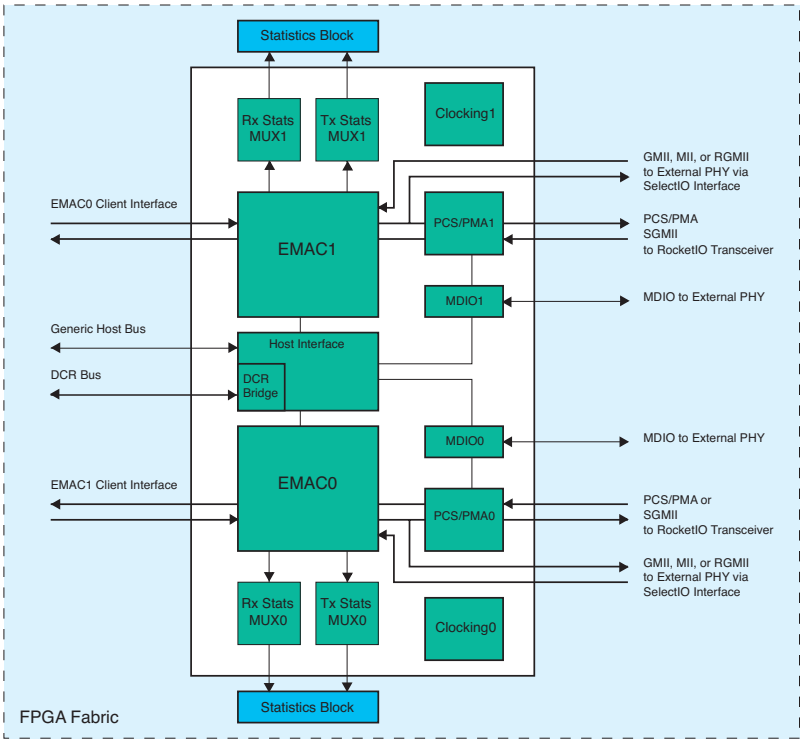

Figure 1 – Block diagram of the Virtex-5 Ethernet MAC

I/O (MDIO) interface. This allows access to the management registers of an external PHY and to the physical interface management registers within the PCS/PMA section of the Ethernet MAC.

Client Interface

Frames are passed to the Ethernet MAC across the transmitter client interface. The transmitter pads the incoming data when it is less than the minimum Ethernet frame length and maintains the minimum interframe gap between frames; however, you can increase the size of the gap. You can also configure the transmitter to add a frame-check sequence to the frame. A separate flow control interface allows you to generate pause frames. In half-duplex mode, the transmitter signals collisions and requests retransmissions for valid collisions.

The receiver interface verifies incoming frames and signals frame errors. Good and bad frame signals are provided. You can also configure the Ethernet MAC to pause and restart frame transmission upon the detection of valid pause frames.

The data on the client interface is 8 or 16 bits wide. The 8-bit interface is used for standard Ethernet applications, giving a 1,000 Mbps data rate with a 125 MHz clock. Using the 16-bit mode, you can increase the data rate to 2,000 Mbps without any increase in clock speed at the client interface.

Each Ethernet MAC outputs statistics vectors containing information about the Ethernet frames seen on its transmit and receive datapaths. An external statistics module is freely available in Xilinx CORE Generator™ software. The statistics module accumulates all of the Tx and Rx datapath statistics of each Ethernet MAC.

New Features in the Virtex-5 Ethernet MAC

In Virtex-4 FPGAs, implementing just the datapath consumes as many as four global clock buffers: one each for the Tx and Rx client interface logic, and one each for the Tx and Rx physical interface logic. For Virtex-5 FPGAs, Xilinx added a clock-enable feature. You can use the clocks derived for the physical interface for all of your client logic. The internally generated clock enable provides a way to maintain the correct data throughput on each of the interfaces. This reduces the number of necessary clock buffers by 50%.

DCR Bus Addressing

The Virtex-5 DCR interface now features an individual base address for each of the Ethernet MACs. This makes the shared DCR bus interface transparent to software drivers. The software no longer needs to know the bit locations for individual Ethernet MACs; the hardware automatically multiplexes in the correct bits depending on the base addresses.

Serial Interface Changes

Xilinx made several changes to the operation of the serial interfaces. Auto-negotiation is now more flexible with the inclusion of a programmable link timer. You can alter the timing of the auto-negotiation process and reduce simulation time.

A newly added unidirectional mode performs the unidirectional enable function from the IEEE802.3ah-2004 specification. When enabled, the Ethernet MAC transmits regardless of whether valid input is present at the receiver.

Finally, loopback can now take place in the Ethernet MAC as well as in the transceiver. This enables the transmission of idles to the link partner while in loopback, ensuring that the link remains active.

Virtex-5 Ethernet MAC Use Models

The versatility of the Virtex-5 Ethernet MAC enables its use in a wide variety of applications. For example, you can:

- Attach the Ethernet MAC to a processor running a protocol stack in network processing or remote monitoring systems, as shown in Figure 2.
- Interface the Ethernet MAC to a packet processing system implemented in the FPGA, such as a checksum offload engine or remote direct memory access design.
- Connect multiple Ethernet MACs to dedicated packet FIFOs and external memory for packet storage, bridging, or switching applications.

Tools and IP Support

Xilinx provides support for the Ethernet MAC through CORE Generator software, LogiCORE™ IP, and reference designs.

Virtex-5 Ethernet MAC Wrappers

Figure 3 shows a block diagram of the HDL wrappers available from the Xilinx CORE Generator tool.

The Ethernet MAC is a complex component with 162 ports and 79 parameters. Wrapper files enable you to easily set the parameters and interface only to those ports required for your application. They also offer benefits in simplifying the use of clocking and physical I/O resources. The different levels of hierarchy enable you to extract the correct wrapper for your application.

- Ethernet MAC Wrapper. In the lowest level, a single or dual Ethernet MAC is instantiated and its attributes are set to your preferred selection in the CORE Generator GUI. All of the unused input ports are tied to ground and the output ports are left open.

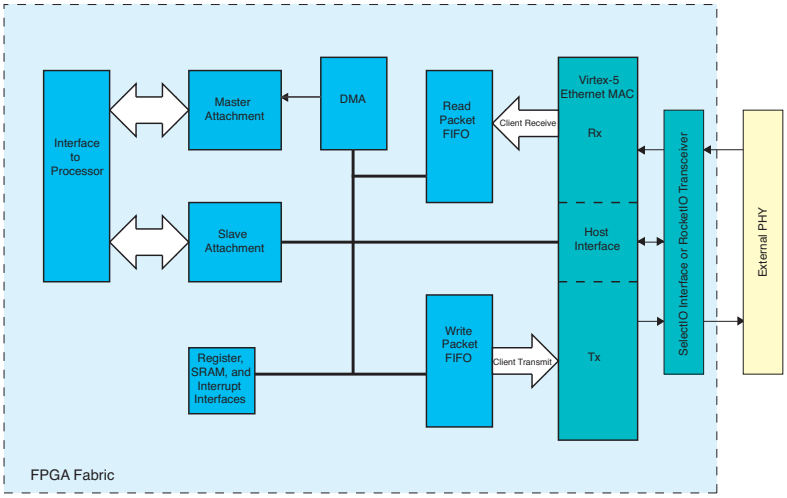

Figure 2 – MAC connected to a processor on the Virtex-5 FPGA

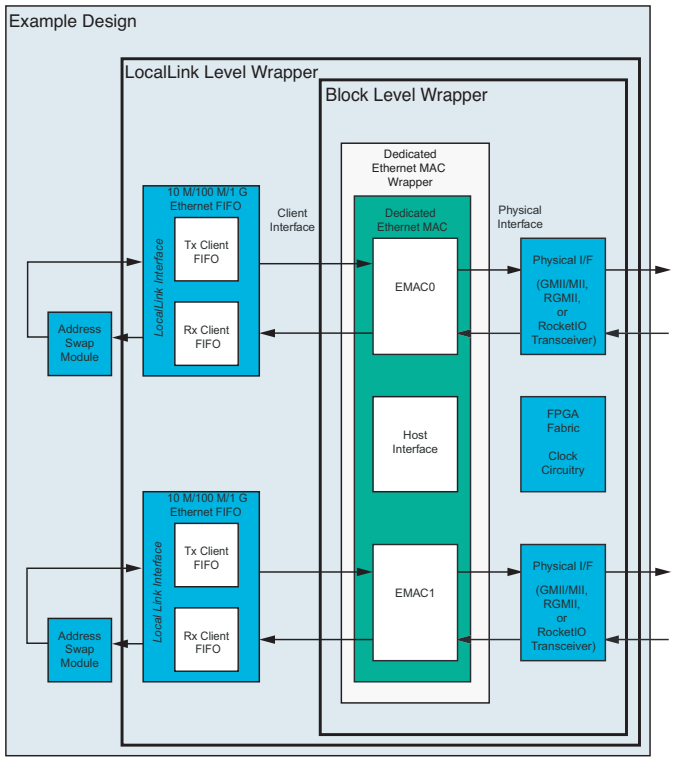

Figure 3 – Block diagram of the Virtex-5 Ethernet MAC wrappers

- Block Level Wrapper. In the next level of hierarchy, the physical interfaces and the required clock resources are instantiated. This includes the RocketIO GTP transceivers for the serial interfaces. Clocking is also optimized for your configuration, and you can clock the output to your design.
- LocalLink Level Wrapper. In this level, FIFOs are added to the client transmitter and receiver interfaces. The FIFOs handle the dropping of bad frames on reception and retransmission of frames in half-duplex mode. LocalLink is used as the backend interface.
- Example Design Wrapper. The top level features a demonstration design where the received data is looped back and sent to the transmitter. You can download this design to a board and stimulate the receiver from a network device to demonstrate the operation of the Ethernet MAC in hardware. Testbenches that stimulate receiver input and monitor the transmitter output of the design are also included in the CORE Generator software.

LogiCORE IP and Reference Designs

Most of the existing Virtex-4 Ethernet MAC documentation is reusable with the Virtex-5 Ethernet MAC. For example, a version of the "Ethernet Cores Hardware Demonstration Platform" (XAPP443, *www.xilinx.com/bvdocs/appnotes/xapp443.pdf*) will be available for the Virtex-5 Ethernet MAC. LogiCORE IP, such as Ethernet statistics, already supports the new architecture.

Conclusion

The Virtex-5 Ethernet MAC provides a cost-effective solution for a wide range of network interfaces, enabling you to connect to BASE-X and BASE-T networks at 10/100/1000 Mbps. Xilinx software tools and IP also allow you to take advantage of the improved feature set of the Ethernet MAC.

For more information, visit the Virtex-5 links on the Xilinx website, *www.xilinx.com/virtex5/*.

Asynchronous Sample-Rate Conversion Between AES Audio Streams

Xilinx Virtex-5 FPGAs provide the perfect platform for implementing AES digital audio sample-rate conversion.

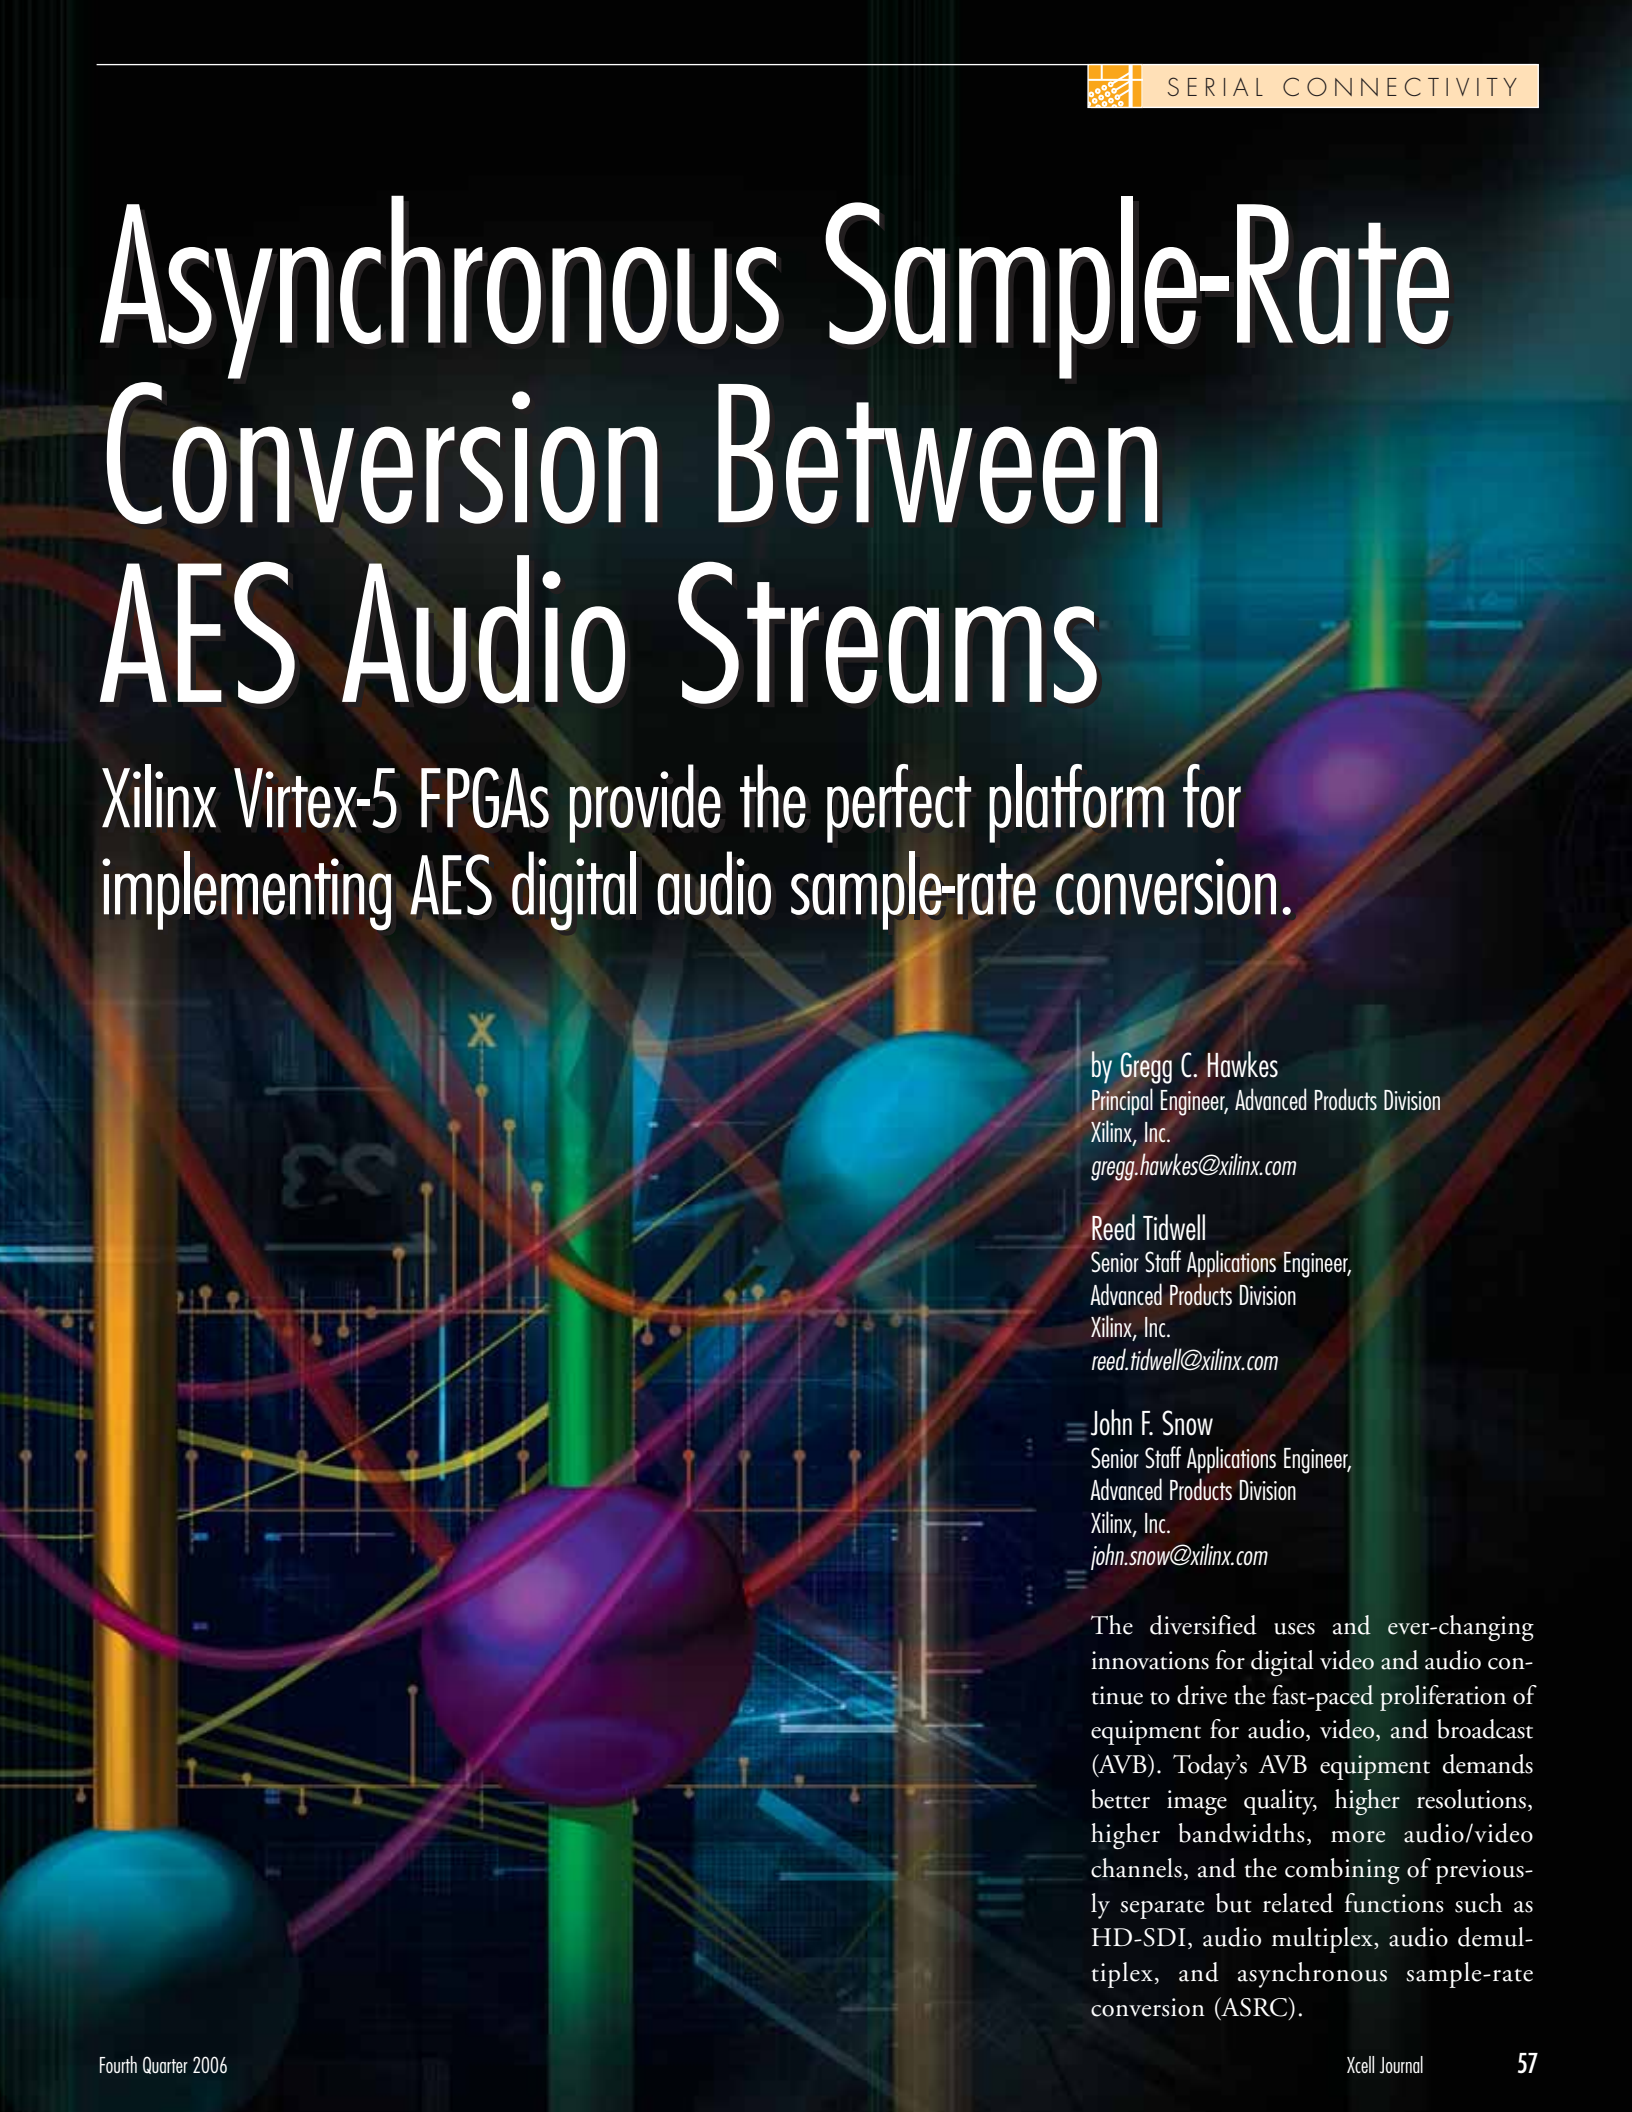

by Gregg C. Hawkes
Principal Engineer, Advanced Products Division
Xilinx, Inc.
gregg.hawkes@xilinx.com

Reed Tidwell
Senior Staff Applications Engineer,
Advanced Products Division
Xilinx, Inc.
reed.tidwell@xilinx.com

John F. Snow
Senior Staff Applications Engineer,
Advanced Products Division
Xilinx, Inc.
john.snow@xilinx.com

The diversified uses and ever-changing innovations for digital video and audio continue to drive the fast-paced proliferation of equipment for audio, video, and broadcast (AVB). Today's AVB equipment demands better image quality, higher resolutions, higher bandwidths, more audio/video channels, and the combining of previously separate but related functions such as HD-SDI, audio multiplex, audio demultiplex, and asynchronous sample-rate conversion (ASRC).

Xilinx® FPGAs have kept pace with customer integration needs by incorporating silicon features, facilitating the absorption of less integrated, complex, and expensive ASSP chips. One such ASSP chip function, ASRC, can be integrated into Xilinx FPGAs by leveraging diffused silicon features known as DSP48E slices and block RAMs to build sophisticated filter functions.

Free Xilinx application notes and reference designs have also kept pace with our customers' need to integrate sophisticated algorithms. The ASRC reference design correctly handles synchronous sample-rate conversion and the far more complex ASRC called for in most audio/video applications.

Simpler "synchronous-only" methods, offered by many ASSP chips and FPGA IP suppliers, can be smaller in terms of utilization per audio channel; however, when applied incorrectly to asynchronous applications, these methods have one or both of the following artifacts:

- The input-to-output latency changes because of accumulating delay
- Artifacts are produced in the audio, such as skipping samples or repeating samples

Both of these cases represent undesirable distortions.

Understanding Sample-Rate Conversion

Before diving into the theory of digital sample-rate conversion, you should look at the basic types of problems audio/video engineers are trying to solve. A few applications exist where you could use a fixed-rate synchronous conversion, such as a 48-KHz input converted to a 44.1-KHz output using the same clock source, or an output clock derived from the input clock. However, more likely is the asynchronous case, where input and output clocks are completely independent, such as two boards communicating audio between them. The different clock oscillators can be the same nominal frequency but several parts-per-million different.

The Xilinx ASRC reference design for the asynchronous case of independent input and output clocks provides two important and difficult design functions:

Figure 1 – ML571 board and frame synchronization demonstration board with an ASRC to match the output digital audio rate to the output digital video rate.

- Automatically and accurately monitoring the input-to-output ratio and sample-rate changes
- Adapting the filter function (filter coefficients) on the fly to provide the highest performance

Supporting ASRC for digital audio with an FPGA means that you can now significantly save costs for every SDI interface in your system – and in many systems, there are many channels.

The Xilinx ASRC IP is very high performance, with a worst-case input-to-output signal-to-noise ratio of -125 dB. It also supports conversion for multiple audio-input frequencies to multiple audio-output frequencies. The rate conversion algorithm adjusts on the fly, maintaining high performance with no special attention needed for input and output clocks. You can verify all of this with the IP running the Xilinx ML571 Serial Digital Video demonstration board shown in Figure 1.

Best of all, the broad functionality and high-performance ASRC IP is free.

Sample-Rate Conversion Theory

Figure 2 illustrates conceptually the general case of up or down conversion. The conversion ratio can vary continuously by rational numbers with fractional values. The diagram shows the up conversion process (creating many more samples and time positions to choose from) followed by down conversion (judiciously choos-

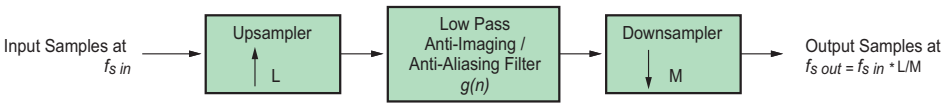

Figure 2 – Classic conceptual data path for sample-rate conversion

The ASRC adjusts the de-embedded audio to match the output video stream clock rate, where it can then be re-embedded into the output SDI video stream.

ing the samples in the output datastream that most closely match the positions of the desired samples). The anti-imaging/anti-aliasing filter in the center of the data path ensures that the spectral content is less than half the Nyquist rate of both the input and output sampling frequencies.

Figures 3 and 4 show that for every output sample location or output phase, a different set of sub-filter coefficients is required because the inputs are in different locations relative to that output phase. The sub-filter, having a set of coefficients that align with the input sample positions, is formed by interpolating the prototype filter coefficients. When this sub-filter is convolved with the corresponding input samples, the output sample of interest is produced. This process repeats, with new sub-filter coefficients interpolated for each output sample.

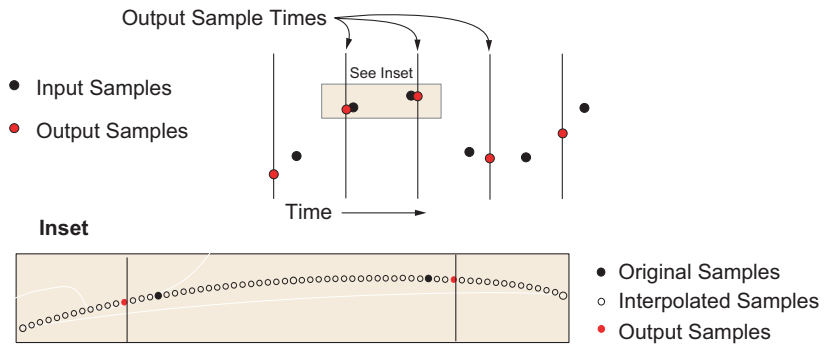

Figure 3 – Output sample position relative to the original sample position dictates which interpolated samples to use.

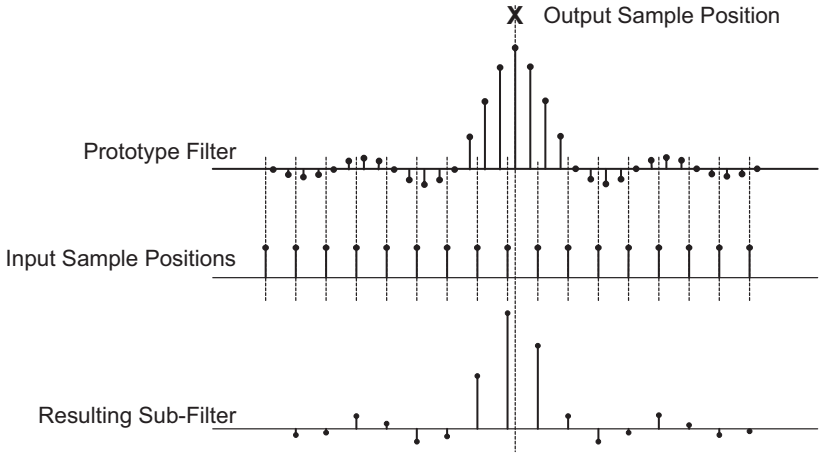

Figure 4 – Prototype filter centered at output sample position

ASRC Example Implemented on the ML571

The simple function known as frame synchronization of video provides a great demonstration of where you might use ASRC. Video can be stored in a frame buffer at some rate and removed at a fractionally different rate. This process can be useful if two pieces of video equipment are not "genlocked" together and operate at different pixel rates.

The result is the occasional need to add or drop a frame of video data. Your eye probably would not notice an added or dropped frame of video on your TV screen, but the human ear is very good at detecting discrepancies in added or dropped audio. The solution is to remove the audio from the starting video stream and reinsert it in the resulting video stream at a fractionally different rate, matching the output audio rate to the new output video rate. The Xilinx ASRC reference design is perfect for this task.

As an example, let's connect two boards with SDI video running at slightly different frequencies because of the different clock oscillators on each board. The receiving board demultiplexes the embedded AES digital audio from the video stream and sends it to the ASRC. The difference in clock frequency between the two boards causes the frame buffer synchronization logic to add or drop video frames. The ASRC adjusts the de-embedded audio to match the output video stream clock rate, where it can then be re-embedded into the output SDI video stream. The difference in clock frequencies between the two boards causes the frame buffer synchronization logic to add or drop video frames. The ASRC adjusts the de-embedded audio to match the output video stream clock rate, where it can then be re-embedded into the output SDI video stream.

For more information about frame buffer synchronization and asynchronous sample-rate conversion techniques, see XAPP514, "Audio/Video Connectivity Solutions for the Broadcast Industry," at *www.xilinx.com/bvdocs/appnotes/xapp514.pdf*.

Block Diagram and Specification Highlights

The simple diagram in Figure 5 illustrates two key design elements required in ASRC. The first element is determining the changes between input sample rate and required output sample rate, labeled "ratio control." The second element within the "re-sampler" is a set of prototype filters that are modified depending on the statistics reported by the ratio control.

The ASRC reference design converts stereo audio from one sample frequency to another. The input and output sample frequencies can be an arbitrary fraction of one another or the same frequency, but based on different clocks. The output is a band-limited version of the input re-sampled to match the output sample timing. The reference design has the following features:

- Fully asynchronous operation
- Expandable to multiple channels
- A -125 dB THD+N worst case with -130 dB THD+N typical
- A 24-bit audio word width in and out, with 31-bit internal math precision and round away from zero
- Automatic input-to-output sample ratio monitoring with continuous filter modification
- Continuous rational/fractional ratio, up conversion, 8:1
- Continuous rational/fractional ratio, down conversion, 1:7.5
- Continuous input-to-output rate monitoring with adaptive filtering
- Input/output rates 8 KHz-192 KHz, continuous
- Low deterministic latency

The reference design has an interpolated coefficient FIR filter coded with Virtex™-5 DSP48E slices as the primary math element and block RAM for input sample buffers and prototype storage.

Conclusion

The need to maintain different input-to-output audio rates for varying numbers of digital audio channels and support new AVB functions is a tremendous challenge. Throw in varying protocols, memory management, different sized payloads, and a variety of different system interfaces, and it is easy to see how these designs require high-performance, cost-effective flexibility that ASSPs and ASICs cannot offer. These challenges open up opportunities for Virtex-5 devices because these devices can enable equipment vendors to provide solutions to the ever-evolving AVB equipment landscape.

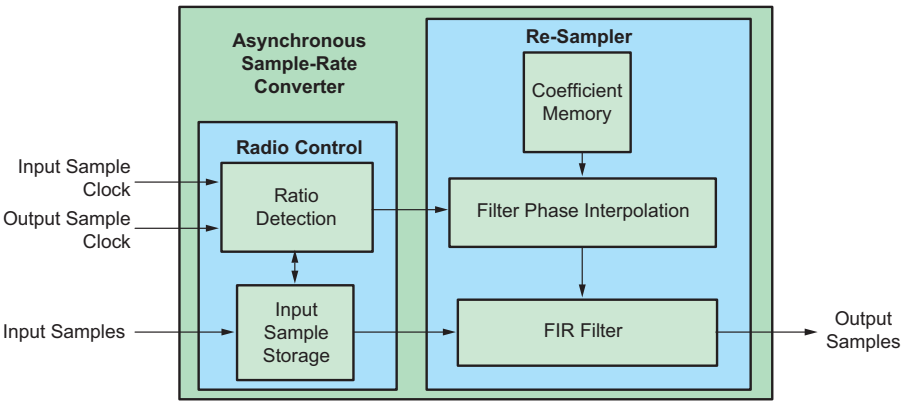

Figure 5 – Xilinx ASRC reference design top-level block diagram

Implementing Integrated Video Connectivity Solutions with Virtex-5 LXT Devices

Xilinx Virtex-5 FPGAs provide the perfect platform for integrating broadcast video solutions inside a single chip.

by Gregg C. Hawkes
Principal Engineer, Advanced Products Division
Xilinx, Inc.
gregg.hawkes@xilinx.com

Reed Tidwell
Senior Staff Applications Engineer,
Advanced Products Division
Xilinx, Inc.
reed.tidwell@xilinx.com

John F. Snow
Senior Staff Applications Engineer,
Advanced Products Division
Xilinx, Inc.
john.snow@xilinx.com

At Xilinx, we understand the challenges that broadcast system designers are facing. The number of emerging new standards for video connectivity creates difficult design challenges and schedules for broadcast products.

With the ever-changing video connectivity landscape prevalent throughout the broadcast chain, our goal is to offer help in the form of free reference designs, forming drop-in building blocks that can solve many system-level video connectivity issues. By providing you with cost-effective and highly integrated solutions compared to ASSP chips, Xilinx hopes to get you to market faster, lower costs, and differentiate your product from the competition.

Our video connectivity IP and reference design book, "Audio/Video Connectivity Solutions for the Broadcast Industry" (*www.xilinx.com/bvdocs/appnotes/xapp514.pdf*), includes chapters about SDI, HD-SDI, DVB-ASI, AES embedded audio, and audio-asynchronous sample rate conversion. Each chapter describes a specific video connectivity topic and links to free reference designs in Verilog and VHDL, providing implementation examples.

Integrating the encoders and decoders for these standards into the FPGA is simple with the clear, concise reference material found within the chapters of XAPP514. The reference design code, offered in both Verilog and VHDL, is clearly documented and illustrated, as shown in Figure 1.

We also offer a suite of validation platforms that can quickly and easily test your video processing algorithms or verify connectivity performance. For example, you can use our new Xilinx® Virtex™-5 ML571 Serial Digital Video (SDV) board (*www.cook-tech.com*) to demonstrate or develop video connectivity with Virtex-5 FPGAs. Figure 2 shows a block diagram; Figure 3 is a photograph of the ML571 board. Many of the free reference designs linked to XAPP514's chapters were verified on the ML571 platform using broadcast industry-standard test equipment.

"The ML571 board is yet another example of how Xilinx provides customers with detailed design assistance for real broadcast industry issues," said Andy DeBaets, senior director, systems and application engineering at Xilinx. "This board demonstrates how engineers can easily implement advanced video networking protocols while greatly increasing system integration, reducing system costs, lowering power, and shortening design schedules."

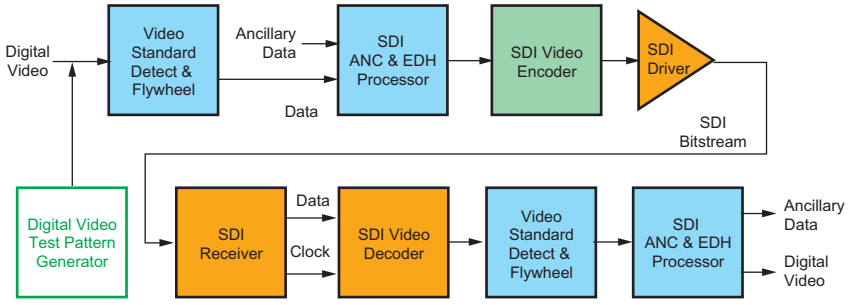

Figure 1 – Example block diagram of free modular Verilog and VHDL reference designs

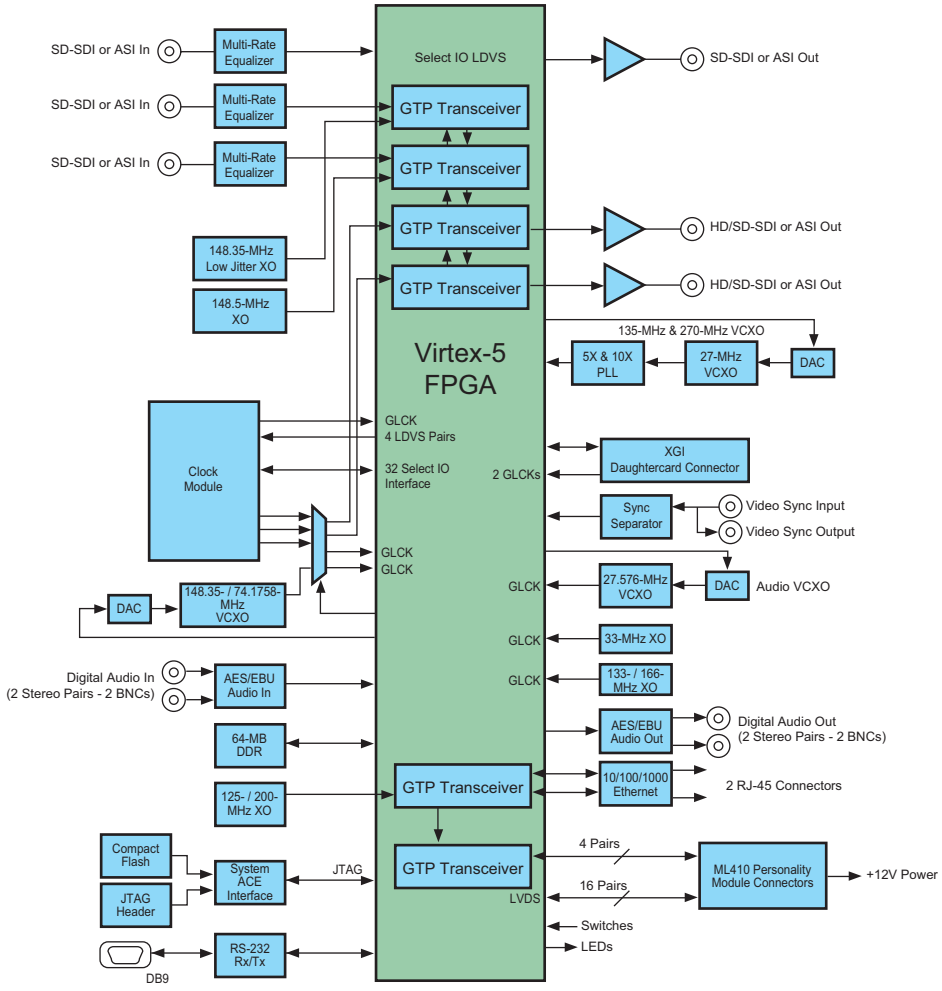

Figure 2 – Xilinx ML571 SDV video connectivity board block diagram

Talk to your Xilinx sales channel about seeing the demonstrations or obtaining one of these boards so that you can test your new algorithms long before your proprietary board is produced. We hope you find this article and the audio/video connectivity book valuable, but it represents just a small sample of the information available about designing with Xilinx programmable logic devices. To access the latest information on these subjects and more, visit *www.xilinx.com/esp/broadcast*.

