ISSUE 56, FIRST QUARTER 2006 XCELL JOURNAL XILINX, INC.

Issue 56
First Quarter 2006

Xcell journal

THE AUTHORITATIVE JOURNAL FOR PROGRAMMABLE LOGIC USERS

Bright Ideas for Better Designs

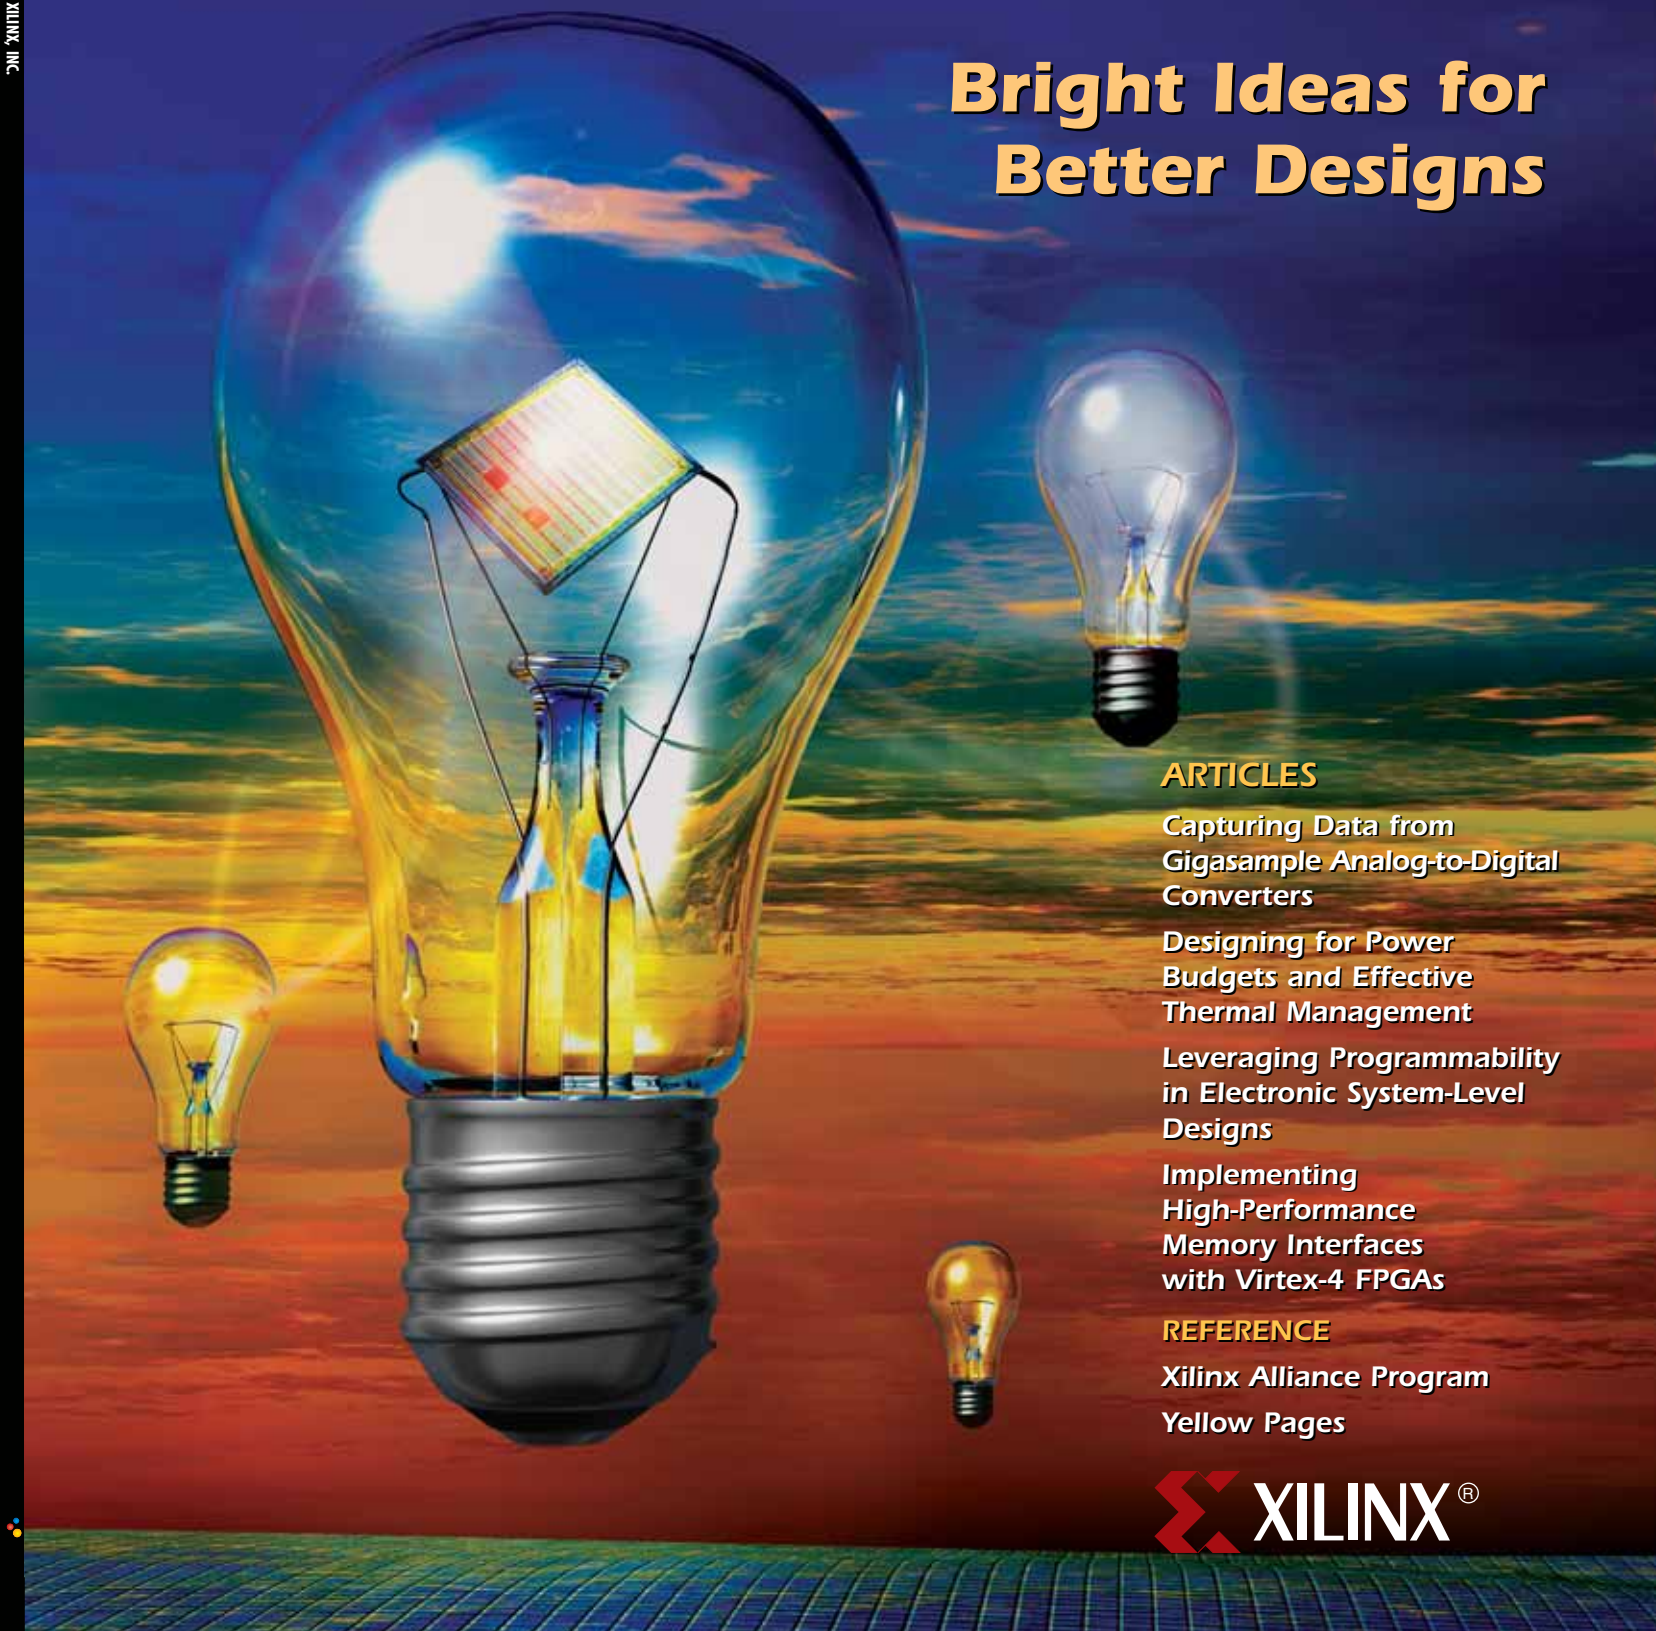

ARTICLES

REFERENCE

XILINX®

Support Across The Board.™

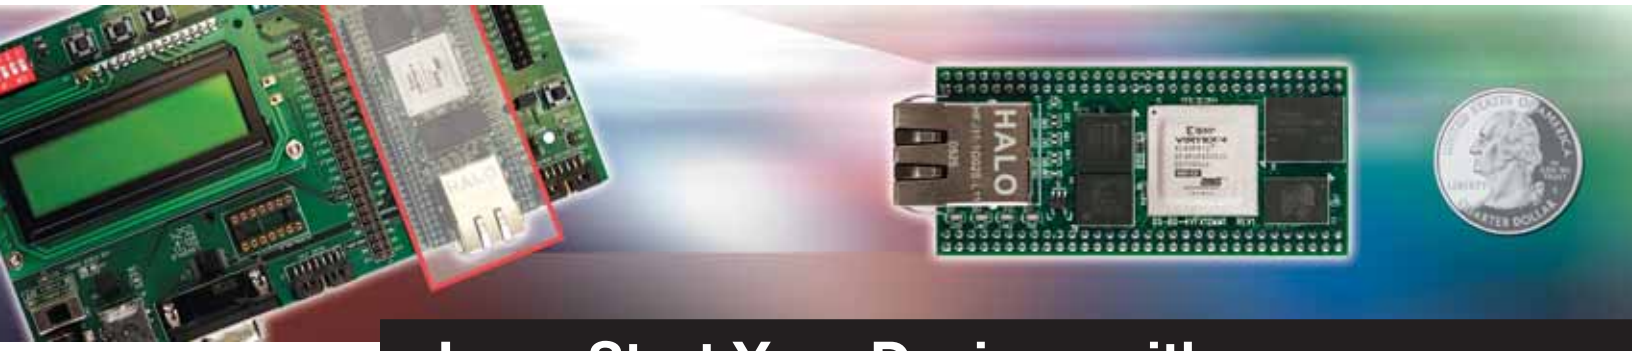

Jump Start Your Designs with Mini-Module Development Platforms

Mini-Module Features

- Small footprint (30 mm x 65.5 mm)
- Complete system-on-module solution with SRAM, Flash, Ethernet port and configuration memory
- Supports MicroBlaze™ or PowerPC® processors
- Based on the Virtex™-4 FX12 and 400K Gate Spartan™-3 FPGAs
- 76 User I/O

Baseboard Features

- Supports mini-module family
- Module socket allows easy connection
- Provides 3.3 V, 2.5 V and 1.2 V
- USB and RS232 ports
- 2 x 16 character LCD

Upcoming SpeedWay Design Workshops

- Developing with MicroBlaze
- Developing with PowerPC
- Xilinx Embedded Software Development
- Embedded Networking with Xilinx FPGAs

The Xilinx® Spartan™-3 and Virtex™-4 FX 12 Mini-Modules are off-the-shelf, fully configurable, complete system-on-module (SOM) embedded solutions capable of supporting MicroBlaze™ or PowerPC® embedded processors from Xilinx.

Offered as stand-alone modules and as development kits with companion baseboards, power supplies, software and reference designs, these modules can be easily integrated into custom applications and can facilitate design migration.

Gain hands-on experience with these development kits and other Xilinx tools by participating in an Avnet SpeedWay Design Workshop™ this spring.

Learn more about these mini-modules and upcoming SpeedWay Design Workshops at:

www.em.avnet.com/spartan3mini
www.em.avnet.com/virtex4mini
www.em.avnet.com/speedway

XILINX®

AVNET® electronics marketing

Avnet Green Initiative

Supplier Authorized Distributor

Enabling success from the center of technology™

1 800 332 8638
em.avnet.com

EDITOR IN CHIEF — Carlis Collins
carlis.collins@xilinx.com
408-879-4519

EXECUTIVE EDITOR — Forrest Couch
forrest.couch@xilinx.com
408-879-5270

MANAGING EDITOR — Charmaine Cooper Hussain

ONLINE EDITOR — Tom Pyles
tom.pyles@xilinx.com
720-652-3883

ART DIRECTOR — Scott Blair

ADVERTISING SALES — Dan Teie
1-800-493-5551

Xcell Publications

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

What's New?

New devices, new software, new services, and new publications to make it all easy to understand.

Our reader surveys tell us that the *Xcell Journal* is the primary source of programmable logic information for most of our customers. Your responses are overwhelmingly positive, and most of you tell us that you want more of the same; more articles, more information, more often. In response, we are launching a new series of publications that focus on the topics in which you are most interested.

We have three new magazines that will initially be produced every six months:

- *Embedded Magazine*, focusing on the use of embedded processors in Xilinx® programmable logic devices. (You'll discover simple and efficient methods of implementing highly optimized application-specific processor platforms more quickly on our devices.)
- *DSP Magazine*, focusing on the high-performance capabilities of our FPGA-based reconfigurable DSPs. (You'll find that you can create extremely high-performance DSP designs with relative ease.)
- *I/O Magazine*, focusing on the wide range of serial and parallel connectivity options available in Xilinx devices. (You'll see that your choices are almost unlimited.)

These magazines will help you explore topics that interest you the most and not waste your time with irrelevant details.

In addition to our new magazines, we've also created Solution Guides, designed to provide useful information on a wide range of hot topics such as Broadcast Engineering, Power Management, and Signal Integrity. As new industry challenges emerge, these Solution Guides will help you find answers quickly.

Until we assess the demand, these new publications are primarily available as PDF downloads (of individual articles and of whole publications). If we find sufficient demand, we plan to offer subscriptions to printed versions of these publications as well.

Please visit our publications website at *www.xilinx.com/xcell*.

If you find these new publications of value to you, send me a note and tell me what you think. For example, are PDF formats sufficient, or do you prefer printed copies mailed to you, like the *Xcell Journal*? Would you prefer to read them online in HTML format?

Thanks again for your feedback; it helps us serve you better.

Forrest Couch
Executive Editor

FIRST QUARTER 2006, ISSUE 56

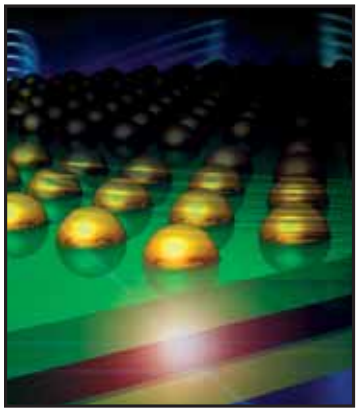

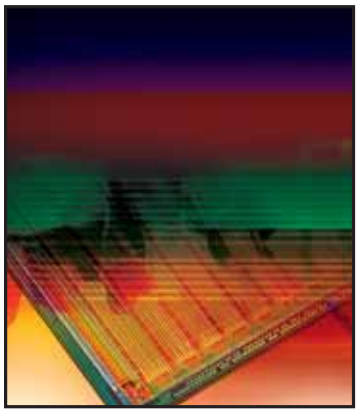

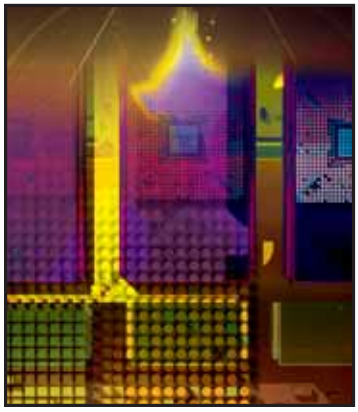

Xcell journal

ARTICLES

Capturing Data from Gigasample Analog-to-Digital Converters

Interfacing National Semiconductor's ADC08D1500 to the Virtex-4 FPGA allows quick-start customer application development.

by Ian King
Application Engineer
National Semiconductor
ian.king@nsc.com

Data conversion within the test and measurement domain and communications industry is moving into the gigasamples per second (GSPS) range. Developing a system capable of processing data at these speeds requires diverse engineering disciplines from the initial system concept through to board design, FPGA logic design, signal processing, and application software.

National Semiconductor has developed a leading-edge analog-to-digital (A/D) converter that can deliver as many as three billion samples per second to an 8-bit resolution. One of the main system design questions from customers regarding this product is how data can be reliably captured and processed at this speed. Therefore, National's applications team designed a development platform to provide a solution to this query and demonstrate a reliable data-capture method. This allows the design focus to shift away from the high-speed front end so that developers can focus on their intended application.

The platform also demonstrates that high clock speeds can be reached while maintaining low power dissipation sufficient for the entire system to be housed in a small enclosure, as would be required for a commercial or industrial system. In this article, I'll explain the techniques and analysis involved in achieving this goal.