Virtex-5 Features Support Broadcast Designs

The Virtex-5 feature set supports many aspects of broadcast solutions by providing high performance, flexibility, and scalability with unique, cost-optimized family members built on the following features:

- High-density, high-speed, reprogrammable ExpressFabric™ technology
- 550-MHz, 36-Kb, dual-port block RAM/FIFO
- 550-MHz, 25 x 18 DSP48E slice
- 550-MHz clock management tile (CMT)
- SelectIO™ technology
- Reduced power consumption
- Sparse chevron package

These features are described throughout the articles in this issue of *Xcell Journal*, with detailed descriptions of the features and performance at *www.xilinx.com/products/silicon_solutions/fpgas/virtex/virtex5/index.htm*.

Overview of the Xilinx ML571

The new serial digital video (SDV) board for demonstrating and testing high-bandwidth video communications channels based on Xilinx Virtex-5 platform FPGAs shows you how to easily implement high-speed serial interfaces to popular industry standards like HD-SDI.

Standards and Functionality Supported

The diffused silicon integration of high-performance and low-power multi-gigabit serial I/O, tri-mode Ethernet MACs, PowerPC™ processor, and PCI Express Endpoint block

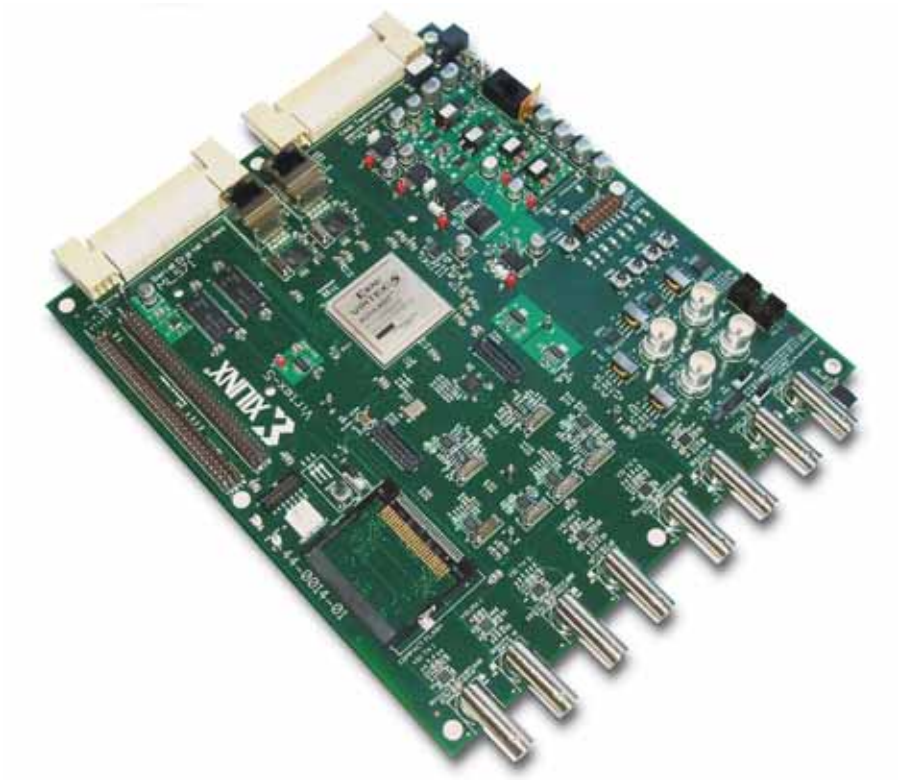

Figure 3 – Xilinx ML571 SDV video connectivity board

into Virtex-5 platforms has enabled the support of many more networking standards than previously possible.

The ML571 board now supports:

- Virtex-5 XC5VLX50T-FF1136 FPGAs (LX110T offered in a pin-compatible package)
- Two RocketIO™ GTP HD/SD-SDI receivers and two RocketIO GTP transmitters. The transmitters have Gennum tri-mode, 3 Gbps-capable cable drivers and the receivers have Gennum tri-mode, 3 Gbps-capable receiver equalizers. The standards supported are:
  - 3 Gbps HD-SDI (SMPTE424M), 2.97 Gbps
  - HD-SDI dual link (SMPTE372M) 1.485 Gbps, 1.4835 Gbps
  - HD-SDI (SMPTE292M) 1.485 Gbps, 1.4835 Gbps
  - SD-SDI (SMPTE 259M), 270 Mbps
  - DVB-ASI (CENELEC EN 50083-9 Annex B), 270 Mbps
- A SelectIO video input and video output providing differential LVDS I/O. This demonstrates the ability of the Virtex-5 SelectIO interface to transmit and receive video bitstreams supporting the following video standards:
  - SD-SDI (SMPTE 259M), 270 Mbps
  - DVB-ASI, 270 Mbps
- Select IO technology, LVDS, AES3 digital audio (AES3id) I/O. Two BNC input connectors provide two stereo pairs of AES3id digital audio in. Two BNC output connectors provide two stereo pairs of AES3id digital audio out. These inputs meet the SMPTE 276M 75-Ohm unbalanced AES3 audio input electrical specifications.
- SDI AES digital audio, embed and de-embed (SMPTE272M-2004)
- AES digital audio, high-performance, asynchronous sample-rate conversion (ASRC)
- DVB-ASI to/from Ethernet for video over IP
- Frame synchronization using external DDR DRAM
- Sync separator and genlock capability. A sync separator can accept a variety of video sync sources including bi-level and tri-level video sync (HD and SD). The separated sync signals from the sync separator go to the FPGA, where they can be used to build genlock PLLs using any of the VCXO clock sources available to the FPGA.
- An XGI-compatible expansion connector set is provided to allow video I/O daughtercards
- Two 10/100/1000 Ethernet interfaces
- Debug RS-232 serial port
- Configuration six-pin JTAG header for connection to a Xilinx download cable
- Xilinx System ACE™ configuration controller with a CompactFlash Type II socket

Conclusion

The need to support new AVB designs and assist you, our customers, with implementations in Xilinx FPGAs is a tremendous challenge. However, at Xilinx, we pride ourselves in striving to keep up with the demand for excellence. With varying protocols and a variety of different system interfaces, it is easy to see how these designs require high-performance, cost-effective flexibility that ASSPs and ASICs cannot offer. These challenges open up opportunities for Virtex-5 devices, for these devices can enable you to provide solutions to the ever-evolving AVB equipment landscape.

The ML571 board is designed and sold by Cook Technologies. The Cook Technologies part number for the ML571 is CTXIL406. There are many clock and connectivity option daughtercards that also plug into the ML571 SDV board. For more information, e-mail *colin@cook-tech.com* or visit their website at *www.cook-tech.com.*

Enhancing System Management and Diagnostics with the Virtex-5 System Monitor

You can use the Virtex-5 System Monitor to greatly increase environmental monitoring coverage of your FPGA design.

by Anthony Collins
Staff Product Marketing Engineer
Xilinx, Inc.
anthony.collins@xilinx.com

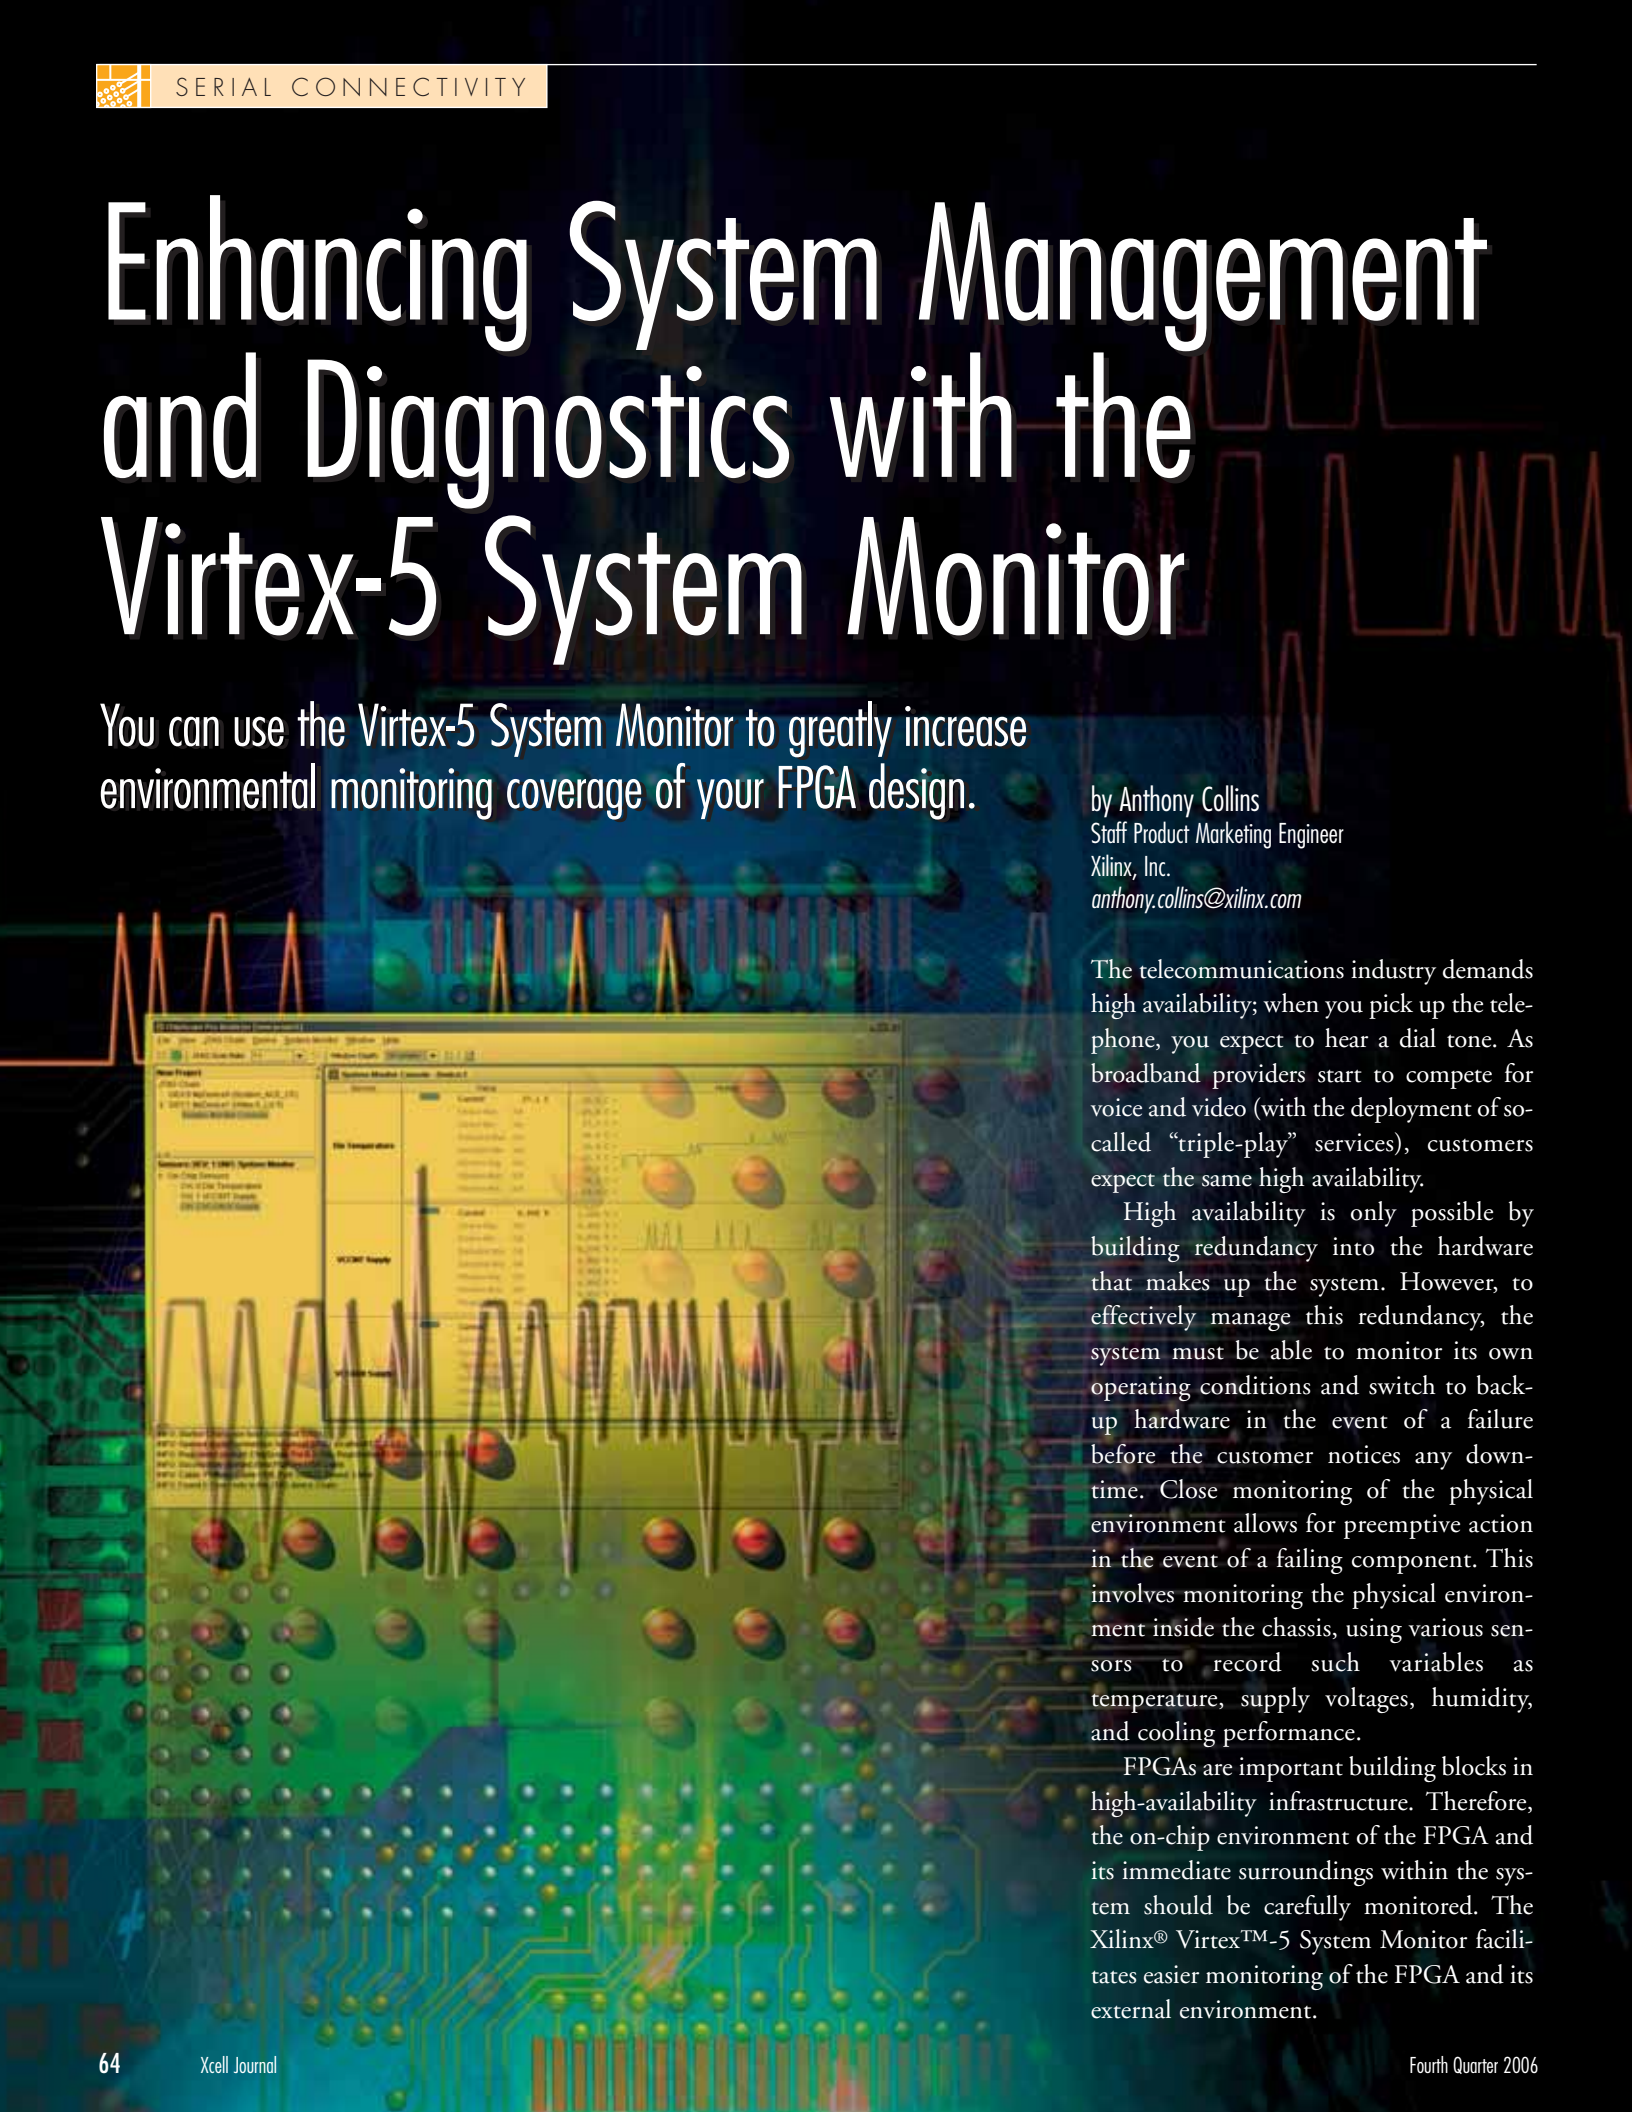

The telecommunications industry demands high availability; when you pick up the telephone, you expect to hear a dial tone. As broadband providers start to compete for voice and video (with the deployment of so-called "triple-play" services), customers expect the same high availability.

High availability is only possible by building redundancy into the hardware that makes up the system. However, to effectively manage this redundancy, the system must be able to monitor its own operating conditions and switch to back-up hardware in the event of a failure before the customer notices any down-time. Close monitoring of the physical environment allows for preemptive action in the event of a failing component. This involves monitoring the physical environment inside the chassis, using various sensors to record such variables as temperature, supply voltages, humidity, and cooling performance.

FPGAs are important building blocks in high-availability infrastructure. Therefore, the on-chip environment of the FPGA and its immediate surroundings within the system should be carefully monitored. The Xilinx® Virtex™-5 System Monitor facilitates easier monitoring of the FPGA and its external environment.

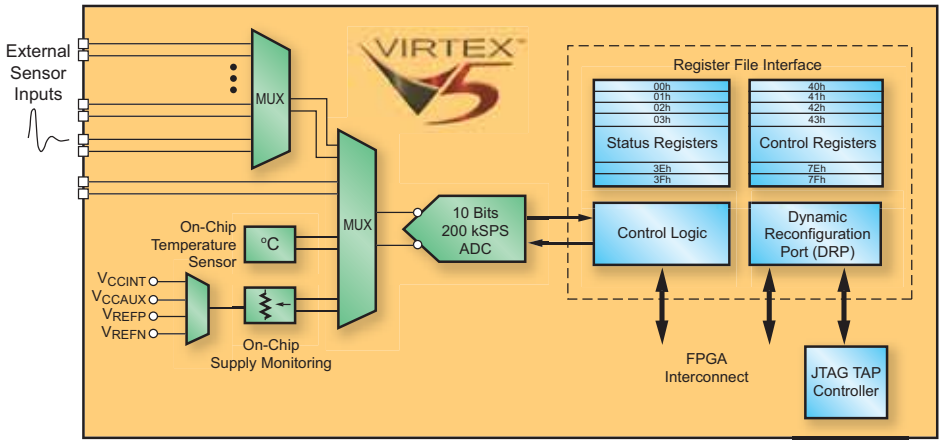

Figure 1 – Virtex-5 System Monitor

Virtex-5 System Monitor

The Virtex-5 System Monitor allows you to easily access information about the FPGA on-chip (die) temperature and power supply conditions. The system monitor also provides access to external sensor information through external analog input channels (monitoring as many as 17 external sensors). Access to this information involves little or no design effort, depending on the required functionality. Common functionality like alarms, automatic channel sequencer, and data averaging are available within the System Monitor block, enabling you to develop a solution easily.

Figure 1 shows a block diagram of the Virtex-5 System Monitor. The system monitor is built around a 10-bit, 200 kilosample-per-second analog-to-digital converter (ADC). The analog input range of the ADC is 0V-1V. At a resolution of 10 bits, the ADC can resolve an input voltage to an accuracy of approximately 1 mV.

As shown in Figure 1, both the on-chip sensors and external analog input channels are connected to the ADC input using analog multiplexers. Therefore, the output voltages from various sensors must be sequentially converted to a digital word by the ADC. These measurement results are written to status registers, where they are easily read using the FPGA fabric, or externally through the FPGA and PC board JTAG infrastructure. The System Monitor control registers can be written or read using the same interfaces. The control registers configure the System Monitor operation (for example, selecting sensor channels for measurement, program alarm limits, and sensor averaging). The System Monitor is fully functional shortly after power-up and does not require the FPGA to be configured for correct operation. By default, only the on-chip sensors are monitored after power-up; however, you can also enable external analog inputs. Measurement information can only be accessed through the JTAG test access port (TAP) before configuration.

User Alarms

One of the useful built-in features of the System Monitor is its ability to generate alarm signals for the on-chip sensors. As a designer, you can specify the threshold limits for these alarm signals. The System Monitor can autonomously monitor the sensors and alert the system only when an alarm condition is detected.

The System Monitor also contains a factory-set alarm condition called over temperature (OT). If you enable this feature, the System Monitor can request a full chip power-down if a die temperature greater than 125° C is detected. Chip power-up is initiated after the die has cooled to a level that you specify. The System Monitor continues to operate and monitor the on-chip sensors during chip power-down.

The OT functionality is disabled by default and must be explicitly enabled.

Checking the Checker

Using the Virtex-5 System Monitor to provide accurate and reliable environmental information requires reliability checks on the measurement data and system monitor operation. The System Monitor has a number of features that help to confirm reliable operation. Built-in auto-calibration of the ADC and sensors correct any drift in the analog measurement system because of the operating environment. Self-check features also allow the system host to monitor the operation of the System Monitor.

Leveraging System Monitor JTAG Access

A novel feature of the Virtex-5 System Monitor is the ability to access the full functionality of the block using the JTAG TAP. By enabling analog testing and access to analog information, you can obtain greater value and efficiencies using the existing JTAG infrastructure in the system. This access is available before configuration of the FPGA for use as part of a PC board test scheme in production, or during normal operation to facilitate a debugging effort.

To facilitate off-chip measurements such as supply voltages and currents on the PC board, you can use special JTAG commands to enable external analog inputs before FPGA configuration. Even after FPGA configuration, the System Monitor does not require an explicit instantiation in your design, thereby allowing full access to its features for debugging work through the JTAG TAP, even at a late stage in the design process. To ensure the availability of the System Monitor, the only requirement is that the correct PC board support must be in place. This involves the connection of an external 2.5V reference IC as described in the System Monitor User Guide (*www.xilinx.com/bvdocs/userguides/ug192.pdf*).

Figure 2 illustrates a typical diagnostic application where the physical operating environment of the FPGA is monitored during normal operation. In the example illustrated in Figure 2, the System Monitor is used to look at the voltage (IR) drop in the power distribution system (PDS) during a period of heavy current demand starting at time t0. The temperature of the FPGA is also monitored during this period

of high activity. Potential issues with the power supply or PC board design can be quickly identified during development. The JTAG access also provides an easy way to confirm that adequate cooling is in place for a particular design. The ChipScope™ Pro Analyzer provides an easy way to access the System Monitor; however, access can easily be incorporated into other JTAG test and programming environments.

System Integration

In addition to convenient access through the JTAG TAP, full access to the System Monitor control and status registers is also provided through the FPGA fabric. These registers can be configured and read at any time from the fabric. Dual access to the System Monitor registers by the JTAG TAP controller and fabric interface is permitted, and an arbitration scheme is available to manage possible contention.

You can also define the contents of these registers when the System Monitor is instantiated in a design and initialized during FPGA configuration. Thus, the System Monitor can be configured to start up in a user-defined mode of operation post-configuration. The fabric interface is known as the dynamic reconfiguration port (DRP). The DRP is a parallel 16-bit synchronous data port (similar to block RAM).

For more advanced applications where greater control over the System Monitor is required, the DRP allows the System Monitor to be easily mapped into the peripheral address space of a hard or soft microprocessor. Figure 3 illustrates a typical system management application where the MicroBlaze™ processor is running a protocol-like intelligent platform management interface (IPMI) and communicating with the system host over management channels like Ethernet or even a simple UART/modem.

The System Monitor also provides an important microprocessor peripheral in the form of a general-purpose ADC. This is the first time analog peripherals like those commonly found in microcontrollers have been integrated into an FPGA. Full control over the ADC operation is supported. The ADC offers a number of sampling modes and can support unipolar, bipolar, and full-differential analog input schemes.

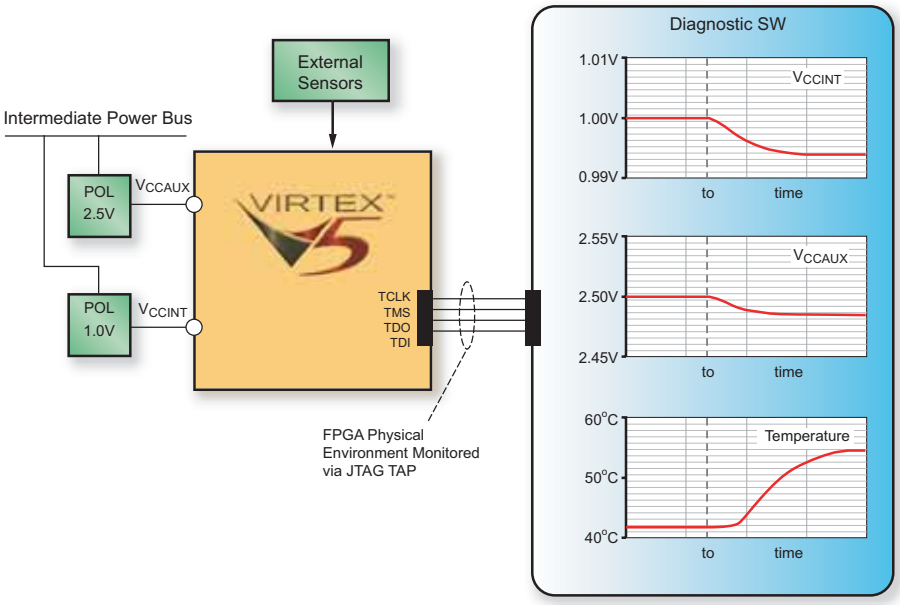

Figure 2 – You can access System Monitor measurements through the JTAG TAP.

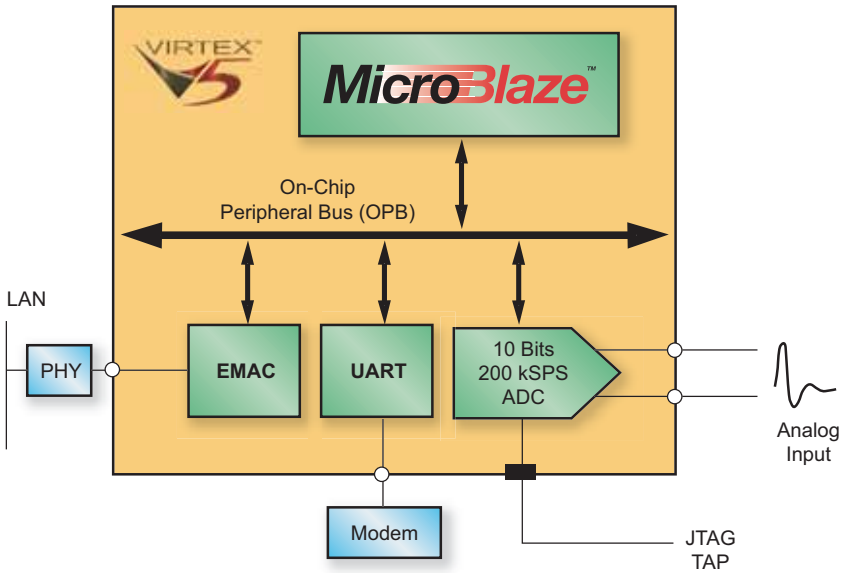

Figure 3 – System Monitor (or ADC) as a microprocessor peripheral

Conclusion

The Virtex-5 System Monitor delivers a greatly simplified solution for common on-chip and external environmental monitoring needs. Minimal development and design effort are required to access the functionality. By interfacing the System Monitor to the JTAG TAP controller, JTAG functionality has been extended into new application areas, thus enabling new test capabilities.

We would like to hear your comments and feedback regarding any topics touched on in this short article; in particular, how our development team can better support your system monitoring and test requirements.

The Ultimate Embedded Solution

for Xilinx MicroBlaze™

With a rich debugging environment and a small, auto-configurable RTOS, Accelerated Technology provides the ultimate embedded solution for MicroBlaze development. Our scalable **Nucleus** RTOS and the Eclipse-based **EDGE** development tools are optimized for MicroBlaze and PowerPC processors, so multi-core development is easier and faster than ever before. Combined with world-renowned customer support, our comprehensive product line is exactly what you need to develop your FPGA-based device and get it to market before the competition!

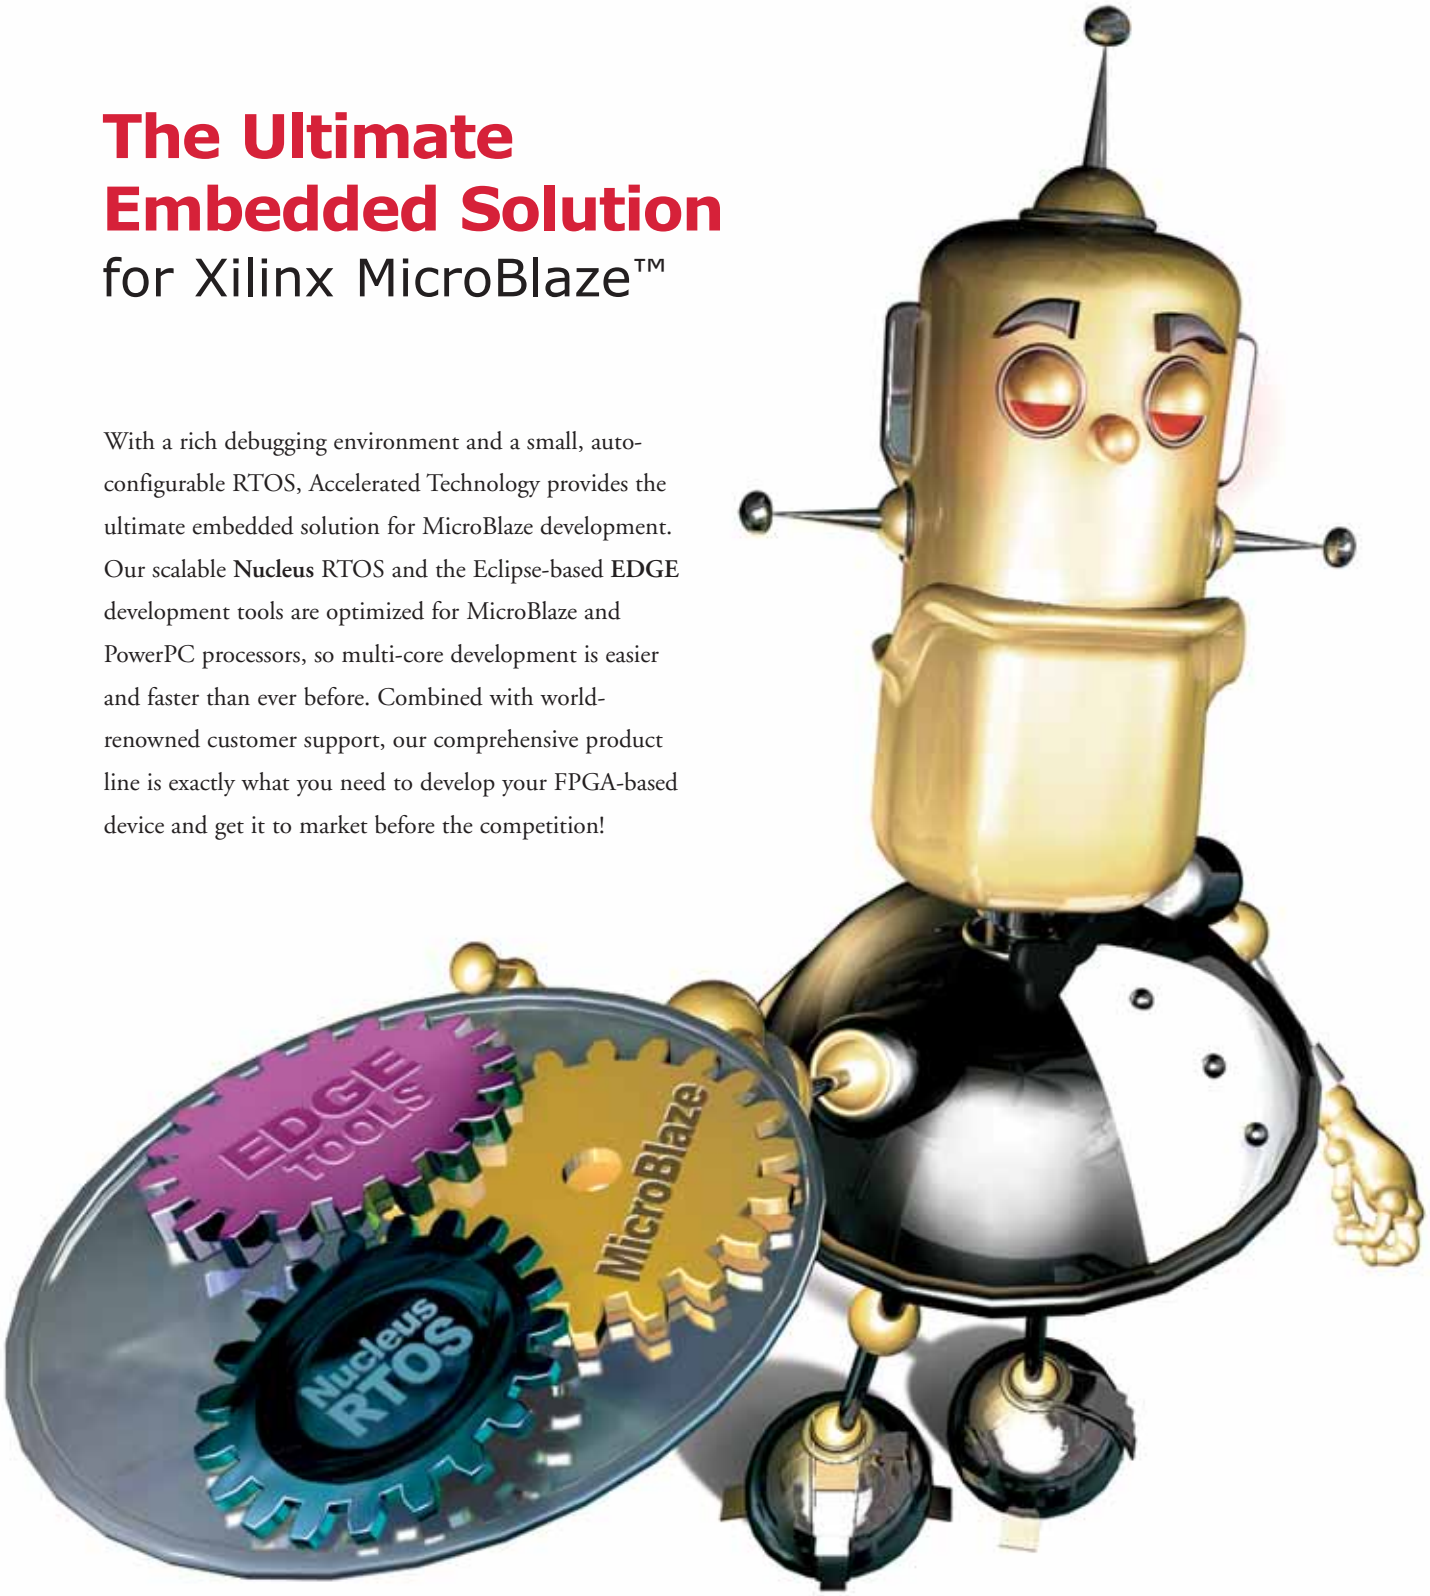

For more information on the Nucleus complete solution, go to
www.acceleratedtechnology.com/xilinx

Accelerated Technology, A Mentor Graphics Division
info@AcceleratedTechnology.com • www.AcceleratedTechnology.com

Accelerated Technology®
A Mentor Graphics Division

Real-Time Debugging for Virtex-5 FPGAs

Version 8.2 of the ChipScope Pro Analyzer delivers verification performance for Xilinx FPGAs.