Power Considerations

When selecting an FPGA for data capture that can achieve low power levels and performance, a 90 nm device is the first choice. In applications where data is captured in bursts (such as oscilloscopes and radar), the static power of the FPGA device becomes an important factor. This is because the high-speed data transfer between devices takes place over a very short time period, so the capture logic will be static while the application consumes the data.

Figure 1 shows a comparison of Xilinx® Virtex™-4 FPGA static power figures over device density. This indicates that the static power is significantly less than the power consumed by the National Semiconductor ADC08D1500 A/D converter, which is typically 1.8W when running from a 1.5 GHz sample clock. Therefore, for systems processing the captured data in bursts, the ADC can be the main source of heat and power dissipation. Having an ADC with low power figures is a key parameter in the design of products, especially those that are required to be small and portable. The design of this development platform confirms that these qualities are achieved by interfacing the ADC08D1500 to the Virtex-4 device.

Data Transmission

The next consideration for systems using the ADC08D1500 and Virtex-4 FPGA is the signaling between these devices. There are two key issues when handling two channels (each providing data at a rate of 1.5 billion (1.5×10^9) conversions per second):

- Signal integrity between the ADC and FPGA
- The rate of data transfer for each clock cycle

The ADC08D1500 uses low voltage differential signaling (LVDS) for each of its data outputs and clock signal. The main advantage of the LVDS signaling method is that you can achieve high data rates with a very low power budget. Two wires are used for each discrete signal that is to be carried across the circuit board, which should be designed to have a characteristic impedance of 100 Ohms (defined by the LVDS standard). These traces are differentially terminated at the receiver with a 100 Ohm resistor to match the transmission line (see Figure 2).

A signal voltage is generated across the terminating resistor by a 3.5 mA current source within the driving output buffer, which provides a 350 mV signal swing for the receiving circuit to detect. The ADC08D1500 has a total of four 8-bit data buses, plus a clock and over-range signal that require an LVDS type connection to the FPGA (Figure 3). This adds up to a total of 34 differential pairs, all of which require 100 Ohm termination.

The Virtex-4 device offers active digitally controlled impedance (DCI) and a simple passive 100 Ohm termination on-chip within the I/O buffers of the device. These on-chip termination methods eliminate the need to place passive resistors on

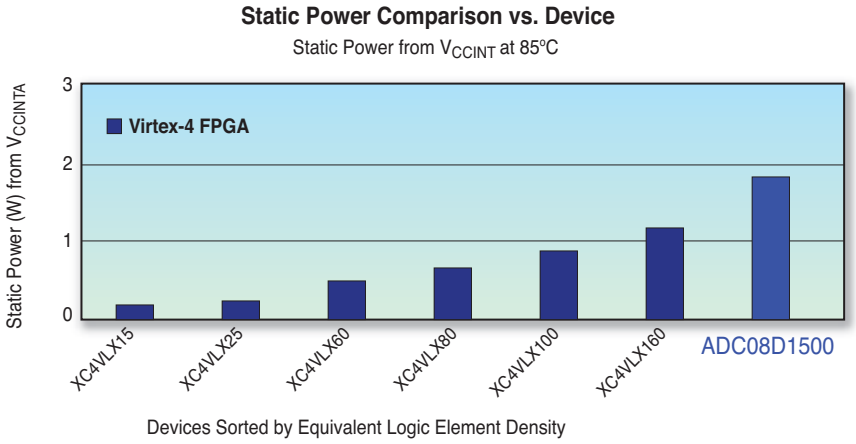

Figure 1 – Comparing the Virtex-4 static power over device density with the operating power of the ADC08D1500

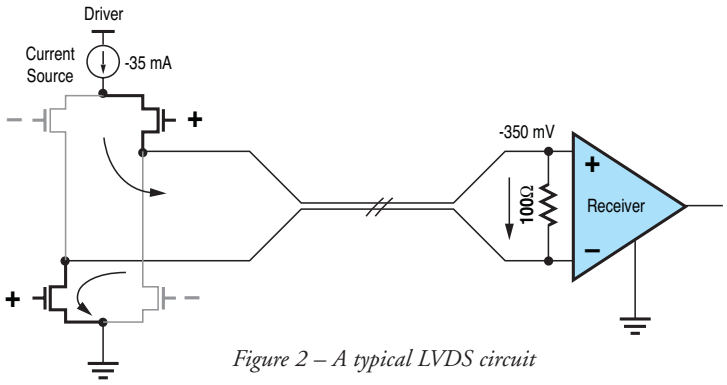

Figure 2 – A typical LVDS circuit

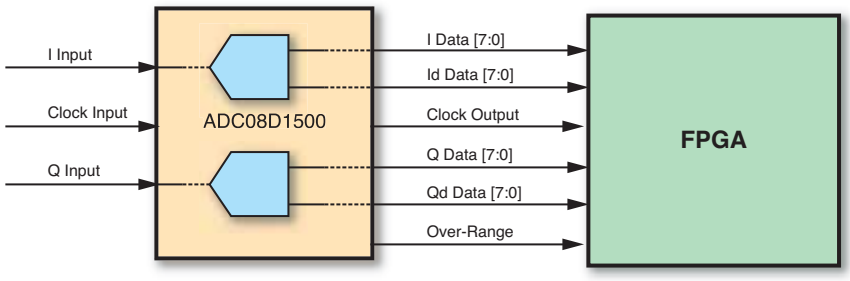

Figure 3 – ADC08D15000 connections to the FPGA

The ADC08D1500 provides a de-multiplexed data output for each of its two channels. Instead of providing a single 8-bit bus running at a data rate equal to the sampling speed, the ADC outputs two consecutive samples simultaneously on two 8-bit data buses (1:2 de-mux).

the circuit board and simplify the routing on the PCB. The DCI option consumes significantly more power than the passive option in this case, simply because of the number of discrete signal lines (68 total) that require termination. Therefore, I would advise turning on the DIFF_TERM feature within each of the IOBs (I/O buffers) to which the ADC signals are connected.

Data Capture

After transmitting data at high speeds using a robust signaling method, it is necessary to store this data into a memory array for post processing. The ADC08D1500 provides a de-multiplexed data output for each of its two channels. Instead of providing a single 8-bit bus running at a data rate equal to the sampling speed, the ADC outputs two consecutive samples simultaneously on two 8-bit data buses (1:2 de-mux).

If the ADC is configured as a single-channel device and put into DES (dual-edge sampling mode), then the sampling speed can be doubled (from 1.5 GSPS to 3.0 GSPS); thus, four consecutive samples are available simultaneously on each of the four buses (1:4 de-mux). This method of de-multiplexing the digital output reduces the data rate to at least half the sampling speed (1:2 de-mux), but increases the number of output data bits from 8 to 16.

For a 1.5 GHz sample rate, the conversion data will be output synchronous to a 750 MHz clock. Even at this reduced speed, FPGA memories and latches would not be able to accept this data directly. It is therefore beneficial to make use of a DDR method, where data is presented to the outputs on the both the rising and falling edges of the clock (Figure 4).

Although the data rate remains the same for DDR signaling, the clock frequency is halved again to a more manageable 375 MHz. This frequency is now in the realms of the FPGA IOB data latches. Before this data can be stored away to memory, a small pipeline constructed from a series of data latches is required. Starting with the inputs, for each data line connected to an IOB pair on the FPGA, two latches will be used to capture the incoming data. One latch is clocked on the rising edge of a phase-locked data clock, while the second latch is clocked using a signal that is 180 degrees out of phase.

The relative position of these clocks should be adjusted so that the edges are aligned with the center of the data eye, taking into account the propagation delay of the signal as it enters the FPGA (Figure 5). To simplify this clocking scheme, the Virtex-4 device is equipped with DCMs that allow these clock signals to be generated internally and can be phase-locked to the incoming data clock.

After latching the incoming data using a DCM, the clock domain must be shifted

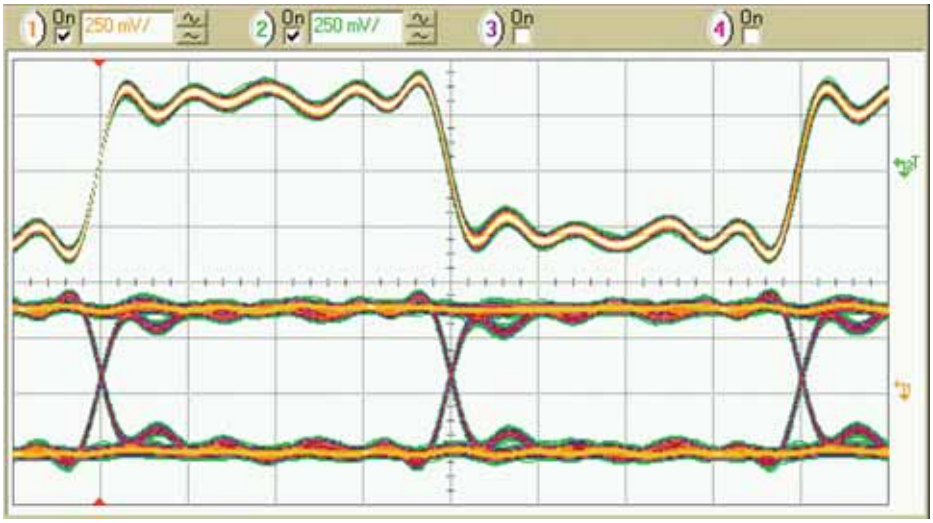

Figure 4 – Oscilloscope plot of clock (top trace) and data from the ADC in DDR mode

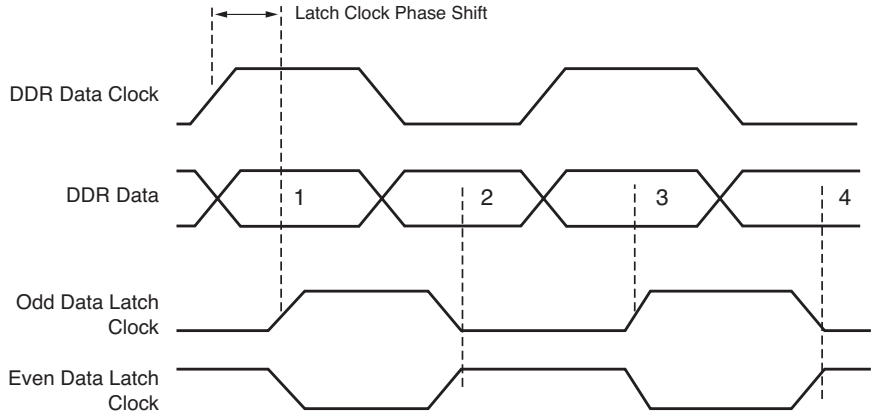

Figure 5 – DDR signaling with DCM-generated data-capture clocks

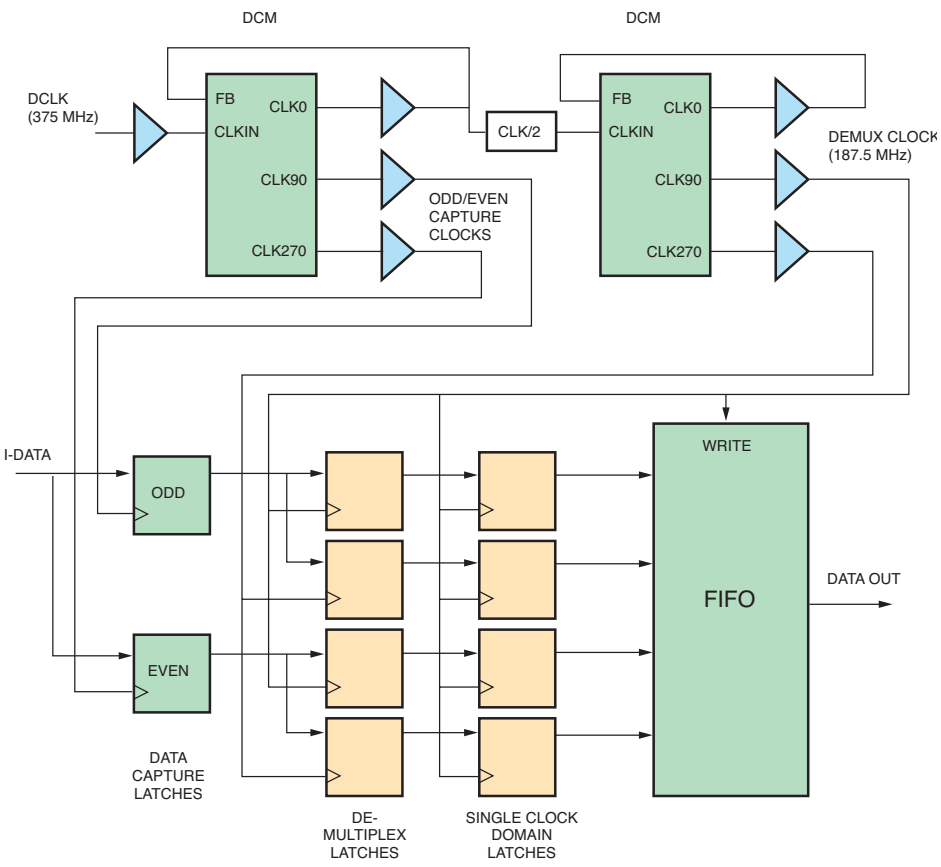

Figure 6 – Data-capture block diagram using two DCMs, latches, and a FIFO memory

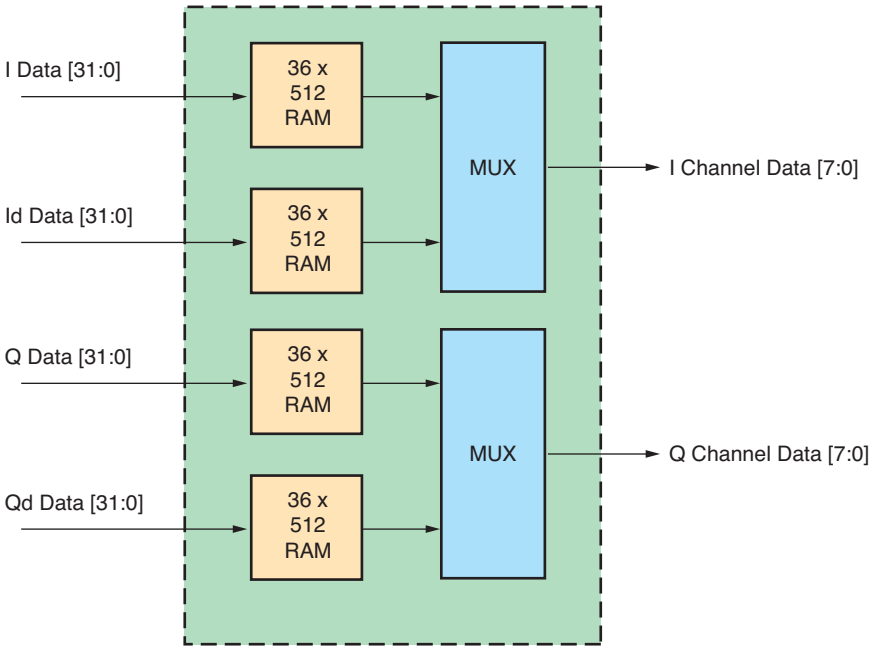

Figure 7 – 128 bit input, 16 bit output, 4 KB deep FIFO

using an intermediate set of latches so that all of the data can be clocked into a memory array on the same clock edge. Because of the speed of the clock, there is not sufficient setup and hold time to re-clock the data; therefore the data must be de-multiplexed again to lower the data rate to 187.5 MHz. Once lowered, the data captured on the out-of-phase clock (even) can be re-captured using the in-phase clock (odd) running at the de-multiplexed rate (see Figure 6).

A second DCM is used to produce the de-mux clock. The clock input frequency is internally divided by two, which produces the 187.5 MHz clock signal. This DCM will provide an output that is phase-locked to the synchronous data clock (DCLK).

Data Storage

As shown in Figure 6, a single 8-bit data bus from the FPGA has been de-multiplexed by four. When all four data buses from the ADC are considered, this method produces a data word 128 bits wide running eight times slower than the sample speed for two-channel operation. The data can now be stored into a FIFO memory buffer.

Creating the custom FIFO for this application is made easy using the Xilinx LogiCORE™ FIFO Generator. Using this software wizard, you can create a FIFO with an input bus width as wide as 256 bits, having an aspect ratio (input-to-output bus width ratio) of 8 to 1. As this design has a 128 bit input bus, the minimum output bus width is 16 bits. This works out well, allowing one 8 bit output bus to be used for I Channel data and the other for the Q channel.

Because the aspect ratio is not 1:1, the FIFO generator will create the memory design using block RAM within the FPGA. A single block RAM can be configured as 36 bits wide by 512 locations deep, so to capture the 128-bit conversion word, the design will use four block RAMs. This gives each channel a 4 KB storage depth without having to cascade FIFO blocks (Figure 7). Having 4K bytes of storage is more than sufficient data for

The low power consumption of the two devices enables systems to operate without forced cooling in small enclosures and does not contribute to a large change in ambient temperature.

a Fast Fourier Transform (see Figure 8) to be applied to the digital conversion of the input signal and represents around 2.7 µS of time-domain information at the 1.5 GHz conversion rate.

Conclusion

When used for the data capture application described, about 85% of the logic fabric inside the Virtex-4 (LX15) device remains available for proprietary firmware development. This leaves space for additional signal processing and data analysis to be performed in hardware, reducing the burden on the software application.