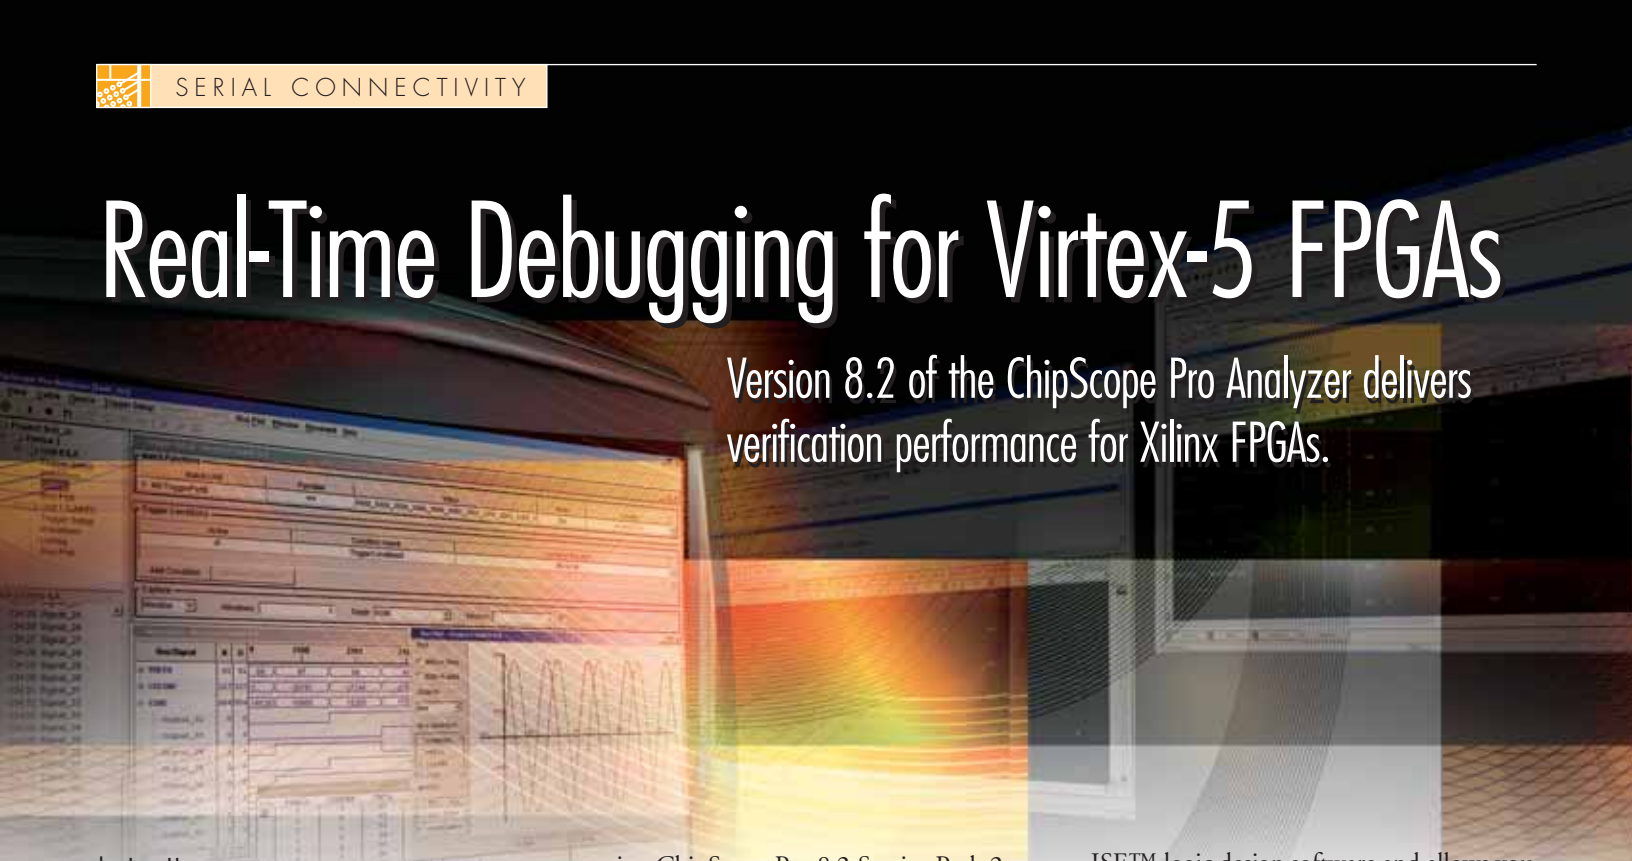

by Lee Hansen
Design Methodologies Sr. Marketing Manager – Horizontal Platform Solutions
Xilinx, Inc.
lee.hansen@xilinx.com

Xilinx® Virtex™-5 devices set a new benchmark in FPGA functionality, with as much as 12 times the logic capacity, 112 times more memory, 2 times the bandwidth, and 2.5 times the performance of the leading FPGA devices of just 8 years ago. Additional dedicated hardware functionality like DCM-based clock management tiles, embedded hard processors, high-speed MGTs, and DSP48E slices extend platform functionality to a broad spectrum of end applications. This extreme functionality places a huge demand on the design cycle and in particular the verification cycle, which tends to be the most time-consuming and time-critical phase of the design flow. The Xilinx ChipScope™ Pro software and analyzer deliver advanced real-time debugging functionality to complex Virtex-5-based designs, moving you through the verification cycle faster than ever before.

New Functionality

The functionality of the ChipScope Pro Analyzer version 8.2 has been enhanced with Virtex-5 performance in mind. All ChipScope Pro-optimized software debugging cores work with Virtex-5 devices when using ChipScope Pro 8.2 Service Pack 2 or later versions. The debugging cores deliver new enhanced performance, supporting higher clock speeds as fast as 500 MHz. You can analyze signals with greater speed and agility through advanced features like wider data capture of up to 1,024 bits, deeper data capture of up to 128K storage samples, and higher density slice packing of trigger match unit and capture control logic.

The resource estimator introduced with ChipScope Pro version 8.1 lets you see how much memory and device space the debugging cores will take up on the chip, useful for project planning.

Another breakthrough feature is remote debugging, first introduced in version 7.1 of ChipScope Pro software. Remote debugging lets you run the ChipScope Pro Analyzer and capture system through a server/client Internet connection. Your board can be running remotely in the lab while you debug from an office on the other side of the building or the other side of the world. You can share a single board or system in the lab with other engineers on your team or allow helpdesk personnel to debug a problem remotely at a customer site, helping to lower field debugging and repair costs.

Optimized Real-Time Debugging

The ChipScope Pro system is available as a separately purchased option to Xilinx ISE™ logic design software and allows you to debug Virtex-5 devices and other Xilinx FPGA-based projects in real time. You can quickly find and analyze design problems while the chip is running on the board, interacting with the rest of the system. Then, leveraging FPGA re-programmability, design changes can be quickly implemented and sent back to the device on board in a matter of minutes or hours through the programming cable. Such changes might take days or weeks using ASIC or competing FPGA offerings.

The ChipScope Pro system also links internal FPGA debugging to Agilent Technologies' bench-top logic analyzers using the included ChipScope Pro ATC2 core. This core synchronizes the ChipScope Pro system to Agilent's FPGA Dynamic Probe software, an optionally purchased plug-in to your Agilent 1680, 1690, or 16900 logic analyzer.

This unique partnership between Xilinx and Agilent delivers deeper trace memory, faster clock speeds, and more trigger options, all using even fewer pins on the FPGA. The advanced technology contained within the ATC2 core and FPGA Dynamic Probe is not available in other FPGA or ASIC real-time verification solutions.

For more information on the ChipScope Pro Analyzer, visit *www.xilinx.com/chipscopepro* or contact your local sales office for ordering information.

X-ray vision for your designs

Agilent Logic Analyzers with FPGA dynamic probe

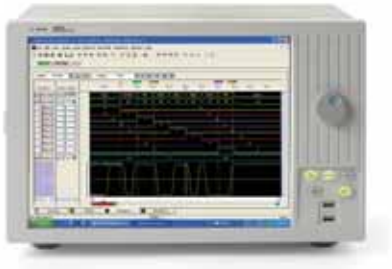

Agilent 16800 Series Portable logic analyzers

- *Increased visibility with FPGA dynamic probe*
- *Portable and modular logic analyzers*
- *Customized protocol analysis with Agilent's exclusive packet viewer software*
- *Low-cost embedded PCI Express packet analysis*
- *Agilent 16800 portable logic analyzers, starting at $9450*

Agilent Direct

Get a quick quote and/or FREE CD-ROM with video demos showing how you can reduce your development time.

www.agilent.com/find/xcell

Now you can see inside your FPGA designs in a way that will save weeks of development time.

The FPGA dynamic probe, when combined with an Agilent Windows®-based logic analyzer, allows you to access different groups of signals inside your FPGA for debug—without requiring design changes. You'll increase visibility into internal FPGA activity by gaining access up to 128 internal signals with each debug pin.

Our newest 16800 Series logic analyzers offer you unprecedented price-performance in a portable family, with up to 204 channels, 32 M memory depth and a pattern generator available.

Agilent's user interface makes our logic analyzers easy to get up and running. A touch screen or mouse makes it simple to use, with prices to fit your budget. Optional soft touch connectorless probing solutions provide excellent reliability, convenience and the smallest probing footprint available. Contact Agilent Direct today to learn more.

Agilent Technologies

Memories are Made of This...

Virtex-5 FPGAs offer a wider range of memories and memory interfaces.

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

All FPGA applications use various amounts of memory for data, parameters, and instructions. To store from a few bits to multiple megabytes, Xilinx® Virtex™-5 devices offer a hierarchy of three different memory implementations:

- LUT-based distributed RAM has a granularity of 64 bits
- Block RAM has a granularity of 18 Kb
- External memory can store practically unlimited amounts of megabytes with the help of on-chip memory interfaces

LUT RAM

Since the early days of the XC4000, Xilinx has made look-up tables (LUTs) available as user RAM. In Virtex-5 devices, the LUT has grown to 64 bits and can be used as either 64 x 1 or 32 x 2 RAM. LUT RAM (see Figure 1) offers very fast (sub-nanosecond) access time, tight integration with the logic fabric, and ultimate design flexibility (see sidebar, "Why 6-LUTs?").

Multiport Option

The four LUTs in a Slice M can share a common write address that does not interfere with the read addressing of the other three LUTs. Together, these four LUTs can thus implement quad-port memory, with one write port and three independent read

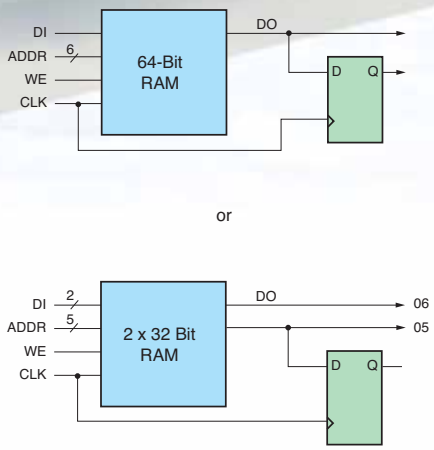

Figure 1 – LUT RAM

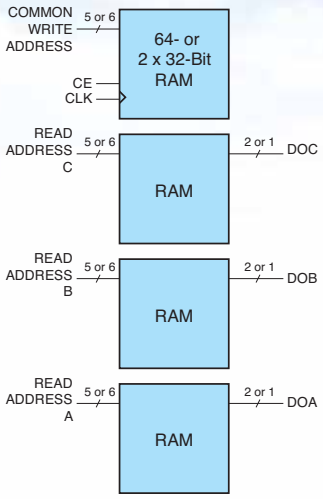

Figure 2 – Slice as quad-port RAM

Why 6-LUTs?

Xilinx invented the use of four-input LUTs in FPGAs 20 years ago. Exhaustive academic and commercial studies had shown that four inputs (16 stored bits) were the optimal size for a LUT that implements random logic.

FPGA evolution has led to ever-smaller transistors as well as ever-more routing and other dedicated structures. As a result, the highly optimized LUT became a much smaller part of the circuitry. For Virtex-5 devices, we re-evaluated the optimal LUT size and found that a four-times-larger six-input LUT (6-LUT) increases the CLB size by only 15%. Extensive benchmarking then showed that, on average, a 6-LUT packs 40% more logic functionality compared to the traditional four-input LUT. The decision was easy: spend 15% more area to gain 40% more logic (or, expressed differently, save roughly 30% in logic area).

The four-times-larger memory capacity of each LUT is an extra, very welcome bonus, as is the ability to make the LUT and the RAM 2 bits wide.

Virtex-5 devices combine four LUTs in a slice. There are two different types of slices: Slice L and Slice M, roughly equal in number on any Virtex-5 device. The LUTs in a Slice L can perform logic and contain a carry chain. The LUTs in a Slice M have the same functionality but can also be used as distributed memory or shift register logic (SRL32) functionality.

ports all accessing the same data. The newest MicroBlaze™ processor uses this feature to reduce its register file from 384 to 44 LUTs. In this kind of application, the new six-input LUT is six times more efficient than a previous-generation four-input LUT (see Figure 2).

Shift Register

You can use any LUT in a Slice M as a serial shift register with addressable length. The LUT is configurable as either a single-bit shift register (a maximum of 32 bits long) or as a 2-bit-wide shift register (a maximum of 16 bits long). Different from earlier SRL16 structures, the Virtex-5 shift register uses a more traditional and scalable design with two latches per shift register bit – hence the maximum 32 bits (not 64 bits) per LUT (Figure 3).

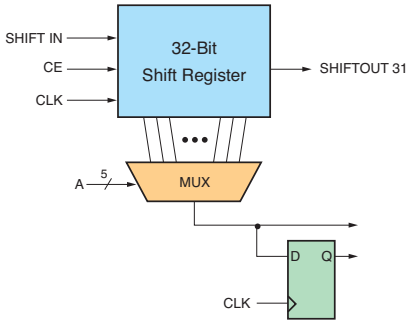

Figure 3 – LUT as shift register

Block RAM

For larger RAM structures, Virtex-5 devices have tens or hundreds of block RAMs, each with a capacity of as much as 36 Kb.

You can structure each block RAM through configuration as:

- 72 bits wide, 512 deep
- 36 bits wide, 1K deep
- 18 bits wide, 2K deep
- 9 bits wide, 4K deep
- 4 bits wide, 8K deep
- 2 bits wide, 16K deep
- 1 bit wide, 32K deep

You can also use the two halves of the 36-Kb block RAM separately as two 18-Kb block RAM, configured as:

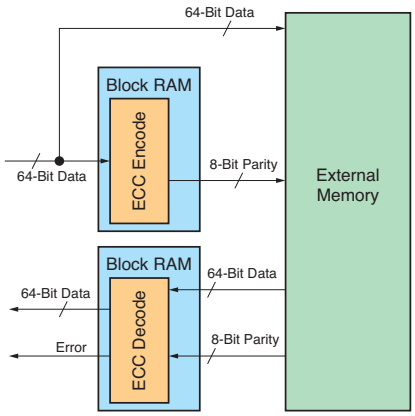

Figure 4 – ECC for external memory

- 36 bits wide, 512 deep
- 18 bits wide, 1K deep
- 9 bits wide, 2K deep
- 4 bits wide, 4K deep
- 2 bits wide, 8K deep
- 1 bit wide, 16K deep

Each block RAM always has two independent access ports and each port can be individually configured. This greatly simplifies data-width conversion.

Read During Write

Each port supports a data-in (DI) bus and a data-out (DO) bus. When writing the data on the DI bus into the memory, the DO bus presents either the previous data at the write address or the new data just being written. A third option keeps DO unchanged from its previous state. These three configuration options offer a design flexibility that is often overlooked.

All block RAM operations require a clock, even for reading data. This requirement is not always desirable, but it is absolute. Nothing happens without an enabled clock. Whenever the clock is enabled, data and address must meet the required setup and hold-time specification. Violating this requirement can contaminate the data content.

ECC

A 72-bit-wide block RAM can provide 64-bit-wide data with error detection and correction (ECC) using Hamming code. The controller is built into each block RAM. It detects single and double errors and corrects all single errors.

The ECC controller can also be used to operate with external memory. In this case, one complete block RAM is necessary for writing and another for reading. The built-in ECC circuit is a great simplification for memory designers who care about the ultimate data integrity (Figure 4).

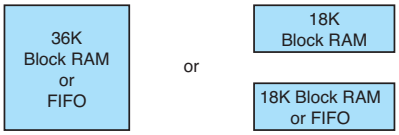

Figure 5 – Dual-ported RAM or FIFO

FIFO

FIFOs are usually implemented using dual-ported SRAMs, with one port used for writing and the other for reading. Many Virtex family block RAMs are traditionally used as FIFOs. That is why Xilinx chose to equip all Virtex-5 block RAMs with a built-in dedicated FIFO controller (Figure 5).

Virtex-5 devices have between 32 and 288 block RAMs, and each can be configured as a 36- or 18-Kb FIFO.

The controller can use the whole block RAM as FIFO with the following configuration options:

- 72 bits wide, 512 deep
- 36 bits wide, 1K deep
- 18 bits wide, 2K deep
- 9 bits wide, 4K deep
- 4 bits wide, 8K deep

But the controller can also use only half of the block RAM and leave the other half to be used as general-purpose block RAM. The FIFO options are then:

- 36 bits wide, 512 deep
- 18 bits wide, 1K deep
- 9 bits wide, 2K deep
- 4 bits wide, 4K deep

In all cases, the FIFO write and read ports have identical width. Unequal width would complicate the interpretation of full/empty flags and is therefore not implemented in Virtex-5 devices.

Soft FIFO controller cores have been available for many years, but the dedicated FIFO controller offers three advantages:

- Higher performance, since dedicated logic is naturally faster than programmable logic
- Smaller size and lower power consumption, as it uses no fabric resources, CLBs, nor additional interconnects
- Guaranteed functionality and performance without any design effort

Write and read clocks can have arbitrary or undefined phase and frequency relationships. But for proper flag operation, both clocks should be free-running.

The challenging aspect of FIFO design is the reliable generation of status flags (FULL, EMPTY, and ALMOST_FULL or ALMOST_EMPTY) when write and read clock frequencies are unrelated. The trailing edges of these flags are inevitably generated by the "wrong" clock domain and must be re-synchronized to the proper clock domain. For a detailed explanation, visit *www.sunburst-design.com/papers/CummingsSNUG2002SJ_FIFO2.pdf*. (See sidebar, "Verifying the EMPTY Flag Synchronization.")

The FIFO controller offers two new options: first-word fall through (FWFT) and synchronous operations.

After a first entry has been written into an empty FIFO, the EMPTY output goes low (inactive), indicating that the read port is allowed to enable its read clock and thus cause the data word to appear at the output. This might be described as a "pull" operation. Data appears at the output after the next enabled read clock.

In FWFT, the newly written data word appears automatically at the output simultaneously with EMPTY going inactive. This might be called a "push" operation.

You can configure the FIFO for either of these modes. The difference is visible only at the read data output of the very first word after the FIFO has been empty. In subsequent operations, no behavioral difference exists between the two modes.

Verifying the EMPTY Flag Synchronization

We tested the EMPTY synchronization logic exhaustively by writing data into the FIFO at 200 MHz and reading it out at 500 MHz, which makes it go EMPTY soon after each write cycle. This exercised the detection logic and re-synchronized the trailing edge of EMPTY 200 million times a second.

More specifically, we wrote an ascending data sequence at 200 MHz and read it out at 500 MHz. We wrote the output data directly into a second FIFO at the same 500 MHz. We then read the second FIFO out at the original 200-MHz rate.

The combined dual FIFO forms a synchronous system but with asynchronous data transfer between the two halves. When we synchronously subtracted the input data from the output data, the difference was constant, indicating flawless transfer at the 500-MHz read/write rate and no flag synchronization problem – even at this high rate.

When two clock frequencies are uncorrelated, each read clock cycle has a different phase relationship with respect to the write clock. During any second, the active read clock edge steps across the ~5 ns write clock period in ~200 million different phase orientations, thus creating a timing granularity of 0.025 femtoseconds. This resolution is millions of times better than any conventional deterministic test methodology can possibly achieve.

We ran this test for several weeks, with more than 10^{14} operations, without any errors.

Asynchronous vs. Synchronous Operation

The main purpose of most FIFOs is to bridge between independent clock domains; most FIFO applications thus use separate and uncorrelated clocks for writing and reading. Because the trailing edge of each flag must be re-synchronized to the opposite clock, there is an unavoidable one-clock-period ambiguity about the delay of the rising flag edge. This can increase the delay for flags going inactive and thus can cause a very small performance loss. The operation is less predictable, but only while the FIFO recovers from the empty or full condition.

In certain applications, there is only one clock domain, and write and read clocks are therefore identical. In that case, you can (optionally) set the mode to "synchronous." This eliminates re-synchronization circuitry, reduces the trailing flag delay, and completely avoids the delay uncertainty. The performance improvement is very small.

External Memory

When a design needs multiple megabytes of memory, it is best implemented in external DRAM devices. High-performance SDRAM controllers can pose a design challenge, but Xilinx offers several application notes and well-documented cores and evaluation boards that implement several different memory controller designs.

Conclusion

Virtex-5 memories and memory interfaces have come a long way since the first LUT RAM appeared in the XC4000 family 16 years ago. Versatile dual-ported block RAMs with FIFO and ECC options simplify system design, while well-documented interface designs allow for unlimited expansion with external DRAMs.

Meeting Memory Interface Design Challenges with Virtex-5 FPGAs

Virtex-5 devices support the latest generation of high-speed memory interfaces.

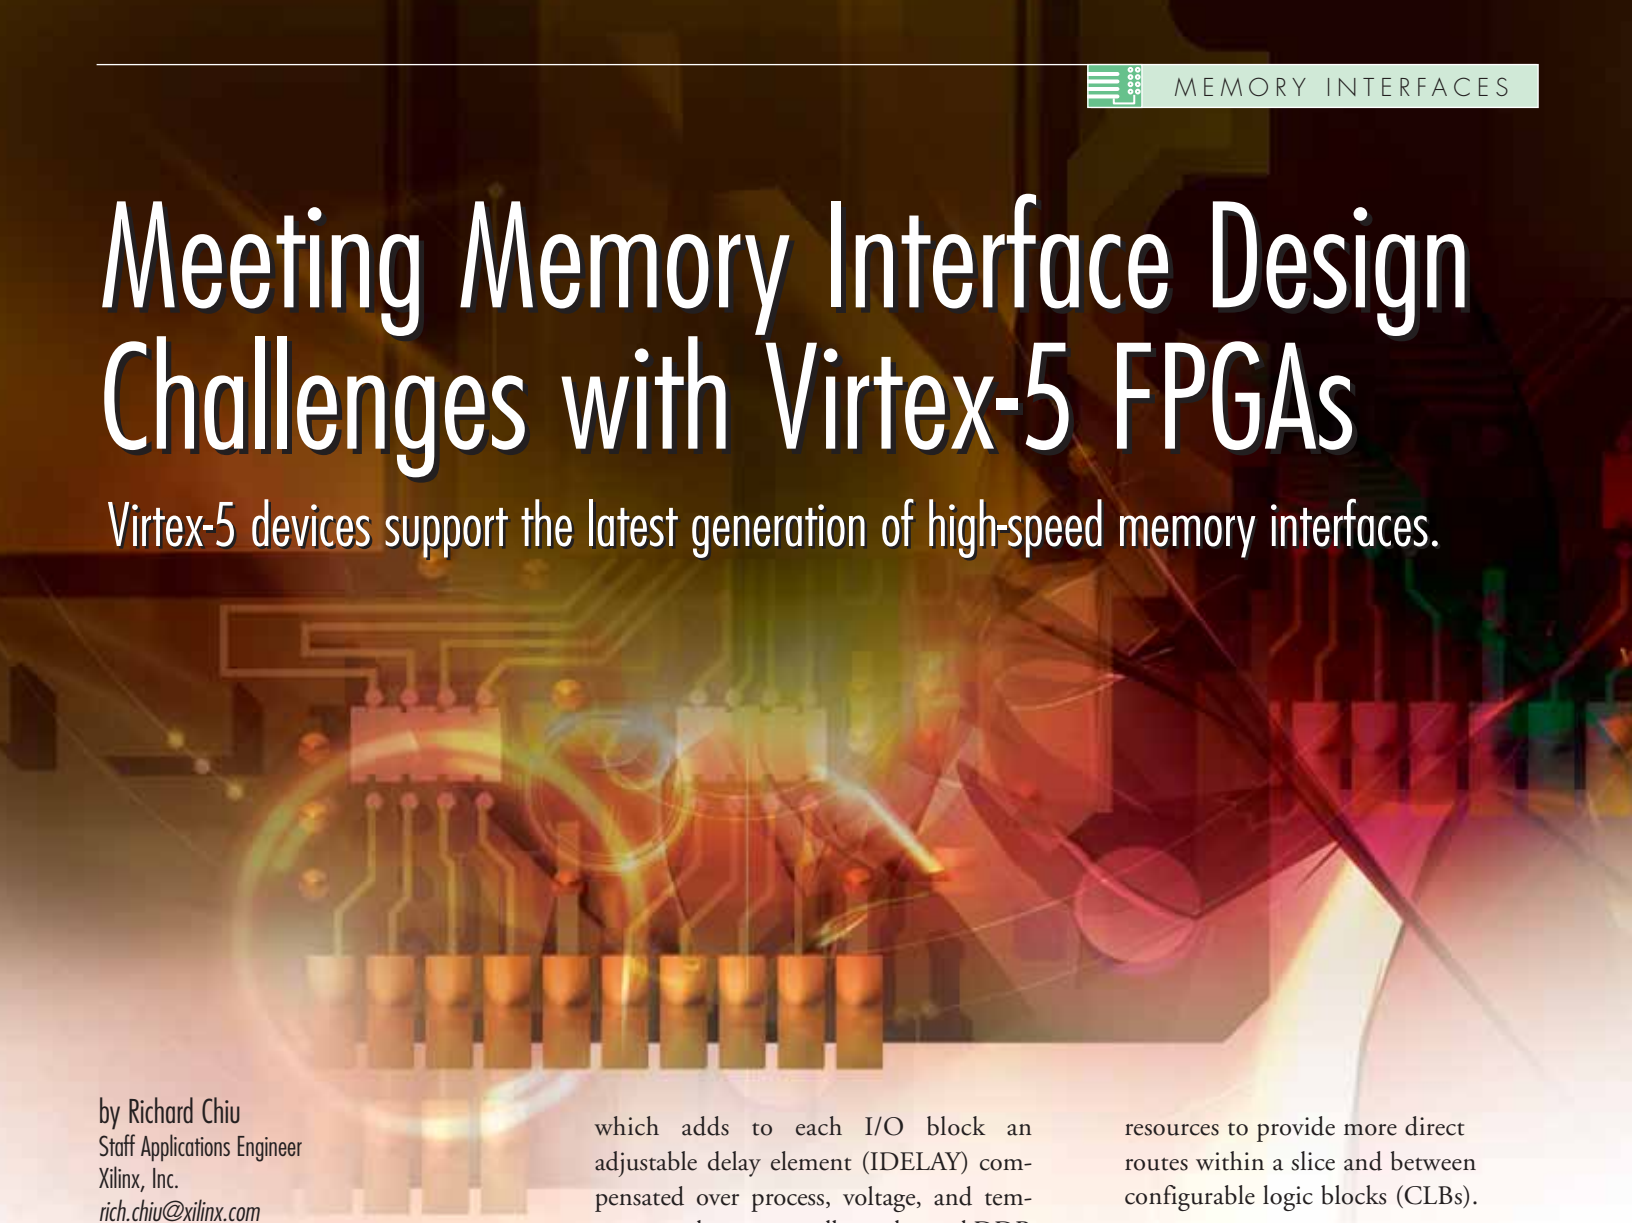

by Richard Chiu
Staff Applications Engineer
Xilinx, Inc.
rich.chiu@xilinx.com

When not supporting new interface protocols, memory interface designers are constantly supporting faster and faster bus speeds for existing interfaces. Today's source-synchronous double-data-rate (DDR) memory devices, such as DDR2 SDRAM, QDR II SRAM, and RLDRAM II, present designers with challenges at chip and PCB levels. Higher clock frequencies result in a rapidly shrinking data valid window. Signal integrity issues, clock jitter, memory uncertainties, varying silicon delays, PCB trace skew mismatch, and other factors now have a proportionally larger impact on meeting timing with a smaller data valid window.

Virtex-5 FPGAs Enhance Memory Interface Design

The Xilinx® Virtex™-4 FPGA family introduced a number of on-chip resources, in particular ChipSync™ technology, which adds to each I/O block an adjustable delay element (IDELAY) compensated over process, voltage, and temperature changes as well as enhanced DDR capture support. These features help meet the challenges of designing with source-synchronous memory interfaces. With Virtex-4 memory interface designs, you can employ calibration algorithms to factor out many of the skews and delays in the timing path and operate your design at higher frequencies.

The Virtex-5 architecture adds additional features that allow you to push the limits of operating frequency. Enhancements to the Virtex-5 device integral to memory interface design include:

- The addition of ExpressFabric™ technology. This architectural enhancement enables internal logic to run at higher clock frequencies. The basic slice look-up table (LUT) has increased from a four- to a six-input LUT (6-LUT), reducing the number of required logic levels. The technology also offers additional routing resources to provide more direct routes within a slice and between configurable logic blocks (CLBs).
- Reduction of the maximum bank size from 64 I/O (or 80 I/O in select Virtex-4 part/package combinations) to 40 I/O, and an increase in the number of banks. This leads to a more efficient implementation of the usual myriad of I/O voltage levels on the same FPGA. More I/O clocking resources have also been added to each bank.
- The availability of phase-locked loop (PLL) blocks as clocking resources in addition to digital clock manager (DCM) blocks. PLLs are useful for low-jitter clock generation and input clock jitter filtering.
- Enhanced block RAM/FIFOs that have doubled in size to 36 Kb and support a maximum width of 72 bits. Applications requiring error-correcting code (ECC) detection and correction can now take advantage of ECC encode/decode logic built into each

block RAM, reducing logic usage and allowing much higher performance over implementing the same functionality in general logic.

- Support for digitally controlled impedance (DCI) on-chip split-Thevenin termination for bidirectional I/O only when the driver is 3-stated. Similar to the on-die termination (ODT) feature implemented in many memory device families, this support is provided for certain HSTL and SSTL I/O standards and can be used to save power when the FPGA is writing to memory.
- The incorporation of low-inductance bypass capacitors directly on the package substrate, simplifying PCB layout by reducing the amount of external bypassing required.

Virtex-5 Data Interface Techniques

Meeting read and write timing for a high-speed source-synchronous bus demands that you keep uncertainties to a minimum. Typically, the capture of read data is the most challenging part of the design.

Write timing for Virtex-5 FPGAs is supported in the same way as in the Virtex-4 device. The DCM (or PLL) generates quadrature phase outputs of the base ("system") clock. The memory strobe is forwarded using an output DDR register clocked by an in-phase copy (CLK0) of the system clock. The write data is clocked by a DCM clock output that is 90 degrees ahead (CLK270) of the system clock. This ensures that the strobe is center-aligned to the data on a write at the outputs of the FPGA.

Both Virtex-4 and Virtex-5 memory interface designs support two kinds of read capture techniques:

- The "direct-clocking" technique delays the read data so that it can be directly registered using the system clock in the input DDR flop of an I/O block. The memory strobe is only used during calibration to determine the optimal time to delay the associated data. Figure 1 shows the direct-clocking read capture path.
- The "strobe-based" technique uses the memory strobe to capture corresponding read data and register it with a delayed version of the strobe distributed through a localized I/O clock buffer (BUFIO). This data is then synchronized to the system clock domain in a second stage of flops. The input serializer/deserializer (ISERDES) feature in the I/O block is used for read capture – the first two levels of flops in the ISERDES transfer the data from the delayed strobe to the system clock domain. Figure 2 shows the read capture path for a Virtex-5 memory interface design.

Most Virtex-4 designs use the direct-clocking method for read data capture. Beginning with the Virtex-4 SERDES DDR2 design and continuing with the new generation of Virtex-5 memory interface designs, the strobe-based method is best to meet the tighter timing requirements at higher clock speeds.

Both techniques involve the use of IDELAY elements that are varied during a calibration routine. This routine is performed during system initialization, delaying both the strobe and data to determine and set the optimal phase between strobe/data and

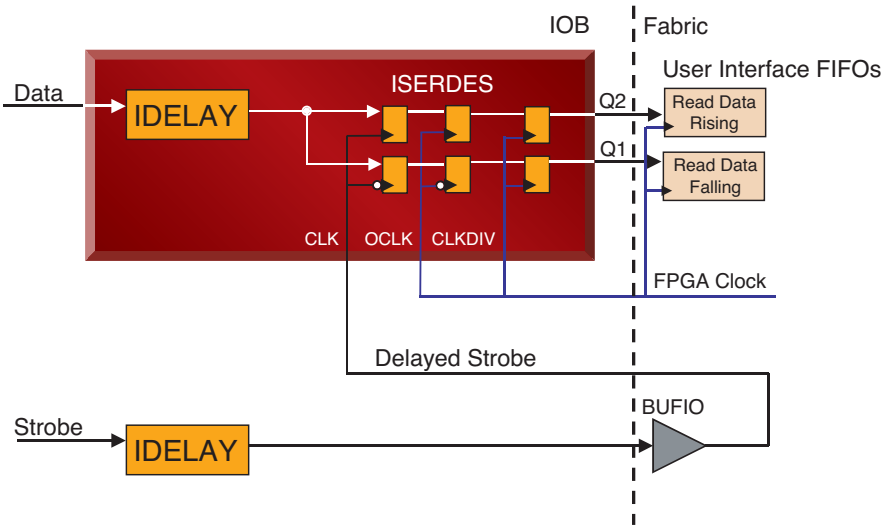

Figure 1 – Virtex-4 direct-clocking read data capture path

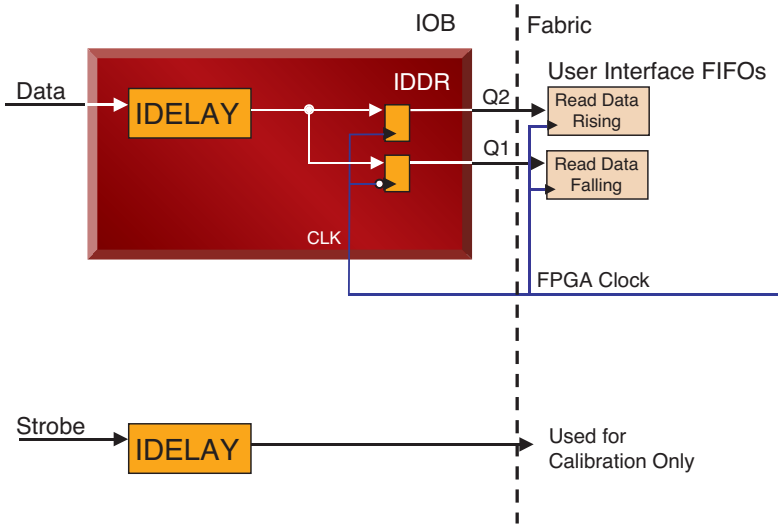

Figure 2 – Virtex-5 strobe-based read data capture path

the system clock to maximize timing margins. Calibration removes any uncertainty caused by process-related delays, compensating for components of the path delay that are static to any one board. These components include PCB trace delays, package delays, and process-related components of propagation delays (both in the memory and FPGA), as well as setup/hold times of capture flops in the FPGA I/O blocks. Calibration accounts for variation in delays that are process-, voltage-, and temperature-dependent at the system initialization stage – you should also factor additional operating temperature and voltage variations separately into your interface timing budget.

During calibration, IDELAY for strobe and data are incremented to perform edge detection by continuously reading back from memory and by sampling either a prewritten training pattern or the memory strobe itself until either the leading edge or both edges of the data valid window are determined. The IDELAY for data or strobe is then set to provide the maximum timing margin. In the case of direct clocking, the optimal delay for the strobe is used to delay the associated data.

For strobe-based capture, the strobe and data can have different delay values because there are essentially two stages of synchronization: one to first capture the data in the strobe domain and another to transfer this data to the system clock domain.

The direct-clocking capture method is simpler in design complexity, and compared to the strobe-based capture method, it has fewer pin-out restrictions. However, the strobe-based capture method becomes necessary at higher clock frequencies. Its two-stage approach offers better capture timing margins for two reasons:

- The DDR portion of the timing is restricted to the first rank of flops in the ISERDES. Because the strobe is used to register the data, timing is limited largely by the strobe-to-data variation; for example, in the case of DDR2, these are given by the tDQSQ and tQHS parameters of the part. For direct clocking, you must consider the data-to-clock variation (for DDR2, this is t_{AC}) because the system clock is used to both drive the memory clock and capture read data. This is a larger uncertainty than the strobe-to-data variation.
- The strobe-to-clock variation is important for the second stage of capture, when the data is transferred from the delayed strobe to the system clock domain. However, by this time the data is split into two separate single-data-rate paths; therefore, aligning the delayed strobe to the system clock can take place over a much larger timing window.

The strobe-based capture method is more pinout-restrictive, as it requires the memory strobes to be placed on clock-capable I/O pins. This can limit the I/O utilization over a given bank. Virtex-5 devices have smaller banks and more I/O clocking resources per bank (for example, the number of BUFIO local clock buffers per bank has increased from two to four), easing this restriction and allowing more strobes and their accompanying I/O (data, mask) to be placed in each bank.

Other significant differences in Virtex-5 memory controllers include:

- Full-speed operation. Both the Virtex-4 SERDES design and Virtex-5 designs use the ISERDES for memory capture. However, Virtex-5 designs do not use the width expansion feature of the ISERDES, and the controller runs at the same speed as the memory clock. The Virtex-4 ISERDES design runs at half the memory clock speed but twice the bus width. Running at the same clock speed as the memory is made possible by the higher performance of the Virtex-5 fabric. This minimizes read-data latency through the ISERDES – as well as controller latency – and simplifies bank-management logic.
- Bank management. The Virtex-5 DDR2 controller employs a least-recently-used (LRU) bank-management algorithm that keeps banks open to reduce the overhead associated with opening and closing banks. In an LRU algorithm, banks are left open at the end of accesses. If a new bank needs to open, the controller closes the bank least recently used. At any time, as many as four banks can be left open.