The low power consumption of the two devices enables systems to operate without forced cooling in small enclosures and does not contribute to a large change in ambient temperature. The ability of the Virtex-4 FPGA to operate with low switching noise and to be placed in very close proximity to a high-bandwidth, high-speed data converter without significantly downgrading the measured performance solved my FPGA design challenge.

The two-channel ADC development board discussed in this article is available to order from National Semiconductor in three speed grades: 500 MHz, 1 GHz, and 1.5 GHz. On-board clocking is provided, so all that is required to get started is to provide an analog signal for sampling, plug in the power supply (included), and connect the USB interface to the host PC.

Single-channel device platforms are also available at 1 GHz and 1.5 GHz sample rates. For more information, visit *www.national.com/xilinx and www.national.com/appinfo/adc/ghz_adc.html.*

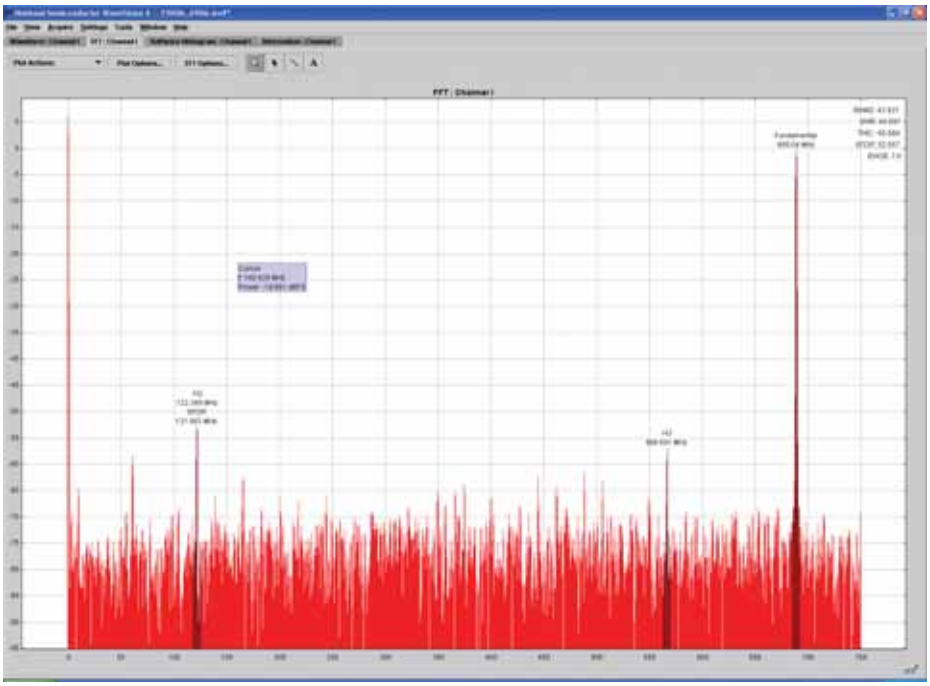

Figure 8 – FFT analysis of 689 MHz input captured by ADC08D1500 and Virtex-4 FPGA

Packet Subsystem on a Chip

Teja's packet-engine technology integrates all of the key aspects of a flexible packet processor.

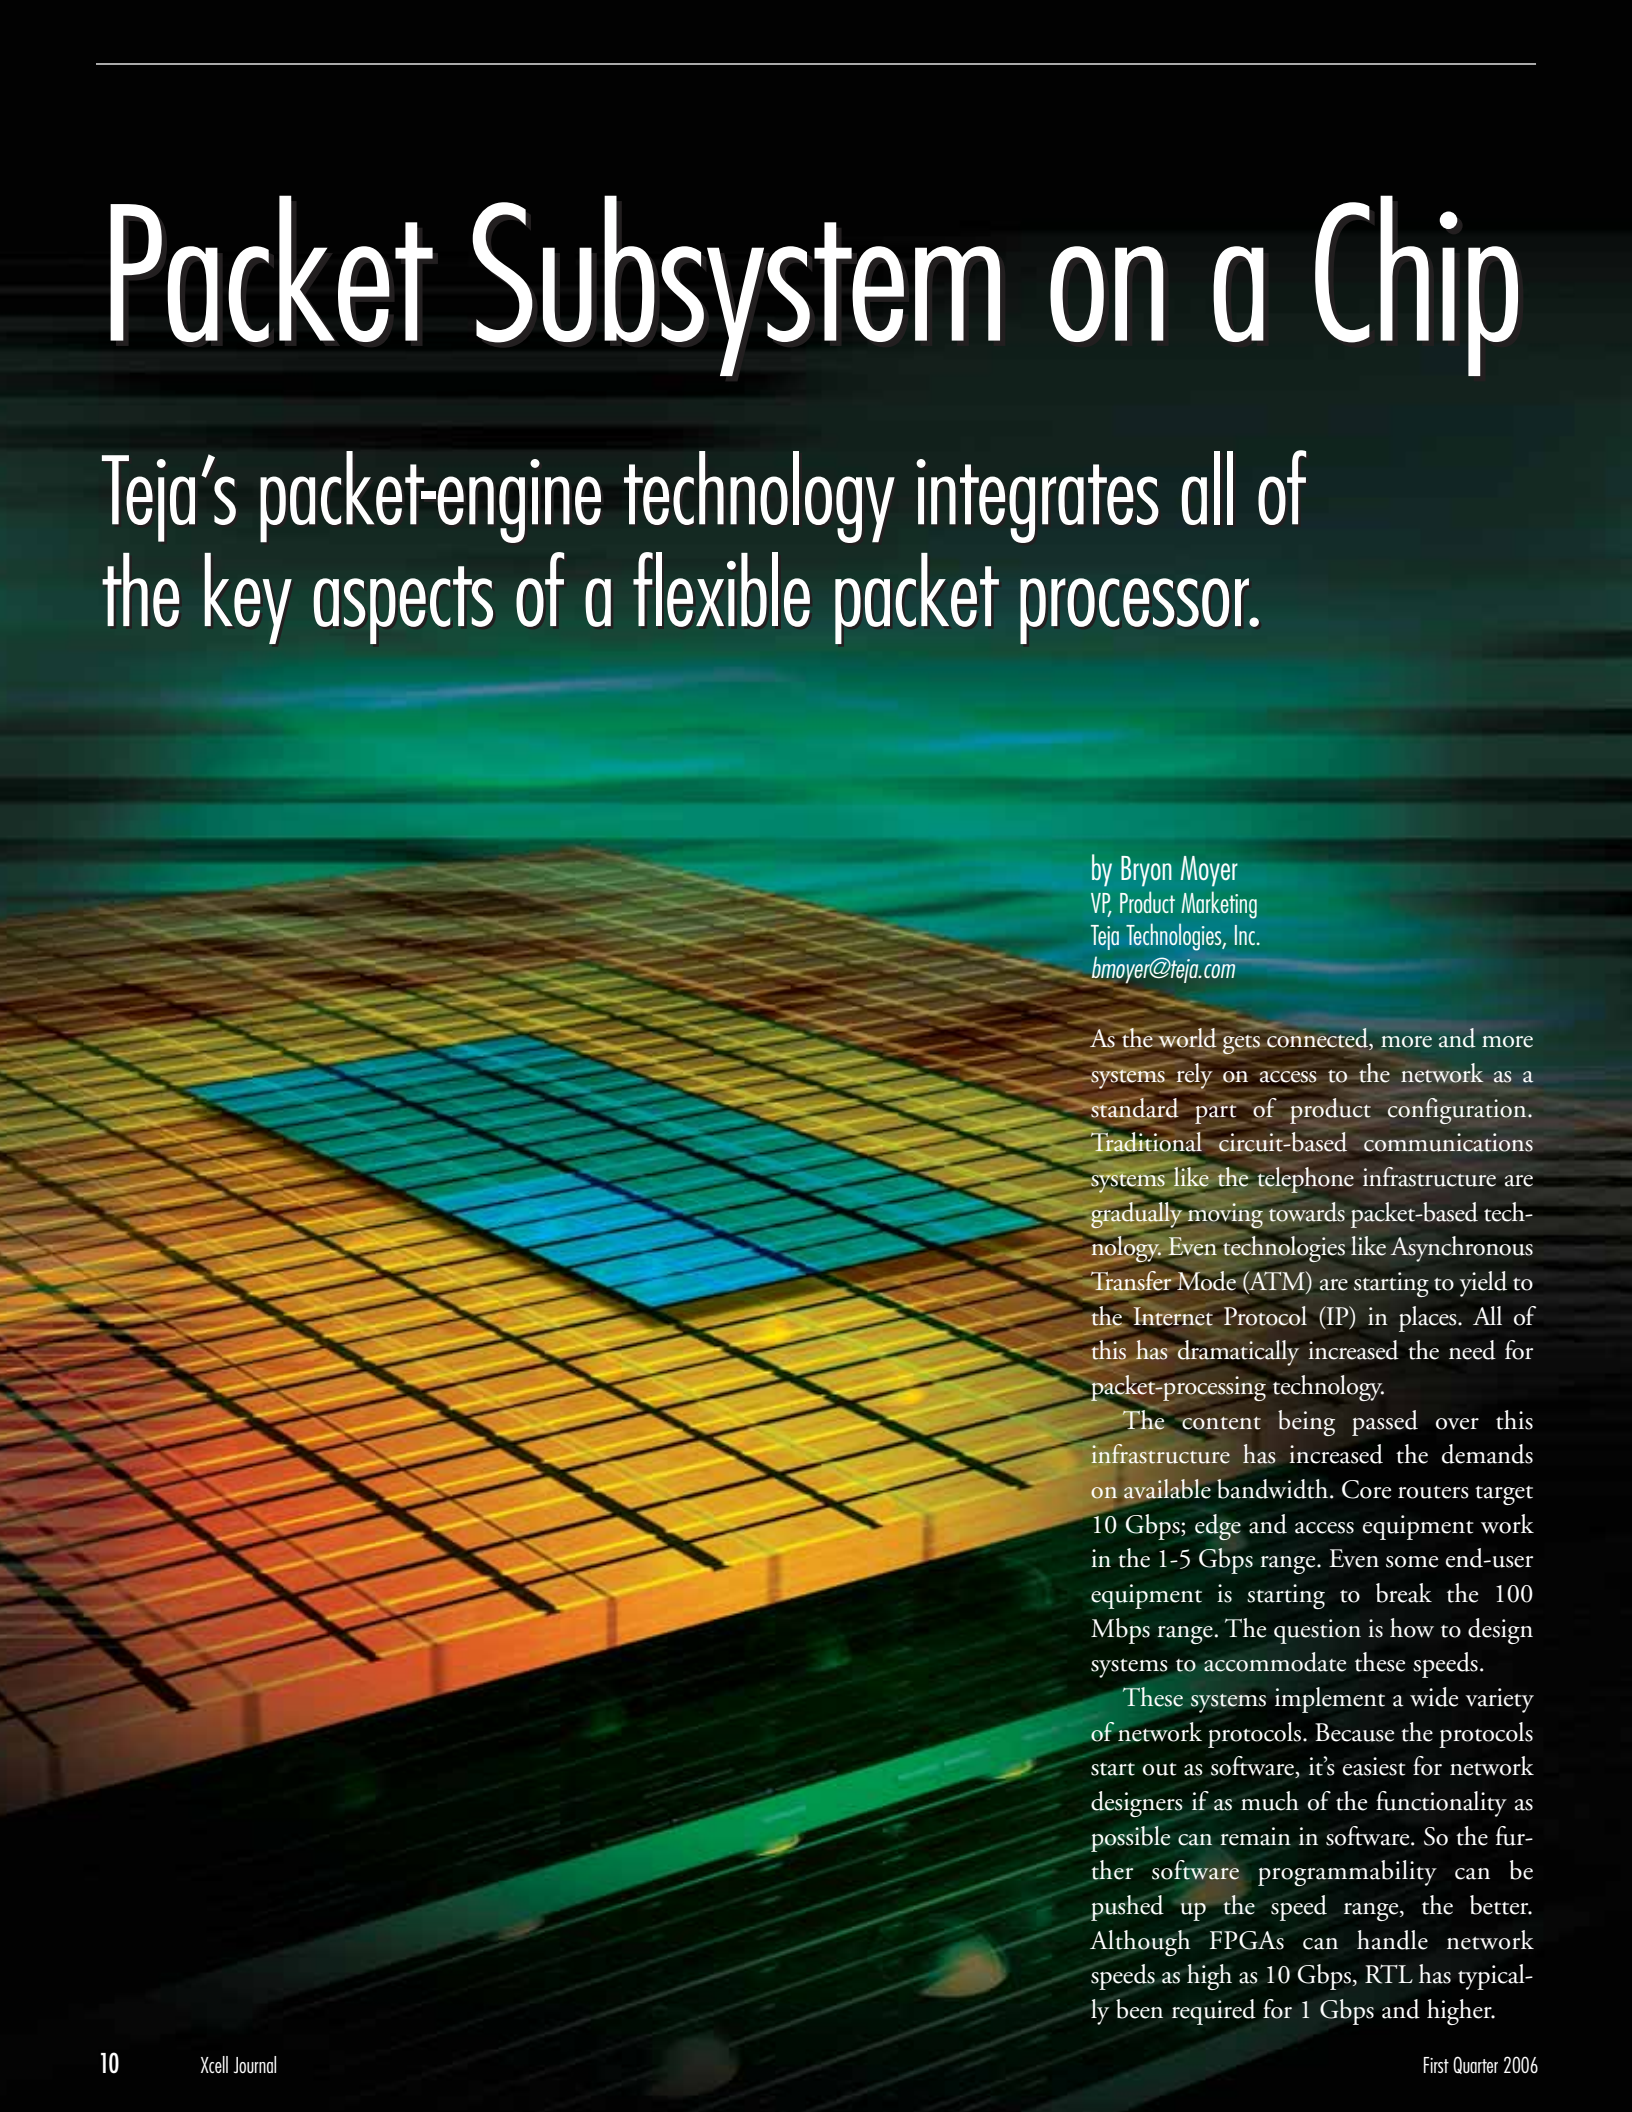

by Bryon Moyer
VP, Product Marketing
Teja Technologies, Inc.
bmoyer@teja.com

As the world gets connected, more and more systems rely on access to the network as a standard part of product configuration. Traditional circuit-based communications systems like the telephone infrastructure are gradually moving towards packet-based technology. Even technologies like Asynchronous Transfer Mode (ATM) are starting to yield to the Internet Protocol (IP) in places. All of this has dramatically increased the need for packet-processing technology.

The content being passed over this infrastructure has increased the demands on available bandwidth. Core routers target 10 Gbps; edge and access equipment work in the 1-5 Gbps range. Even some end-user equipment is starting to break the 100 Mbps range. The question is how to design systems to accommodate these speeds.

These systems implement a wide variety of network protocols. Because the protocols start out as software, it's easiest for network designers if as much of the functionality as possible can remain in software. So the further software programmability can be pushed up the speed range, the better. Although FPGAs can handle network speeds as high as 10 Gbps, RTL has typically been required for 1 Gbps and higher.

Most traffic that goes through the system looks alike, and processors can be optimized for that kind of traffic.

Teja Technologies specializes in packet-processing technologies implemented in high-level software on multi-core environments. Teja has adapted its technology to Xilinx® Virtex™-4 FPGAs, allowing high-level software programmability of a packet-processing engine built out of multiple MicroBlaze™ soft-processor cores. This combination of high-level packet technology and Xilinx silicon and core technology – using Virtex-4 devices with on-board MACs, PHYs, PowerPC™ hard-core processors, and ample memory – provides a complete packet-processing subsystem that can process more than 1 Gbps in network traffic.

The Typical Packet Subsystem

The network "stack" shown in Figure 1 is typically divided between the "control plane" and the "data plane." All of the packets are handled in the data plane; the control plane makes decisions on how the packets should be processed. The lowest layer sees every packet; higher layers will see fewer packets.

The control plane comprises a huge amount of sophisticated software. The data-plane software is simpler, but must operate at very high speed at the lowest layers because it has such a high volume of packets. Packet-processing acceleration usually focuses on layers one to three of the network stack, and sometimes layer four.

Most traffic that goes through the system looks alike, and processors can be optimized for that kind of traffic. For this reason, data-plane systems are often divided into the "fast path," which handles average traffic, and the "slow path," which handles exceptions. Although the slow path can be managed by a standard RISC processor like a PowerPC, the fast path usually uses a dedicated structure like a network processor or an ASIC. The focus of the fast path is typically IP, ATM, VLAN, and similar protocols in layers two and three. Layer four protocols like TCP and UDP are also often accelerated.

Of course, to process packets, there must be a way to deliver the packets to and from the fast-path processor. Coming off an Ethernet port, the packets must first traverse the physical layer logic (layer one of the stack, often a dedicated chip) and then the MAC (part of layer two, also often its own dedicated chip).

One of the most critical elements in getting performance is the memory.

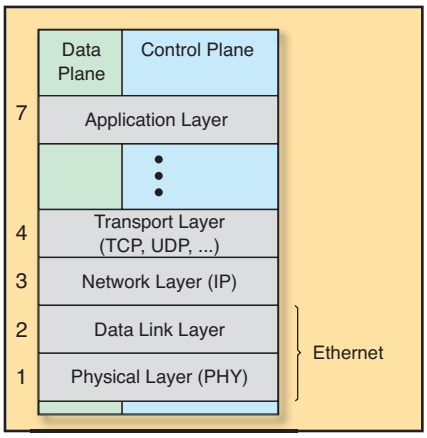

Figure 1 – The network protocol stack

Memory is required for packet storage, table storage, and for program and data storage for both the fast and slow paths. Memory latency has a dramatic impact on speed, so careful construction of the memory architecture is paramount.