Generating Virtex-5 Memory Designs

You can generate a custom memory controller by using the Memory Interface Generator (MIG) tool. The MIG tool is accessed through CORE Generator™ software and outputs HDL source (Verilog or VHDL) design files, along with accompanying constraint and build scripts.

The latest version of the MIG tool (1.6) supports DDR2 SDRAM-registered DIMM and QDR II SRAM component interfaces for Virtex-5 devices. The DDR2 controller supports operation of bus clock speeds as fast as 333 MHz (667 Mbps). The QDR II supports operation of bus clock speeds as fast as 300 MHz (600 Mbps).

Virtex-5 designs generated by the MIG tool also allow the physical layer interface portion of the design to be easily separated from the controller portion. You can then incorporate your own specific controller but retain the memory initialization and high-performance source-synchronous calibration logic.

Conclusion

The Virtex-5 device family builds on the Virtex-4 FPGA, with additional features to ease memory interface design and meet the challenges of supporting ever-increasing bus speeds.

To download the MIG tool and for more information about the implementation and design details of Virtex-5 memory controller reference designs, visit the Xilinx Memory Corner at *www.xilinx.com/memory/*.

Virtex-5 memory controllers are also available as reference designs for downloading from the Memory Corner:

- XAPP858 (DDR2 SDRAM)
- XAPP853 (QDR II SRAM)
- XAPP852 (RLDRAM II)
- XAPP851 (DDR SDRAM)

Implementing Memory Controllers Using the Memory Interface Generator Tool

The Memory Interface Generator tool simplifies designing memory controllers for Xilinx FPGAs.

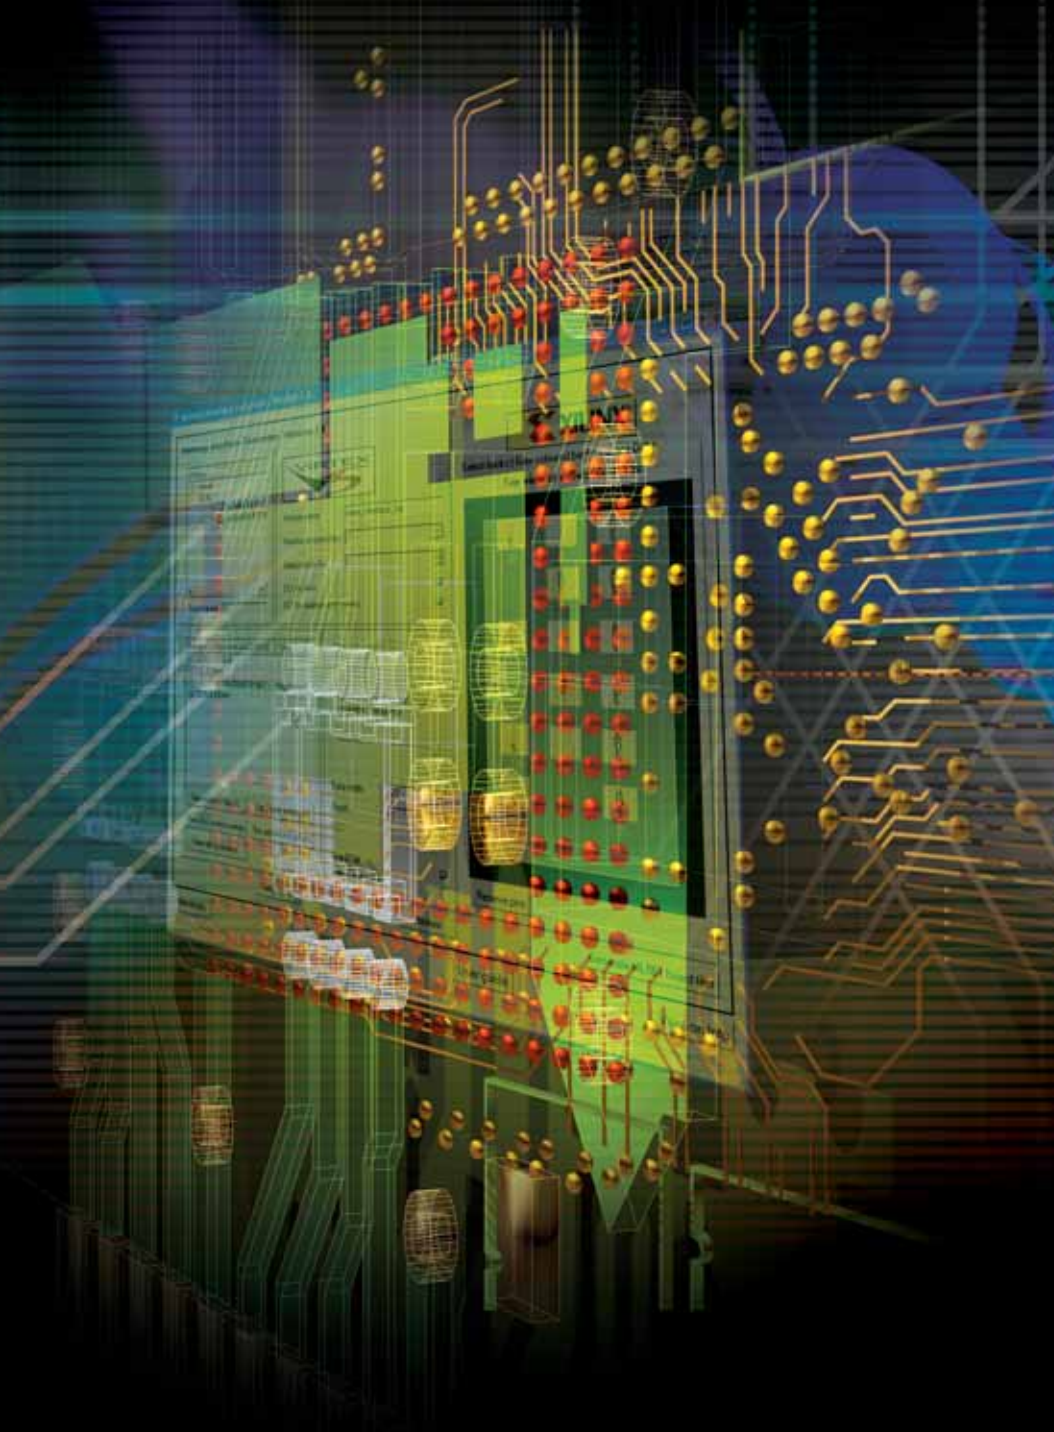

by Nagesh Gupta
Founder & CEO
Taray Incorporated
nagesh@tarayinc.com

The Memory Interface Generator (MIG) tool is a comprehensive tool used to simplify the design of memory controllers for Xilinx® FPGAs. Memories are part of a majority of Xilinx applications. The goal of the MIG tool is to simplify memory interfaces, thus enabling FPGA users to focus on the rest of the system design.

The MIG tool was first introduced in 2002 as a memory controller pin selection utility for Virtex™-II and Virtex-II Pro FPGAs. Since then, the MIG tool has progressed significantly; it now supports all Xilinx FPGA devices, including Virtex-4, Virtex-5, Spartan™-3, and Spartan-3E FPGAs.

The MIG tool dynamically generates HDL in Verilog or VHDL formats based on user inputs. Additionally, the MIG tool generates .ucf pin constraints, any slice and logic placement constraints, and any other constraints required to create high-performance designs with minimal user changes.

MIG outputs are fully available in non-encrypted formats. This enables you to modify the designs.

The Time-to-Market Advantage

High speed memories are complex to design. Conservatively, you can save more than six months by using the targeted reference designs provided by the MIG tool.

Fully verified MIG reference designs enable you to focus on other design activities, thus reducing overall time to market.

MIG Controller Architecture

The MIG tool produces everything required to fully implement a memory controller. MIG controllers are implemented in logical layers comprising:

1. The physical layer, or PHY, which captures the read data, transfers it to a convenient clock domain, and stores it. The PHY also transmits the write data and command/control signals.
2. The controller generates the required commands based on user requests. The controller also implements the state machine for reading, writing, and refreshing the memory.
3. The user interface enables exchange of data and commands to and from your application.

This layered approach allows you to modify the required portions of the design. In Virtex-5 devices, Xilinx has further simplified the layering compared to previous designs. For example, some designers want to use their own controllers, which is possible by replacing the controller that the MIG tool generates. This is easily achieved in Virtex-5 designs.

Hardware Verification

The designs generated by the MIG tool are thoroughly verified to ensure high quality. These quality checks have increased significantly over time as we at Taray learn more from the field.

For a given FPGA family, we verify at least one set of designs in hardware. Hardware verification is usually performed on a Xilinx memory reference board, such as the ML461 or ML561 boards.

Hardware verification starts with a point test, such as a read/write data match, at a particular frequency for a given memory part. We then perform frequency sweeps and ensure that the designs work ±10% in the required frequency range. We also verify all the possible parameters such as column address strobe (CAS) latencies, burst lengths, and data widths, as well as all supported synthesis tools.

Simulations

Taray simulates MIG designs using ModelSim from Mentor Graphics. We simulate a large number of combina-

Hardware-Tested Configurations	
HDL	Verilog and VHDL
Synthesis Tools	XST and Synplicity
Board and FPGAs	ML 461 –> XC4VLX25-FF668-10 and ML 462 –> XC4VLX25-FF668-11
Burst Lengths	4 and 8
CAS Latencies	3 and 4
Additive Latencies	0, 1, and 2
ODT (in Ohms) Verified	0, 75, and 150
Depth Verified for Components	1
Depth Verified for DIMMs	1, 2, 3, and 4
Component Verified	MT47H32M16BT-37E
DIMM Verified	MT18HTF6472G-53E (Registered DIMM)
Component Data Width Verified	16
DIMM Data Width Verified	72 and 144
ECC Verified	72 and 144
Frequency Range	100 MHz to 280 MHz for 16-bit component 100 MHz to 250 MHz for 72-bit DIMM (with and without ECC) 100 MHz to 250 MHz for 144-bit DIMM (with and without ECC)
Simulation-Tested Configurations	
HDL	Verilog and VHDL
Burst Lengths	4 and 8
CAS Latencies	3 and 4
Additive Latencies	0, 1, and 2
ODT (in Ohms) Verified	0, 75, and 150
Depth Verified	1, 2, 3, and 4 (for both components and DIMMs)
Components Verified	All supported by the MIG tool (X4, X8, and X16)
DIMMs Verified	All supported by the MIG tool (registered, unbuffered, and SODIMMs)
Component Data Width Verified	8, 16, 24, 32, 40, 48, 56, 64, 72, 128, and 144
DIMM Data Width Verified	64, 72, 128, and 144
ECC Verified	40, 72, and 144
Frequencies Verified	200 MHz and 267 MHz (for both components and DIMMs)
Initialization	As per both Micron and JEDEC specifications
Multicontroller	1 to 8

Table 1 – Hardware and simulation test summary for Virtex-4 DDR2 SDRAM designs

tions and ensure that every memory listed in the MIG tool is verified with at least one of the test cases. Table 1 is a summary of the different simulation test cases for Virtex-4 DDR2 SDRAM designs. Below are some parameters to generate the test cases:

- All possible data widths
- All of the supported memory components/DIMMs
- Different values for CAS latencies, burst lengths, and additive latencies, depending on the memory type
- Simulated Verilog and VHDL RTL files
- RTL with and without testbench
- RTL with and without DCM
- Use memory models with different frequencies

Key Features

The MIG tool is part of Xilinx ISE™ software and is invoked through the CORE Generator™ tool. Figure 1 is a screen shot of the MIG GUI. The key features of the MIG tool v1.6 are:

- Virtex-5 FPGAs:
  - DDR2 SDRAM, Verilog
  - QDR II SRAM, Verilog
- Support for Virtex-4 FPGAs (and the following designs):
  - DDR2 SDRAM, Verilog and VHDL, direct clocking
  - DDR SDRAM, Verilog and VHDL, direct clocking
  - QDR II SRAM, DDR II SRAM, Verilog and VHDL, direct clocking
  - RLDRAM II, Verilog and VHDL, direct clocking
  - DDR2 SDRAM, Verilog and VHDL, SERDES clocking
  - All Virtex-4 designs support both XST and Synplicity
- Spartan-3 FPGAs:
  - DDR SDRAM, Verilog and VHDL
  - DDR2 SDRAM, Verilog and VHDL
- Spartan-3E FPGAs:
  - DDR SDRAM, Verilog and VHDL
  - All Spartan-3 and Spartan-3E designs support XST, Synplicity, and Precision Synthesis
- Support for many different memory components and DIMMs
- Pins picked are based on the selected memory part and user inputs
- Generates RTL and bit files for Xilinx reference boards containing memories
- Basic I/O design rule check (DRC) engine ensures that signals are allocated correctly
- Verifying a modified MIG .ucf file ensures that MIG pin-out rules are valid

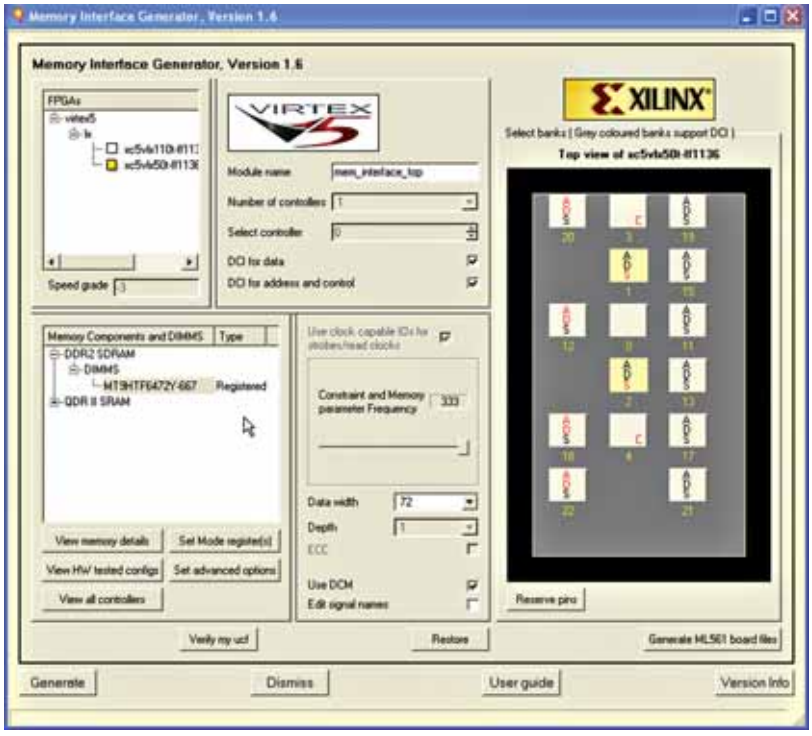

Figure 1 – The MIG tool 1.6 GUI

Using the Outputs of the MIG tool

The MIG tool generates everything required to create a memory interface:

- The RTL (Verilog or VHDL) design files
- Synthesis scripts
- ISE scripts for build, map, and place and route
- A .ucf file for pin locations, RLOCs, and any other constraints

After generating the design RTL, you can execute a batch file to synthesize, map, and place the design. The MIG tool generates two designs – one with a testbench and another without. The MIG scripts work on the version with the synthesizable testbench. However, you can integrate your applications to the version without the testbench.

Conclusion

The MIG tool significantly reduces design burden and improves time to market. It has been used successfully by many customers.

For a copy of the Memory Interface Generator or for additional information, visit *www.xilinx.com/memory*.

Micron Memory Interface

RLDRAM offers a complete low-latency memory interface for networking and communication solutions.

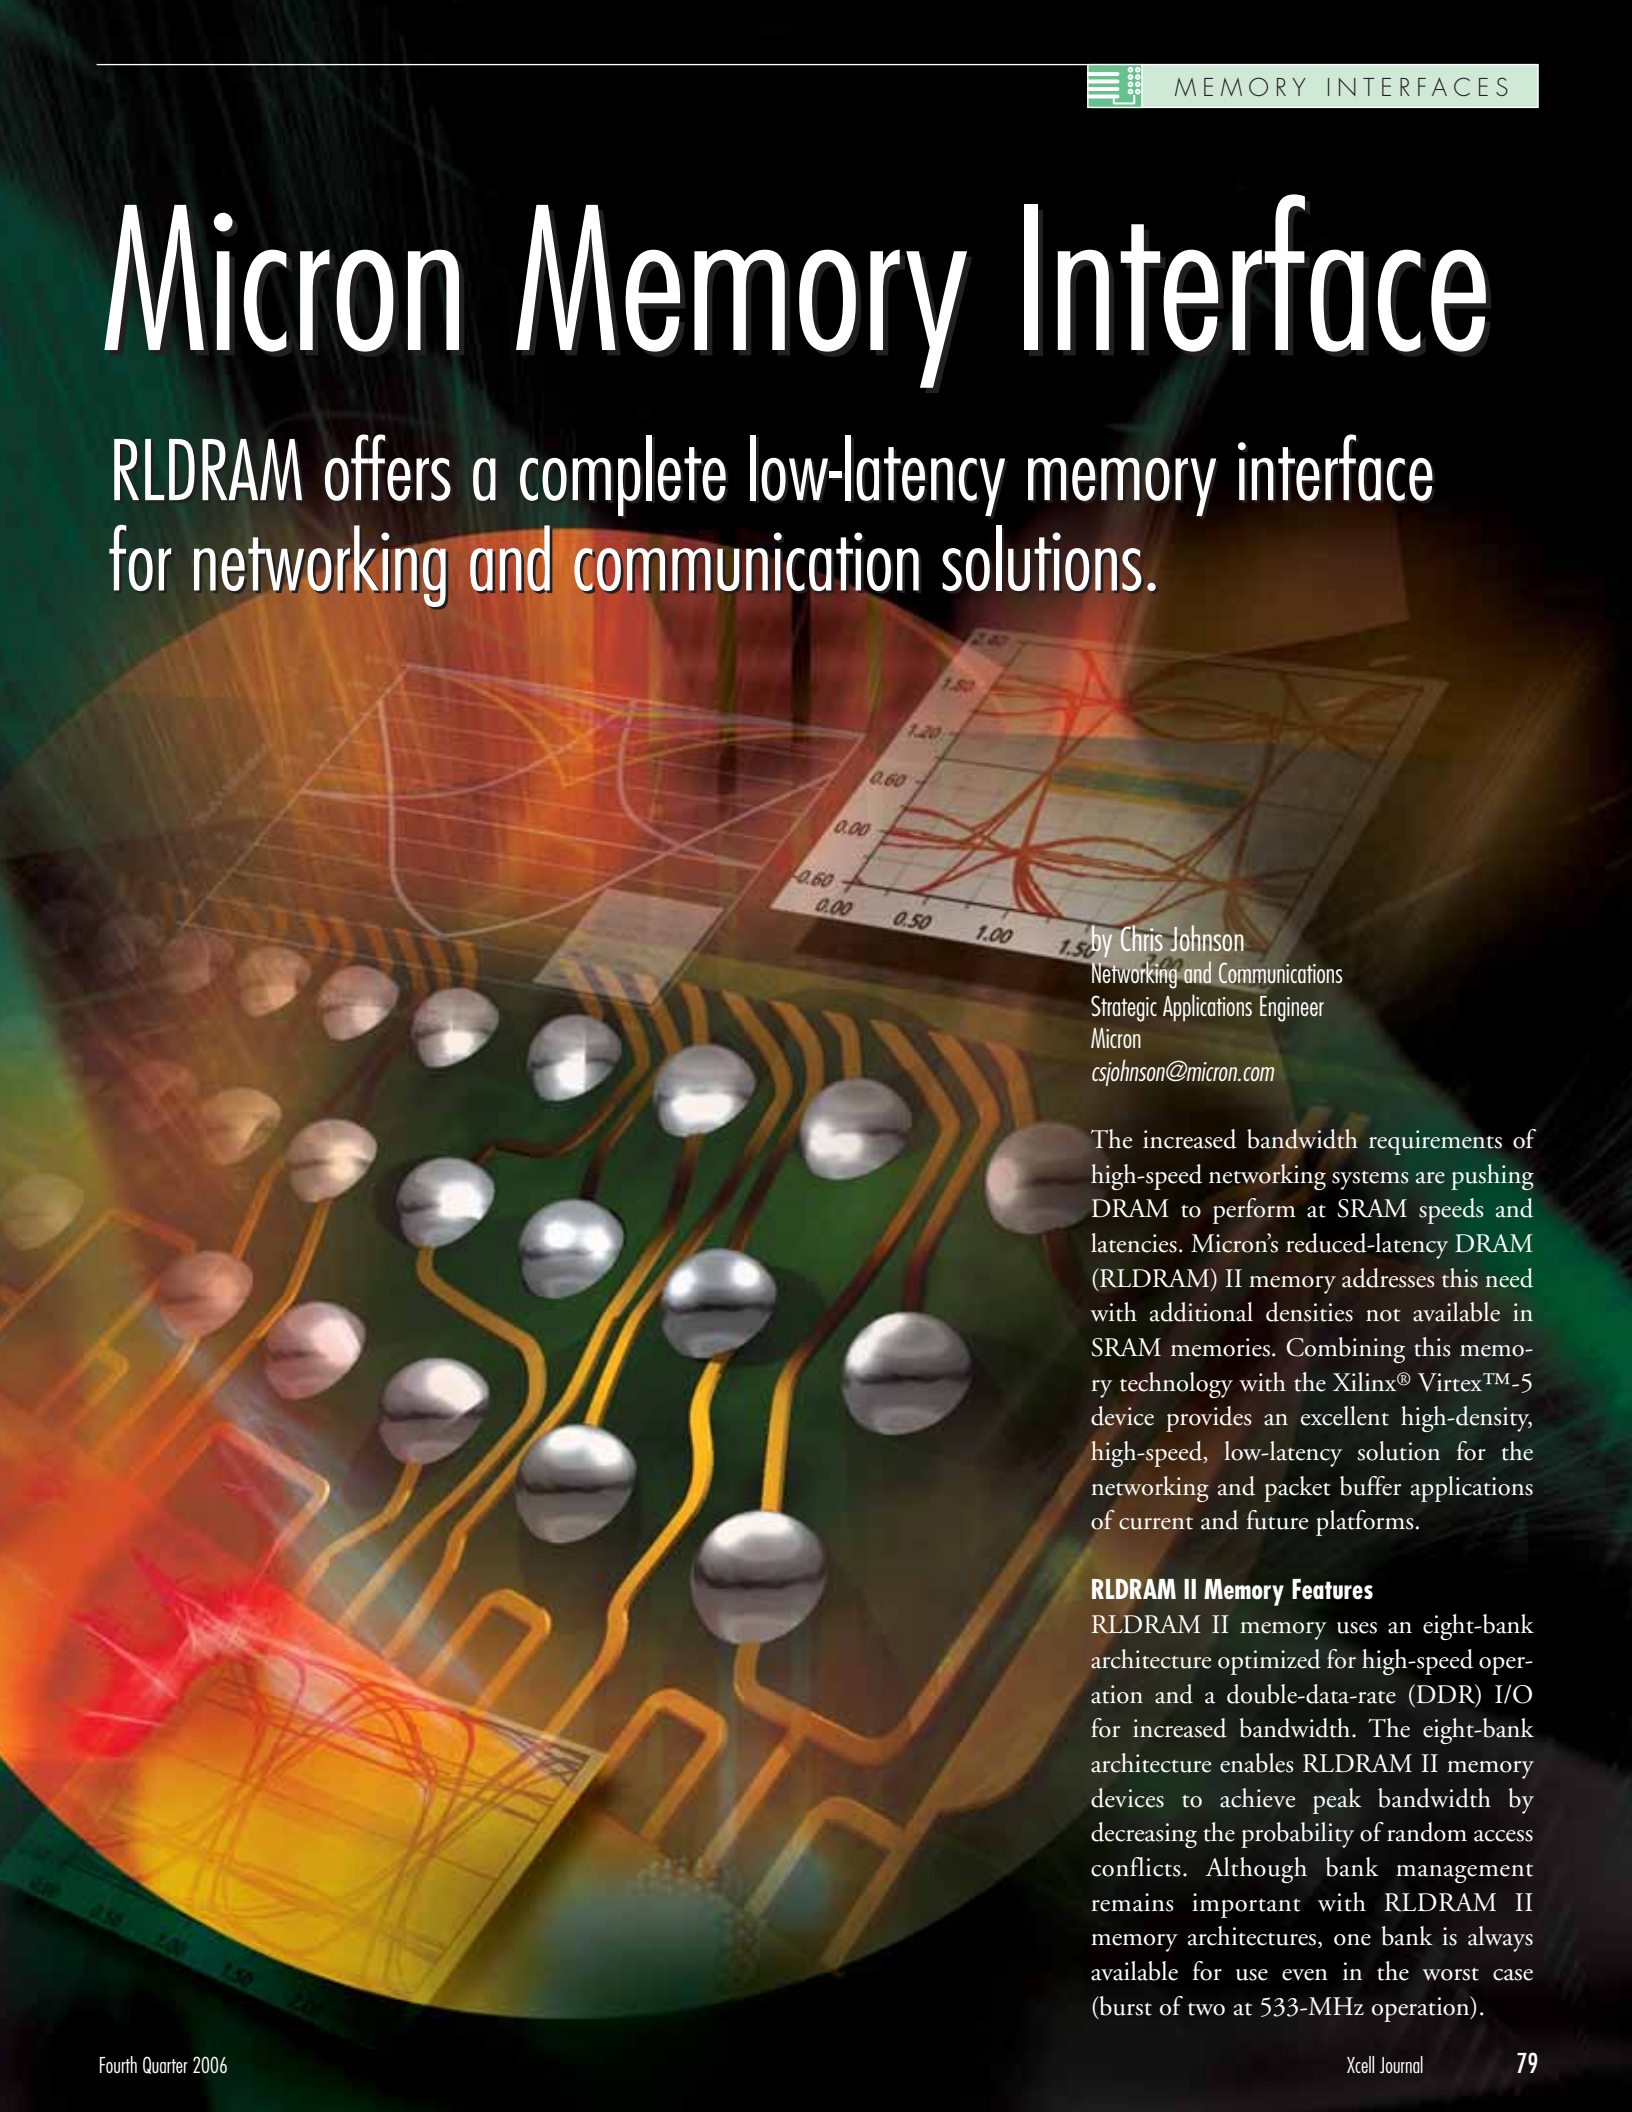

by Chris Johnson
Networking and Communications
Strategic Applications Engineer
Micron
csjohnson@micron.com

The increased bandwidth requirements of high-speed networking systems are pushing DRAM to perform at SRAM speeds and latencies. Micron's reduced-latency DRAM (RLDRAM) II memory addresses this need with additional densities not available in SRAM memories. Combining this memory technology with the Xilinx® Virtex™-5 device provides an excellent high-density, high-speed, low-latency solution for the networking and packet buffer applications of current and future platforms.

RLDRAM II Memory Features

RLDRAM II memory uses an eight-bank architecture optimized for high-speed operation and a double-data-rate (DDR) I/O for increased bandwidth. The eight-bank architecture enables RLDRAM II memory devices to achieve peak bandwidth by decreasing the probability of random access conflicts. Although bank management remains important with RLDRAM II memory architectures, one bank is always available for use even in the worst case (burst of two at 533-MHz operation).

One of the key features added in the RLDRAM II memory architecture is reduced row cycle latency time. Row cycle latency (tRC) is the amount of time that must elapse before a recently accessed bank can be accessed again. Table 1 shows a direct comparison between RLDRAM II memory, DDR2, and DDR at device densities of 576 Mb, 512 Mb, and 512 Mb, respectively.

Latency	RLDRAM II Memory	DDR2	DDR1	Units
^{t}RC	15	55	55	ns

Table 1 – Row cycle time DRAM comparison

I/O Options

RLDRAM II memory offers separate I/O (SIO) and common I/O (CIO) options. The SIO devices have separate read and write ports to eliminate bus turnaround cycles and contention. CIO devices have a shared read/write port that requires one additional cycle to turn the bus around. RLDRAM II memory CIO architecture is optimized for data streaming, where the near-term bus operation is either 100% read or 100% write, independent of the long-term balance.

Figure 1 illustrates the performance variations between the versions at different read-to-write ratios. The reduced latency and eight-bank architecture achievable with RLDRAM II memory allow faster random access from the memory array, increasing sustainable bus utilization.

SRAM-Style Interface

The command bus protocol used in RLDRAM memory is a bit simpler than other DRAM devices. RLDRAM II memory incorporates an SRAM-style interface with read and write commands, replacing the activate and precharge commands used in DDR2 and other similar computer-based memory technologies. This reduction of commands frees up the command bus, helping reduce dead cycles during short burst lengths.

Additional Features for RLDRAM II Memory

RLDRAM II memory also deviates from the refresh requirements of current DRAM technologies. Because ordinary DRAM devices refresh a row in all banks, they require dead clock cycles on the bus after a refresh command. This requires a period of inactivity on the DQ bus, typically 66 ns.

RLDRAM II memory devices have incorporated a bank-based refresh scheme to hide the refresh recovery periods required by other DRAM technologies. The refresh process for RLDRAM memory requires the bank address of the bank that needs to be refreshed, still allowing bus activity during the refresh process and eliminating the dead time seen in other DRAM technologies. Bank-dependent refresh helps increase the percentage of bus utilization, increasing the overall bandwidth of the system using the Virtex-5 device.

Computer-based memory technologies such as DDR2 have relied on modules to support error-correcting code (ECC) technologies to remedy soft errors in the transition of data from the processor to the DRAM. Multiple parts are placed on the DIMM, adding extra data lines for the error-correcting schemes used to eliminate soft errors in the memory channel. RLDRAM II memory is the first DRAM-based technology to add the ECC DQ pins to the devices. RLDRAM II memory is offered in x9, x18, and x36 configurations to provide a single-chip ECC solution without adding unwanted components, reducing board layout space.

Manufacturability in large-component-count systems is a major problem, but it has been ignored in the DRAM industry, largely because the applications that use commodity DRAM devices do not need continuity testing. The module-based business uses so few components that the extra pins are not desirable for JTAG support. Also, discarding manufacturing errors is less expensive than repairing them in small component-based systems.

The target market for RLDRAM memory is an entirely different scenario. Systems based on RLDRAM memory are typically point-to-point applications, where the DRAM is soldered directly to the main board alongside countless other components. Placing all of these components can be a manufacturing challenge. To address this issue, RLDRAM II memory

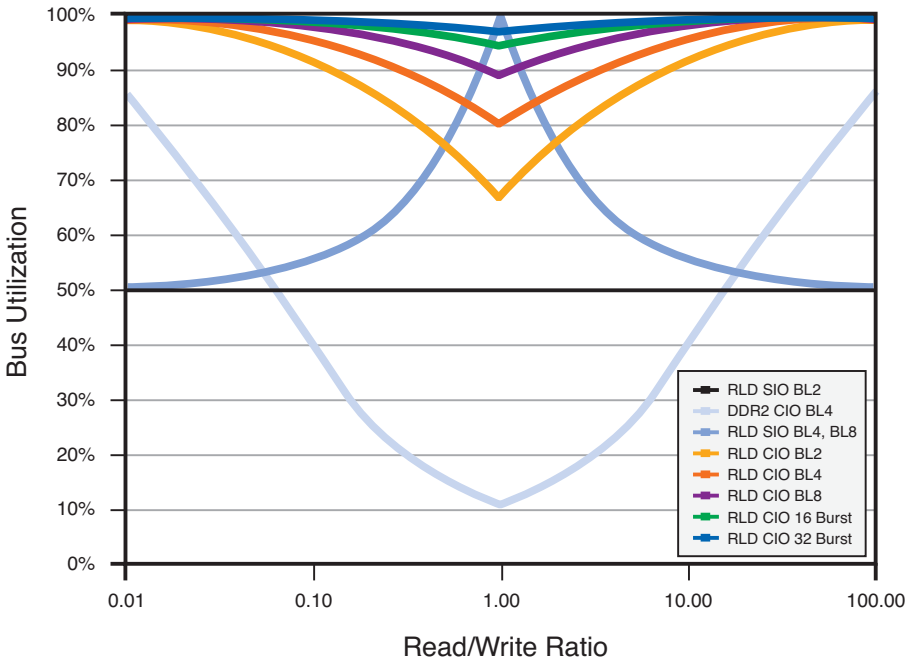

Figure 1 – Peak bandwidth comparison with different DRAM technologies

The density requirements of today's applications have presented a challenge for SRAM systems. The memory cell for SRAM memory devices is approximately five times larger than a DRAM memory cell.

incorporates JTAG continuity technology, which helps with manufacturing issues in the large system-based boards used in the networking and communication industry.

The density requirements of today's applications have presented a challenge for SRAM systems. The memory cell for SRAM memory devices is approximately five times larger than a DRAM memory cell. Figure 2 shows a comparison of the two memory technologies. The DRAM memory cell is substantially smaller, allowing for significantly higher density memory subsystems that still approach SRAM latency speeds.

In the first quarter of 2007, Micron is introducing a 576-Mb RLDRAM II memory device compatible with the high-volume 288-Mb RLDRAM memory device. The 576-Mb device will be offered in multiple configurations that are pin-for-pin compatible, but will have additional address pins to accommodate the higher density. RLDRAM II memory offers one of the highest density, lowest latency DRAM-based solutions available on the market today.

Additional I/O Interface Options

The RLDRAM II memory I/O interface provides other features and options, including support for both 1.5V and 1.8V I/O levels and a programmable output impedance driver that enables compatibility with both HSTL and SSTL I/O schemes.

RLDRAM II memory requires an external one-percent precision resistor (RQ) tied to VSS in order to calibrate the driver to a known value and eliminate the process variation that can be introduced during manufacturing. The calibration process requires the external resistor to operate at five times the desired driver impedance. The programmable impedance control (PIC) circuit calibrates the output impedance to the desired value, eliminating variation introduced during the manufacturing process. Updates are made continuously during device operation without interrupting data transfer, ensuring consistent operation of the RLDRAM memory system independent of temperature and voltage.

Micron's RLDRAM II memory is also equipped with on-die termination (ODT) to enable more stable operation at high speeds without the use of an external termination resistor. ODT provides simplicity and flexibility for high-speed designs by bringing termination resistors on-die, eliminating some of the on-board termination.

At high-frequency operation, however, it is important you analyze the signal driver, receiver, printed circuit board network, and terminations to obtain good signal integrity and the best possible voltage and timing margins. Without proper terminations, the system can suffer from excessive signal attenuation, leading to reduced voltage and timing margins. This, in turn, can lead to marginal designs and cause random soft errors that are difficult to debug. Micron's RLDRAM II memory provides you with simple, effective, and flexible termination options for high-speed memory designs.

Conclusion

RLDRAM II memory combines several performance-critical features to provide flexibility and simplicity for a wide range of high-speed applications. The speed and latency requirements for high-speed applications continue to grow, demanding new and more innovative solutions to meet market requirements. As RLDRAM II memory helps address current market requirements, the demand continues for increased performance. Micron will address this demand with future low-latency devices such as RLDRAM III memory.

Micron continues to innovate and deliver solutions to meet the needs of today's and tomorrow's markets. The joint efforts of Micron and Xilinx help enable our customers to quickly deliver next-generation networking, video, and imaging systems.

For more information, please contact Ray Fontayne at *rfontayne@micron.com.*

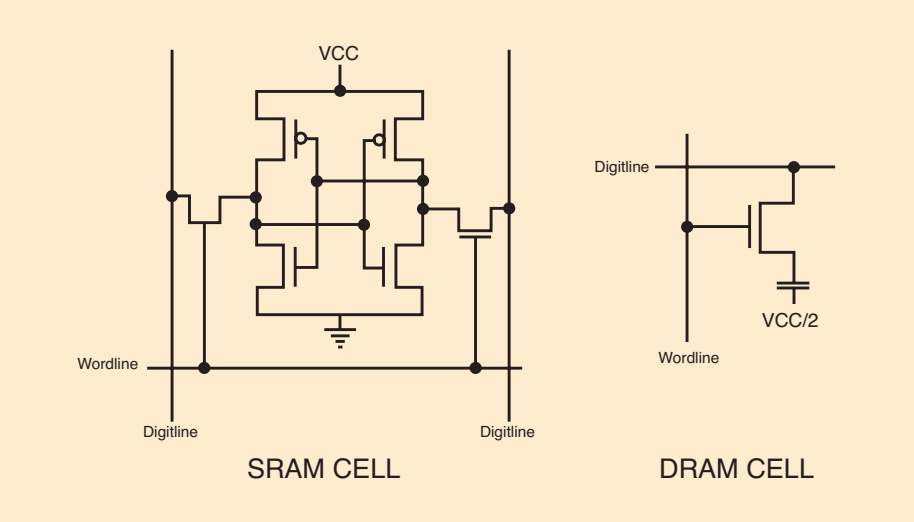

Figure 2 – SRAM cell compared to a DRAM cell

FREE on-line tutorials with Demos On Demand

A series of compelling, highly technical product demonstrations, presented by Xilinx experts, is now available on-line. These comprehensive videos provide excellent, step-by-step tutorials and quick refreshers on a wide array of key topics. The videos are segmented into short chapters to respect your time and make for easy viewing.

Ready for viewing, anytime you are

Offering live demonstrations of powerful tools, the videos enable you to achieve complex design requirements and save time. A complete on-line archive is easily accessible at your fingertips. Also, a free DVD containing all the video demos is available at **www.xilinx.com/dod**. Order yours today!

XILINX®
The Programmable Logic Company℠

Designing Virtex-5 DDR2 Memory Interfaces for Signal Integrity

Follow these guidelines to make your next Virtex-5 DDR2 design experience a success.