Finally, there must be a way for the control plane to access the subsystem. This is important for initialization, making table changes, diagnostics, and other control functions. Such access is typically accomplished through a combination of serial connections and dedicated Ethernet connections, each requiring logic to implement.

A diagram of this subsystem is shown in Figure 2; all of the pieces of this subsystem are critical to achieving the highest performance.

The Teja Packet Pipeline

One effective way to accelerate processing is to use a multi-core pipeline. This allows you to divide the functionality into stages and add parallel elements as needed to hit performance. If you were to try to assemble such a structure manually, you would immediate-

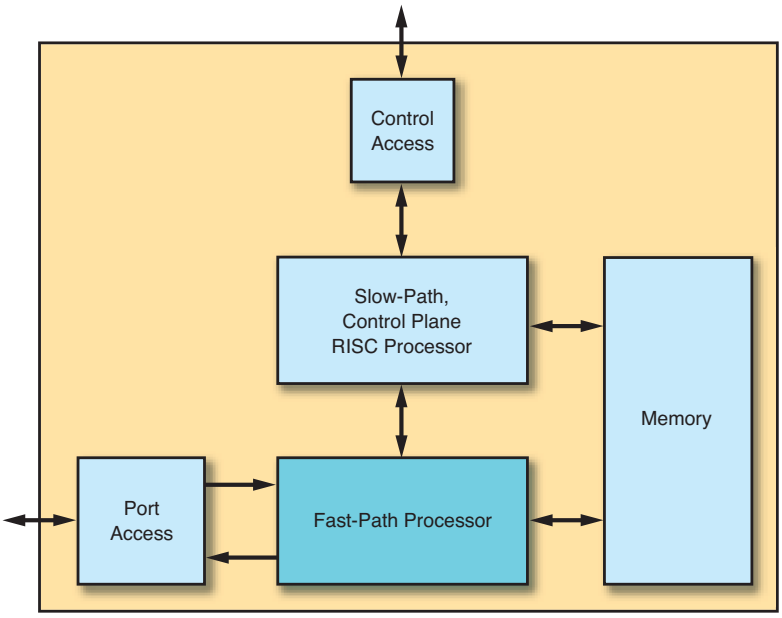

Figure 2 – Typical packet-processing subsystem

ly encounter the kinds of challenges faced by experienced multi-core designers: how to structure communication between stages, scheduling, and shared resource access.

Teja has developed a pipeline structure by creating its own blocks that implement the necessary functions for efficient processing and inter-communication. By taking advantage of this existing infrastructure, you can assemble pipelines easily in a scalable fashion.

The pipeline comprises processing engines connected by communication blocks and accessed through packet access blocks. Figure 3 illustrates this arrangement.

The engine consists primarily of a MicroBlaze processor and some private block RAM on the FPGA. In addition, if a stage has a particularly compute-intensive function like a checksum, or a longer-lead function like an external memory read or write, an offload can be included to accelerate that function. Because the offload can be created as asynchronous if desired, the MicroBlaze processor is free to work on something else while the offload is operating.

The communication blocks manage the transition from stage to stage. As packet information moves forward, the communication block can perform load balancing or route a packet to a particular engine. Although the direction of progress is usually "forward" (left to right, as shown in Figure 3), there are times when a packet must move backwards. An example of this is with IPv4/v6 forwarding, when an IPv6 packet is tunneled on IPv4. Once the IPv4 packet is decapsulated, an internal IPv6 packet is found, and it must go back for IPv6 decapsulation.

Access to the pipeline is provided by a block that takes each packet and delivers the critical parts to the pipeline. Because this block is in the critical path for every packet, it must be very fast, and has been designed by Teja for very high performance.

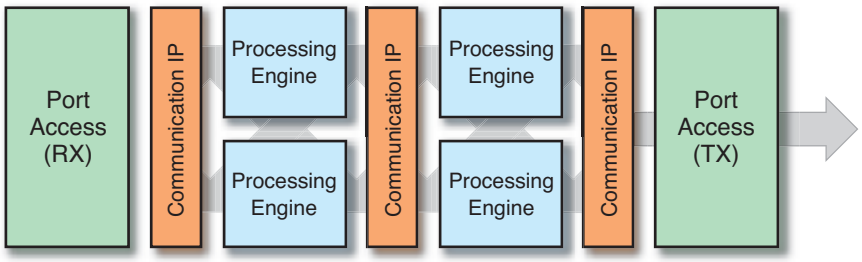

Figure 3 – Teja packet-processing pipeline

The result of this structure is that each MicroBlaze processor and offload can be working on a different packet at any given time. High performance is achieved because many in-flight packets are being handled at once.

The key to this structure is its scalability. Anytime additional performance is needed, you can add more parallel processing, or create another pipeline stage. The reverse is also true: if a given pipeline provides more performance than the target system requires, you can remove engines, making the subsystem more economical.

The Rest of the Subsystem

What is so powerful about the combination of Teja's data-plane engine and the Virtex-4 FX devices is that most of the rest of the subsystem can be moved on-chip. Much of the external memory can now be moved into internal block RAM. Some external memory will still be required, but high-speed DRAM can be directly accessed by the Virtex-4 family, so no intervening glue is required. The chips have built-in Ethernet MACs which, combined with the available PHY IP and RocketIO™ technology, allow direct access from Ethernet ports onto the chip.

The integrated PowerPC cores (as many as two) allow you to implement the slow path and even the entire control plane on the same chip over an embedded operating system such as Linux. You can also provide

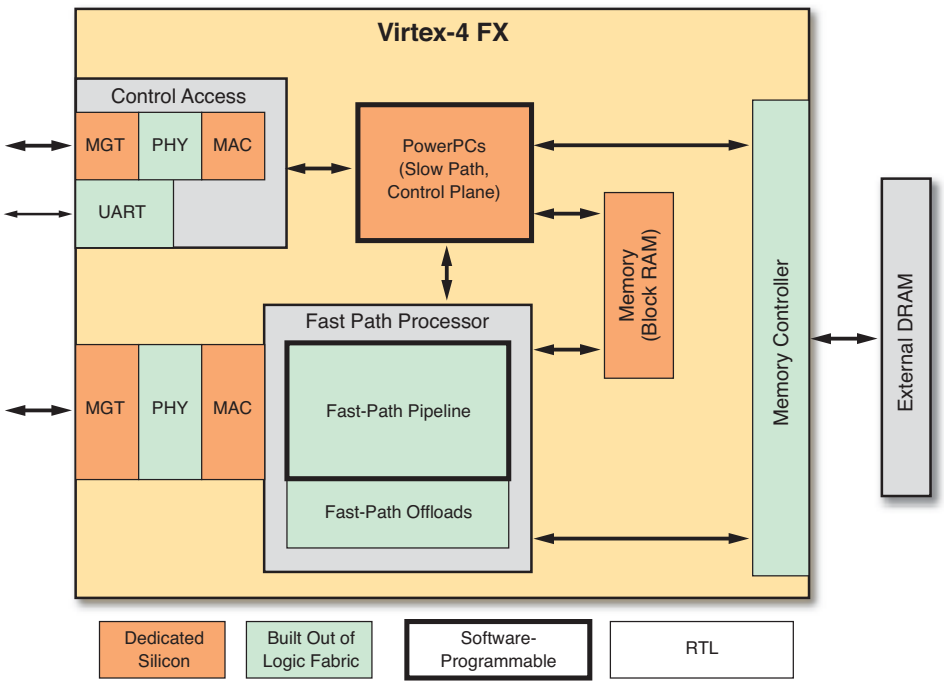

Figure 4 – Single-chip packet-processing subsystem

control access through serial and Ethernet ports using existing IP.

As a result, the entire subsystem shown in Figure 2 (with the exception of some external memory) can be implemented on a single chip, as illustrated in Figure 4.

Flexibility: Customizing, Resizing, Upgrading

Teja's packet-processing infrastructure provides access to our company's real strength: providing data-plane applications that you can customize. We deliver applications such as packet forwarding, TCP, secure gateways, and others with source code. The reason for delivering source code is that if you need to customize the operation of the application, you can alter the delivered application using straight ANSI C. Even though you are using an FPGA, it is still software-programmable, and you can design using standard software methods.

An application as delivered by Teja is guaranteed to operate at a given line rate. When you modify that application, however, the performance may change. Teja's scalable infrastructure allows you to tailor the processor architecture to accommodate the performance requirements in light of changed functionality.

In a non-FPGA implementation, if you cannot meet performance, then you typically have to go to a much larger device, which will most likely be under-utilized (but cost full price). The beauty of FPGA implementation is that the pipeline can be tweaked to be just the right configuration, and only the amount of hardware required is used. The rest is available for other functions.

One of the most important aspects of software programmability is field upgrades. With a software upgrade, you can change your code – as long as you stay within the amount of code store available. As the Teja FPGA packet engine is software-programmable, you can perform software upgrades. But because it uses an FPGA, you can also upgrade the underlying hardware in the field. For example, if a software upgrade requires more code store than is available, you can make a hardware change to make more code store available, and then the software upgrade will be successful. Only an FPGA provides this flexibility.

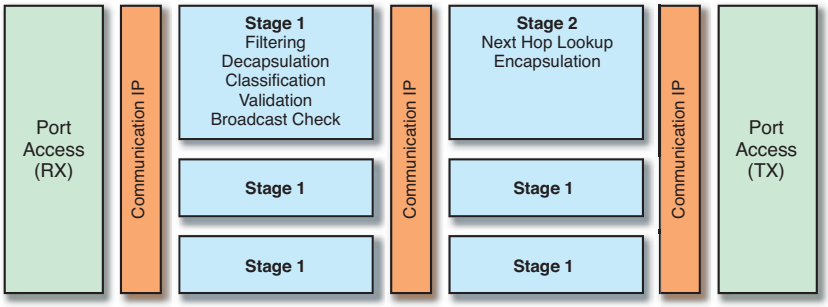

Figure 5 – IPv4 forwarding engine

Because a structure like this is typically designed by high-level system designers and architects, it is important that ANSI C is the primary language. At the lowest level, the hardware infrastructure, the mappings between software and hardware, and the software programs themselves are expressed in C. Teja has created an extensive set of APIs that allow both compile-time and real-time access from the software to the various hardware resources. Additional tools will simplify the task of implementing programs on the pipeline.

IPv4 Forwarding Provides Proof

Teja provides IPv4 and IPv6 forwarding as a complete data-plane application. IPv4 is a relatively simple application that can illustrate the power of this packet engine. It is the workhorse application under most of the Internet today. IPv6 is gradually gaining some ground, with its promise of plenty of IP addresses for the future, but for now IPv4 still dominates. At its most basic, IPv4 comprises the following functions:

- Filtering
- Decapsulation
- Classification
- Validation
- Broadcast check
- Lookup/Next Hop calculation
- Encapsulation

Teja has implemented these in a two-stage pipeline, as shown in Figure 5. Offloads are used for the following functions:

- Checksum calculation
- Hash lookup
- Longest-prefix match
- Memory access

This arrangement provides full gigabit line-rate processing of a continuous stream of 64-byte packets, which is the most stringent Ethernet load.

Conclusion

Teja Technologies has adapted its packet-processing technology to the Virtex-4 FX family, creating an infrastructure of IP blocks and APIs that take advantage of Virtex-4 FX features. The high-level customizable applications that Teja offers can be implemented using software methodologies on a MicroBlaze multicore fabric while achieving speeds higher than a gigabit per second. Software programmability adds to the flexibility and ease of design already inherent in the Virtex family.

The flexibility of the high-level source code algorithms is bolstered by the fact that the underlying hardware utilization can be specifically tuned to the performance requirements of the system. And once deployed, both software and hardware upgrades are possible, dramatically extending the potential life of the system in the field.

Teja Technologies, the Virtex-4 FX family, and the MicroBlaze core provide a single-chip customizable, resizable, and upgradable packet-processing solution.

Learn to increase Design Performance *using the new PlanAhead Software Tool.*

www.xilinx.com/education

The new PlanAhead software tool from Xilinx allows you to quickly analyze, floorplan, and improve placement and timing of even the most complex designs. Use your Xilinx Training Credits or pay as you go for this in-depth, hands-on course:

Designing with PlanAhead **(2 Days)**

Perfect for FPGA designers, system architects, and system engineers interested in analyzing and driving the physical implementation of their designs to maximize performance and capacity. With this comprehensive training, you will develop the necessary skills to import designs into the PlanAhead software, analyze design statistics, partition and floorplan, improve performance, run Design Rule Checks (DRC) and implement a floorplanned design, and more.

Visit ***www.xilinx.com/education*** for more information and registration. Or call us at 1-877-XLX-CLASS.

XILINX®

Designing Reconfigurable Computing Solutions

The Virtex family of FPGAs is the foundation for Cray XD1 high-performance co-processing solutions.