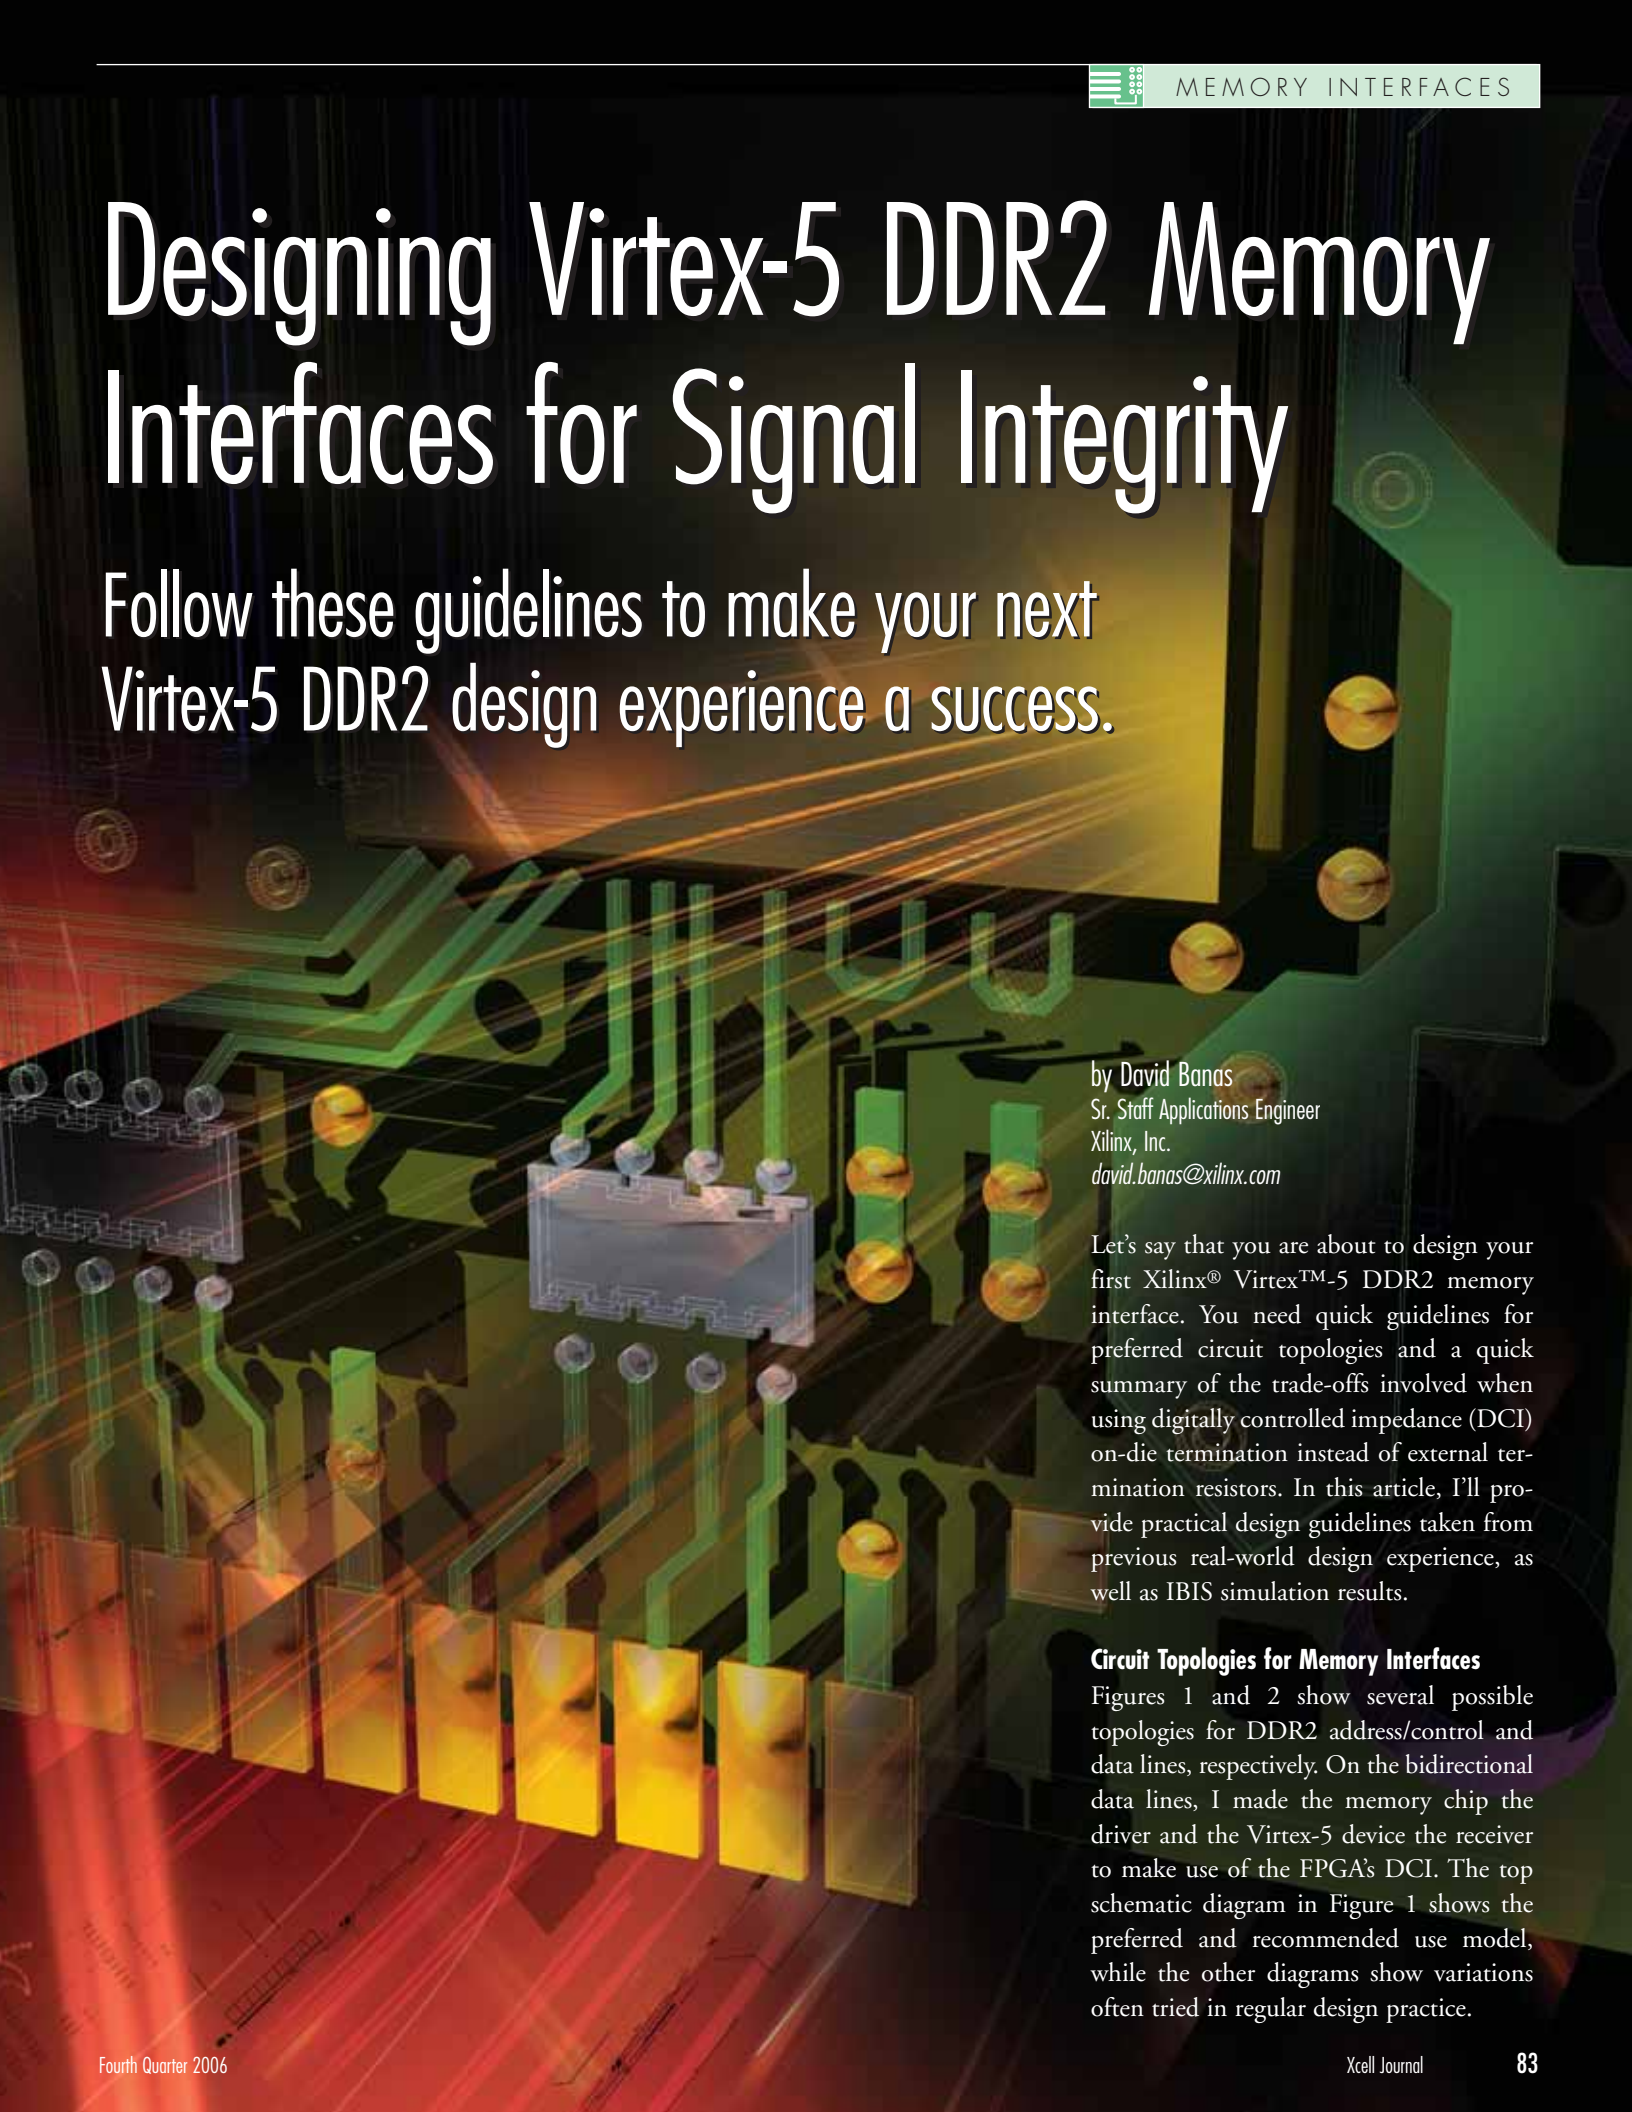

by David Banas
Sr. Staff Applications Engineer
Xilinx, Inc.
david.banas@xilinx.com

Let's say that you are about to design your first Xilinx® Virtex™-5 DDR2 memory interface. You need quick guidelines for preferred circuit topologies and a quick summary of the trade-offs involved when using digitally controlled impedance (DCI) on-die termination instead of external termination resistors. In this article, I'll provide practical design guidelines taken from previous real-world design experience, as well as IBIS simulation results.

Circuit Topologies for Memory Interfaces

Figures 1 and 2 show several possible topologies for DDR2 address/control and data lines, respectively. On the bidirectional data lines, I made the memory chip the driver and the Virtex-5 device the receiver to make use of the FPGA's DCI. The top schematic diagram in Figure 1 shows the preferred and recommended use model, while the other diagrams show variations often tried in regular design practice.

Figures 3 and 4 show typical receiver eye diagrams corresponding to the topologies shown in Figures 1 and 2, respectively. The input switching thresholds of the receiver are shown as horizontal dashed blue lines for reference. The color of the "probe" arrows in Figures 1 and 2 correspond to the colors of the associated traces in Figures 3 and 4, respectively. I used Mentor Graphics's HyperLynx software to generate these eye diagrams with the following parameter settings:

- Pseudo-random binary sequence (PRBS) with bit order 7 (a sequence length of 127)
- Bit interval = 1.5 ns (667 Mbps)
- One sequence repetition
- First 50 bits skipped
- Zero added jitter

When looking at the traces in Figure 3, it should be obvious that of the three topologies shown, the recommended use model gives by far the cleanest eye.

The middle schematic in Figure 1 shows a typical mistake made by novice DCI users, which is to assume that using SSTL18_I_DCI drivers eliminates the need for any external termination components. Some DCI users often incorrectly assume that the "_DCI" versions of the SSTL (stub series terminated logic) driver family adjust their output impedance to match the DCI calibration resistors and can therefore be used as matched impedance drivers of the transmission line.

But this is not true. The SSTL18_I_DCI output driver, for instance, has a fixed output impedance of approximately 20Ω, as per the SSTL18 specification. The disastrous results of this erroneous assumption are clearly visible in the yellow trace shown in Figure 3. Not only has the eye been drastically narrowed, but problematic overshoot/undershoot has also been introduced at the receiver input.

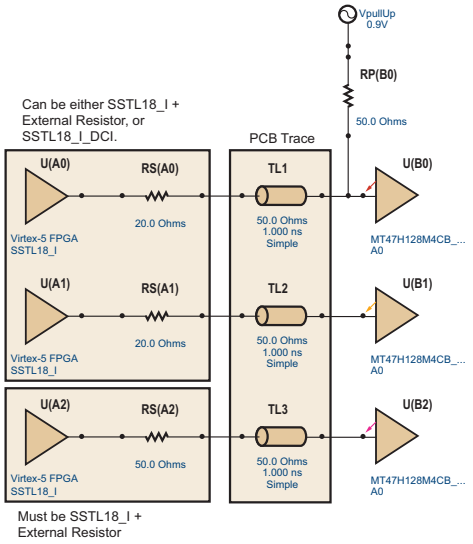

Figure 1 – Typical address/control circuit topology

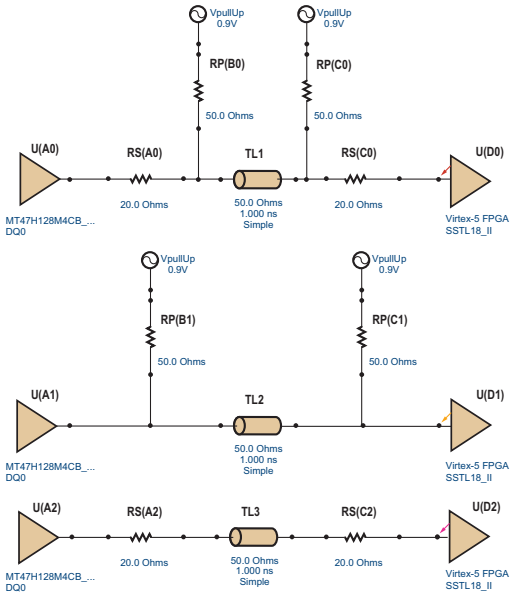

Figure 2 – Typical data circuit topology

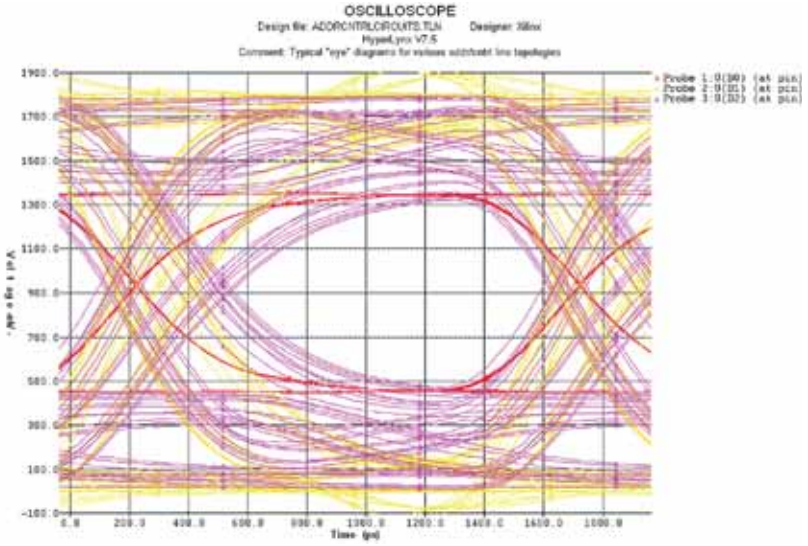

Figure 3 – Typical eye patterns for address/control

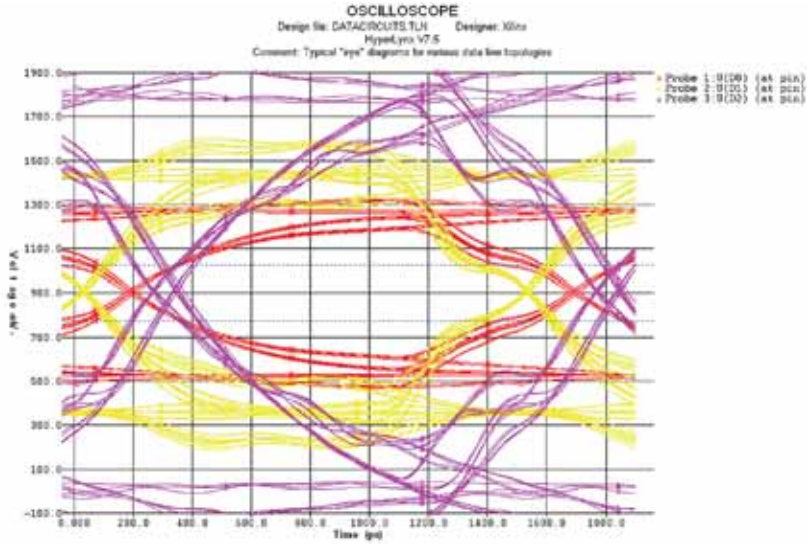

Figure 4 – Typical eye patterns for data

...blindly following recommended use models, rules of thumb, or general guidelines is never a good substitute for simulating your design.

Increasing the series termination to 50Ω, as shown in the bottom schematic in Figure 1, successfully eliminates overshoot/undershoot but does nothing to restore the eye to its original width. Therefore, you should always use parallel termination at the end of all address/control lines.

If an appropriate termination voltage source is not available, you can form a Thevenin-equivalent termination using two resistors connected in series between the V_{CCO} supply and ground, where each resistor has a value twice that of the desired impedance. In this case, simply terminate the line by connecting its end point to the net that connects the two resistors. Note that your circuit consumes more power when terminated in this fashion because of the constant load on the V_{CCO} supply formed by the two resistors.

The data-line eyes in Figure 4 also illustrate that removing the parallel termination from the ends of the line causes unacceptable overshoot/undershoot. However, in this case, removing the series terminations appears to have improved the eye, making it slightly wider and providing more "head room" against noise without introducing overshoot/undershoot. This serves as an excellent reminder that blindly following recommended use models, rules of thumb, or general guidelines is never a good substitute for simulating your design.

Keep in mind that before approving an engineering change order (ECO) for the removal of the series termination resistors, you should reverse the direction of the line with the Virtex-5 device driving data to the memory chip and check the eye for good signal integrity (SI).

Why Use DCI?

The benefits of using DCI, as opposed to equivalent external termination, are numerous, including:

- Better SI at receiver inputs
- Reduced PCB size
- Reduced bill of materials (BOM) parts count

SI at receiver inputs improves when using DCI because the termination lies closer to the inputs than when you use an external termination resistor.

PCB size and BOM parts count are both reduced when using DCI because of the elimination of external termination components.

Caveats when using DCI include:

- Impedance variation over process/voltage/temperature
- Greater power consumption

The termination impedance when using DCI is provided by CMOS transistors; therefore, the value of that impedance can vary along with variations in the fabrication process, supply voltage, and operating temperature (PVT) of the FPGA. You should always perform system-level SI simulations twice, using the high and low extremes for the value of the termination impedance to ensure correct system operation across all possible combinations of PVT.

Using DCI for parallel termination at the end of a transmission line results in higher power consumption than using an external resistor, assuming that an appropriate termination voltage source is available. In this case, the end of the line can be connected to the voltage source through a resistor with a value equal to the characteristic impedance of the line, and no load will be placed across the supply rails.

Conversely, when DCI is used to terminate the line, two pass transistors connect the receiver input to V_{CCO} and ground, respectively. Each transistor is adjusted to have an effective resistance equal to twice the characteristic impedance of the line, thus producing a Thevenin-equivalent termination impedance of Z0 to $V_{CCO}/2$.

A side effect of this termination scheme is that an additional load of $4Z_0$ appears across the supply rails, consequently increasing overall system power consumption. If the system architectural design specifications do not provide for a voltage supply at $V_{CCO}/2$, then no power penalty is incurred for using DCI because the termination scheme has to be Thevenin-equivalent in either case.

Table 1 gives the worst-case output power dissipation of a single line for three termination styles: source series, external parallel (assuming availability of $V_{CCO}/2$), and internal parallel (DCI).

Termination Type	Power Dissipation
External Source Series Termination	41 mW
External Parallel Termination	49 mW
Internal (DCI) Parallel Termination	57 mW

Table 1 – Power dissipation versus termination type (for a discussion of output power calculations, see "High-Speed Digital Design: A Handbook of Black Magic" by Howard Johnson and Martin Graham)

Conclusion

In this article, I've shown how various digressions from the Xilinx recommended use model for circuit topology of DDR2 memory interfaces affect the eye at the receiver. I hope you are convinced that system-level simulation with an IBIS simulator such as HyperLynx is a necessity when designing DDR2 memory interfaces. And I've given you some pros and cons for using DCI as an alternative to external termination resistors in your next DDR2 design.

Going forward, as memory interface speeds continue to increase and I/O voltage levels continue to decrease, you will be able to apply the general principles learned here to the design of more complex memory interfaces as those standards emerge.

For more information, visit *www.jedec.com.*

Honey, will you please tell Alex
to stop programming the FPGA!

Viva! Absurdly Easy.

Viva! Software

The Easiest Development Tool for Programming FPGAs

Viva may just be the easiest and most intuitive development tool ever created for FPGAs. That's because only Viva delivers a true data set recursive and polymorphic object oriented programming language with a graphical, drag-and-drop interface. In fact, Viva is so easy you can get the hang of it in a just few minutes. Whether you're working with FPGAs in high-performance or embedded environments, Viva is the right solution for you. And Viva is perfectly suited for use with hardware from SGI, Nallatech and Digilent, including Xilinx Starter Kits. Find out why organizations like NASA and the US Air Force use Viva! software to program FPGAs. Visit www.starbridgesystems.com/viva or call (801) 984-4444 today.

www.starbridgesystems.com/viva

starbridge

Viva! Absurdly Easy.

Improve System Reliability with SPI-4.2 LogiCORE Solutions and Virtex-5 FPGAs

Virtex-5 devices provide an ideal platform for source-synchronous designs like the widely adopted SPI-4.2 interface.

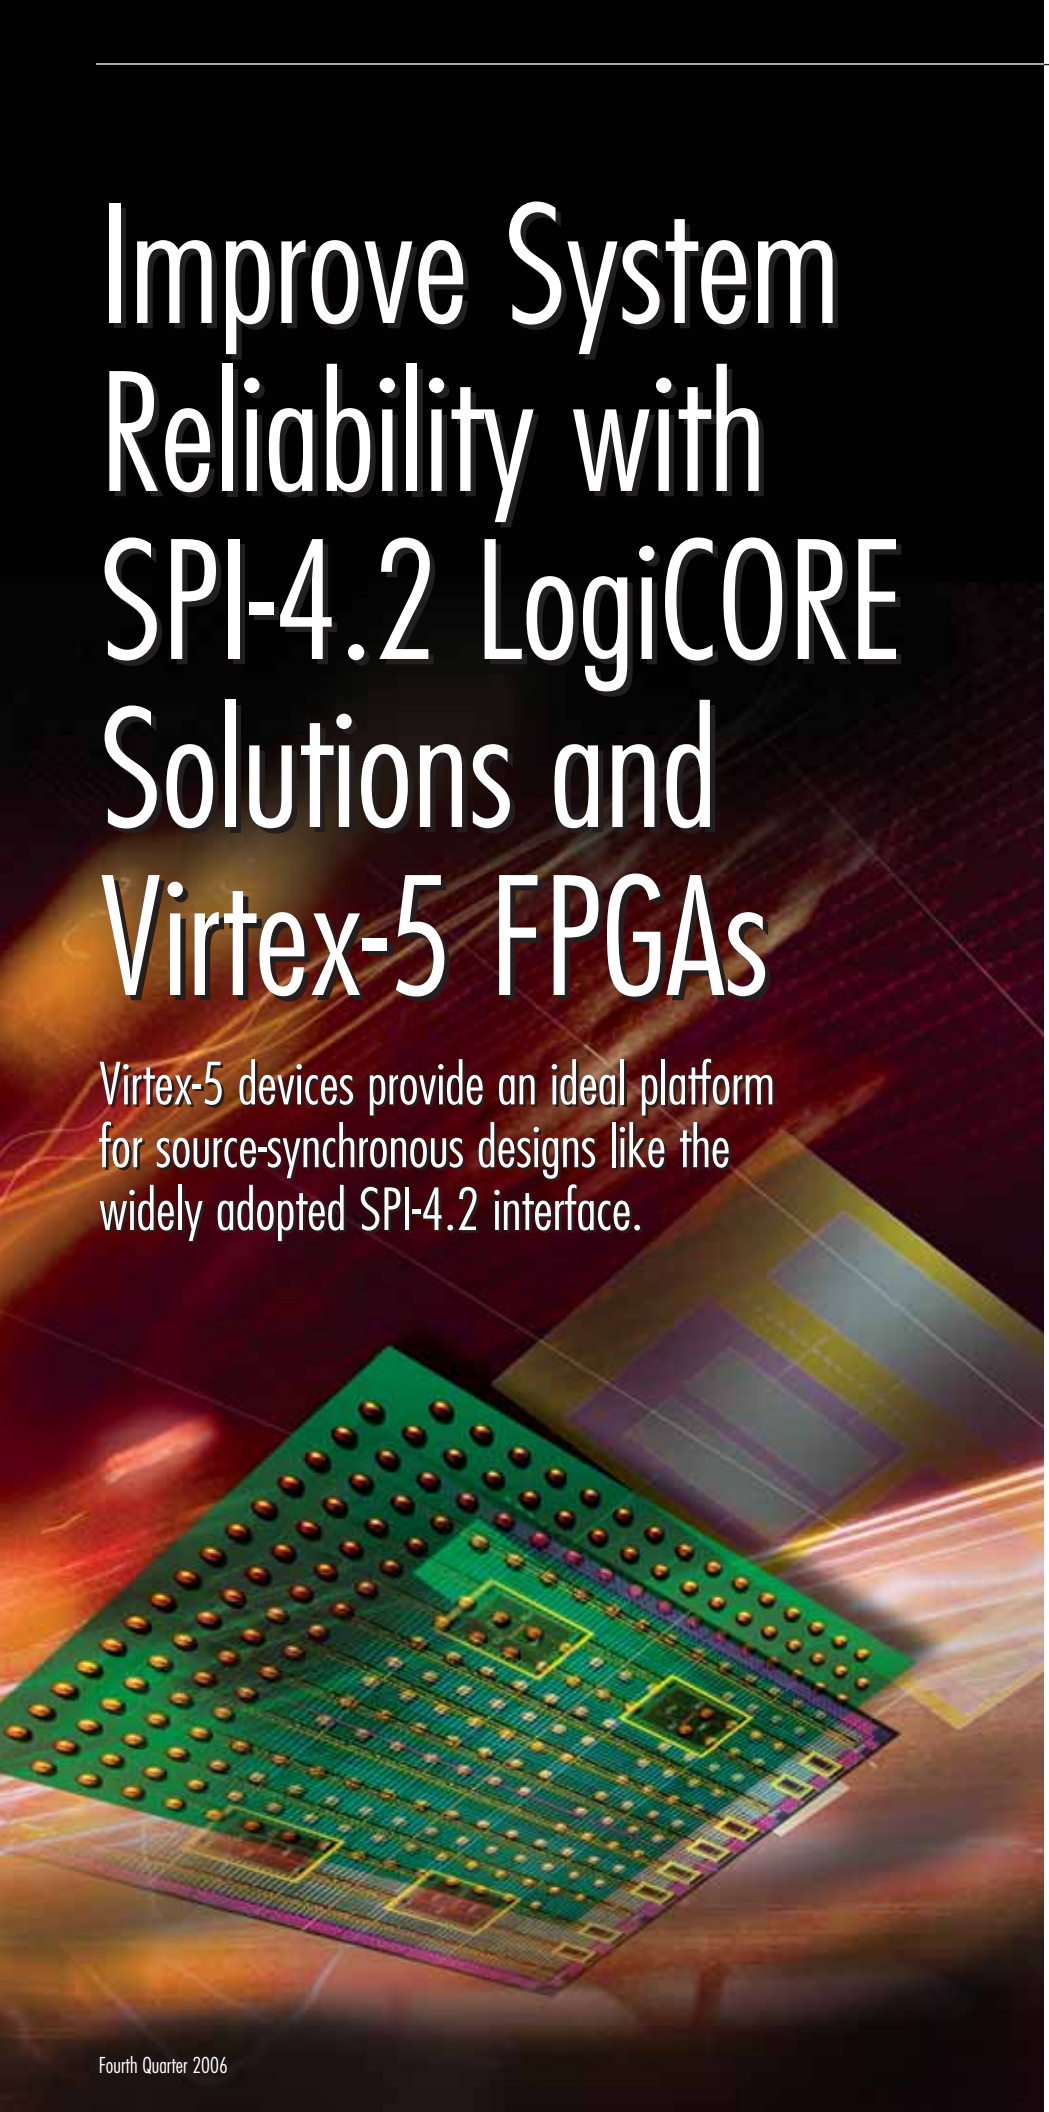

by Dean Armintrout
Product Marketing Engineer, IP Solutions Division
Xilinx Inc.
dean.armintrout@xilinx.com

Chris Ebeling
Principal Engineer, IP Solutions Division
Xilinx, Inc.
chris.ebeling@xilinx.com

System Packet Interface Level 4 Phase 2 (SPI-4.2) is the Optical Internetworking Forum's recommended interface for the interconnection of devices for aggregate bandwidths of OC-192 (ATM and POS) and 10 Gbps (Ethernet), as illustrated in Figure 1.

The SPI-4.2 interface has become the standard for interconnecting leading-edge 10 Gbps framers, traffic managers, network processors, and switch fabrics. SPI-4.2 is popular because of its efficient interface, which offers high bandwidth and low pin count, along with seamless handling of typical system requirements such as flow control, error detection, synchronization, and bus realignment.

The Xilinx® Virtex™-5 architecture provides an ideal platform for implementing SPI-4.2. The Xilinx SPI-4.2 LogiCORE™ IP targeting Virtex-5 devices provides a significantly smaller solution with dramatic power savings, 1.2 Gbps LVDS DDR I/O, and complete pin assignment flexibility.

SPI-4.2 LogiCORE IP

Continually improving on its SPI-4.2 solution, Xilinx has made the latest implementation 25% smaller than previous versions by leveraging the 65-nm ExpressFabric™ technology and real six-input look-up tables (LUTs) of Virtex-5 FPGAs.

Enhanced ChipSync™ technology is supported on every pin of the Virtex-5 device family, allowing you to target the SPI-4.2 LogiCORE solution to any device pinout to meet your system and PCB requirements. High-performance interfaces are supported by 1.2 Gbps LVDS data rates.

For applications requiring multiple SPI-4.2 interfaces, the Virtex-5 FPGA's logic density, high pin count, and extensive clocking resources support four or more full-duplex cores in a single device.

ChipSync Source-Synchronous Technology

Virtex-5 devices build on ChipSync technology to ensure reliable high-speed data transfer for source-synchronous applications like SPI-4.2 with these features:

- Built-in serializer/deserializer (SERDES) logic enables the fabric to interface to the I/O at a fraction of the source-synchronous clock rate. The included bitslip function allows shifting of the deserialized data to achieve word alignment when linking multiple pins (bus deskew).
- Input delay (IDELAY) components allow the dynamic phase alignment (DPA) logic to independently adjust the delay of each bit of a bus in 75-ps increments, providing a mechanism for tuning the interface timing to the system environment.
- DDR registers integrated into the I/O pins simplify the interface between the FPGA fabric and the I/O blocks by supporting data transfer on a single clock edge.

SPI-4.2 and ChipSync Technology

The SPI-4.2 interface has a DDR source-synchronous data bus that comprises 18 LVDS pairs (16 data, 1 control, and 1 clock), operating at a minimum rate of 311 MHz.

The SPI-4.2 core uses ChipSync technology to serialize/deserialize bus data to a four-word SPI-4.2 datastream at a lower clock rate; thus, you can implement high-frequency SPI-4.2 interfaces in slower speed grade Virtex-5 devices.

The SERDES functions allow the core logic to transfer these four words to and from the I/O logic without using any CLB logic resources and operate at half the source-synchronous DDR clock rate. For example, a SPI-4.2 interface with a 500-MHz DDR reference clock only requires an FPGA fabric clock of 250 MHz – easily achievable in the Virtex-5 architecture.

As the frequency of the source-synchronous clock increases, data recovery at the receiving (sink) device becomes more challenging. The SPI-4.2 protocol provides a training pattern that permits a receiving device to adjust its data sampling to the system interface timing – a process referred to as dynamic phase alignment (DPA).

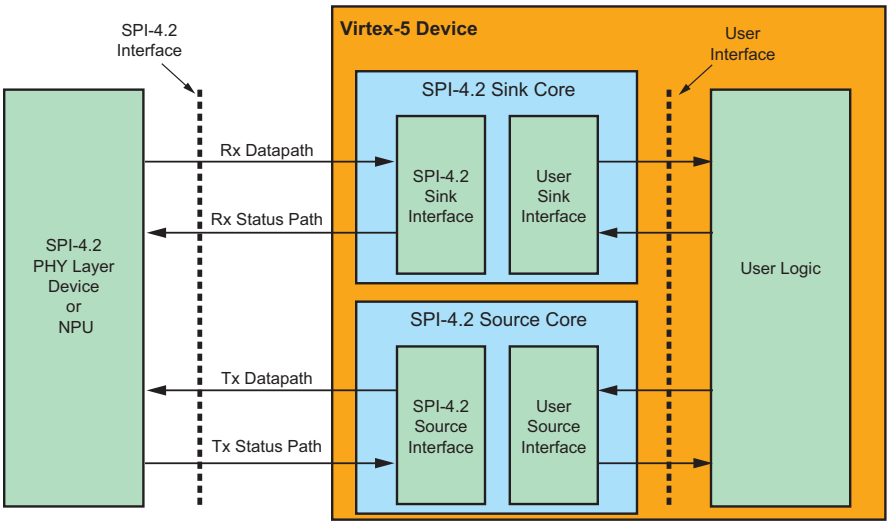

Figure 1 – Typical SPI-4.2 application

In Virtex-5 FPGAs, the IDELAY feature present in every I/O is ideally suited to adjust the clock-data phase relationship for maximum I/O timing margin. This has two primary benefits for the SPI-4.2 core:

- Integrating the IDELAY feature into the input pin (ILOGIC) reduces the FPGA resources required for DPA to less than 350 slices.
- The IDELAY function's ability to adjust the data sampling point enables DPA to be implemented in the I/O – except for a small control state machine implemented in the fabric. The state machine portion is fully synchronous and does not require a complex macro. Thus, there are no restrictions on SPI-4.2 pin assignments.

Continuous DPA

The Xilinx SPI-4.2 LogiCORE solution enhances communication system reliability with continuous DPA, which monitors the clock-data alignment during operation and constantly adjusts the data sampling points to adapt to system timing changes.

Following the initial clock-data alignment phase, the sampling point of each data bit is aligned to the middle of the data valid window (Figure 2a). This window can shift with changes in operating conditions, such as voltage and temperature, as well as other variations (Figure 2b). Continuous DPA addresses this by constantly monitoring the ingress data and adjusting the sample point of each data bit to provide the maximum timing margin (Figure 2c).

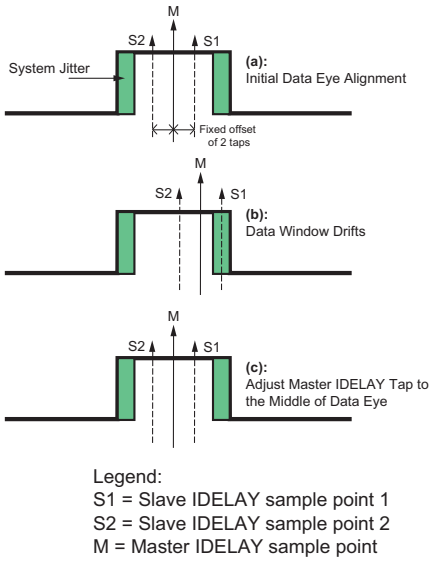

Figure 2 – Continuous DPA operation

Although the OIF SPI-4.2 Implementation Agreement calls for the insertion of periodic training patterns to maintain clock-data alignment over time, Xilinx continuous DPA does not depend on the presence of training patterns. By

reducing/eliminating the need for periodic training patterns, continuous DPA enables the maximum data bandwidth in your system while maintaining the optimal clock-data alignment at each pin.

DPA Diagnostics

If your hardware operation encounters alignment issues, the Xilinx SPI-4.2 core includes DPA diagnostic ports to aid with debugging. The DPA diagnostic data monitors the data eye and final sampling point of the initial alignment process, as well as a second sweep of the data valid window to determine if any changes have occurred.

You can connect the diagnostic ports to the ChipScope™ analyzer or other logic probes to analyze alignment conditions while the FPGA is on the board, interacting with the rest of the system

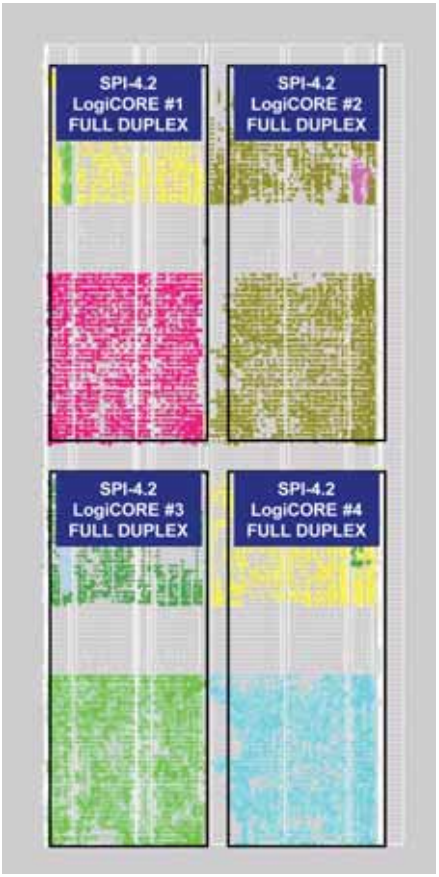

Figure 3 – Illustration of four instances of SPI-4.2 LogiCORE IP implemented on a Virtex-5 XC5VLX110 device

Clocking Resources

Virtex-5 FPGAs provide an unprecedented number of clock resources for implementing multiple SPI-4.2 interfaces in a single device. The abundance and flexibility of clock distribution in the Virtex-5 family solves this challenge, supporting as many SPI-4.2 interfaces as the device logic and I/O will accommodate.

In the Virtex-5 family, all devices have 32 global clock resources, with any 10 of the 32 total global buffers available in each clock region. The global clock trees and associated buffers are implemented differentially for best duty-cycle fidelity and greater common-mode noise rejection.

In addition, each region in the device has four regional clock nets, which are ideal for source-synchronous interface clocking at rates above 1 Gbps. You can configure the SPI-4.2 LogiCORE IP to use either global or regional clock resources.

These high-performance clock resources support as many as four SPI-4.2 interfaces in a mid-range device (LX85/LX110) and more than four SPI-4.2 interfaces in the larger devices (Figure 3). The Virtex-5 clocking capability enables a whole new class of SPI-4.2 applications and provides an ideal platform for applications such as multiplexing and demultiplexing, bridges, and switches.

Higher Performance at Lower Power

Virtex-5 silicon is manufactured with a 65-nm triple-oxide process that reduces power consumption by as much as 35%. This has a positive impact for all designs, including the SPI-4.2 interface; the power savings are summarized in Table 1.

With Virtex-5 devices, SPI-4.2 uses significantly less power than its predecessors, both because of the enhanced 65-nm process and because the LogiCORE solution uses 25% less fabric resources. At the same time, Virtex-5 FPGAs support 20% higher performance for SPI-4.2, with high-speed 1.2 Gbps LVDS data rates on every I/O of the device.

This means that not only can you place multiple SPI-4.2 interfaces anywhere on the device, but for each interface, you can realize an aggregate bandwidth as high as 19 Gbps. Designs not requiring this level of performance (such as more typical framer interfaces running at 10-12 Gbps) automatically get additional performance overhead, ensuring ease of design integration and timing closure.

	Virtex-4 FPGA	Virtex-5 FPGA
Power: Static Alignment @ 700 Mbps per LVDS Pair	1.55W	1.42W
Power: Dynamic Alignment Performance per LVDS Pair	2.0W @ 1 Gbps	1.66W @ 1 Gbps
Speed Grades Supporting 800 Mbps per LVDS Pair	-10, -11, -12	-1, -2, -3

Table 1 – SPI-4.2 power estimates for Virtex-4 and Virtex-5 FPGAs

Conclusion

Xilinx SPI-4.2 LogiCORE IP coupled with Virtex-5 features provides a highly efficient and reliable SPI-4.2 solution. We developed ChipSync technology and continuous DPA specifically for source-synchronous interfaces like SPI-4.2.

This technology allows you to design the most efficient and reliable SPI-4.2 solutions, which use significantly less resources (25% less), allow fully flexible device pin assignments (you choose the pinout), and support extremely high interface speeds (1.2 Gbps LVDS DDR I/O).

The higher performance is even more remarkable because Virtex-5 FPGAs achieve this while consuming significantly less power. The wealth of Virtex-5 clocking resources, combined with full pin assignment flexibility, enables a new class of applications with multiple SPI-4.2 interfaces.

For more information about the SPI-4.2 LogiCORE IP targeting Virtex-5 devices, visit the Xilinx IP Center at *www.xilinx.com/systemio/spi-4.2*. A hardware demonstration is also available; for more information, contact your Xilinx representative.

Using Virtex-5 FPGAs in COTS Board-Level Products

Optimize your COTS designs with the many improvements in the Virtex-5 family.

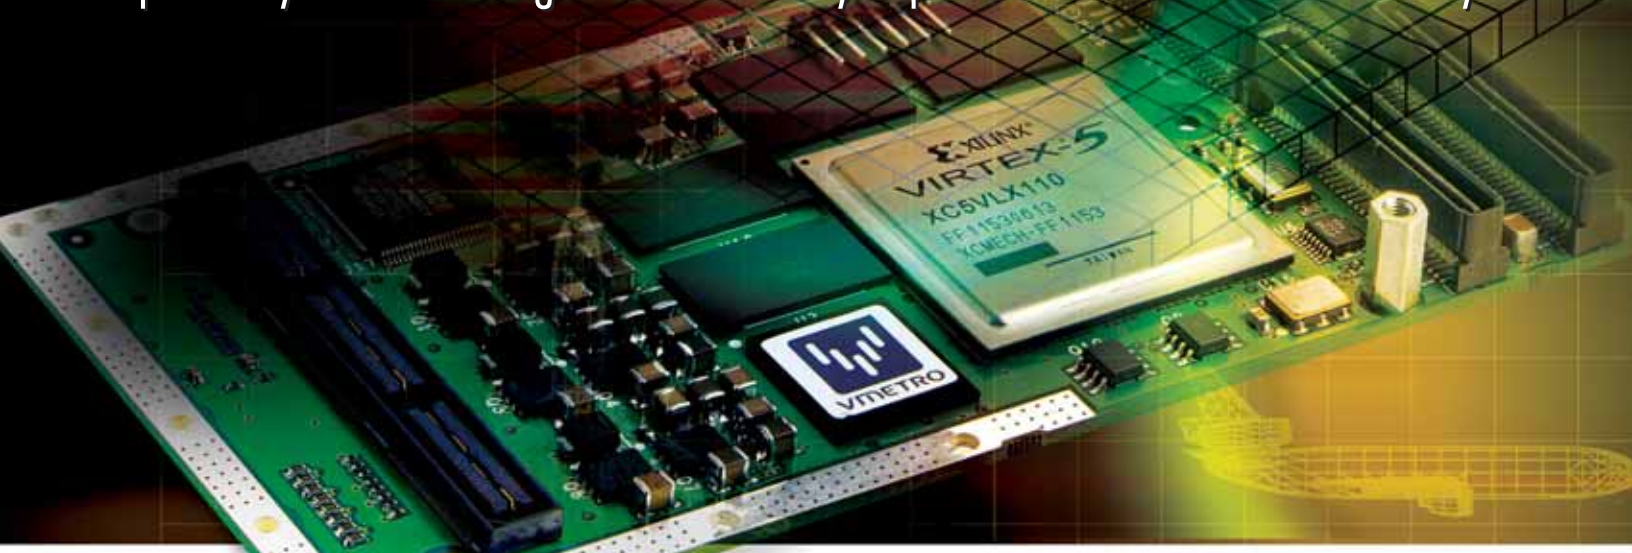

by Craig Davies
Firmware Engineer
VMETRO Ltd. (High Wycombe, UK)
cdavies@vmetro.com

Jeff Bateman
Senior Systems Engineer
VMETRO Inc. (Ithaca, NY)
jbateman@vmetro.com

In the fast-paced world of FPGA development, Xilinx has struck again with its second-generation ASMBL™ architecture devices, the Virtex™-5 family. This device family has many upgrades from its predecessor, the Virtex-4 family, and likewise continues the evolution of the ASMBL architecture, with scalable FPGAs catering to the application-specific marketplace. For commercial off-the-shelf (COTS) developers, this means a platform that is low cost, light on power consumption, and optimized for high performance.

When creating the Virtex-4 family, Xilinx harnessed the flexibility of the ASMBL architecture to build the first multiplatform FPGA family. Xilinx continues this approach with the Virtex-5 family. The initial offering is the Virtex-5 LX platform, optimized for high-performance logic.

Seasoned FPGA users expect new FPGA generations to deliver more and the Virtex-5 family certainly delivers, all while consuming less power. Compared to Virtex-4 LX devices, Virtex-5 LX FPGAs offer:

- 65% higher logic capacity with as many as 330,000 logic cells
- 70% more block RAM
- 100% more DSP slices
- 25% more SelectIO™ pins

For COTS board vendors, these features enable powerful products capable of handling the very high data rates and processing complexity required of modern real-time DSP systems. In this article, we'll examine those Virtex-5 architecture components that enable COTS designers to deliver more bang for the buck.

COTS FPGA Backgrounder

Freed from the need to design hardware and IP from scratch, COTS board-level users can focus their energies on implementing their specialist algorithms. COTS products incorporating user-programmable FPGAs target a variety of applications, from simple customizable digital I/O to RADAR, video, and signals intelligence (SigInt).

Typically, the FPGA requires hardware connections to a real-world data source or destination, plus a standardized interface to a host processor. COTS products must usually follow an industry-standard form factor (such as PMC, VME, VXS, and CompactPCI), enabling end users to integrate products from a range of vendors.

With its parallel architecture and high-speed I/O capabilities, the Virtex-5 FPGA is capable of streaming and processing data at the gigabyte-per-second rates typically required for today's applications. It is well suited to algorithms where a core "inner loop" can be parallelized to speed up operation, employing the resources available in modern devices. Many DSP algorithms dovetail with this architecture. Conversely, even the fastest CPUs cannot easily process data at gigabyte-per-second rates; they are, however, well suited to decision making and user interaction tasks.

Given these trade-offs, FPGA-based DSP systems often employ a hybrid approach, as illustrated in Figure 1. Here, a wide-bandwidth RADAR or video source is digitized at gigasample-per-second rates and fed to an FPGA. The FPGA performs some heavy-duty number crunching to eliminate unwanted data, focusing in on the key area of interest. Pre-processed data is fed at a more manageable rate to a general-purpose CPU for post-processing control and display.

Key COTS FPGA board requirements are:

- Large, reconfigurable FPGAs with ample room for customer-programmable application logic
- Regular air-cooled and rugged conduction-cooled options
- High-speed interface for efficient transfers to and from a host processor
- Flexible, fast I/O to and from a variety of real-world interfaces
- Local memory interfaced directly to the FPGA for I/O buffering as well as temporary storage during algorithm operation
- Wide range of I/O and signal-processing IP cores to speed end-user development cycle times
- Flexible FPGA development tools covering both budget-conscious and extreme-performance users
- Debugging interface capable of in-FPGA logic analysis
- Comprehensive board support firmware and software

With a proven track record in the high-end FPGA DSP arena and comprehensive tool and IP support from a variety of sources, the Virtex-5 FPGA family is a natural choice for COTS board-level vendors.

Optimization of Soft Components

A great addition to the Virtex-5 architecture is the replacement of traditional four-input look-up tables (LUTs) with new six-input LUTs (6-LUTs) for more efficient mapping of wider functions. Because 6-LUTs are also configurable as dual five-input LUTs, design software tools can achieve greater efficiency in logic mapping when six-input functions are not required.

Most FPGA devices these days base their soft fabric components – those components configured to implement logic equations – on LUTs. Previously, the common choice was the four-input LUT, as this was a nice binary base and was relatively easy to work with for optimizing a logic function. A given equation can be optimized to contain a sum of products of four inputs. For many of today's applications, especially those in DSP, this optimization reduces significantly as system-level algorithms increase in complexity.

Configurable logic block storage density improvements increase the shift register LUT (SRL) length from 16 bits to 32 bits (SRL32), while retaining a dual SRL16 option. Distributed RAM now offers a 64-bit option, up from 16 bits. With improved reduced-hop routing and more logic per slice (four LUTs/four flip-flops versus two LUTs/two flip-flops), speed improvements of as much as 45% are possible.

Improved FIR Efficiency

Let's consider a finite impulse response filter implemented in distributed logic. Distributed arithmetic filters are often selected because their operating frequency is not tied to the length of the tap vector. This characteristic is highly desirable because increasing the tap vector length is fundamental to improving the overall filter response. However, these types of filters are

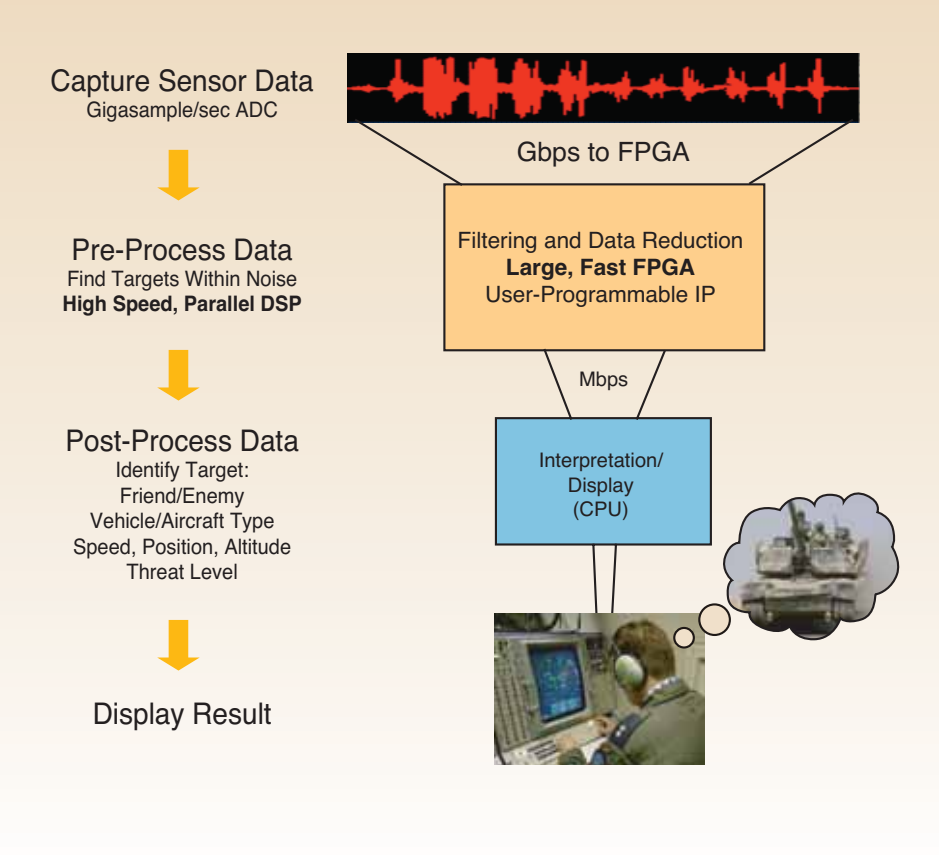

Figure 1 – Processing chain

serial in nature; most applications need valid output at a marginally decimated rate with respect to the sampling frequency. Thus, a fully parallel architecture is required.

Parallel Distributed Arithmetic FIR filters (DAFIRs) utilize significantly more logic versus other FIR implementations to perform the many partial products on a clock-to-clock basis (even when decimating the sampling rate). In the distributed arithmetic architecture, a corresponding output product y(n) is produced by summing the products of a time-delayed series of input x(n) and coefficients a(m), where m, an integer between 1 and N, is the filter length.

For the sake of simplicity, let's say that each tap or filter coefficient is two bits wide and that the input vector is six bits. In total, our filter is 96 taps in length. If we calculate this product using the partial products method, we need 6 x 2 partial products using four-input LUTs. Each LUT is capable of a 2 x 2 multiplication, which means using three 4-input LUTs. Using 6-LUTs, we can reduce this to just two LUTs. For 96 taps, we have saved 96 LUTs of a possible total of 288. This is just the savings when producing the partial product.

LUTs and SRLs are also used for shift registers in the input delay pipe and for the scaling accumulator responsible for summation and normalization of the output. Expanding our example input and tap widths to a more applicable precision of 16 bits increases the depth of our partial product multiplication tree, requiring even more LUTs. Using 6-LUTs as opposed to four-input LUTs results in a LUT logic reduction of more than 33%.

Wider LUTs Improve Efficiency and Speed

Switched fabric developers will also benefit from the 6-LUTs, as these are often used to implement multiplexers. 6-LUTs mean a reduction in the overall depth of a logic equation. For implementing multiplexers, this means an effective increase in speed for an equivalent multiplexer implemented using 6-LUTs, as opposed to four-input LUTs.

Depending on the application, changing to 6-LUTs can make as much as a 1.6x improvement in logic utilization over previous generations of the Virtex device family.

Multiplying Computational Power

Not to be outdone by the soft-logic components, the hard-logic dedicated multipliers have also been optimized for the Virtex-5 FPGA. The 18 x 18-bit multipliers present in the Virtex-II and Virtex-4 families have been upgraded to 25 x 18-bit multipliers in the new family. Application developers who implement beam-forming arrays or other advanced computations will benefit from this enhancement.

Large multiplication arrays that require a high degree of precision traditionally required a large tree structure of multipliers. As output is carried between intermediary stages of a large multiplication, the maximum allowable output value increases with each subsequent stage. To handle this bit-width increase, typical solutions involve a precision reduction by truncation or some other intelligent scheme such as convergent rounding or (less often) by breaking down the multiplication into smaller stages and then rebuilding the final product by summation. Utilizing 25 x 18-bit multipliers, more precision is carried through intermediary stages of a multiplication and thus reduces the impact of intermediary truncation/rounding errors while improving on overall speed and minimizing pipeline latency.

Suppose convergent rounding is employed to reduce the precision at each stage of multiplication within an FFT. If we implement an 8K FFT using a mixed-radix base of radix-4 and radix-2, that gives us six radix-4 butterfly stages and one radix-2 butterfly stage. In an FFT, at each subsequent stage we perform calculations that produce partial products. These outputs are fed into the multipliers of the next stage until the time-domain data is transformed completely to the frequency domain. However, each stage must employ a scheme to reduce the precision of the output so that subsequent stages can accept them as inputs. After each stage of multiplication, scaling is employed to reduce the precision. Each stage of scaling introduces quantization errors.

But what if more precision is carried into the input side of each butterfly stage? Using 25 x 18-bit multipliers, we can carry more precision from our partial products when multiplying them to new sample data and in turn introduce less rounding errors into our results.

Improved Source-Synchronous Memory Access

Much to the delight of many designers using Virtex-4 FPGAs, Xilinx introduced a primitive called IDELAY capable of synchronizing data and strobes to a source clock off the FPGA. This feature meant that high-speed DDR and DDR2 SDRAM and QDR and QDR II SRAM memories could be accessed through controllers inside the Virtex-4 device at high data rates.

COTS developers are increasingly finding applications that require fast and deep onboard memory. For example, data recording applications benefit greatly from fast onboard memory to implement the sizeable buffers needed to sustain high-speed data transfers over PCI/PCI-X buses. Video processing applications also require large, fast external memories to store the very-high-resolution, high-frame-rate images produced by today's leading camera equipment.

The introduction of the IDELAY primitive also benefits ruggedized application developers, as the IDELAY taps can be constantly monitored by logic to perform run-time resynchronization to the source clock; this technique is known as dynamic clock-to-data centering.

Now, with the Virtex-5 family, Xilinx has expanded the primitive to add ODELAY, enabling delay control on both input and output signals. The key component of the IDELAY primitive is to delay the input data relative to the clock such that the internal FPGA version of the source clock edge is centered with the input data. The ODELAY enables variable delays per output data line to better match trace-length differences.

Improving High-Speed I/O Communication

As the COTS marketplace moves more and more into high-speed serial implementations, clock and data recovery techniques

become more in demand. When implementing general-purpose high-speed serial links, transmission errors and data loss become a reality, especially when targeting data rates beyond 1 Gbps.

For developers using previous generations of Virtex devices, the choices for clock and data recovery (CDR) implementation to de-serialize incoming streams without using multi-gigabit transceivers (MGTs) were limited. Although the delay-locked loops (DLLs) used for clock generation in previous families are very stable in nature because of their first-order loop architecture and digital implementation, they are not able to filter input jitter or handle phase alignment beyond their discrete range.

With the phase-locked loop (PLL) blocks introduced with Virtex-5 family, jitter reduction is an intrinsic feature, resulting in large improvements in higher data-rate sustainability. Filtering input jitter to produce stable internal versions of source clocks is critically important to correctly sample and store incoming data at the FPGA I/O boundary. Using these new blocks, implementing SERDES components using regular SelectIO pins becomes practical even at 1 Gbps and above.

Together with SelectIO performance of as much as 800 Mbps per pin single-ended and 1.25 Gbps differential, the Virtex-5 device is able to input, process, and output the high data rates generated by current real-world interfaces. For example, interfacing directly with ADCs and DACs running in the gigasample-per-second range is now perfectly feasible.

Looking to the future, new high-speed serial fabric interfaces will be a natural fit for interfacing between FPGAs, external devices, and host systems.

65-nm Copper CMOS

COTS developers will greatly benefit from the move into the 65-nm copper CMOS process. One of the consequences of process shrinks is that density and performance increase with the next generation. This is true in the case of the Virtex-5 LX platform, which has increased the amount of CLBs by 65% over the Virtex-4 LX platform, block RAM by 70%, DSP slices by 100%, and SelectIO pins by 25%.

Figure 2 – VMETRO PMC-FPGA05

Figure 3 – Standard PCI option

With increased logic density in the overall package, power consumption has been reduced significantly. While the Virtex-4 FPGA operates at 1.2V core voltage, the Virtex-5 FPGA improves power efficiency with a core voltage of 1.0V. You can achieve further power savings by optimizing the soft and hard components, such as the 6-LUTs and 25 x 18-bit multipliers.

When implementing large designs – such as in software-defined radio (SDR) applications where multi-channel digital filters consume significant CLB space – the dynamic power dissipation is quite high because of the large amount of switching activity that occurs. This is in part caused by the extensive signal routing required to implement these designs. With the new components in the Virtex-5 family, existing designs implement in a smaller number of primitives, reducing the overall switching activity. In addition, the Virtex-5 architecture includes the enhancement of diagonally symmetric routing for more efficient design implementation.

COTS developers often complain about violating power specifications when developing mezzanine cards in existing designs. With the Virtex-5 device, their mezzanine cards will be less power-hungry and more desirable to end users. In other words: less power required, simplified cooling, and greater reliability.

COTS products are often employed in environments where power consumption is a significant challenge – for example, high ambient temperatures may limit a cooling system's effectiveness, so reducing heat output is an important motivation. Other applications such as unmanned airborne vehicles (UAVs) have limited electrical power availability, so using every Watt available effectively is of paramount importance.

The VMETRO PMC-FPGA05

A good example of a current COTS product implementing the advanced Virtex-5 65-nm technology is the VMETRO PMC-FPGA05, a general-purpose high-end FPGA PCI mezzanine card (PMC) pictured in Figures 2 and 3 and illustrated in the block diagram shown in Figure 4.

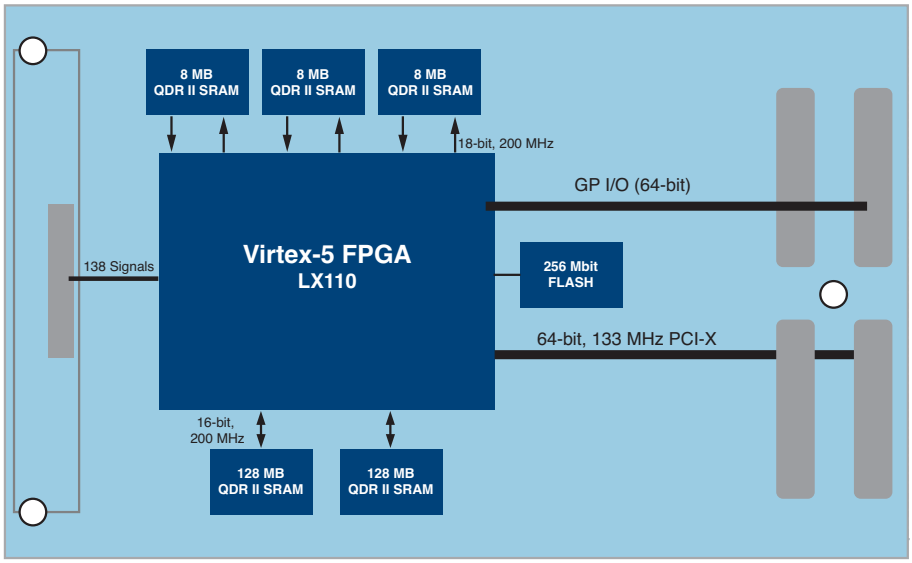

Figure 4 – PMC-FPGA05 block diagram

A Virtex-5 FPGA at the heart of a PMC-FPGA05 means that it is a highly integrated design. There is no need for external bridges and controllers because the Virtex-5 device is more than capable of handling these functions directly (see Figure 5) while using only a small amount of resources. This leaves plenty of space in the PMC-FPGA05 to include IP such as digital receivers, FFTs, or other DSP functions.

Implementing Flash and PCI-X interfaces inside the FPGA introduces some design challenges. How is the FPGA configured and debugged without these interfaces already in place – especially if the IP is complex? This is where the ChipScope™ Pro Analyzer (version 8.2) comes in – it embeds a logic analyzer inside the FPGA and connects the user interface to the FPGA through JTAG. It provides a debugging portal into the FPGA that can be inserted through HDL entry.

Once built, a ChipScope Pro ILA (integrated logic analyzer) port assignment can be changed in FPGA Editor. Therefore, changes to the ILA do not require another full run through the ISE™ flow when debugging, as shown in Figure 6. Instead, you can make changes at the map stage of the design flow, reducing the time between generating bitstreams and improving efficiency during the debugging cycle. The ChipScope Pro Analyzer consumes a limited amount of the FPGA's resources, though this is largely a function of the depth of the analyzer sample memory. Features such as the analyzer memory depth and triggering functions are parameterizable through the ChipScope Pro inserter tool, eliminating unnecessary resource waste.

The ChipScope Pro Analyzer is a valuable asset when debugging designs such as the VMETRO PMC-FPGA05. The PCI-X interface in particular represents a challenge. In simple terms, this might be to power up the PCI bus, configure the IP onto the FPGA, and re-initialize the PCI bus to detect the FPGA's PCI controller. This low-level development requires a tool like the ChipScope Pro Analyzer to reduce the number of bus initialization cycles. Without the ChipScope Pro tool, debugging would require a sophisticated bus analyzer – and even then the internal FPGA functionality would be inaccessible. Utilizing the ChipScope Pro Analyzer saves the design team considerable time by reducing the number of iterations required to debug the PCI-X interface.

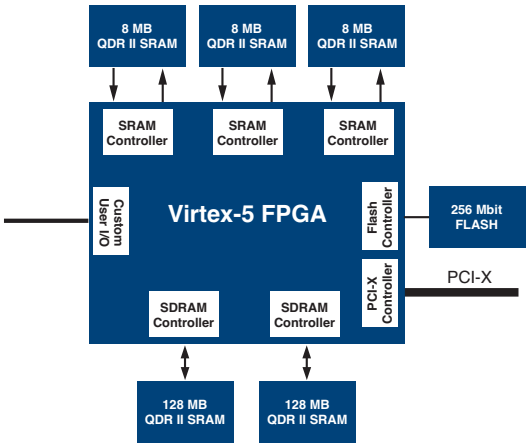

Figure 5 – PMC-FPGA05 controller IP components

The ChipScope Pro tool is not only useful for VMETRO board development, but also for customers integrating their IP and external interfaces. We designed the PMC-FPGA05 with a high-density parallel interface that you can use with a range of I/O modules including analog I/O, RS485, LVDS, FPDP, and Camera Link. Using the ChipScope Pro Analyzer makes these interesting projects manageable.

The PMC-FPGA05 also offers application developers a platform that can sustain high-bandwidth data transfers and implement sophisticated DSP and processing algorithms at a fraction of the power of previous generations.

FPGA Resources

More than any other FPGA family, the Virtex series leads the way in IP availability. Through the Xilinx® Alliance Program, IP is available from both Xilinx directly as well as third-party IP suppliers. With a strong worldwide network of IP vendors accessible directly from *www.xilinx.com*, COTS FPGA board users can choose IP cores from suppliers experienced with as many as five generations of the Virtex family. This is key to success when developing projects to aggressive timescales, and therefore is a high-priority reason for selecting the Virtex-5 family.

FPGA Firmware Development Process

With the Xilinx ISE tool chain offering an easy-to-use GUI development environment and an included VHDL/Verilog synthesis tool (XST), many end users need only purchase ISE software for a complete cost-

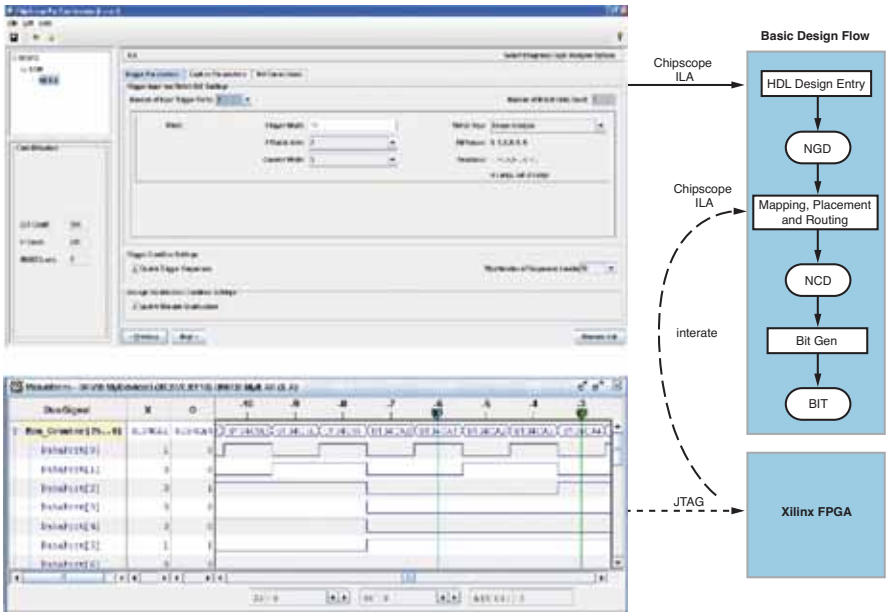

Figure 6 – ChipScope debugging cycle

effective development solution. Those implementing the most complex algorithms may benefit from high-end synthesis tools from third-party vendors. These integrate directly with ISE software to maintain a simple project management process.

Simulation may not reveal all errors in the design, particularly for complex projects. When an implementation does not function as expected, you must connect up to the actual hardware to see what is going on. But with dense packaging covering many thousands of pins, there's no practical way to connect a traditional logic analyzer.

VMETRO engineers have many years of experience developing firmware for Xilinx FPGAs. Using ISE software as the primary synthesis/place and route tool, Model Technology's ModelSim PE for simulation, and the ChipScope Pro tool for in-circuit debugging, VMETRO developed the IP necessary for interfacing with the board hardware. This includes interfaces for the SRAM and SDRAM memory devices, to which users simply connect their address and data signals, and a high-performance bus mastering PCI-X interface core supporting customizable registers and simple FIFO-based DMA transfers for streaming data.

Another good reason to choose the Virtex-5 family is a commitment to education and support: from introductory courses in VHDL to DSP logic design courses to development laboratories equipped with the latest gear for testing high-speed serial interfaces, Xilinx offers the resources necessary for successful FPGA deployment.

As the Virtex-5 device is brand new, the PMC-FPGA05 is still in development. Stay tuned for an update on the challenges, solutions, and lessons learned as VMETRO works hard to bring you the world's first Virtex-5 COTS product.

Conclusion

To meet the growing demands being placed on the COTS marketplace, you must adapt and implement platforms with the right tools for the application. Today, this means integrating high-speed serial communication, fast access memory, and plenty of optimized logic space for advanced algorithm development. Equally important are efficient development and debug tools, IP resources, and a commitment to high-speed DSP development.

The highest logic density available, the lowest power consumption, and the best performance are what COTS developers need to meet the needs of their customers. Virtex-5 FPGAs, with their ASMBL architecture and 65-nm process, deliver on these demands.

Xilinx Worldwide Events

Xilinx participates in numerous trade shows and events throughout the year. This is a perfect opportunity to meet our silicon and software experts, ask questions, see demonstrations of new products, and hear other customer success stories.

For more information and the current schedule, visit *www.xilinx.com/events/*.

January 06 - 10, 2007	Int'l Conference on VLSI Design & Int'l Conference on Embedded Systems	Bangalore, India
January 08 - 11, 2007	International Consumer Electronics Show	Las Vegas, NV
February 12 - 15, 2007	3GSM	Spain
February 13 - 15, 2007	Embedded World	Germany
February 21 - 23, 2007	IDGA Software Radio Summit	Vienna, VA
April 03 - 05, 2007	Embedded Systems Conference - Silicon Valley	San Jose, CA
April 16 - 19, 2007	NAB	Las Vegas, NV

Tackling Serial Backplane Interface Design Challenges

The Virtex-5 LXT FPGA enables robust, high-performance, and high-integration serial backplane interface solutions.

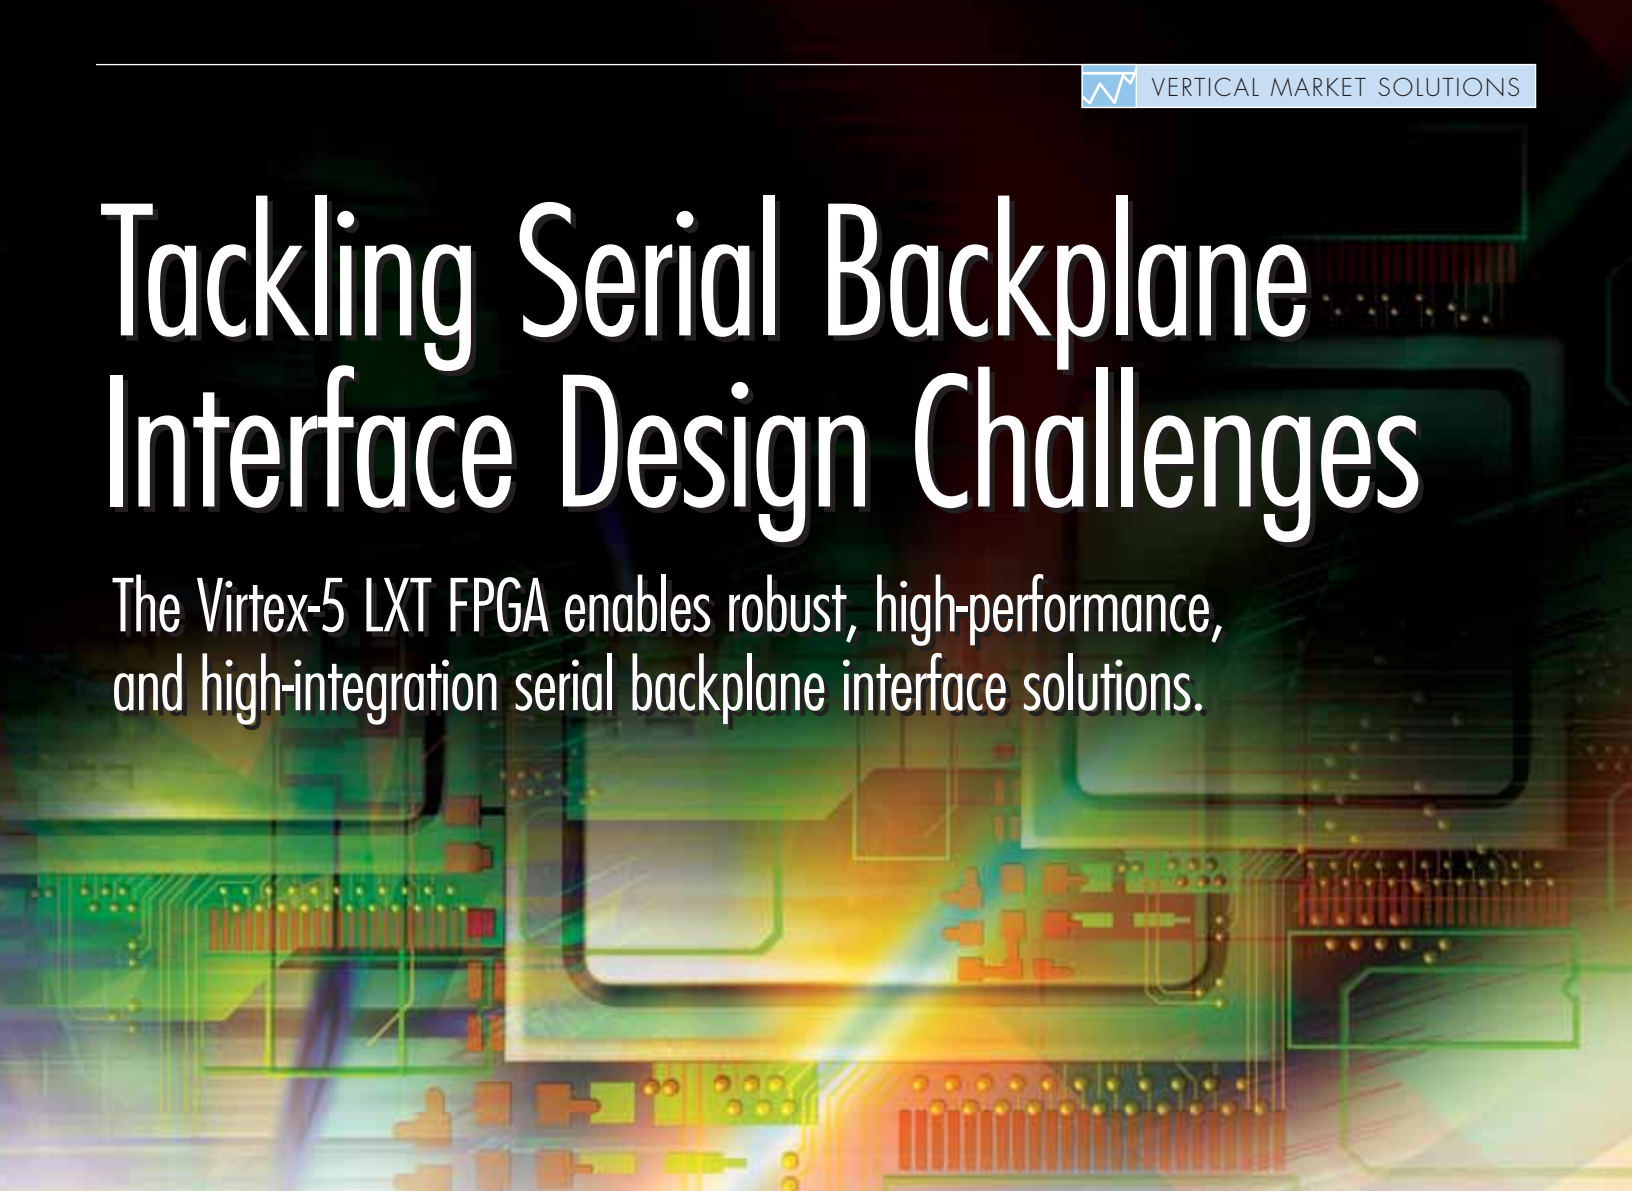

by Delfin Rodillas
Senior Manager, Wired Communications
Xilinx, Inc.
delfin.rodillas@xilinx.com

The rate of adoption of serial technology in high-end system design has reached critical mass. As shown in Figure 1, 92% of respondents in a recent *EE Times* survey answered "yes" when asked if they were designing serial I/O systems in 2006, compared to 64% serial design activity in 2005.