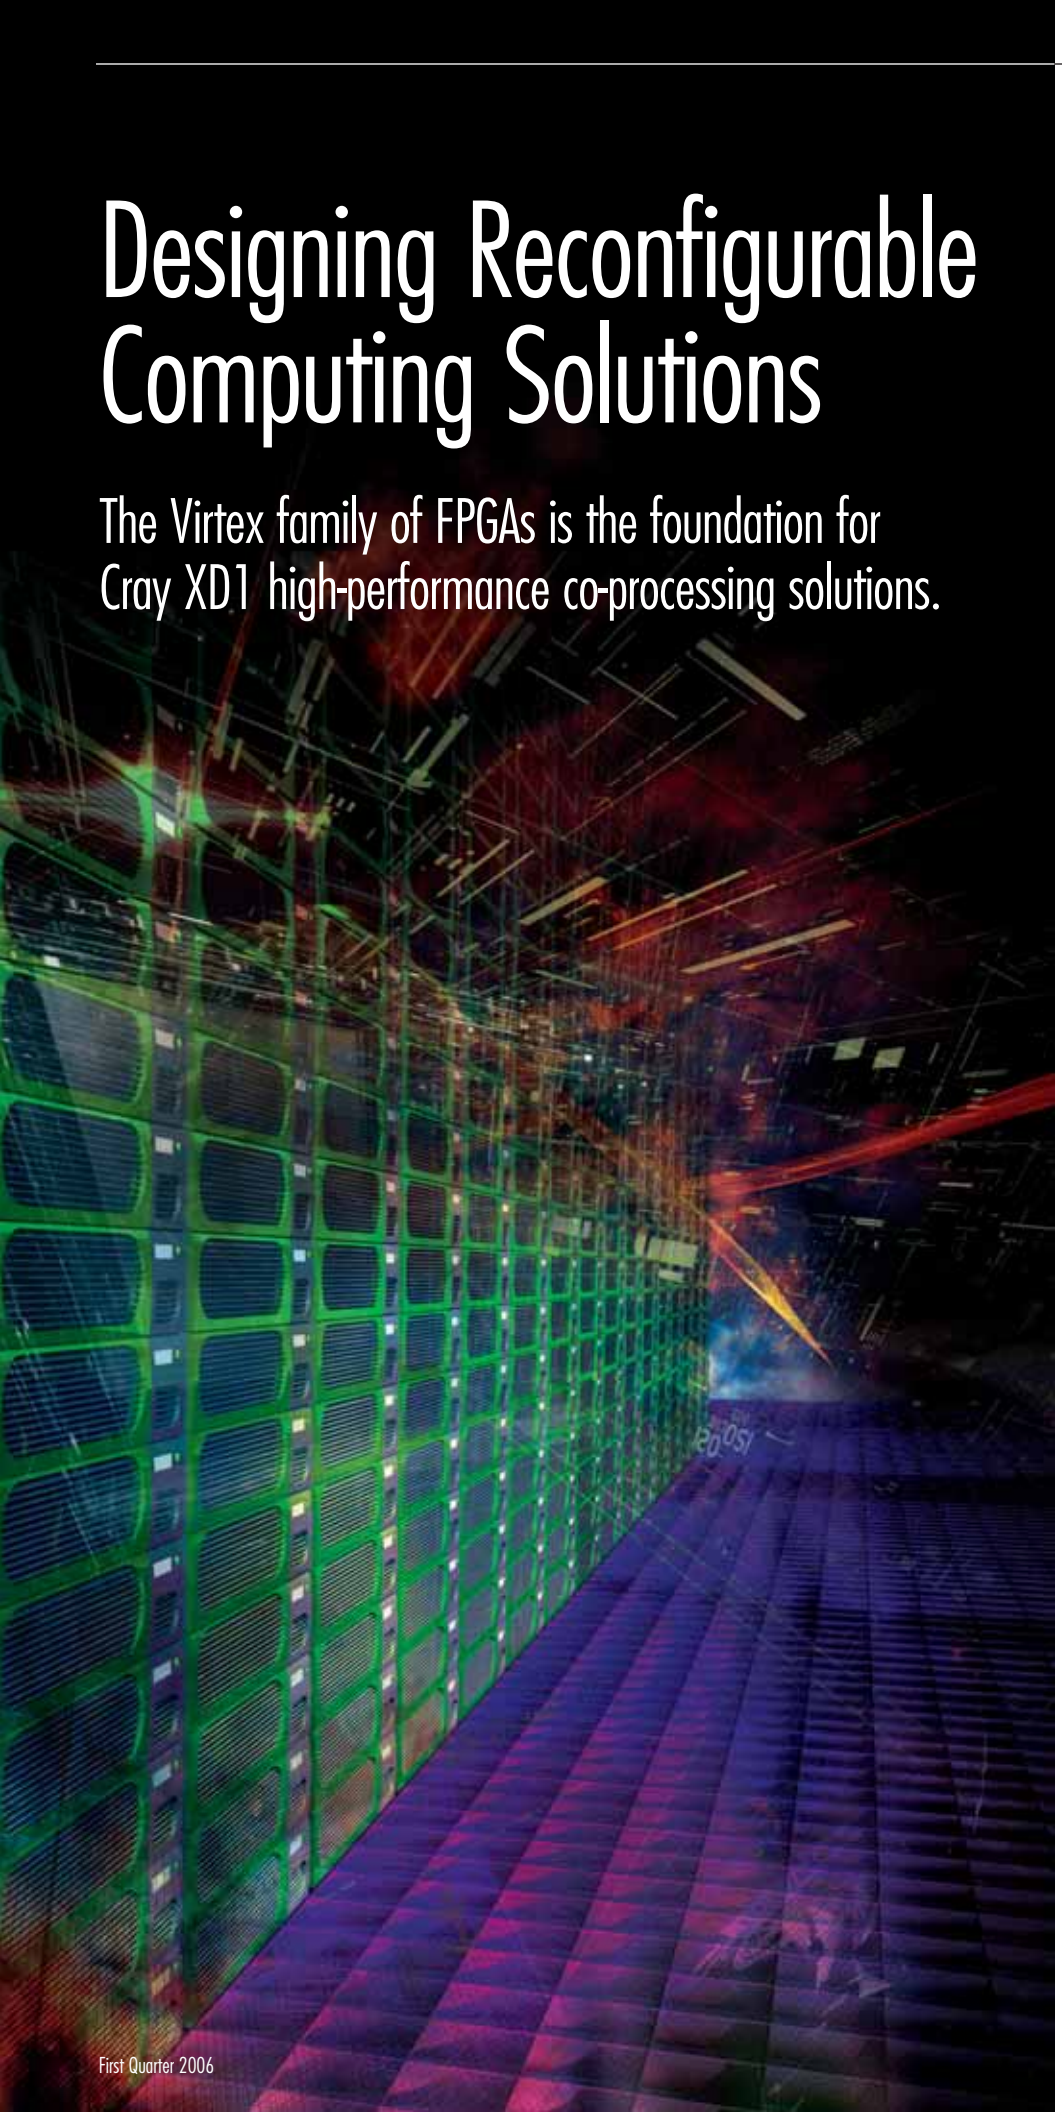

by Geert C. Wenes, Ph.D.
HPC Architect
Cray, Inc.
geertw@cray.com

Steve Margerm
Senior Hardware Designer
Cray, Inc.
smargerm@cray.com

Sriram R. Chelluri
Marketing Manager
Xilinx, Inc.
sriram.chelluri@xilinx.com

The reconfigurable computing (RC) architecture enables software logic that can be reconfigured or reprogrammed to implement specific functionalities on tunable hardware rather than on a general-purpose processor (GPP). RC can achieve orders-of-magnitude performance improvements on selected applications. An RC solution in a high-performance computing environment should include tight coupling and a transparent interface to a general-purpose processor and its auxiliary resources (storage, I/O, networking).

Because FPGAs are tunable, high-density logic cores that provide high-performance and low-latency features, they are well suited for RC solutions. In this article, we'll describe what RC is, how it fits into the Cray, Inc. XD1 high-performance architecture, and why Cray selected Xilinx as the foundation for high-performance application acceleration.

What's Different about RC?

A GPP is primarily designed to execute sequential instructions. System designers have introduced parallelism in two different ways: off-chip, by simply clustering many GPPs and distributing the workload among these multiple processors; or on-chip, either by duplicating the independent computational units on the chip or (more recently) by doubling the cores on a processor.

In contrast, FPGAs can perform many operations in parallel. Xilinx® FPGAs with embedded PowerPC™ processors can perform fixed-point arithmetic and embedded functions, offloading the main processor to other tasks. Furthermore, standard HDLs

for logic devices (such as VHDL and Verilog) have at their core the notions of parallel execution of statements and event-driven simulation.

FPGAs, on the other hand, do not suffer from serial execution model restrictions found in standard languages such as C for generic CPUs. FPGAs are increasingly used to provide application performance. Industries such as manufacturing, government, research, media, and biosciences are deploying FPGAs as hardware application accelerators that can provide orders-of-magnitude (10 to 100x) performance improvements on selected applications over generic microprocessors. The advantages of such an RC solution are:

- Dramatic overall system performance gains at only incremental co-processing board costs
- The ability to reprogram, utilizing a consistent API to the hardware
- Easily upgrade or retarget hardware

Cray XD1 Architecture

The Cray XD1 high-performance computer is based on the directly connected processor (DCP) architecture. The DCP architecture views the system as a pool of processing, logic, and memory resources interconnected by a high-bandwidth low-latency network. This innovative new computer unifies as many as hundreds of processors into a single, resilient computer. The Cray XD1 system combines both on-chip and off-chip parallelism.

The Cray XD1 architecture includes three key subsystems:

- Compute environment. The Cray XD1 compute subsystem comprises single- or dual-core AMD Opteron 64-bit processors integrated on a single board and six blades constituting a chassis (Figure 1). The operating system is Linux, supporting 32- and 64-bit x86-compatible software.
- Interconnect. The Cray XD1 RapidArray interconnect directly connects blades over high-speed, low-latency pathways (Figure 2). Each fully configured chassis includes two interconnect components:
  - An FPGA-based RapidArray communications processor (RAP) is tightly coupled to the AMD Opteron processors and switching fabric to offload and accelerate communications functions from the Opteron processors, freeing the latter to perform core compute tasks and enabling concurrent computing and communication. The FPGA enables interconnect bandwidth on par with memory bandwidth, solving a major system performance bottleneck.
  - The RapidArray embedded switching fabric is a 96 GB/s, non-blocking, crossbar switching fabric in each chassis that provides four 2 GB/s links to each node and twenty-four 2 GB/s inter-chassis links.
- Acceleration application modules. The application acceleration subsystem incorporates RC capabilities to deliver substantial performance increases for targeted applications (Figure 3). Each Cray XD1 chassis can be configured with six application acceleration processors (one per blade), originally developed using the Xilinx Virtex™-II Pro device but soon shipping with Virtex-4 FPGAs that you can program to accelerate key algorithms. The application acceleration processors are tightly integrated with Linux and the AMD Opteron processors and use standard software programming APIs, removing a major obstacle to application development.

Figure 1 – Single Cray XD1 chassis

The FPGA as Co-Processor Model

Cray views the FPGA as a very tightly coupled application accelerator platform to speed up computations on demanding applications. The FPGAs are designed into an expansion module that connects to the high-speed, low-latency, point-to-point HyperTransport subsystem.

There are three main advantages that an FPGA has over a microprocessor:

- FPGAs have a flexible architecture. You can customize and optimize the logic and manipulate variable length data.
- FPGAs are inherently parallel devices.
- FPGAs can be reprogrammed to execute new applications without having to change the hardware.

Xilinx Solution

Cray selected the Virtex-II Pro and Virtex-4 series of FPGAs as the foundation for its RC solutions because of their industry-leading technical features and support for third-party development.

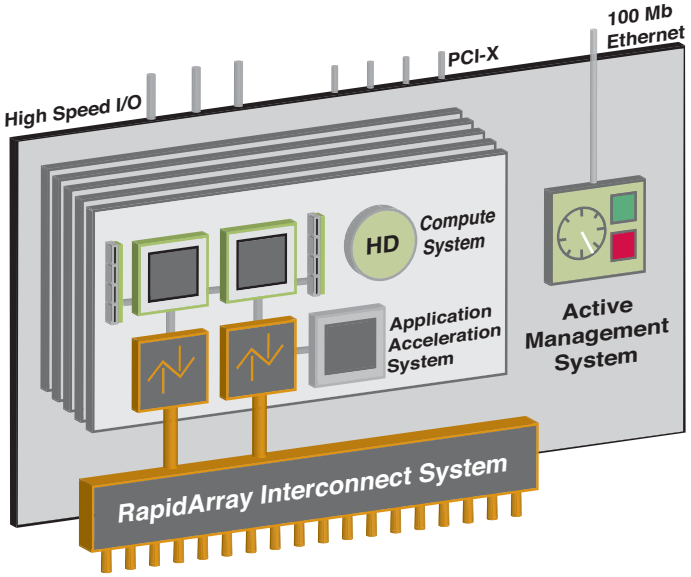

Figure 2 – Cray XD1 compute blade

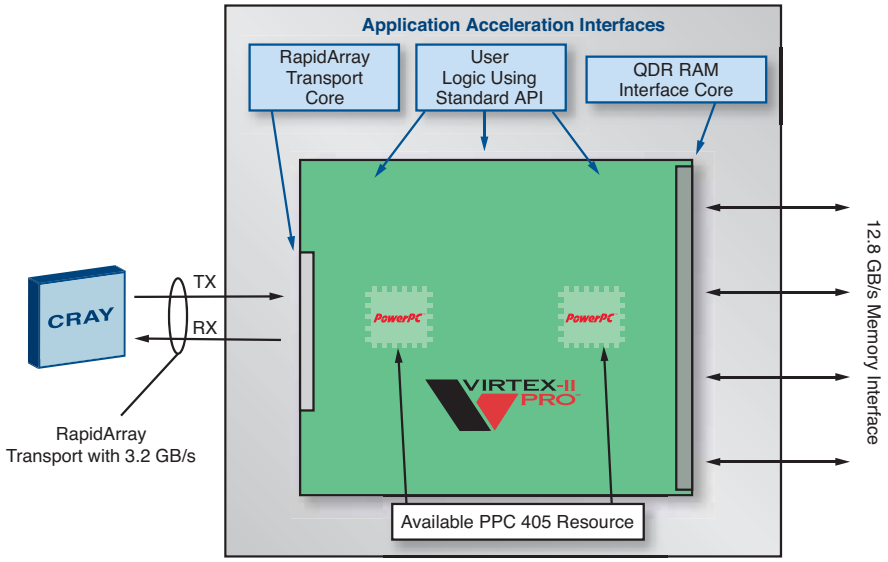

Figure 3 – Xilinx-based application acceleration module

Advantages of Xilinx Virtex FPGAs include:

- Low latency and a high throughput range of 622 Mbps-10.3125 Gbps
- High-speed 500 MHz internal I/O clock
- Low power consumption
- Optional embedded hardware PowerPC processor or a software co-processor
- Built-in scalable RAM

Cray chose the Virtex family of FPGAs because it was developed for high performance, from low- to high-density designs based on IP cores and customized modules.

Development Environment

Many Xilinx and third-party tools simplify the development of co-processing solutions. Software developers can interface with software APIs provided by Cray and use a rich variety of third-party design tools to create their own applications using C and C++.

For instance, Celoxica Ltd. provides a C-based design and synthesis tool (DK Design Suite) for customers who want to use a software design flow to accelerate their applications using FPGAs integrated into the Cray XD1 supercomputer (*www.celoxica.com/products/dk/default.asp*).

As another example, the Cray XD1 system equipped with Mitrionics' Mitrion Virtual Processor and Mitrion Software Development Kit make it possible for supercomputer users to program FPGAs integrated into the Cray XD1 system on a software level, reducing the time and effort required to take advantage of FPGA-based computation (*www.mitrionics.com/technology.shtml*).

For DSP solutions, Xilinx System Generator for DSP enables you to design DSP blocks using commercial tools like the MATLAB package from The MathWorks.

Applications Acceleration

The XD1 has proven to be a successful product. For example, the Naval Research Laboratory (NRL) facility in Washington, D.C., is home to one of the largest Cray XD1 supercomputers ever installed and also employs the largest known number of application acceleration modules in the world. Equipped with 288 AMD Opteron dual-core processors and 144 Virtex-II Pro FPGAs, the 24-chassis machine will provide peak performance of 2.5 teraflops.

The use of FPGAs as application accelerators has been successfully demonstrated in many fields. For instance:

- In encryption/decryption applications, the RC5 cipher-breaking application runs 1000x faster than on a 2.4 GHz Pentium 4, while for elliptic curve cryptography, speedups of 895 to 1300x compared to (a relatively slow) 1 GHz Pentium III are possible. Encryption algorithms like 3DES have been shown to run at more than 16 Gbps throughput from a high-level abstraction language like Mobius.
- In bioinformatics applications, the well-known Smith-Waterman code performs about 26 times faster than on the AMD Opteron, while applications in proteomics such as thin-spline algorithms for comparing 2D gel contents run more than 20 times faster, reducing analysis times from days to hours.
- Complex, realistic vehicular traffic simulation codes perform 34x faster on Virtex-II Pro devices than on a 2.2 GHz Opteron, and remarkable sustained bandwidths from the FPGA to AMD processor have been observed at more than 1 GB/s, considerably higher than any PCI bandwidth.

Conclusion

Cray's FPGA-based RC solutions are ideal for applications in industries such as manufacturing, government, research, media, and biosciences. Cray selected the Virtex-II Pro and Virtex-4 series because of their high performance, low latency, scalable memory, and an array of software tools and support.

With the Cray XD1 high-performance computer, reconfigurable computing has taken a major step forward by breaking down performance barriers at substantially lowered cost by using off-the-shelf components from Xilinx to solve difficult computational problems.

For more information about Cray RC solutions' HPC architectures, contact *geertw@cray.com* or visit *www.cray.com/products/xd1/index.html*.

For more information about Xilinx FPGA-based co-processing solutions, contact *sriramc@xilinx.com* or visit *www.xilinx.com/products/design_resources/dsp_central/resource/coprocessing.htm*.

Integrating Security Services into Network Equipment

Using Sensory Networks' NodalCore technology, Virtex-4 FPGAs can solve the design challenge of integrating high-performance security acceleration to cope with continually evolving network threats.

by Ryan Junee
Product Manager
Sensory Networks
ryan.junee@sensorynetworks.com

Sriram R. Chelluri
Vertical Marketing Manager
Xilinx, Inc.
sriram.chelluri@xilinx.com

In today's network environment, security threats are everywhere – from Web- and e-mail-borne viruses and worms to productivity-hampering spam and spyware to malicious break-in attempts. Security functionality in network equipment is no longer "nice to have" – it is now a fundamental design requirement.

Designers of network equipment must be able to integrate applications such as antivirus, anti-spam, anti-spyware, intrusion detection and prevention, content filtering, and sometimes all of these in one box (a universal threat management or UTM appliance), without sacrificing performance throughput or detection accuracy. With today's rapidly increasing network speeds and the growing taxonomy of network threats, this is a difficult design challenge.

In this article, we'll present an FPGA-based solution to the problem and describe how to integrate security functions using an architecture that can be easily upgraded over time.

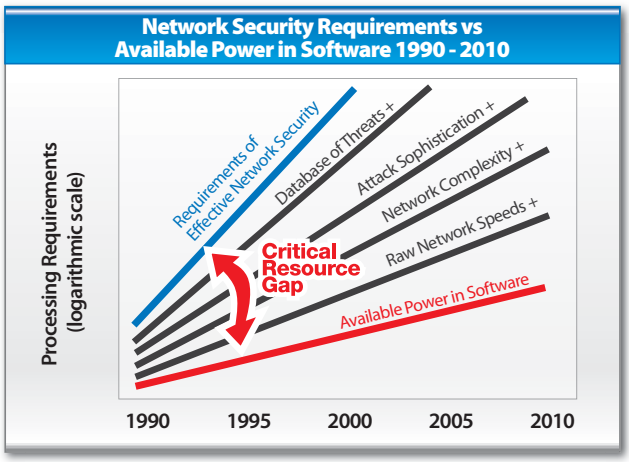

Figure 1 – Critical resource gap in traditional security appliances

Current Technical Challenges

Today's security devices provide a "best effort" service but fall dramatically short of providing comprehensive protection. The growing number and sophistication of attacks and the exponential growth in raw network traffic leads to increases in both complexity and processing speed requirements in network security applications (Figure 1). This sort of processing power is simply not available in traditional security appliances built from standard PC hardware packaged into a rack-mount chassis.

Figure 2 shows an example of the reduced accuracy in Snort (the industry-standard open-source intrusion detection system) as the number of rules (types of attacks detected) and the traffic rate are increased. Any IDS that cannot scale and maintain an accuracy rate of more than 90% is not worth implementing. Security processing has quickly become the major bottleneck in today's network environments, and is often dealt with by reducing the number of attack signatures that are detected or by dropping packets, leading to missed attacks and security holes in both cases.

There are two forms of data inspection for network security appliances:

- Header (or packet) processing. Processing is restricted to read and write operations on the packet header; thus the processing complexity is independent of the size of the packet. Examples include TCP flow reassembly, IP forwarding, and some Snort header rules.
- Content processing. Read and write operations are applied to all of the data in the packet; in particular, the payload of the packet. Examples include pattern matching (searching packets for viruses), content decoding (MIME), and decompression/encryption (VPNs).

Today's security appliances perform these types of data inspection using either general-purpose processors, network processing units (NPUs), or fixed-function ASICs. These architectures all have drawbacks in performance, scalability, adaptability, and cost.

Under a general-purpose architecture, network traffic is captured using a network interface card (NIC) that connects to the host processor over a bus such as PCI. The function of reassembling data streams is handled by the operating system, while content decode and inspection operations are handled by a user-space software application. Although functionally separate, these operations all share the same CPU, memory, buses, and other system resources. The most computationally expensive operations will therefore consume most of the resources, leaving little processing time for other operations and causing critical bottlenecks.