A good portion of this dramatic adoption rate is caused by the penetration of serial technology in backplane applications. As system throughput requirements increase, the parallel backplane technologies of old will be displaced by SerDes-based backplane subsystems that provide higher bandwidth, better signal integrity, lower EMI and power, and simpler PCB designs.

Further promoting this growth is the emergence of standard serial protocols such as XAUI and Gigabit Ethernet (GbE), which allow reduced engineering efforts and interoperability. Standardization efforts for serial backplane form factors such as AdvancedTCA and MicroTCA in the PCI Industrial Computer Manufacturers Group (PICMG) have also contributed to the accelerated adoption. The benefits of serial backplanes are so compelling that they have been used as the backbone of not only communications, compute, and storage systems but also broadcast, medical, defense, and industrial/test systems.

Persistent Design Challenges

Regardless of the increased rate of adoption, many design challenges still exist. Because the backplane subsystem is the heart of the system, it must be able to pass signals from card to card reliably. Thus, designing backplanes with high signal integrity (SI) is of primary importance.

Also significant is the use of proper silicon ICs with SerDes technology, capable of driving backplanes with very low bit-error rates. Silicon-based approaches to mitigating SI issues are particularly important in "legacy upgrade" scenarios, in which designers re-use older backplanes with legacy components and design rules.

There are also challenges in developing serial backplane protocols and fabric interfaces. The majority of backplane designs leverage legacy ASICs, which have proprietary protocols. Even some newer backplane designs require a proprietary backplane protocol. Silicon solutions must therefore be flexible and provide the necessary customizability. Although an ASIC allows this, it can often be costly and risky, with unproven product demand/volume and the possibility of design bugs and specification changes.

An approach that has recently gained traction is the use of off-the-shelf standards-based switch fabrics. This saves development time, but you must have silicon solutions that conform to the standard

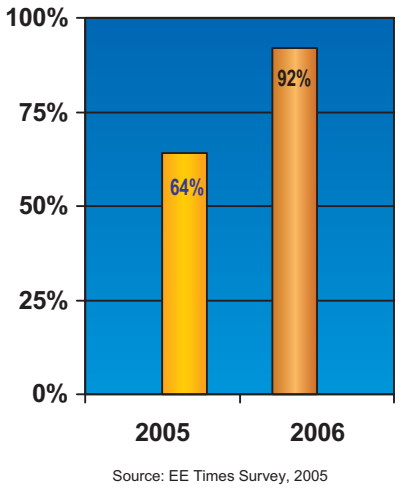

Figure 1 – Percentage of engineers designing Serial I/O systems

protocol, as well as the flexibility to customize the end product and make it unique.

And of course, there are the ever-present challenges of cost, power, and time to market. To meet the challenges of serial backplane design, Xilinx provides the Virtex™-5 LXT platform of FPGAs as well as IP solutions.

Xilinx Solutions for Serial Backplanes

The key technology that enables the application of Xilinx® Virtex-5 LXT FPGAs in serial backplane applications is the embedded RocketIO™ GTP low-power serial transceiver. There are as many as 24 serial tranceivers in the largest Virtex-5 LXT FPGA; each serial tranceiver is capable of running from 100 Mbps to 3.2 Gbps. Coupled with programmable fabric, the FPGA is capable of supporting virtually any serial protocol – proprietary or standard – up to 3.2 Gbps.

More important for serial backplane applications are built-in signal conditioning features, including transmit pre-emphasis and receive equalization. These features enable transmission of multi-gigabit signals over long distances, often reaching 40 inches or longer. Both equalization methods minimize the impact of inter-symbol interference (ISI) by boosting high-frequency signal components and attenuating low-frequency components. The difference is that pre-emphasis is performed on the transmitted signal as it goes out of the line driver, while equalization occurs on the received signal after it enters the IC package. Both pre-emphasis and equalization features are programmable to different states to allow for optimum signal compensation.

Besides signal conditioning features, the serial tranceivers also provide additional features beneficial for backplanes, such as programmable output swings that allow interfacing to a variety of other current mode logic (CML)-based devices and built-in AC coupling capacitors that simplify transmission line design and reduce ISI.

IP Cores

Proprietary protocols still make up most serial backplane implementations. However, some newer designs have used standards-based protocols such as XAUI and GbE. This growing acceptance has been driven primarily by the maturity of these standards and the emergence of switch fabric ASSPs utilizing these protocols. Using ASSPs for switching applications saves tremendous development time, but designers realize that they need to differentiate their products by adding value-added capabilities, primarily on the line card.

FPGAs are the ideal platform for providing customizability, as the serial tranceivers are designed to support a majority of standard serial backplane protocols. Together, the serial tranceivers and fabric allow for standards-compliant designs with value-added functions – all in a single silicon device.

To reduce design time, Xilinx offers off-the-shelf available IP cores for key serial I/O interface standards such as XAUI, GbE, SRIO, and PCIe. To ensure interoperability, these IP cores are tested through consortia plug-fests and independent third-party verification. To facilitate the creation of light-weight serial protocol designs, Xilinx also created the Aurora protocol, which is ideal for simpler designs requiring minimal overhead and optimized slice/resource utilization.

With increased usage of Ethernet and PCIe, Virtex-5 LXT FPGAs also include embedded tri-mode Ethernet MACs and PCIe Endpoint blocks. These allow significant savings of FPGA slice resources for customers needing interfaces in control plane applications, for example.

Because many chips with parallel interfaces are still used even in newer systems, Xilinx also offers IP cores for popular parallel interfaces such as SPI-4.2, SPI-3, and PCI. These allow you to rapidly create serial-to-parallel bridges, which are still required in many applications.

Besides serial and parallel interface IP, Xilinx offers more complete IP solutions that further reduce development time and time to market. These solutions include a Traffic Manager for prioritizing traffic flows across backplanes, as well as a Mesh Fabric Reference Design that allows "every-to-every" connectivity between cards. Lastly, the ChipScope™ Pro Serial I/O Tool Kit enables rapid serial tranceiver setup and debugging as well as BERT testing. Table 1 summarizes the serial backplane-related IP available from Xilinx.

Application Examples

Let's look at how you could integrate all of the solution components to create a complete serial backplane fabric interface FPGA for both a star and mesh system.

IP Category	Available IP
Serial Interfaces	XAUI, GbE, PCI Express, Serial RapidIO, Aurora, CPRI, OBSAI
Parallel Interfaces	SPI-4.2, SPI-3, Utopia, PCI, CSIX
System-Level Solutions	10G Traffic Manager, Mesh Fabric Reference Design
Serial Backplane Test Solutions	ChipScope Pro Serial I/O Tool Kit

Table 1 – Xilinx IP for serial backplanes

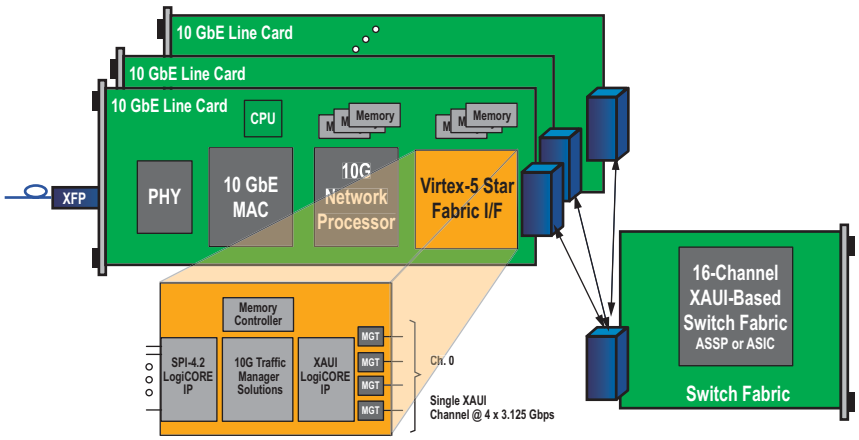

Figure 2 – Star fabric I/F FPGA in a 10 GbE line card

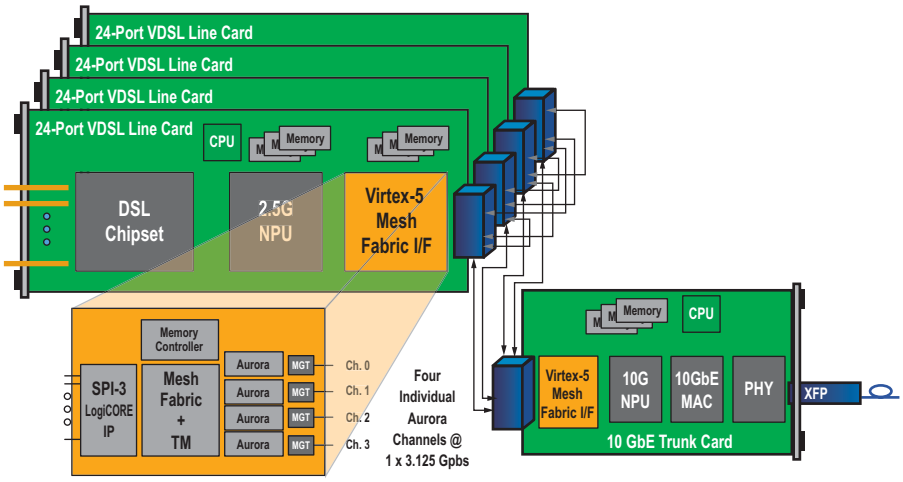

Figure 3 – Mesh fabric I/F FPGA in a VDSL line card

Star Backplane Topology Application

Star fabric topologies prevail in high-end infrastructure equipment because of their cost-effectiveness, particularly in systems with a high number of cards. Figure 2 is an example of a 10 GbE line card that implements an FPGA-based star fabric interface. This FPGA instantiates the XAUI LogiCORE™ IP and uses four serial tranceivers to connect to the 16-channel XAUI switch fabric card. A LogiCORE SPI-4.2 core is also realized in the FPGA to interface to the 10 Gbps network processing unit.

Between the serial and parallel interface is the Traffic Manager IP, which performs QoS-related functions on ingress and egress traffic. A memory controller controls the external memory, which is used primarily as packet buffers. The benefits of this architecture include increased integration of SerDes and logic functions and quick time to market through the use of IP, while allowing an implementation that meets your exact system specifications. It also provides solid SI as well as low SerDes power consumption (~400 mW total). You can implement all of this in the lowest cost speed grade XC5VLX50T device.

Mesh Fabric Architectures

Star topologies prevail in most cases, but in some smaller systems, a mesh topology is required. Take the case of the five-slot IP DSL access multiplexer shown in Figure 3, which requires full connectivity between four 24-port VDSL line cards and a 10 GbE backhaul card that connects to a metro Ethernet network. Each card uses a Virtex-5 LXT device and four embedded serial tranceivers to realize the four independent channels of the mesh fabric physical layer. Implementing the four link layers is the Aurora protocol, which runs at approximately 3 Gbps to transport the 2.4 Gbps payload – plus additional overhead such as the encoding.

A SPI-4.2 and SPI-3 LogiCORE IP is used on the trunk card and line cards, respectively, providing connectivity to the network processor. The Mesh Fabric Reference Design and Traffic Manager solution provide the distributed switching and QoS functions required on all of the line cards.

The line card fabric interface could easily fit in an XC5VLX30T device, while the trunk card fabric could fit in an XC5VLX50T device. Similar to the star example, you can realize significant benefits in integration, time-to-market reduction, feature optimization, and power and cost reduction by using the Virtex-5 LXT solution.

Conclusion

Serial backplane technology is now mainstream; its adoption will only continue to increase with the rapidly growing demand for bandwidth. Evolution of backplane system requirements in terms of rates and protocols is inevitable and designers will face new challenges.

However, with Xilinx Virtex-5 LXT FPGAs and off-the-shelf-available IP for serial backplanes, system architects have an option that can accommodate legacy as well as newer backplane designs. Virtex-5 LXT FPGAs with embedded SerDes have the critical SI-improving features and integration required to provide high reliability and area- and cost-optimized designs.

Furthermore, Xilinx off-the-shelf IP reduces development time and time to market. Together, the powerful silicon and IP cores are what make the Virtex-5 solution the ideal vehicle for tackling even the toughest serial backplane design challenge.

For more information, visit *www.xilinx.com/backplanes* and *www.xilinx.com/qos*.

Where Great Ideas Go to Work

A career with Xilinx® puts you at the Leading edge of technology.

The world leader in programmable systems, Xilinx solutions are found in numerous applications including wireless, networking, storage, automotive, aerospace, and much more.

Visit our website today, ***www.xilinx.com/hr/ncg/index.htm,*** and let's talk about putting your ideas to work.

XILINX®
The programmable Logic Company℠

Enabling Multi-Port 1 Gbps and 10 Gbps TCP/iSCSI Protocol Offload Solutions

The Virtex-5 LXT platform enables low-footprint, system-level, multi-port 1 Gbps and 10 Gbps TOE solutions.

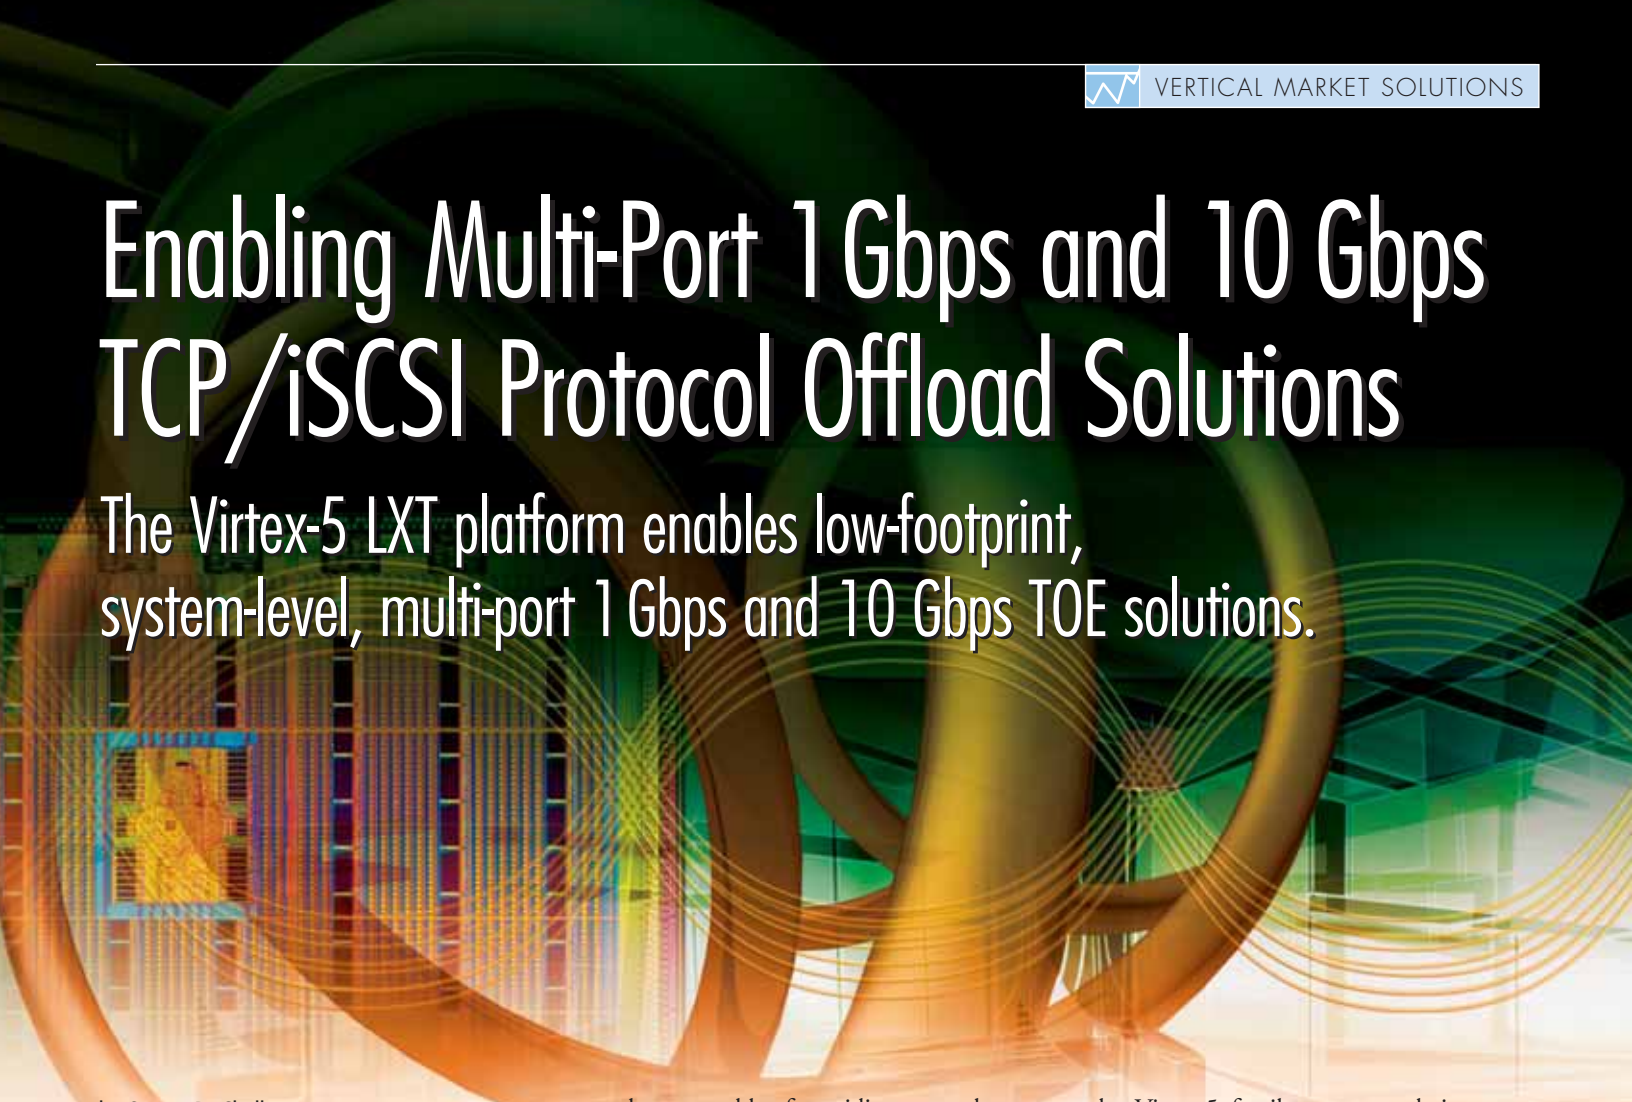

by Sriram R. Chelluri
Senior Manager, Storage and Servers
Xilinx, Inc.
sriram.chelluri@xilinx.com

As the data center network infrastructure migrates to 10 Gbps, moving data traffic to an Ethernet-based solution becomes economically viable without sacrificing performance and latency. Hardware-based host interfaces like PCI Express and multi-Gigabit Ethernet (GbE) support open up design possibilities for low-cost, high-performance products in the computer and data-processing markets. The Xilinx® Virtex™-5 family of FPGAs sets the stage for designing system-on-chip (SoC) solutions with higher functionality and low power.

The Virtex-5 architecture brings to market critical features that make SoC designs easy to implement for TCP and iSCSI offload engines:

- Built-in PCI Express (PCIe) block – An integrated standards-compliant PCIe endpoint for supporting one to eight lanes, capable of providing as much as 32 Gbps full-duplex performance
- Built-in Gigabit Ethernet MAC (GEMAC) – four hardcore GEMACs enable multi-port gigabit solutions, reducing total real estate requirements for SoC designs
- Real six-input LUT (6-LUT) technology – improves slice utilization and reduces routing latency for high performance
- 36-Kb dual-port block RAM – higher memory density with error-correction circuitry enables support for reliable computational logic structures and increased on-chip TCP sessions for simultaneous transmit and receive operations
- DSP48E slices – enable massively parallel computations for image processing and multimedia applications

Because the Virtex family is a programmable platform, you can adapt your designs to changing standards and market requirements. Leveraging the resources available in the Virtex-5 family, you can design cost-effective TCP and iSCSI offload solutions for the server, storage, multi-protocol switch, and wireless base station markets with extended product life cycles.

TCP Offload Engine (TOE) Overview

Current TCP offload solutions rely on a complete software stack or on special network interface cards (NICs) based on ASICs for handling TCP/IP processing. An all-software solution is acceptable for low-bandwidth applications, but high-performance applications would consume all of the CPU resources, creating a system bottleneck for critical applications.

ASIC-based solutions are primarily from start-ups looking to capitalize on the high-performance 10 Gbps market. These solutions are still expensive and prone to vendor lock-in with an uncertain financial future.

Xilinx and its third-party IP partners provide fully standards-compliant TCP/iSCSI offload solutions that you can implement as is or customize for functionality, size, speed, or the target application.

FPGA-Based TCP/iSCSI Engine

With standards-compliant built-in GEMACs, a PCIe core, and increased block RAM, the Virtex-5 LXT device is a programmable platform chip that system architects can exploit for TCP and iSCSI protocol processing without worrying about serial connectivity issues on the network or host interface side. Some of the protocol offload design challenges are:

- The number of TCP connections to support
- TCP reassembly/reorder
- IP fragmentation and reassembly
- Latency
- On-chip versus off-chip TCP session management

These issues can be mitigated with the unique features of Virtex-5 devices and available IP cores. With built-in GEMAC and PCIe interfaces, you can implement direct memory access solutions with minimal FPGA resources, reducing memory transfer latencies and enabling TCP reassembly without using temporary memory. Virtex-5 FPGAs also feature a 36-Kb dual-port block RAM, allowing you to support twice as many TCP connections than previous generations. With the Xilinx LogiCORE™ high-speed memory controller, you can use external DDR2 memory to scale TCP session management. Let's look at the resources you could save in an FPGA-based network interface solution.

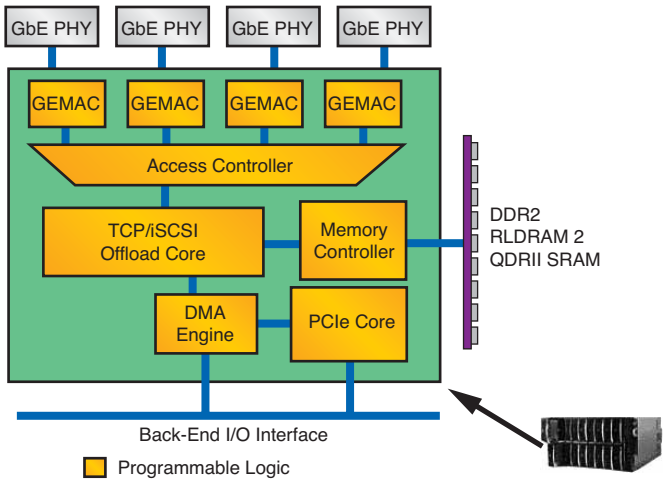

Figure 1 – Designing a TCP offload solution with traditional FPGAs

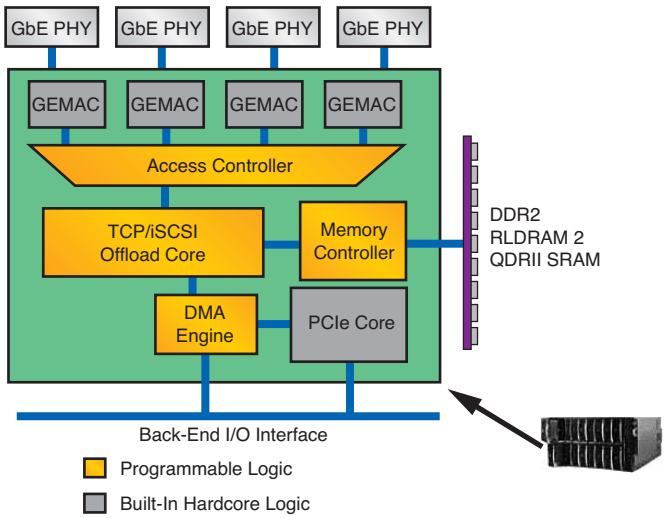

Figure 2 – Designing a TCP offload solution with Virtex-5 LXT FPGAs

1 Gbps and 10 Gbps NIC Solution

An integrated multi-port 1 Gbps and 10 Gbps TCP offload NIC for IP storage and bladed servers enables companies to leverage network infrastructure for storage traffic. Figure 1 shows a typical FPGA-based NIC design.

Depending on the IP cores used, this design can take as many as 20,000 slices to implement. The Virtex-5 LXT platform can reduce resource utilization by 50%, enabling you to develop lower cost solutions without sacrificing performance. Besides hardware efficiency, system architects can also reduce NRE costs because they are not required to implement high-speed I/O interfaces for GbE and PCIe. Figure 2 shows a redesign of the TCP offload NIC leveraging the built-in resources of the Virtex-5 family.

Conclusion

With standards-compliant TCP and iSCSI offload IP cores from third-party vendors implemented on Xilinx FPGAs, you can now design a drop-in or custom SoC at a much lower total cost of development. The Virtex-5 LXT platform – with hardened GEMACs and PCIe Endpoint blocks, larger block RAMs, and 6-LUTs – uses fewer FPGA resources to implement complex solutions for the server, storage, multi-protocol switch, and wireless base station markets.

To learn more about Virtex-5 LXT FPGAs, visit *www.xilinx.com/virtex5*. To learn more about protocol offload solutions, visit *www.xilinx.com/esp/storage/*. And to learn how Xilinx FPGAs can help you in other applications, visit *www.xilinx.com/esp*.

Implementing Encryption Algorithms with the Virtex-5 LXT Platform

The Virtex-5 LXT platform makes encryption product development easy.

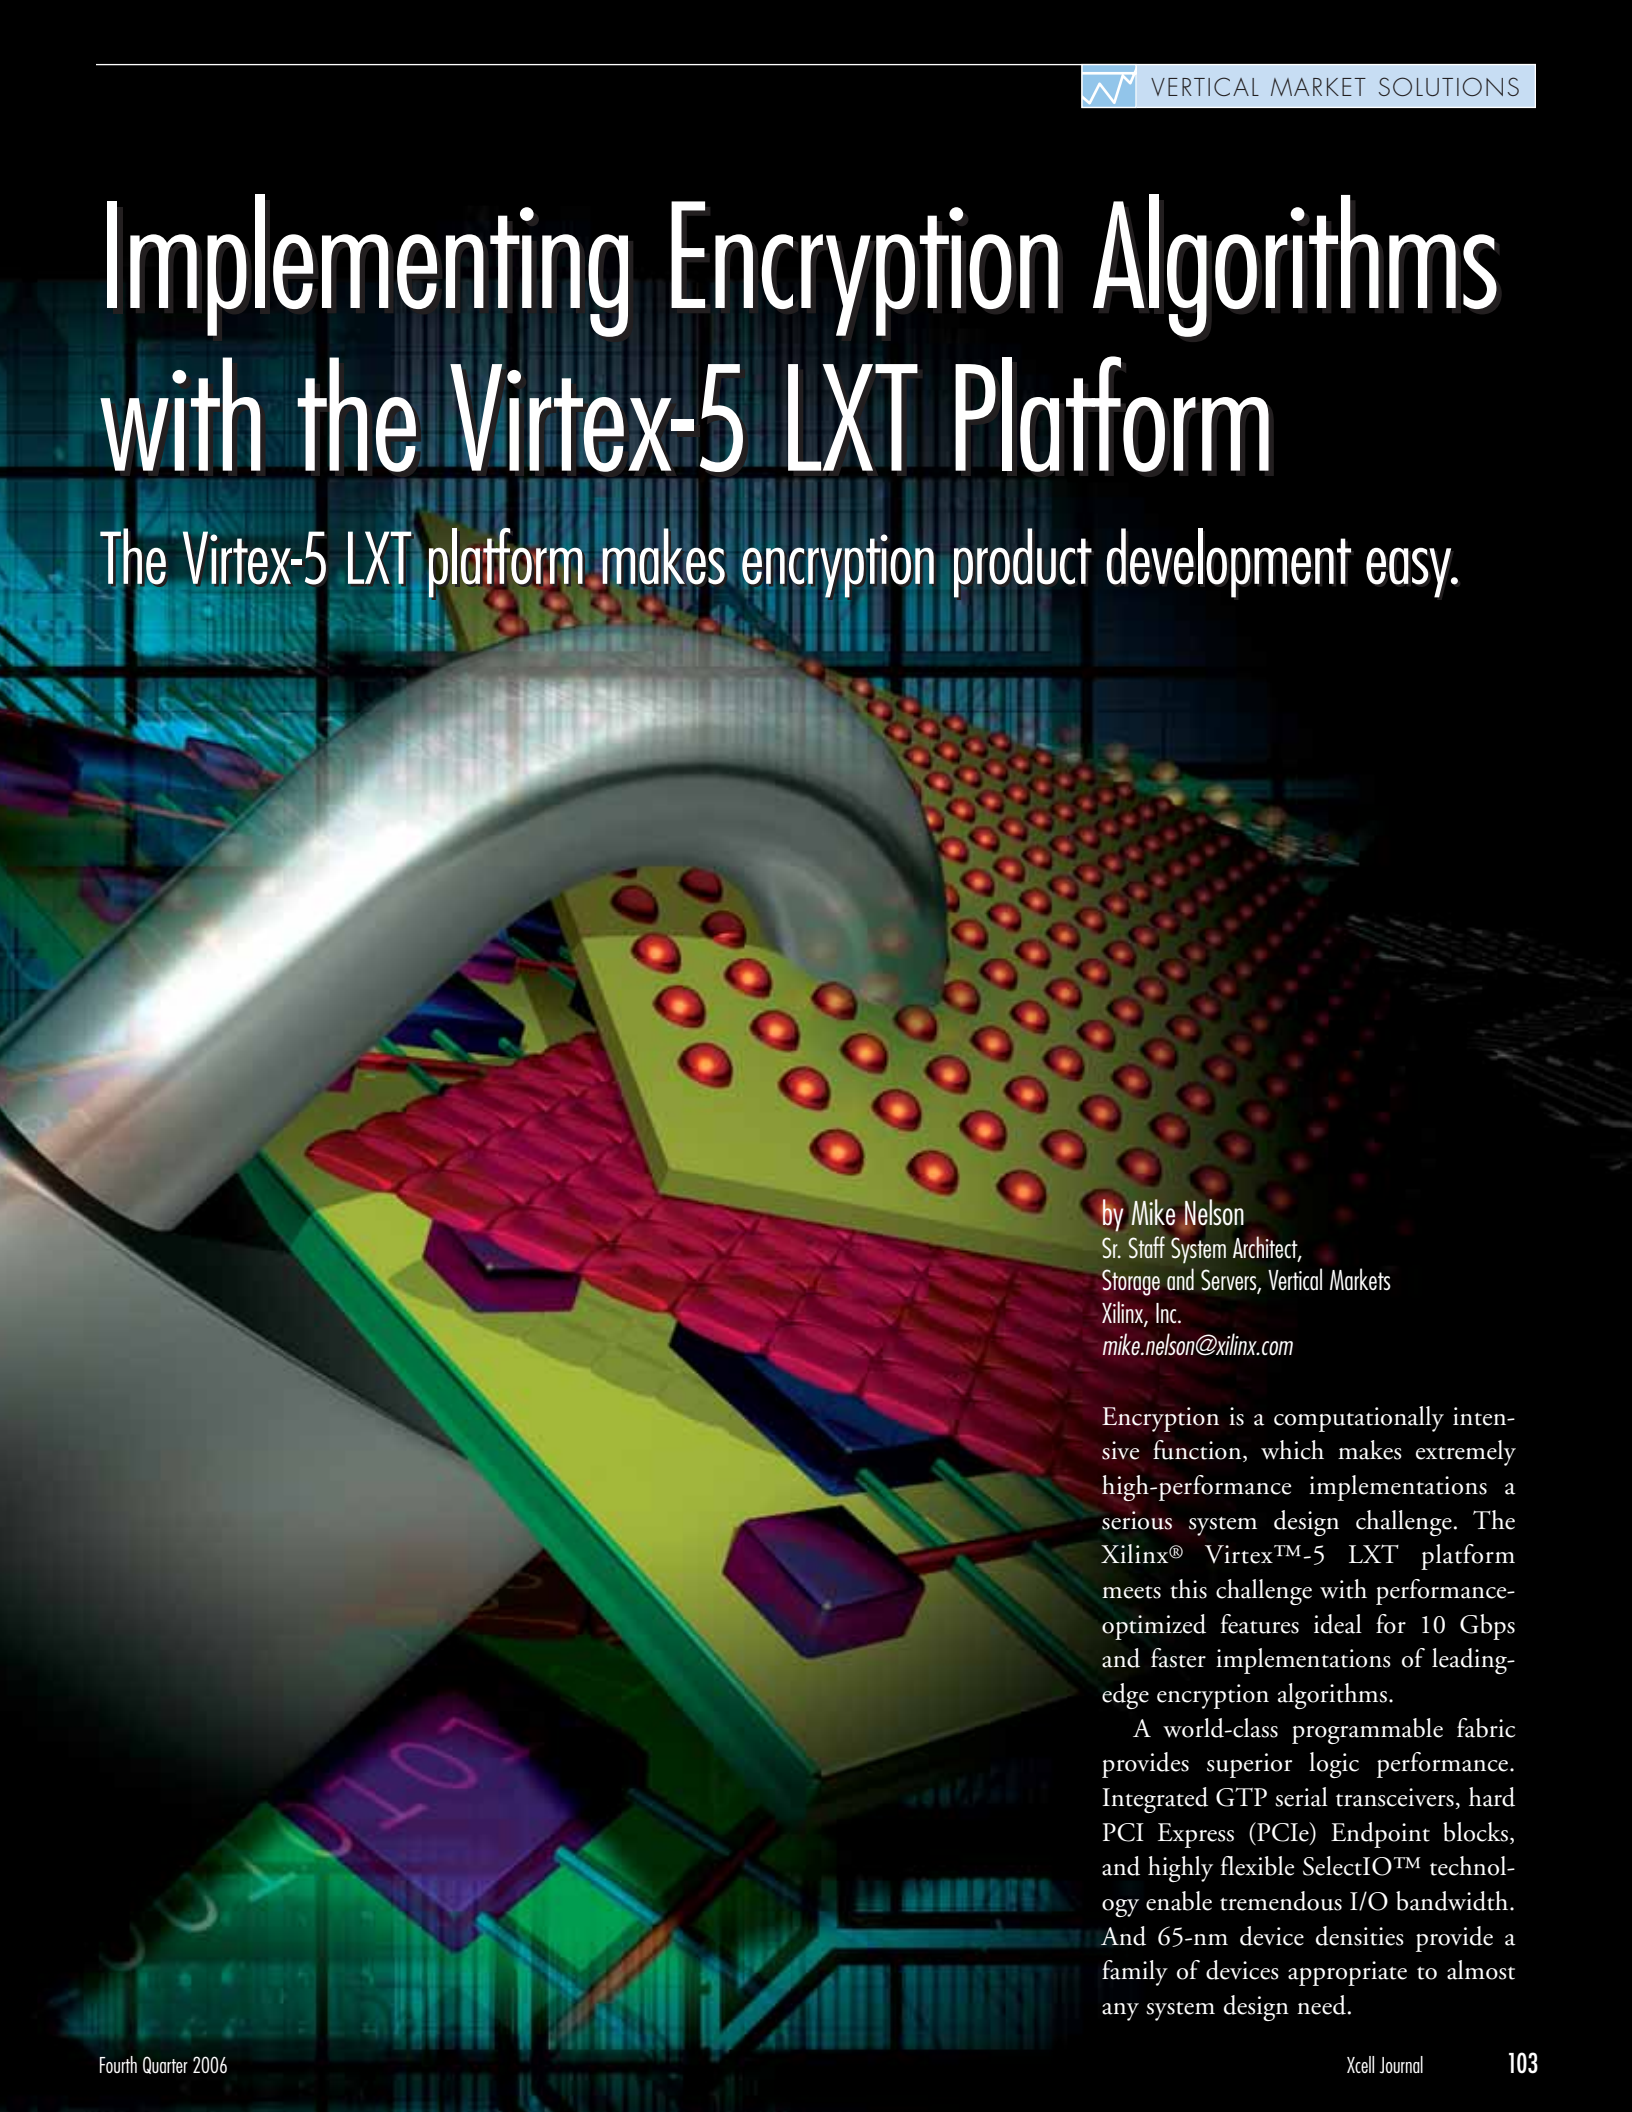

by Mike Nelson
Sr. Staff System Architect,
Storage and Servers, Vertical Markets
Xilinx, Inc.
mike.nelson@xilinx.com

Encryption is a computationally intensive function, which makes extremely high-performance implementations a serious system design challenge. The Xilinx® Virtex™-5 LXT platform meets this challenge with performance-optimized features ideal for 10 Gbps and faster implementations of leading-edge encryption algorithms.

A world-class programmable fabric provides superior logic performance. Integrated GTP serial transceivers, hard PCI Express (PCIe) Endpoint blocks, and highly flexible SelectIO™ technology enable tremendous I/O bandwidth. And 65-nm device densities provide a family of devices appropriate to almost any system design need.