NPU-based solutions are limited when it comes to content decode and inspection operations. Although NPUs are very good at processing packets, they break down when scanning content across multiple packets. Even multi-core NPUs designed for content processing often cannot scale to wire speed for deep packet inspection.

ASIC-based security appliances are expensive, with long lead times and high risk. They are not adaptable to changing threat environments and cannot be upgraded with new performance and features. This leads to a solution that is not future-proof.

What is needed for the modern-day dynamic environment of network security is an adaptable co-processing platform, with the ability to:

- Accelerate deep packet inspection
- Deal with large databases
- Adapt to new security threats
- Adapt to new content types
- Adapt to new security applications

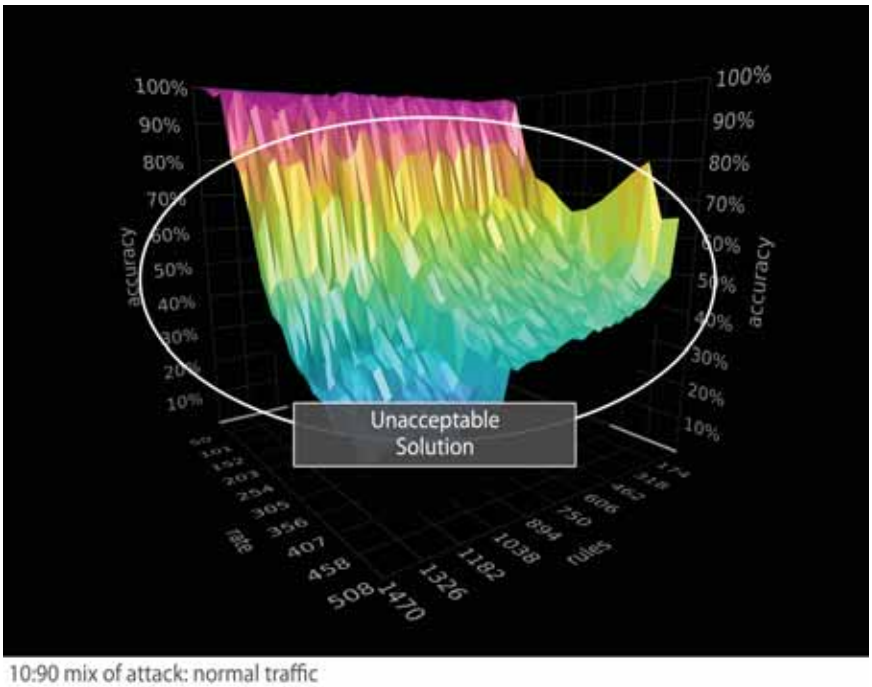

Figure 2 – Snort 2.0 performance on a dual 3 GHz Xeon

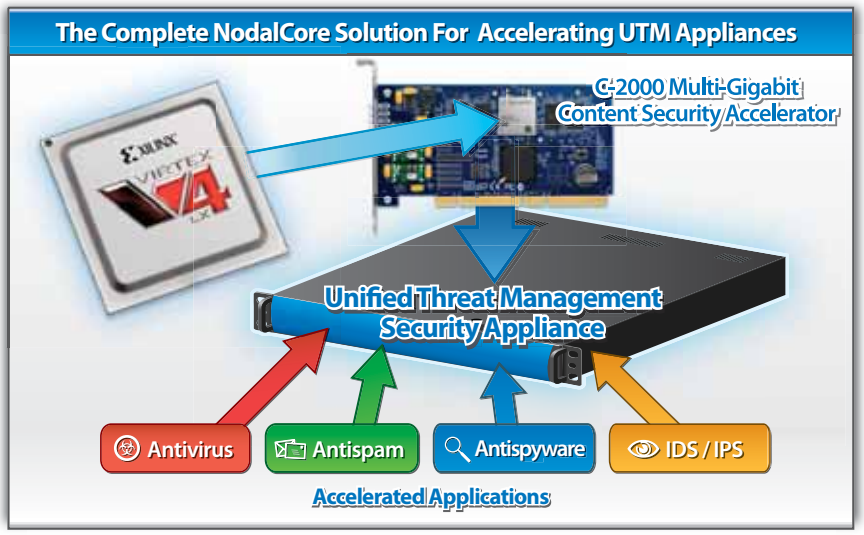

Figure 3 – Building a security appliance with the NodalCore SPU

To meet these design criteria, a security co-processor based on reconfigurable technology is required, and must meet the following guidelines to be market-viable:

- Interfaces with standard architectures (Pentium, NPU, multi-core MIPS)
- Designed for deterministic processing in the data plane (pattern matching, decoding)
- Allows in-field upgrades of functionality to adapt to new threats, new applications, and new content types
- Scalable number of parallel execution units

The NodalCore Co-Processing Solution

Sensory Networks has designed a co-processing solution called the NodalCore Security Processing Unit (SPU) based on Xilinx® Virtex™-II Pro and Virtex-4 FPGAs (Figure 3). The NodalCore architecture can be used in conjunction with a CPU- or NPU-based system to bridge the performance and scalability gap by providing the computational power necessary for the next generation of adaptable security products.

The NodalCore SPU is designed to offload and accelerate bottleneck operations associated with high-performance network security applications such as antivirus, anti-spam, anti-spyware, intrusion detection, and content classification (Figure 4). The NodalCore SPU's dedicated regular expression pattern-matching engine operates with compressed signature databases stored in external SRAM or RLDRAM memory. Using patent-pending high-performance matching algorithms and a scalable architecture, the SPU can achieve data throughput as high as 1.6 Gbps per bank of RLDRAM memory.

Sensory Networks' CorePAKT massive memory architecture (MMA) compression technology ensures that compiled signature databases occupy a memory footprint smaller than any other competing technology on the market. Compression rates of more than 90% are possible, meaning that databases of more than 10 million signatures are easily supported. The ability to store multiple databases in a compressed format while still achieving gigabit throughputs makes NodalCore technology ideal for applications demanding high performance on large signature databases. The NodalCore SPU is one of the only products on the market that can provide wire-speed performance on full commercial antivirus databases that often contain hundreds of thousands of signatures.

The SPU is a packet-based architecture with standard PCI-64/66, QDR, and soon-to-be-released PCI-X and PCI-Express interfaces, allowing seamless integration with standard host systems. The ability to process packets in parallel and rapidly adapt to hostile network security threats is a key differentiator of Sensory Networks' technology. The NodalCore SPU architecture allows hardware modules to be upgraded in the field to deal with new attacks, content types, and application demands.

Heart and Soul

Virtex-II Pro and Virtex-4 FPGAs provide the foundation for NodalCore technology. The dynamic environment of network threats requires a reconfigurable solution with the processing power to handle deep packet inspection – something that cannot be solved by pure general-purpose CPUs or NPUs without encountering performance and scalability issues. A FPGA-based configurable co-processing solution enables security devices to:

- Adapt to changing threats
- Meet performance and scalability requirements

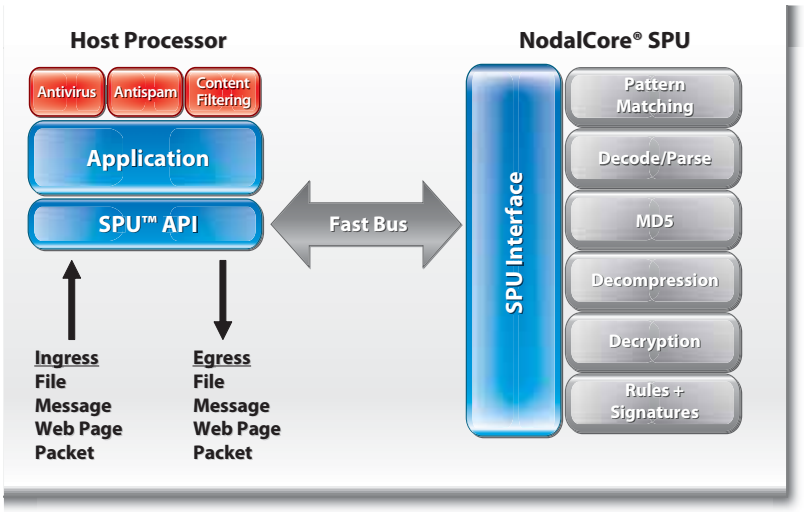

Figure 4 – SPU architecture

- Support a broad range of features and performance at various price points
- Provide phased product releases
- Achieve fast time to market

The Virtex product line comprises low-power and high-performance FPGAs capable of speeds as high as 10.3125 Gbps with built-in multi-gigabit transceivers. Both the Virtex-II and Virtex-4 devices feature embedded PowerPC™ 405 hardware cores that run at speeds as fast as 450 MHz, soft IP cores such as the 32-bit MicroBlaze™ RISC processor, and up to 200,000 cells of configurable logic. The Virtex-4 series features a scalable Advanced Silicon Modular Block (ASMBL) architecture that has allowed Xilinx to produce a programmable silicon platform with unprecedented capabilities and configurations to build co-processing solutions.

The NodalCore SPU was initially designed using the Virtex-II Pro FPGA to meet its high-performance and field-upgradability requirements. The dynamic nature of the network security landscape meant that an FPGA-based solution was a much better fit than an ASIC-based solution.

The ability for appliance vendors to offer after-sale performance upgrades is another important benefit that increases the marketability of Sensory Networks' range of products. The latest NodalCore SPU is designed using the Virtex-4 FPGA because of its higher performance, faster I/O, and attractive cost. It also offers an easy upgrade path to newer and faster FPGA families.

Porting Security Applications to NodalCore

Developers can easily port new or existing security applications to the NodalCore platform and gain immediate improvements in throughput and accuracy. Sensory Networks provides a comprehensive suite of drivers, libraries, and development tools that abstract the hardware layers behind fully-featured C and C++ APIs. You can push bottleneck functionality down to the hardware by making simple function calls to pass the data.

To demonstrate this process, Sensory Networks has ported Clam Antivirus to the NodalCore platform and shown as much as a 50x performance improvement. Clam AV is one of the world's largest open-source antivirus projects, and is used by many large telecommunications companies and managed service providers (see *www.clamav.net*). An analysis of Clam AV showed that the pattern-matching portion of the code was a severe bottleneck, so a NodalCore C-2000 accelerator card was added to the appliance.

Using Sensory Networks' C++ API and Linux device drivers (Clam AV runs under Linux), the pattern-matching code was pushed down to the NodalCore hardware. This meant that the virus signatures also needed to be loaded onto the C-2000. Thanks to Sensory Networks' CorePAKT memory compression technology, the entire ClamAV virus database (more than 40,000 virus signatures) was compiled and loaded onto the card. The resulting improvement in throughput is illustrated in Figure 5, using different base platforms for comparison.

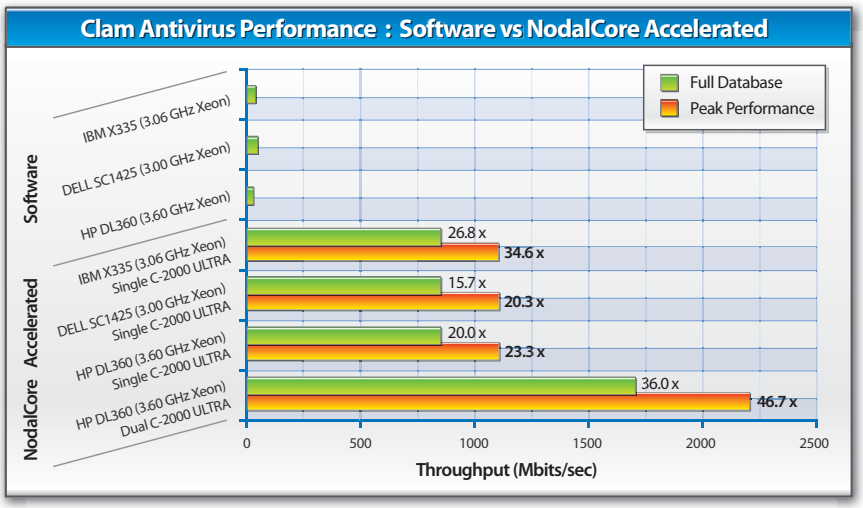

Figure 5 – Antivirus throughput comparison

Integrating Commercial Security Applications

If you do not wish to develop your own security applications, you can select from a wide variety of commercial plug-in applications that run on the NodalCore platform. Sensory Networks has a range of industry-leading application partners that provide best-of-breed antivirus, anti-spam, anti-spyware, intrusion detection, and content filtering applications. In addition, Sensory Networks supports a wide range of open-source security applications such as Clam AV, Spam Assassin, DSpam, and Snort that are fully ported and tested.

These applications can be quickly and easily integrated without the need to write code, providing "security on tap." You simply need to design a Xilinx Virtex-4 FPGA onto the motherboard (or use a separate PCI card) and install the software to have a security product with market-leading performance. Sensory Networks has a range of reference designs available, and an expert consulting services team that can assist with aspects of this development.

Conclusion

Because of the growing number and sophistication of network attacks, designers of network appliances must integrate specialized security processing hardware. This hardware needs to provide gigabit throughput, support very large signature databases, and adapt to new types of threats.

Sensory Networks provides a range of solutions based on the Virtex FPGA family that allow you to quickly and easily integrate security applications into network appliances.

To download white papers, datasheets, and API reference documentation for the NodalCore SPU, visit *www.sensorynetworks.com* or contact *sales@sensorynetworks.com*. Development kits can be ordered online. To learn more about implementing reconfigurable computing solutions, contact *sriram.chelluri@xilinx.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.

Designing for Power Budgets and Effective Thermal Management

Xilinx provides the critical edge with power-optimized FPGAs, thermal-efficient packaging, and power analysis tools.

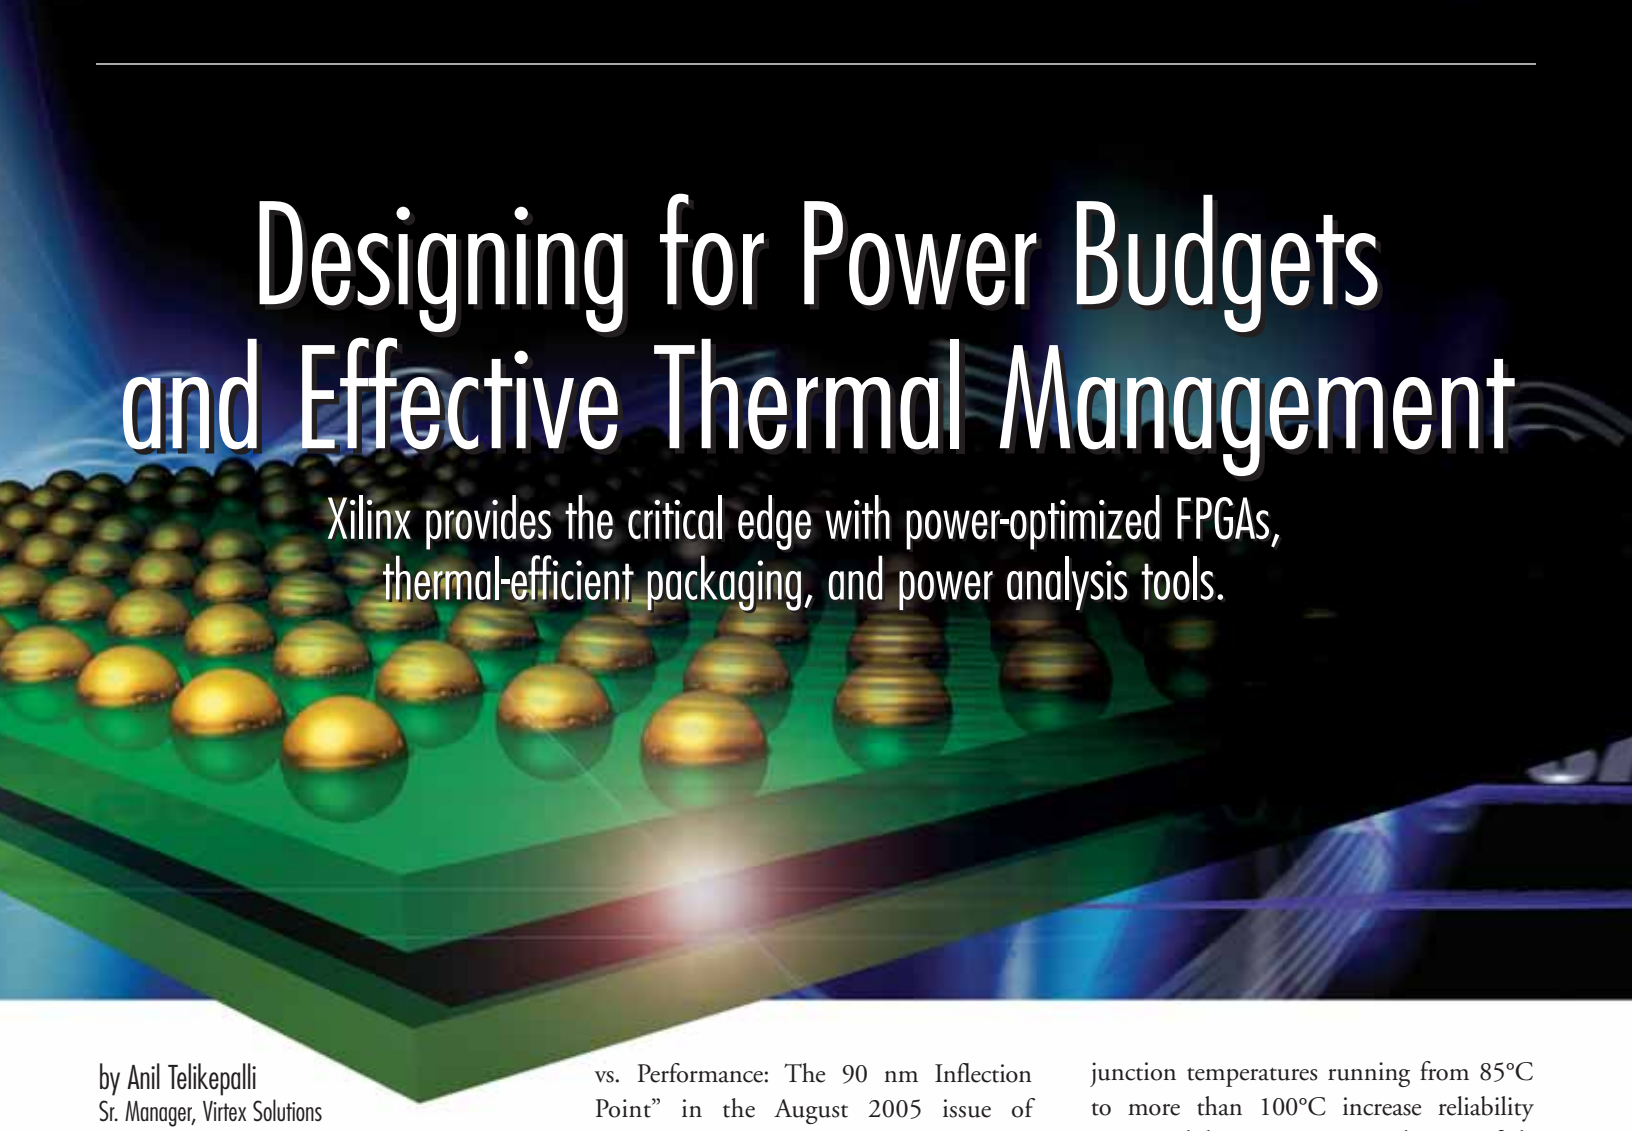

by Anil Telikepalli
Sr. Manager, Virtex Solutions
Xilinx, Inc.
anil.telikepalli@xilinx.com

Meeting power and thermal budgets is one of the essential criteria by which our customers measure the success of their FPGA-based system designs. Constrained both by the density (bandwidth, logic density, and functionality) of the applications they are creating and the environments in which these systems are deployed, FPGA designers need every advantage as they painstakingly balance performance, cost, reliability, and power.

Xilinx demonstrated its commitment to empower its customers' success with the introduction of the Virtex™-4 family of FPGAs. Delivering the lowest total power in any 90 nm high-performance FPGA without compromising performance (1 to 5W lower than any competing 90 nm device), Virtex-4 devices employ a variety of power-saving design techniques to provide the lowest inrush current, static power, and dynamic power. (For more details about the Virtex-4 power advantage, please see "Power vs. Performance: The 90 nm Inflection Point" in the August 2005 issue of the *Power Management Solution Guide*, *www.xilinx.com/solguides*.)

This substantial power advantage provided by Virtex-4 devices translates directly to additional measurable benefits in performance, cost, and reliability. Reducing the FPGA's contribution to total power in a given design alleviates a host of problems incurred by excessive power.

What Cost Power?

Excessive power is expensive in many ways. It creates the need for special system design and operational considerations such as electricity changes and battery back-up costs. Increased power requires more of everything, including more area on the PCB, a larger chassis, more floor space, and larger air-conditioning systems. And because dynamic power is a function of clock frequency, you may be forced to run a design at a lower performance just to keep total power consumption low.

Perhaps the most critical issue is the effect excessive power can have on reliability. Continuously operating systems with junction temperatures running from 85°C to more than 100°C increase reliability issues and decrease mean time between failures (MTBF).

Power and Thermal Budgets

Every system has a power and thermal budget driven by a standard such as NEBS, ETSI, or other requirements (see sidebar, "NEBS"). A look at various types of electronic equipment can provide valuable insight into the importance of power budgets and demonstrate how a seemingly small power savings for each FPGA can translate into enormous value within overall system operation.

Infrastructure Equipment

Metro aggregation systems (such as enterprise Ethernet/layer 2 switches and multi-service provisioning platforms) are typically deployed in temperature-controlled rooms to ensure acceptable transistor junction temperatures. A typical multi-service provisioning platform (MSPP) is populated with multiple line cards supporting a range of data and voice transmission standards. Each line card usually contains numerous ASSPs, ASICs, and FPGAs.

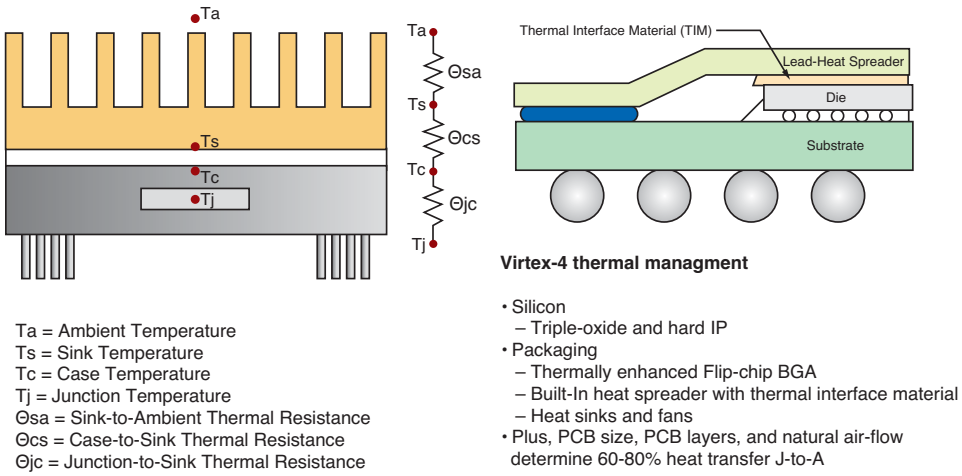

Figure 1 – Heat transfer in devices and how Virtex-4 FPGAs address them

For an FPGA that performs connectivity and traffic management functions, reducing power is critical to system reliability and operational costs. Using a power-optimized FPGA that saves 1-5W per device will dramatically impact the MSPP's total power, potentially allowing the use of smaller power supplies and reducing the size and number of heat sinks.

Now consider wireless base stations that are deployed outdoors in harsh environmental conditions with ambient temperatures exceeding 80°C and transistor junction temperatures in excess of 100°C. A typical wireless network may comprise as many as 35,000 base stations, with each base station utilizing 16 or more line cards.

Wireless base station power budgets are 1.2-2.5 KW; each line card (excluding backplane, fans, and system control) is about 20W. Using an FPGA for connectivity and signal processing, which saves as much as a third of the typical 6W power budget, can in turn save as much as $1 million in operational expenditures.

System Reliability Case Study

As stated previously, many applications are restricted in power by standards. Consider the networking OEMs who sell their equipment to operators and service providers. In the U.S., the operators require OEMs to meet NEBS (National Equipment Building Systems) standards. In particular, the GR-63-CORE and GR-1089-CORE standards specify equipment reliability criteria, on which the power, signal integrity, and thermal management of the equipment (and hence the components used) have a direct impact.

For example, consider an RPR (Resilient Packet Ring) MSPP. Such a system is limited to 4 KW/rack and must remain operational for 99.999% of the time, which equates to no more than five minutes of downtime per year.

Because FPGAs are used extensively in these systems within the various line cards, it is critical to keep the FPGA static and dynamic power consumption to an absolute minimum to help OEMs meet the GR-63-CORE standard.

The typical line card power budgets and FPGA power consumption illustrate the value Virtex-4 FPGAs offer by saving a just couple of watts per FPGA:

- 12-port DS3 card: 30W; FPGA = 4-5W
- 4-port OC-12 card: 28W; FPGA = 4-5W
- 12-port 10/100 Base-T card: 50W; FPGA = 4-5W
- 32-port T1/E1: 9W; FPGA = 2-3W

Thermal Management

Given the increasing criticality with which system vendors are scrutinizing their power and thermal budgets, Xilinx continues to strive to improve upon the edge it already provides with the Virtex-4 family.

Figure 1 shows heat transfer at various locations on a device with definitions of temperature and thermal resistance at each point. In addition to the effective thermal management Virtex-4 FPGAs provide at the silicon level, Xilinx also offers thermally enhanced flip-chip packages with built-in heat spreaders.

The new Xilinx flip-chip BGA packages are the latest package offering for Xilinx high-performance FPGA products. Unlike traditional packaging in which the die is attached to the substrate face-up and the connection is made by using wire, the solder bumped die in the flip-chip BGA is flipped over and placed face down, with the conductive bumps connecting directly to the matching metal pads on the laminate substrate.

The flip-chip package accommodates more I/O pins than traditional wire-bond packages by using the internal chip area for package connections. Furthermore, the cavity-up nature of flip-chip packages allows for superior thermal dissipation through the top of the package.

Virtex-4 pinout provides abundant power and ground pins (Figure 2) – one adjacent to every user pin, providing even distribution for better thermal dissipation when compared to the clustered power and ground pins found in typical flip-chip packages.

To facilitate user design, Xilinx provides

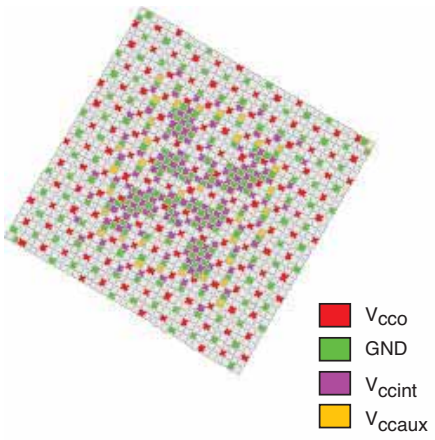

Figure 2 – Virtex-4 packaging, with uniform distribution of power and ground pins, enables better thermal dissipation.

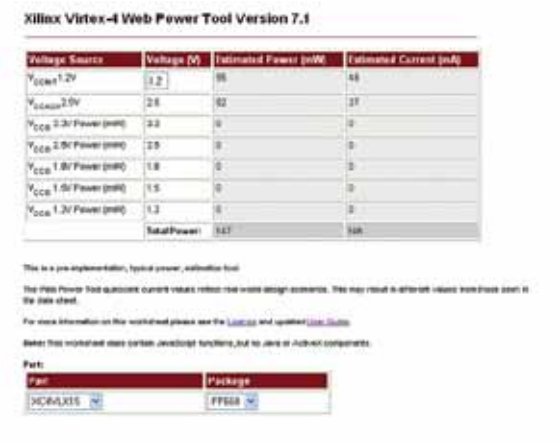

Figure 3a – The Web Power tool allows you to set the operating voltage for the design; when feasible, this enables you to lower the voltage to save power.

Figure 3b – You select the particular utilization of slices, flip-flops, shift register and LUT RAMs, their toggle rates, and routing density to calculate the power used by the configurable logic blocks (CLBs).

you to obtain more precise power-consumption estimates. Of course, the actual power consumption must ultimately be determined in circuit under the appropriate operating conditions.

Web Power

Xilinx Web Power tools help you estimate total power consumption accurately before design implementation. System architects can estimate power using high-level design details and make intelligent design choices on clock frequencies, implementing a function using hard IP or logic, type of I/O standard, and other factors. The Web Power tool relies exclusively on your estimates of FPGA design parameters such as utilization, toggle rates, and operating conditions (Figure 3a/3b).

To calculate a power estimate for a given design, simply enter the estimated design parameters for resource utilization, operating environment, clock rates, and toggle rates, push a button to calculate power, and save the settings. A brief look reveals the intuitively simple nature of this process:

1. Set the operating voltage of the design.
2. Enter the settings for the part and package, typical or worst case process, and thermal information (ambient temperature and airflow).
3. Select the particular utilization of slices, flip-flops, shift register and RAM LUTs, their toggle rates and routing density, to calculate the power used by the configurable logic blocks (CLBs).
4. Specify the number of DSP48s, their architectural usage level, and toggle rates.
5. Select the digital clock manager (DCM) settings.
6. Specify input/output settings.
7. Press the "calculate" button to generate the pre-implementation power estimate with totals at every level.
8. Press the "save settings" button to save the settings for future analysis and modifications.

XPower

Xilinx XPower is a post-implementation power analysis and power estimation software. Included in all configurations of the Xilinx Integrated Software Environment (ISE™ software), XPower enables you to analyze total device power.

JEDEC JESD-51 thermal characteristics (*http://direct.xilinx.com/bvdocs/userguides/ug075.pdf*) for a four-layer board. Although a four-layer board is quite conservative for most applications, it does provide a uniform standard-based comparison. For simulating your unique board setup, Xilinx can provide compact resistor models for the Virtex-4 device package combination.

Power Analysis Tools

Another edge Xilinx offers its customers is the ability to calculate power estimates both before and after design implementation. Xilinx Web Power tools provide pre-implementation estimates of a design's power consumption based on the expected utilization of device resources, operating frequencies, and toggle rates.

Post-implementation, XPower enables

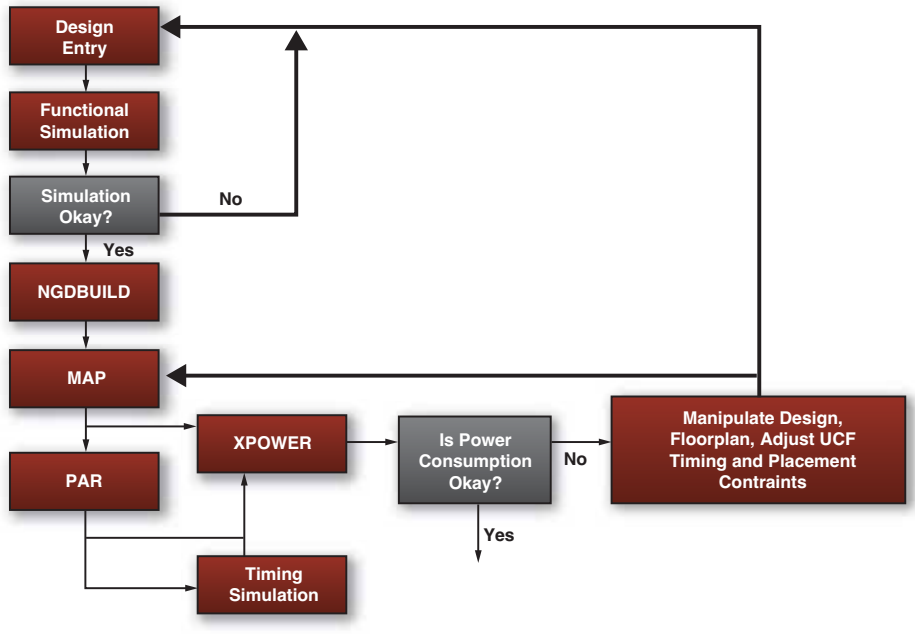

Figure 4 – You can run XPower after map, place and route, or timing simulation.

XPower can read HDL simulation data to estimate toggle rates. XPower uses proprietary algorithms to estimate toggle rates even if the simulation vectors are missing, which is useful for iterative design. Figure 4 shows XPower in the design flow.

XPower includes a design wizard that provides a step-by-step dialog box format to help you import design and simulation data and set loading and default data. The design wizard also helps you get accurate power estimates quickly and easily.

The step-by-step process for using XPower is straightforward and intuitive:

1. Invoke XPower in ISE software.
2. Open design.
3. Load design (NCD).
4. Load optional PCF and simulation data files.
5. Launch the new design wizard.
6. Set voltage, ambient temperature, and airflow.
7. Set frequencies and activity rates.
8. Set capacitive loads for outputs.
9. Set/verify DC loads for outputs.
10. Set/verify enable rates for bi-directional I/Os if simulation data is not loaded.
11. Set global default activity to estimated value.
12. Set estimated activity for any single or groups of signals.

Corroborative Results

Given the value of these power estimation and analysis tools and the essential need for accuracy both before and after the design has been implemented, you might wonder just how much of an advantage these tools provide.

Of course, a great deal hinges on the accuracy of the settings relative to actual design parameters and operating conditions. Nonetheless, Xilinx offers the results of this simple, yet revealing test example as evidence of the veracity of these power tools.

Using a Virtex-4 4VLX60 device, we tested for a correlation of logic, register, and interconnect power across a large design. We ran the design through ISE 7.1i software and compared the power data to estimates developed by the Web Power tool and measurements in a lab. The results in Figure 5 speak for themselves.

Virtex-4 4VLX60 FF1148
Design Parameters
100 MHz, 200 MHz
21,849 LUTs, 25421 FFs; 36% toggle rate

Power Data Comparison – Dynamic			
Frequency (MHz)	Web Power (mW)	XPower (mW)	Measured (mW)
0	0	0	0
100	1,864	1,750	1,742
200	3,727	3,429	3,482

Figure 5 – When compared with actual measured results, the close approximation of the pre- and post-implementation power estimation tools demonstrates the value these tools provide to system architects and FPGA designers.

Conclusion

Designing high-performance FPGA-based systems constrained by exacting power and thermal budgets will never cease to present a daunting challenge to the design engineers and architects who must regularly build them.

As Xilinx customers continue to place their trust in our ability to provide them with the essential edge in power-optimized FPGAs, tools, and package technology, we have and will continue to strive to reduce the challenge of effective power and thermal management. For more information, visit *www.xilinx.com/virtex4/lowpower.*

NEBS

The National Equipment Building Systems (NEBS) standards prevent the target system from interfering with other equipment, while ensuring scalability and safety. The two major standards, GR-63-CORE and GR-1089-CORE, were originally developed by Bellcore/Telcordia in the 1970s and deployed in 1985.