As the world of cryptography continuously evolves with additional modes and algorithmic refinements, your design can evolve with it…

The Virtex-5 architecture features several advances that enable the very high-performance logic necessary for high-bandwidth encryption applications:

- Real six-input LUT-based fabric means that you can map circuits into denser structures with fewer levels of logic, increasing device utilization and performance
- Improved routing architecture increases the reach of low-latency logic interconnection, providing more flexibility to synthesis tools and also increasing device utilization and performance
- 36-Kb dual-port block RAMs with integrated ECC allow extremely high-performance on-chip memory resources for creating FIFOs and computational logic structures

Combined, these resources enable very cost-effective 10 Gbps and faster implementations of IPsec AES-CBC/AES-XCBC-MAC-96, 802.1ae MACSec, LRW-AES, AES-GCM, SHA-256/384/512, and many other cryptographic algorithms. Furthermore, as the world of cryptography continuously evolves with additional modes and algorithmic refinements to these algorithms, your design can evolve with it – because the Virtex-5 family is a programmable logic platform.

I/O Bandwidth and Flexibility

Computationally intensive core logic requires high-bandwidth I/O. But the nature of that I/O will vary based on your system architecture. Figure 1 shows two common architectures for implementing encryption processing.

Look-aside co-processing is an attractive option widely used in x86-based system appliances. This model leverages the excellent value of the commodity x86 platform to implement the application framework and selectively "looks aside" to an optimized accelerator to achieve high performance for the target application. FPGAs have always been well suited for this role, but scaling to approach to very high performance has been problematic.

PCI and PCI-X solutions require modest soft logic but have limited performance and must share what bandwidth they do have. PCI Express can implement a very high-performance non-blocking switched fabric, but traditionally requires extensive soft-logic resources to implement the controller, and possibly an external PHY for the electrical connection.

Virtex-5 LXT platform FPGAs address these limitations by combining embedded RocketIO™ GTP transceivers and a hardened PCI Express Endpoint block in every device. With the LXT platform, extremely high-performance co-processor I/O is easy and efficient, as shown in Figure 2.

Virtex-5 LXT platform FPGAs are also ideal for in-line applications, as illustrated in Figure 3. A key requirement for in-line encryption applications is flexibility. They may require identical – or different – input and output ports, port aggregation,

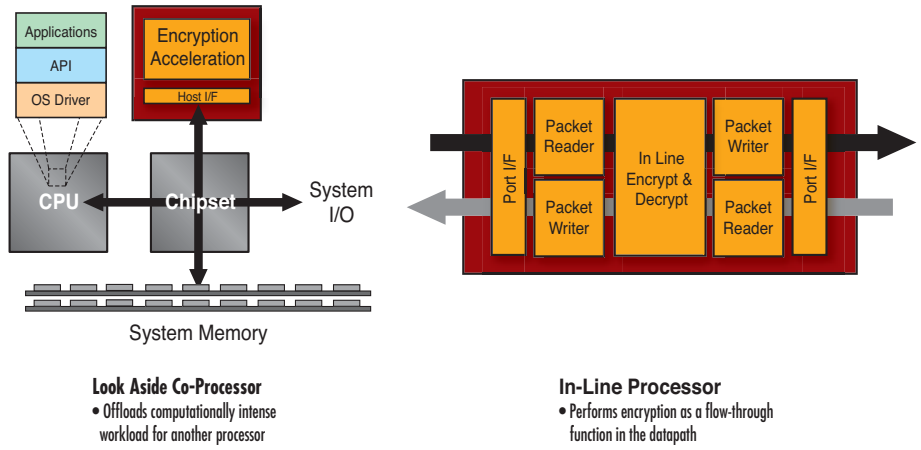

Figure 1 – Look-aside and in-line encryption processing

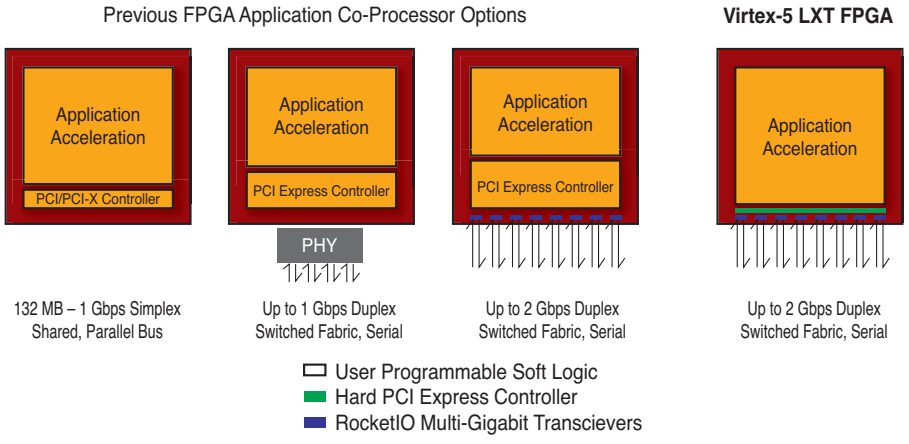

Figure 2 – I/O bandwidth and soft logic progression for FPGA co-processor options

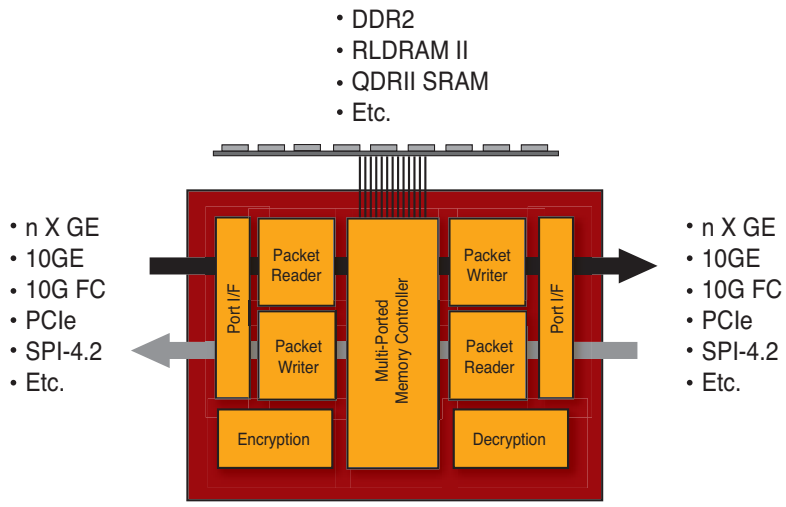

Figure 3 – Virtex-5 LXT device in-line encryption platform flexibility

or local subsystem memory. The Virtex-5 LXT platform meets this challenge with a wide range of capabilities:

- Gigabit Ethernet (GbE) – Each device in the Virtex-5 LXT platform includes four independent hardened GbE MACs, making multi-port Ethernet a very efficient I/O option. You can add additional ports as necessary with 100% form-, fit-, and function-equivalent soft LogiCORE™ IP.
- 10 Gigabit Ethernet – A Xilinx soft LogicCORE function is available that can be connected to four RocketIO MGTs for a XAUI interface or to a SelectIO pinout for an XGMII interface.
- 10 Gbps Fibre Channel (FC) – A XAUI-like Fibre Channel standard uses four RocketIO MGTs operating at 3.1875 Gbps in parallel to create a 10.2 Gbps FC channel.
- PCI Express – Available to interface to a variety of industry-standard PCIe-based port controllers.
- SPI-4.2 – Soft LogicCORE IP supports this networking industry standard for chip-to-chip connectivity over high-performance SelectIO technology.
- Memory – In addition to port I/O standards, Virtex-5 SelectIO technology also supports a wide range of memory interface technologies including DDR2, RLDRAM II, and QDR II SRAM. These capabilities enable virtually any local memory subsystem that an in-line processing engine might require.

These features allow you to create in-line solutions that will connect to the ports you need with the integrated encryption technology you want.

Conclusion

The Virtex-5 LXT platform expands the capabilities of the Virtex-5 FPGA architecture with the addition of RocketIO GTP transceivers, plus hard PCI Express Endpoint and tri-mode Ethernet MAC blocks. The result is a platform ideally suited to support very high-performance look-aside and in-line encryption functions.

Other applications where LXT platform devices will excel include high-performance packet handling and deep content inspection for networking; high-speed data mining for databases; time-critical computational processing for industrial, scientific, and medical applications; and real-time image processing for aerospace/defense and video graphic applications.

To learn more about Virtex-5 LXT platform FPGAs, visit *www.xilinx.com/virtex5*. To learn more about Xilinx in encryption, visit *www.xilinx.com/esp/security/data_security/index.htm*. And to learn how Xilinx FPGAs can help you in other applications, visit *www.xilinx.com/esp*.

Great Test, Less Filling

SystemBIST™ enables FPGA Configuration that is less filling for your PCB area, less filling for your BOM budget and less filling for your prototype schedule. All the things you want less of and more of the things you do want for your PCB – like embedded JTAG tests and CPLD reconfiguration.

Typical FPGA configuration devices blindly "throw bits" at your FPGAs at power-up. SystemBIST is different – so different it has three US patents granted and more pending. SystemBIST's associated software tools enable you to develop a complex power-up FPGA strategy and validate it. Using an interactive GUI, you determine what SystemBIST does in the event of a failure, what to program into the FPGA when that daughterboard is missing, or which FPGA bitstreams should be locked from further updates. You can easily add PCB 1149.1/JTAG tests to lower your down-stream production costs and enable in-the-field self-test. Some capabilities:

- User defined FPGA configuration/CPLD re-configuration
- Run Anytime-Anywhere embedded JTAG tests
- Add new FPGA designs to your products in the field
- "Failsafe" configuration – in the field FPGA updates without risk
- Small memory footprint offers lowest cost per bit FPGA configuration
- Smaller PCB real-estate, lower parts cost compared to other methods
- Industry proven software tools enable you to get-it-right before you embed
- FLASH memory locking and fast re-programming
- New: At-speed DDR and RocketIO™ MGT tests for V4/V2

If your design team is using PROMS, CPLD & FLASH or CPU and in-house software to configure FPGAs please visit our website at http://www.intellitech.com/xcell.asp to learn more.

Intellitech®

Virtex-5 Configuration Options Offer Designers a Choice

Xilinx provides a host of flexible choices in configuration memory to help you make the best decision for your design.

by Frank Toth
EasyPath FPGAs and Configuration Solutions
Xilinx, Inc.
frank.toth@xilinx.com

System designers are always making trade-offs between alternative requirements. Considerations include time to market, ease of use, total ownership cost, and system speed.

Every alternative offers different total-cost-of-ownership (TCO) considerations that you should examine when designing a configuration system, including design time, prototyping, manufacturing and test costs, and the per-bit costs of the configuration device. The trade-offs of all these factors should enter into your decision of which configuration method to use.

Designers using Xilinx® Virtex™-5 devices have many additional choices for configuring the new FPGA family, including new configuration modes built right into the chips; support for 32-bit-wide high-performance parallel SelectMAP, which offers the ultimate in speed; and both BPI (byte parallel interface) and SPI (serial peripheral interface) using industry-standard SPI and parallel flash memory devices (see Figure 1). The easy-to-use, full-featured, configuration-engineered Platform Flash PROMs offer a pre-engineered way to flexibly configure Virtex-5 devices and manage multiple bitstreams.

Platform Flash PROMs

Dropping a Platform Flash PROM into a design provides a seamless solution with a low TCO, including minimum board space, high configuration speed, a guaranteed source of supply, and value-added features like bitstream de-compression, design revision management, JTAG Boundary Scan for test and configuration, and additional storage for boot and scratch-pad memory.

Figure 1 – Virtex-5 configuration modes

Platform Flash features on-board decompression that can result in as much as 50% more configuration data into the same overall memory space. Design revisioning allows you to switch between memory blocks for various configurations: for instance, using the board and system in different geographical regions or loading a diagnostic followed by a mission load of configuration. In addition, unused Platform Flash memory can be allocated to boot code or scratch-pad memory. Both of these features work without any additional glue logic or special software. The Boundary Scan JTAG port enables configuration and includes the Platform Flash in overall board tests.

SPI Flash PROMs

Virtex-5 FPGAs support direct connection to SPI PROMs using the industry-standard four-wire SPI interface. Many systems currently use SPI PROMs; you can now easily take advantage of the on-board SPI PROM without any additional circuitry or software. Designers should think about design trade-offs, including different features offered by SPI PROM manufacturers and the slower configuration speed compared to parallel SelectMAP, Platform Flash, and BPI.

BPI Flash PROMs

Virtex-5 devices include on-board circuitry to directly connect – without any additional glue logic or software – to industry-standard parallel flash devices. The parallel flash interface can be directly connected to the FPGA and the memory shared by the system bus.

Conclusion

Virtex-5 FPGAs offer you the widest variety of configuration alternatives in the industry, including Platform Flash, 32-bit SelectMAP, and interfaces that directly connect to SPI and BPI PROM devices. You should understand these alternatives and make an informed decision on which one meets your needs based on the trade-offs between speed, complexity, and features.

For more information, please see Xilinx Application Note 483, "Multiple-Boot with Platform Flash PROMs," at *www.xilinx.com/bvdocs/appnotes/xapp483.pdf.*

Introducing Virtex-5 EasyPath FPGAs

The world's first 65-nm FPGA cost-reduction solution.

by Gokul Krishnan, Ph.D
EasyPath Marketing
Xilinx, Inc.
gokul.krishnan@xilinx.com

Derek Johnson
APD Marketing
Xilinx, Inc.
derek.johnson@xilinx.com

With increasing competition in many different market segments, many companies must drive down their product development costs while at the same time adding more and more complexity and features. In addition, companies must react to fast-changing market requirements and heightened time-to-market pressures.

Xilinx® FPGAs can help you face these challenges by continually innovating to provide increasingly complex functions at a lower cost per logic cell. The recently introduced Virtex™-5 FPGAs, in combination with Virtex-5 EasyPath™ FPGAs, are the latest generation of 65-nm devices that provide higher performance, lower system cost, and greater embedded functionality than ever before.

Virtex-5 EasyPath FPGAs are the industry's only 65-nm customer-specific FPGA cost-reduction solution, providing the lowest total cost of ownership (TCO) when compared to other solutions. EasyPath FPGAs are identical to standard Xilinx FPGA offerings but use patented testing techniques and customer-specific test patterns to significantly improve FPGA yields for designs that no longer require the full programmability of a standard FPGA. You can reap the benefits of these improved yields in the form of lower costs. EasyPath technology is available across multiple platforms, different product families, and 28 different devices over a range of gate and memory counts.

Lowest TCO with Virtex-5 EasyPath FPGAs

Virtex-5 EasyPath FPGAs devices are manufactured using a 65-nm process, which intrinsically offers a cost advantage (see Table 1). In addition to a low unit price, EasyPath FPGAs provide many other cost advantages, such as:

- Low NRE
- No re-qualification required
- No engineering resources required
- Shorter lead times (12-16 weeks)

With Virtex-5 EasyPath FPGAs, you can realize a 30%-75% price reduction when moving to high volume as compared to standard FPGAs. EasyPath FPGAs are identical to their standard FPGA counterparts, effectively eliminating any conversion work. This has two important implications. The first is that you pay less yet incur very little risk, because every single feature in a standard FPGA is supported and will work in an EasyPath FPGA. Second, you do not need to re-qualify your boards or systems when you move to EasyPath FPGAs. This saves valuable engineering time and resources and provides cost savings of $500K or more.

Unlike structured ASICs, where customers have to go through multiple reviews with the vendor and spend many months of valuable engineering resources, EasyPath FPGAs demand almost no resources from you. Once you have finalized the design and handed off the relevant files to Xilinx, you can get to full production directly in 8-12 weeks. No intermediate prototyping is required because the design has already been finalized (prototyped) in a standard FPGA. The lead time to get to production is at least three to four months less than with structured ASICs.

Alternatively, those of you in fast-moving markets can postpone the design freeze milestone by three to four months to better address dynamic market conditions. Getting to market faster can have a significant impact on the market share a product can capture. Studies have indicated that just a three-month delay in time to market can reduce market share by as much as 15%, according to research from International Business Strategies, Inc.

Total Cost Driver	Structured ASICs	Xilinx EasyPath
Time to Cost Reduction	20 to 24 Weeks	Yes
NRE Costs	$100K to $400K	No
Cost of Requalification	$100K to $500K	Yes
Engineering Costs	$250K to $300K	Identical
Cost of Design Tools	$100K to $200K	Identical
Unit Costs	Lowest	Low
Cosy of Respin	High	High

EasyPath Total Cost of Ownership Advantage
Only Cost-Reduction Path Supporting Complex IP
100% FPGA Feature Support
100% Package Support

Table 1 – EasyPath total cost of ownership advantage

Difficulties with 65-nm ASICs

One of the industry trends in place for some time now is that ASIC design starts are decreasing every year. Part of the reason for this is that FPGAs have been able to provide lower unit costs for higher functionality. The other driver has been the rising cost of mask sets, design, and verification. By some recent estimates, the cost of developing a new 90-nm ASIC design is in excess of $10 million (International Business Strategies, Inc). A significant portion of this cost occurs in the verification phase, which has become longer and longer with increases in chip complexity. At 65 nm, operating voltages and transistor sizes are so small that very small process variations can have a big effect on the functionality of a design.

What this means to you is that you must now factor in DFM (design for manufacturability) rules during the physical design phase – something previously assumed to have been embedded in the library itself. Furthermore, because the interconnects are very closely spaced, signal integrity becomes even more important; you must take care to reduce parasitic capacitance issues when signals are being transmitted simultaneously on adjoining metal lines.

Another major problem with 65-nm design is the issue of power consumption. Although Virtex-5 FPGAs take advantage of triple-oxide technology to reduce leakage power consumption significantly, each ASIC design must factor in power consumption and use techniques such as clock gating and selective transistors to mitigate leakage current. So although Virtex-5 EasyPath FPGAs retain the simplicity they had at 130 nm and 90 nm, 65-nm ASICs offer unique challenges that could significantly reduce the chance of first-time success with a design.

Conclusion

The Virtex-5 family is one of the most cost-effective, high-performance FPGA families in the industry. With advanced features such as a higher utilization logic fabric, more integrated block memory, higher precision DSP slices, as well as advanced connection and embedded processing blocks, you can reduce your overall system cost by fitting into a smaller FPGA or replacing external discrete devices on your boards. This complements the natural cost reduction that comes from fabricating devices on a 65-nm process.

Virtex-5 FPGAs are a faster time-to-market alternative to ASICs and other custom logic solutions, and enable a lower total system cost. In addition, Virtex-5 FPGAs are designed for the lowest overall power consumption, highest signal integrity, and highest performance. All of these attributes can lead to lower overall system cost: lower power consumption and high signal integrity can cut design and debugging costs, and high performance can save device costs by allowing the design to be done in a lower, less-expensive speed grade.

Connectivity Solutions

Realize the full potential of the solutions in our silicon with Xilinx application notes.

Memory Interfaces

XAPP851 – DDR SDRAM Controller Using Virtex-5 FPGA Devices

By Toshihiko Moriyama and Rich Chiu

This application note describes a 200-MHz DDR SDRAM memory controller implemented in a Virtex™-5 device. This reference design uses the Virtex-5 ChipSync™ features to calibrate and adjust read data timing.

A straightforward back-end user interface is provided to allow integration into a complete FPGA design.

On the Web at ***www.xilinx.com/bvdocs/appnotes/xapp851.pdf***

XAPP852 – Synthesizable CIO DDR RLDRAM II Controller for Virtex-5 FPGAs

By Benoit Payette and Rodrigo Angel

This application note describes how to use a Virtex-5 device to interface to common I/O (CIO) double data rate (DDR) reduced latency DRAM (RLDRAM II) devices. The reference design targets two CIO DDR RLDRAM II devices at a clock rate of 200/300 MHz, with data transfers at 400/600 Mbps per pin.

On the Web at ***www.xilinx.com/bvdocs/appnotes/xapp852.pdf***

XAPP853 – QDR II SRAM Interface for Virtex-5 Devices

By Lakshmi Gopalakrishnan

This application note describes the implementation and timing details of a four-word-burst quad data rate (QDR II) SRAM interface for Virtex-5 devices. The synthesizable reference design leverages the unique I/O and clocking capabilities of the Virtex-5 family to achieve performance levels of 300 MHz (600 Mbps), resulting in an aggregate throughput for each 36-bit memory interface of 43.2 Gbps.

The design greatly simplifies the task of read data capture within the FPGA while minimizing the number of resources used. A straightforward user interface is provided to allow simple integration into a complete FPGA design utilizing one or more QDR II interfaces.

On the Web at ***www.xilinx.com/bvdocs/appnotes/xapp853.pdf***

XAPP858 – High-Performance DDR2 SDRAM Interface in Virtex-5 Devices

by Karthi Palanisamy and Maria George

This application note describes the controller and data capture technique for high-performance DDR2 SDRAM interfaces. This data capture technique uses the input serializer/deserializer (ISERDES) and output double data rate (ODDR) features available in every Virtex-5 I/O.

On the Web at ***www.xilinx.com/bvdocs/appnotes/xapp858.pdf***

Source-Synchronous Interfaces

XAPP855 – 16-Channel DDR LVDS Interface with Per-Channel Alignment

by Greg Burton

This application note describes a 16-channel source-synchronous DDR LVDS interface. The design takes advantage of the Virtex-5 I/O ChipSync feature's ability to adjust the delay of the receiver datapaths, creating dynamic setup and hold timing for each device at initialization and compensating for skews associated with the manufacturing process. The receiver operates at 1:8 deserialization on each of the 16 data channels.

On the Web at ***www.xilinx.com/bvdocs/appnotes/xapp855.pdf***

XAPP860 – 16-Channel DDR LVDS Interface with Real-Time Window Monitoring

by Greg Burton

This application note describes a 16-channel source-synchronous DDR LVDS interface. The receiver operates at 1:6 deserialization on each of the 16 data channels. Similar to XAPP855, the design also includes a real-time window monitoring circuit for added performance. This reference design calibrates and compensates for skews associated with process, voltage, and temperature (PVT) at initialization and also dynamically during operation.

On the Web at ***www.xilinx.com/bvdocs/appnotes/xapp860.pdf***

Serial Connectivity

XAPP861 – Efficient 8x Oversampling Asynchronous Serial Data Recovery Using IDELAY

by John Snow

Virtex-5 devices a have a high-precision programmable delay element (IDELAY) associated with every input pin. This application note shows how to implement 8x oversampling of many data streams using a single DCM, two global clock resources, and minimal FPGA logic resources. This solution provides better jitter tolerance than techniques using multiple DCMs. When paired with a suitable data recovery scheme, this oversampling technique can be used with many different data protocols up to 550 Mbps. A reference design is included that implements a SD-SDI (SMPTE 259M) receiver running at 270 Mbps.

On the Web at ***www.xilinx.com/bvdocs/appnotes/xapp861.pdf***

Intellectual Property Offerings

The Xilinx IP Center on the Web allows you to search for IP by function, type, vendor, or keywords.

The IP Center includes intellectual property from Xilinx and its third-party vendors. To access the IP solutions highlighted here, visit ***www.xilinx.com/ipcenter*** and type the keywords listed in the search box.

Source-Synchronous Interfaces

SPI-4 Phase 2 Interface Solutions (DO-DI-POSL4MC) Xilinx IP Core

The Xilinx® SPI-4 Phase 2 core provides a fully compliant packet-over-SONET/SDH (POS) solution, which can be quickly integrated into networking systems.

Through user-configurable options, the Xilinx SPI-4.2 core provides ultimate flexibility while seamlessly interoperating with industry-leading ASSPs to maximize the data transfer bandwidth. The Xilinx SPI-4.2 core is fully compliant with the OIF's System Packet Interface Level 4 (SPI-4) Phase 2 standard, as well as the Saturn Development Group's POS-PHY Level 4 (PL4) interface specification.

Type search keywords: **SPI-4 Phase 2**

Serial Connectivity

Virtex-5 RocketIO GTP Wizard

The Virtex™-5 RocketIO™ GTP Wizard automates the task of creating HDL wrappers to configure Virtex-5 RocketIO GTP transceivers. The wizard's customization GUI allows you to configure one or more GTP transceivers using pre-defined templates to support popular industry standards, or from scratch to support a wide variety of custom protocols.

Type search keywords: GTP Wizard

Virtex-5 Embedded Tri-Mode Ethernet MAC Wrapper

The CORE Generator™ tool supports the Virtex-5 Tri-Mode Ethernet Media Access Controller (MAC) Wrapper to automate the generation of HDL wrapper files for the tri-mode Ethernet MAC in Virtex-5 LXT devices. Preconfigured HDL wrappers, testbenches, and implement and simulation scripts are generated automatically based on user-defined options.

Type search keywords: **Virtex-5 Ethernet MAC**

Virtex-5 PCI Express Endpoint Block Wrapper

The Xilinx PCI Express Endpoint block wrapper integrates and interfaces to the on-chip PCI Express Endpoint block, supporting 1-lane, 2-lane, 4-lane, and 8-lane complete endpoint core implementations. In addition, a PCI Express Endpoint block development kit is also available. This solution is used in communication, multimedia, server, storage, and mobile platforms and enables applications such as high-end medical imaging, graphics-intensive video games, DVD quality streaming video on the desktop, and 10 Gigabit Ethernet interface cards.

Type search keywords: **PCI Express Block**

Virtex-5 LXT PCI Express Block Plus LogiCORE Xilinx IP Core

The Xilinx PCI Express Plus LogiCORE IP integrates and interfaces to the PCI Express Endpoint block, supporting 1-lane, 4-lane, and 8-lane complete endpoint core implementations. In addition, a PCI Express development kit is also available. This solution is used in communication, multimedia, server, storage, and mobile platforms and enables applications such as high-end medical imaging, graphics-intensive video games, DVD quality streaming video on the desktop, and 10 Gigabit Ethernet interface cards.

Type search keywords: **PCI Express Plus**

Virtex-5 Boards and Kits

Jumpstart your Virtex-5 designs with these development platforms and tool kits.

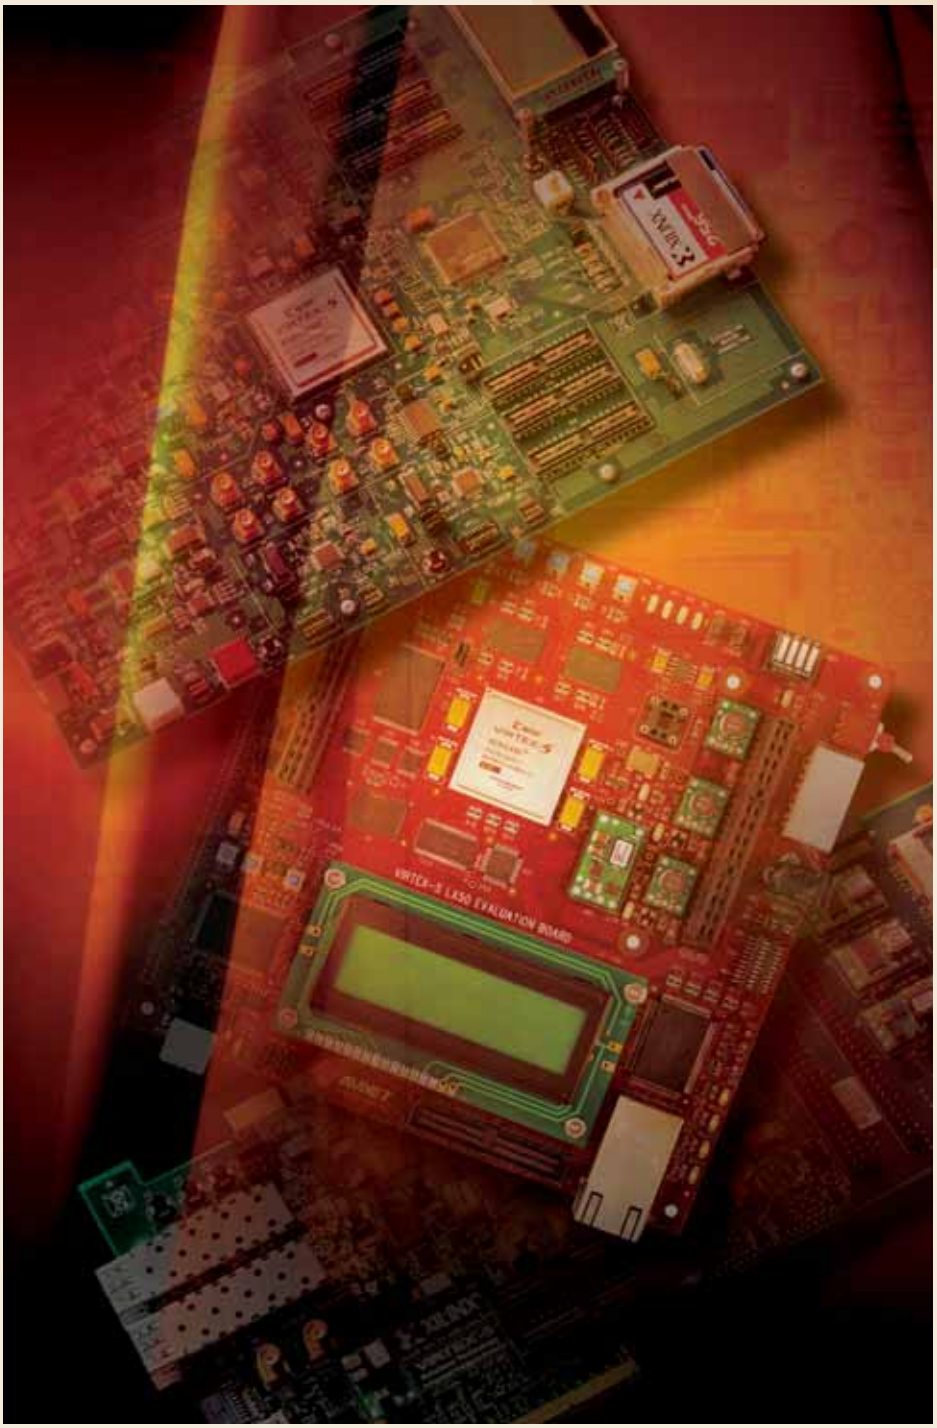

Nu Horizons Virtex-5 LXT Evaluation Kit

Nu Horizons creates a low-cost evaluation kit for Virtex-5 LXT platform FPGAs.

Nu Horizons's newest evaluation kit is designed for customers interested in evaluating Virtex-5 LXT FPGAs. This kit differs from other Nu Horizons kits in that it has the added ability for high-speed serial communication.

On the Web at ***www.nuhorizons.com***

Xilinx Virtex-5 ML505 Evaluation and Development Platform

A low-cost embedded system and RocketIO GTP transceiver development platform.

The Xilinx Virtex-5 ML505, based on RocketIO™ technology, is a feature-rich, low-cost evaluation/development platform that provides easy and practical access to the resources available in the on-board Virtex-5 LXT FPGA.

Supported by industry-standard interfaces/connectors, generous memory resources, and companion chipsets, the ML505 evaluation platform is a versatile development platform for multiple applications including embedded systems.

On the Web at ***www.xilinx.com/XOB***

Avnet Virtex-5 LX Development Kit

A complete development platform for designing and verifying applications based on the Xilinx Virtex-5 LX FPGA family.

Available with the Xilinx® Virtex™-5 XC5VLX50-1FF676 device, the Avnet Virtex-5 Development Kit allows you to prototype high-performance designs with ease, while providing expandability and customization through the EXP expansion slot.

The system board includes DDR2 SDRAM, flash memory, a 10/100/1000 Ethernet PHY, and a serial port, making it an ideal platform for MicroBlaze™ development. Other board features include a USB port, programmable LVDS clock, 10-bit Tx/Rx high-speed LVDS interface, user switches and LEDs, and a 2 x 16-character LCD panel.

The board also provides a full EXP expansion slot, providing a total of 168 high-speed, single-ended, and differential user I/O. You can easily add EXP modules to the board for additional application-specific functions.

On the Web at *www.avnet.com*

Xilinx Virtex-5 ML501 Evaluation and Development Platform

An ideal general-purpose, low-cost development platform.

The Xilinx Virtex-5 ML501 evaluation and development platform is a feature-rich, low-cost evaluation/development platform that provides easy and practical access to the resources available in the on-board Virtex-5 LX FPGA. Supported by industry-standard interfaces and connectors, the ML501 is a versatile development platform for multiple applications. Video, audio, and communication ports as well as generous memory resources extend the functionality and flexibility of the ML501 evaluation platform beyond a typical FPGA development platform.

On the Web at *www.xilinx.com/ML501*

HiTech Global Virtex-5 PCI Express Development Platform

Seamless serial interface connectivity enabled by the Virtex-5 LXT FPGA.

Powered by a Xilinx Virtex-5 LXT FPGA, supported by mainstream peripherals, and designed with excellent signal integrity performance, the HiTech Global HTG-V5PCIE is the ideal platform for serial interface/connectivity developments, including PCI Express subsystems, Serial ATA (SATA), Fibre Channel, RapidIO, and XAUI.

On the Web at *www.hitechglobal.com*

Xilinx Virtex-5 ML550 Networking Interfaces Tool Kit

Designing networking, telecom, servers, and computing systems with Virtex-5 FPGAs.

Many of today's telecom and networking systems use high-bandwidth interfaces based on LVDS or other differential I/O standards. Differential I/O standards simplify system design by improving system performance and signal integrity.

Protocols based on source-synchronous I/Os such as SPI-4.2 and SFI-4 are central to leading-edge system design. To take advantage of these technologies, you have to work through multiple challenges to ensure device interoperability and standards compliance. Xilinx provides the Virtex-5 network interface board, as well as standards-compliant IP cores and free reference designs, to help you tackle these high-speed, source-synchronous interface challenges. This allows you to focus on user application design and not worry about interoperability and standards compliance.

On the Web at *www.xilinx.com/XOB*

Xilinx Virtex-5 ML555 PCI Express Development Tool Kit

A highly configurable pre-verified development solution.

The Xilinx ML555 RoHS-compliant PCIe/PCI-X/PCI development board provides a pre-verified solution to parallel and serial PCI interface design challenges. Using an established development environment can dramatically shorten the design cycle. By using proven Xilinx dedicated blocks, you can focus your efforts on specific application development and avoid time-consuming PCIe or PCI development.

On the Web at *www.xilinx.com/XOB*

Xilinx Virtex-5 ML561 Advanced Memory Development System

Achieve your performance targets in the shortest development time.

Building interfaces to high-performance memory devices presents challenges such as high-speed synchronous data capture, along with implementing complex physical-layer interfaces and control logic. The ML561 advanced memory development system offers an excellent platform to develop and verify high-performance memory interfaces using Virtex-5 FPGAs.

On the Web at *www.xilinx.com/XOB*

ASIC PROTOTYPING

HAPS – HARDI ASIC Prototyping System
a high performance, high capacity FPGA platform for ASIC prototyping and emulation composed of multi-FPGA boards and standard or custom-made daughter boards

HapsTrak
a set of rules for pinout and mechanical characteristics, which guarantees compatibility with previous and future generation HAPS motherboards and daughter boards

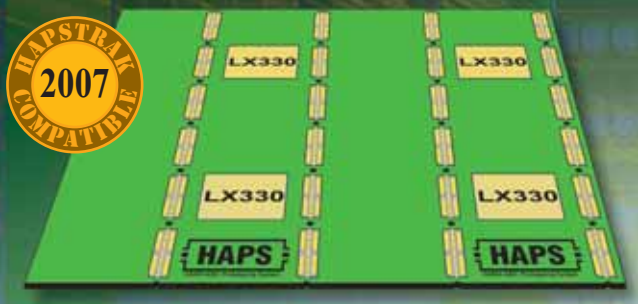

HAPS-40

VIRTEX-5

HAPS-30

VIRTEX-4

HAPS-20

VIRTEX-II PRO

HAPS-10

VIRTEX-II

HARDI Electronics

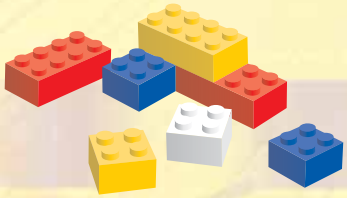

www.hardi.com, haps@hardi.com

HARDI Electronics Inc., 26831 Magdalena Lane, Mission Viejo, CA 92691, (949) 202-5572

IS YOUR CURRENT FPGA DESIGN SOLUTION HOLDING YOU BACK?

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

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

Mentor Graphics®

THE EDA TECHNOLOGY LEADER

Low-Power Transceivers

Ultimate Connectivity . . .

VIRTEX™ 5

Reduce serial I/O power, cost and complexity with the world's first 65nm FPGAs.

With a unique combination of up to 24 low-power transceivers, and built-in PCIe™ and Ethernet MAC blocks, Virtex-5 LXT FPGAs get your system running fast. Whether you are an expert or just starting out, only Xilinx delivers this complete solution to simplify high-speed serial design.

Power consumption and area required to implement a typical design including 8-lane PCIe endpoint

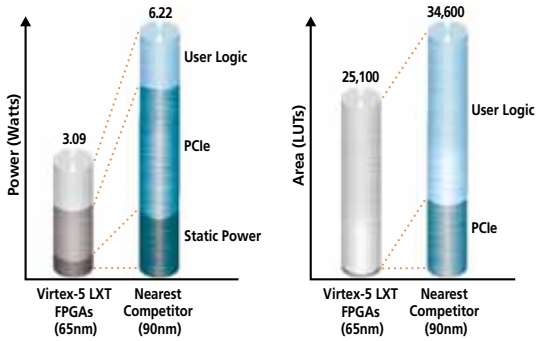

5VLX30T vs. 2SGX60D. Target Frequency = 200 MHz. Worst-case process. 25K LUTs, 17K Flip-Flops,1 Mbit On-Chip RAM, 64 DSP Blocks, 128 2.5V I/Os. Based on Xilinx tool v8.2 and competitor tool v6.0.1

Lowest-power, most area-efficient serial I/O solution

RocketIO™ GTP transceivers deliver up to 3.2 Gbps connectivity at less than 100 mW to help you beat your power budget. The embedded PCI Express® endpoint block ensures easy implementation and reduced development time. Embedded Ethernet MAC blocks enable a single-chip UNH-verified implementation. And the Xilinx solution is fully supported by development tools, design kits, IP, characterization reports, and more.

Visit our website today, view the Webcast, and order your free eval CD to give your next design the ultimate in connectivity.

XILINX®

The Programmable Logic Company℠

www.xilinx.com/virtex5

The Ultimate System Integration Platform

PN 0010999