These standards establish physical requirements such as space planning, temperature, and humidity environmental specifications. They also specify acceptable preventative measures to guard against potential hazards such as fire, earthquake, vibration, transportation, acoustics, poor air quality, and illumination.

The standards also specify the electrical requirements that safeguard the equipment against electrostatic discharge (ESD) and electromagnetic impulse (EMI), as well as establishing resilience to lightning and AC power faults.

Other preventative measures established by the standards address issues such as steady state power induction, corrosion, DC potential differential, electrical safety, bonding, and grounding.

With Competing FPGAs, Power Has Become a Burning Issue.

LX SX FX

Design Example:

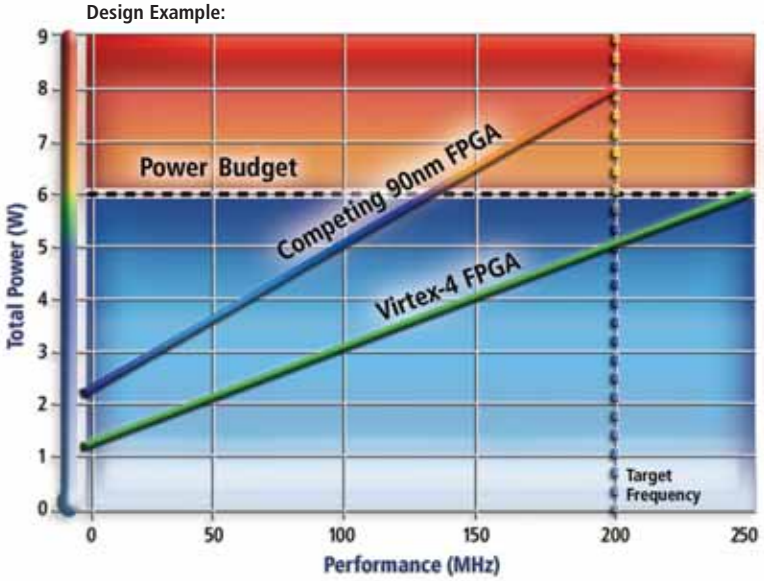

Design Example: LX60 vs. 2S60. Target Frequency = 200 MHz. Worst-case process.
20K LUTs, 20K Flip-Flops, 1Mbit On-Chip RAM, 64 DSP Blocks, 128 2.5V I/Os
Based on Xilinx tool v4.1 and competitor tool v2.1
For higher density devices, achieve up to 5W lower power

VIRTEX V4

Get 1-5W lower power per FPGA, only with the Virtex-4 family

Check the specs for yourself at realistic operating temperatures (T_j = 85°C). Different logic architecture or dielectric just won't do it. No competing FPGA comes close to Virtex-4 for total power savings–take it to the lab and see.

- **73% lower static power**
- **Up to 86% lower dynamic power**

Unique Triple-Oxide Technology & Embedded IP

At the 90nm technology node, power is the next big challenge for system level designers. An inferior device can suffer leakage, dramatic surges in static power, and thermal runaway. That's why we designed our Virtex-4 FPGAs with Triple-Oxide Technology™, embedded IP, and power-saving configuration circuitry. Now you can meet your performance goals, while staying within the power budget.

Visit ***www.xilinx.com/virtex4/lowpower*** today, and get the right solution on board before your power issues start heating up.

FPGA and Structured ASIC JOURNAL Reader's Choice Awards

Highest customer satisfaction:
- Virtex & Spartan FPGAs
- ISE design tools
- Support

XILINX®
The Programmable Logic Company℠

www.xilinx.com/virtex4/lowpower

View The TechOnLine Seminar Today

Breakthrough Performance At The Lowest Cost

Leveraging Programmability in Electronic System-Level Designs

You gain productivity and accuracy with platform-based design and programmable platforms.

by Douglas Densmore
Graduate Student Researcher
University of California, Berkeley
densmore@eecs.berkeley.edu

Alberto Sangiovanni-Vincentelli
Professor
University of California, Berkeley
alberto@eecs.berkeley.edu

Adam Donlin
Research Engineer
Xilinx Research Labs
adam.donlin@xilinx.com

Platform-based design (PBD) is an electronic system-level (ESL) design methodology forged in the inferno of ASIC design complexity: it tackles shrinking design time pressures, the growing complexity of applications, nanometer-era design effects, and heterogeneity between chip and system components. Individually, the challenges are formidable. Combined, they are almost insurmountable: the methods ASIC designers use to address one challenge are often at odds with the methods that solve another. Programmable platforms can utilize PBD and be a welcome new home to refugees fleeing ASIC design complexity.

In this article, we'll show you how to effectively combine the benefits of FPGAs with lessons learned from ASIC design. You will see how the unique properties of programmable platforms leverage innovative techniques to make PBD a powerful FPGA design methodology. We will demonstrate how programmable platforms and PBD allow designers to build complex systems not just faster, but also of a higher quality than non-programmable solutions using traditional design techniques.

Platform-Based Design

At its core, PBD (Figure 1) attempts to:

- Maximize the level of design re-use available to system architects and chip designers
- Enable design verification, analysis, and synthesis at all levels of abstraction
- Provide the basis for design chain integration
- Make architectural exploration, including implementation platform selection, a fast and rigorous process

The basic tenets of the methodology are:

1. Orthogonalization of concerns. Function (what the system does), architecture (how it does it), communication, and computation are modeled independently with a unified mathematical framework.
2. Design formalization. Incorporate a mathematical framework with mapping of functional components to architectural elements. This process can be made (partially) automatic by providing common semantics between two consecutive layers of abstraction, similar to what was achieved when logic synthesis was introduced in the traditional design flow.
3. Successive refinement. The design progresses through a refinement process that links the levels of abstraction defined with the help of the design formalization.

The structure of the platform-based design methodology is illustrated in Figure 1. However, these techniques are severely limited if, during the development of systems with PBD, the results of simulation to predict performance before system creation are inaccurate. Thus, a critical element of PBD is the characterization of architectural components in terms of physical quantities such as time, power, and size. The choice to map functional components to architectural elements is guided by constraints and cost functions expressed in terms of these quantities. To evaluate the quantities for the entire design in terms of the quantities associated with the components, you would use simulation.

Architectural components at high levels of abstraction are characterized in terms of these quantities (time, power, and size), also through estimation or abstraction. If the abstraction or estimation is inaccurate, the results of simulation to predict performance before system creation are inaccurate.

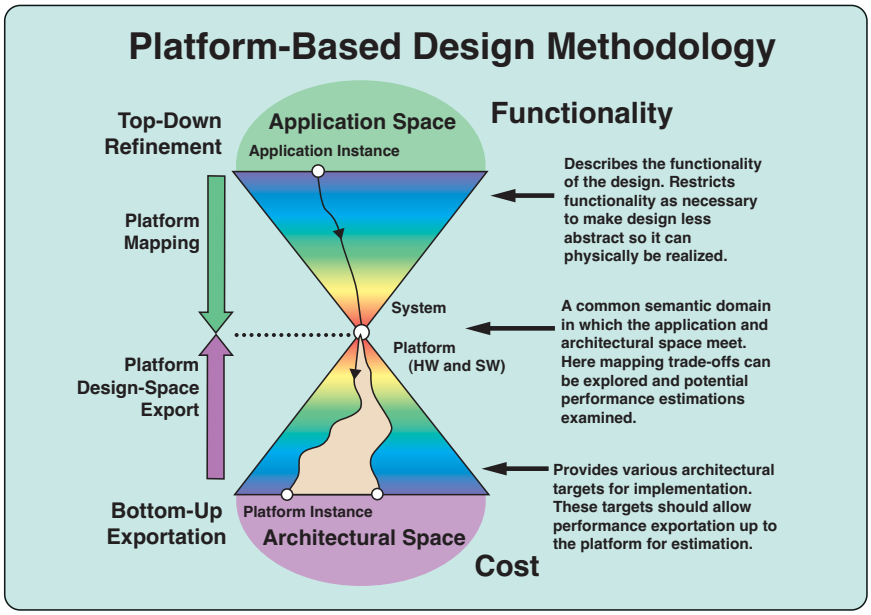

Figure 1 – Platform-based design

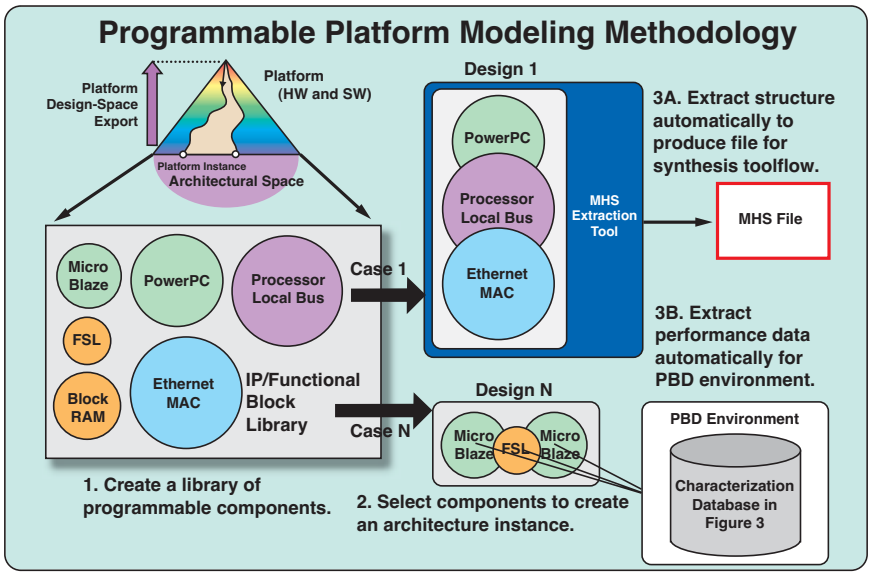

Figure 2 – Programmable platform modeling

Designers require performance models that are abstract and modular while still having a close correlation (accuracy) to actual devices. In addition, these models need to be convenient to use and efficient to simulate. Although this is certainly possible for all design styles and is one of the essential values of the methodology, it is easier to carry out if the implementation platform has fewer degrees of freedom.

Programmable platform FPGAs such as the Xilinx® Virtex™-4 device are excellent substrates for PBD: they are easily reconfigurable using the FPGA fabric and contain powerful high-performance pre-designed processors whose performance can be accurately predicted before simulation.

Programmable Platform Modeling

In our approach, a platform is a collection of IPs that you can use to implement your designs. Each IP has to be given a functional model (what it can do) and a performance model (the cost of doing it). In the Virtex-4 platform, a functional model of an element corresponds to a number of different implementations regarding actual selection of the basic building blocks. Each of these implementations is characterized by a different performance model.

Functional Modeling

The architectural space in Figure 2 contains functional models of FPGA blocks that you can combine to form a programmable system. The key requirement of the models is abstraction to increase designer productivity; however, they must still remain efficient in their ability to yield a good design.

There are three aspects of programmable platforms that facilitate this task as compared to other design styles (such as ASICs):

1. Programmable platforms allow one device to represent many possible architectural topologies. Therefore, one model set for a device allows many potential designs tailored to various levels of concurrency and application specificity.
2. Programmable platforms are naturally partitioned into functional or IP-based blocks. This is the case with the Xilinx implementation of the IBM CoreConnect system in EDK. You can create models at this granularity with a transaction-level set of models.
3. Programmable models are a collection of configured elements. This configuration can be explicitly captured by the models and provided directly to the tool flow. In the case of Xilinx FPGAs, this could be the microprocessor hardware specification (MHS) file.

Figure 2 shows an example of the entire flow. Notice that these three techniques increase efficiency by facilitating library creation, transaction-level modeling, and design synthesis, all while maintaining the required abstraction level.

Performance Models

Accurate programmable platform performance models are easier to obtain than other implementation styles such as ASICs because the physical characteristics of the implementation fabric are known at design time. Figure 3 illustrates a flow based on a pre-characterization process.

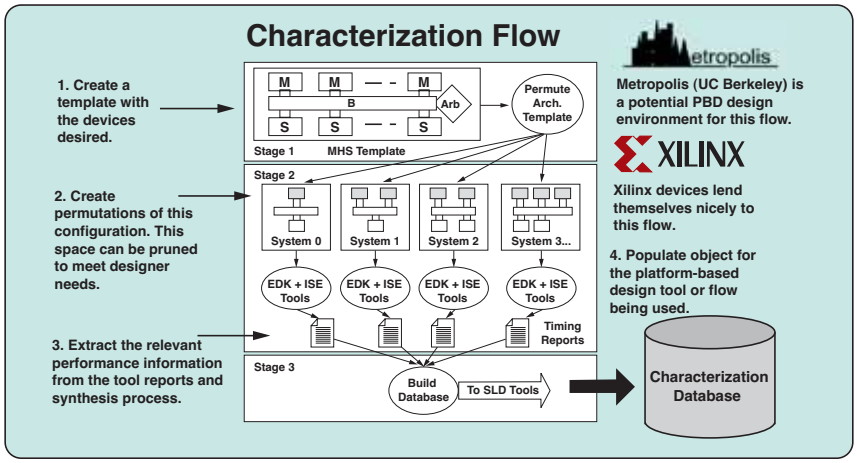

Figure 3 – Programmable platform characterization

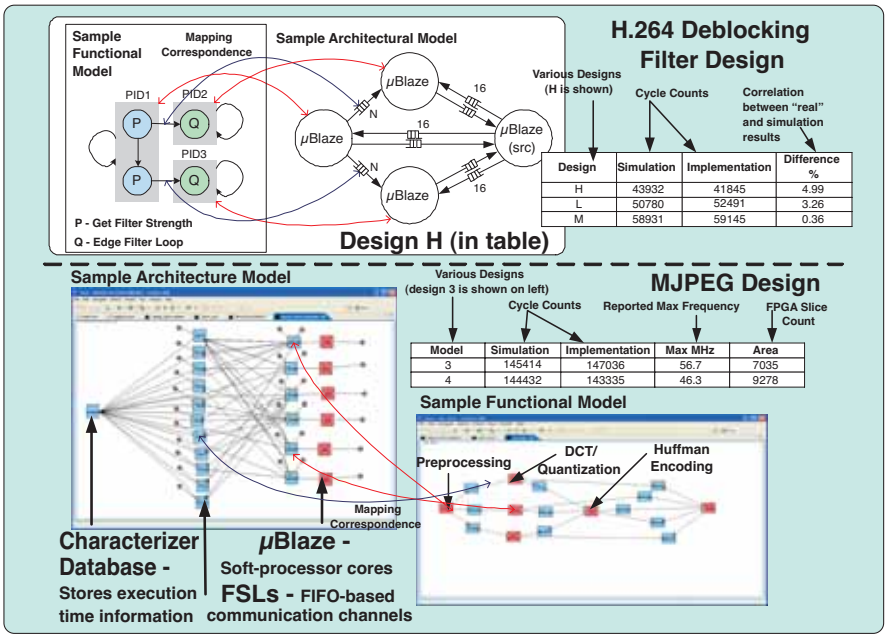

Figure 4 – Example methodology results

- Permutations of various architecture topologies are created in the first stages. These permutations result from an initial "seed" topology tailored to your specific needs.
- Permutation generation is followed by actual synthesis, place and route of these designs for the device being characterized. You use the exact tool flow that the device will eventually use for programming. This is not easily available in an ASIC flow.
- The data from synthesis, place and route is captured in a database structure that can be used by a tool supporting PBD during simulation. Instead of relying on static estimations, simulation can annotate transactions directly with costs that have been realized on an actual system corresponding to the simulation. This database is independent of the actual system, thus yielding a great deal of modularity. You can reuse the database for other designs or add to it on an independent, individual basis. And because the models are created to have a very specific topology, you can use the topology information to index the database during simulation.

Example Designs Using this Approach

We have two example designs using our approach: a motion JPEG encoder and an H.264 video standard deblocking filter. Both of these designs were first specified in the Metropolis Design Environment from the University of California, Berkeley. (Metropolis is a design environment that fully supports the PBD methodology.)

The Metropolis models were of Xilinx Virtex-II Pro CoreConnect components. Specifically, these designs were of MicroBlaze™ and Fast Simplex Link (FSL) networks. Figure 4 shows the designs and a sample of the results we obtained. Key insights include:

- The accuracy of simulation versus actual implementation performance is very high. In the cases shown, the worst case is a mere 5% difference.
- Characterization allowed us to observe the actual system clock speed and design area values. This is important because not only does clock speed combine with the clock cycle requirements to form an actual execution time, but it also demonstrates the potentially detrimental effect of increasing the size of the design. This effect would be more difficult to obtain quickly in ASIC-based design flows.

Conclusion

Platform-based design is an ideal methodology to convert system design ideas into implementation realities, especially when coupled with pre-characterized programmable platforms. High-level abstract and modular models supported by the methodology allow fast, accurate, and efficient design space exploration and early validation.

The link to implementation offered by a pre-design fabric, as offered by the Xilinx Virtex family, yields rapid design turns. You can accurately predict the performance of your designs during the early design phases, further reducing the need to re-design.

We believe that the fast and efficient mapping of systems into programmable platforms offered by PBD can ease the pain that today's designers must endure to carry out their tasks. For more information about PBD and supporting tools, see *http://embedded.eecs.berkeley.edu/metropolis/index.html.*

Successful DDR2 Design

Mentor Graphics highlights design issues and solutions for DDR2, the latest trend in memory design.

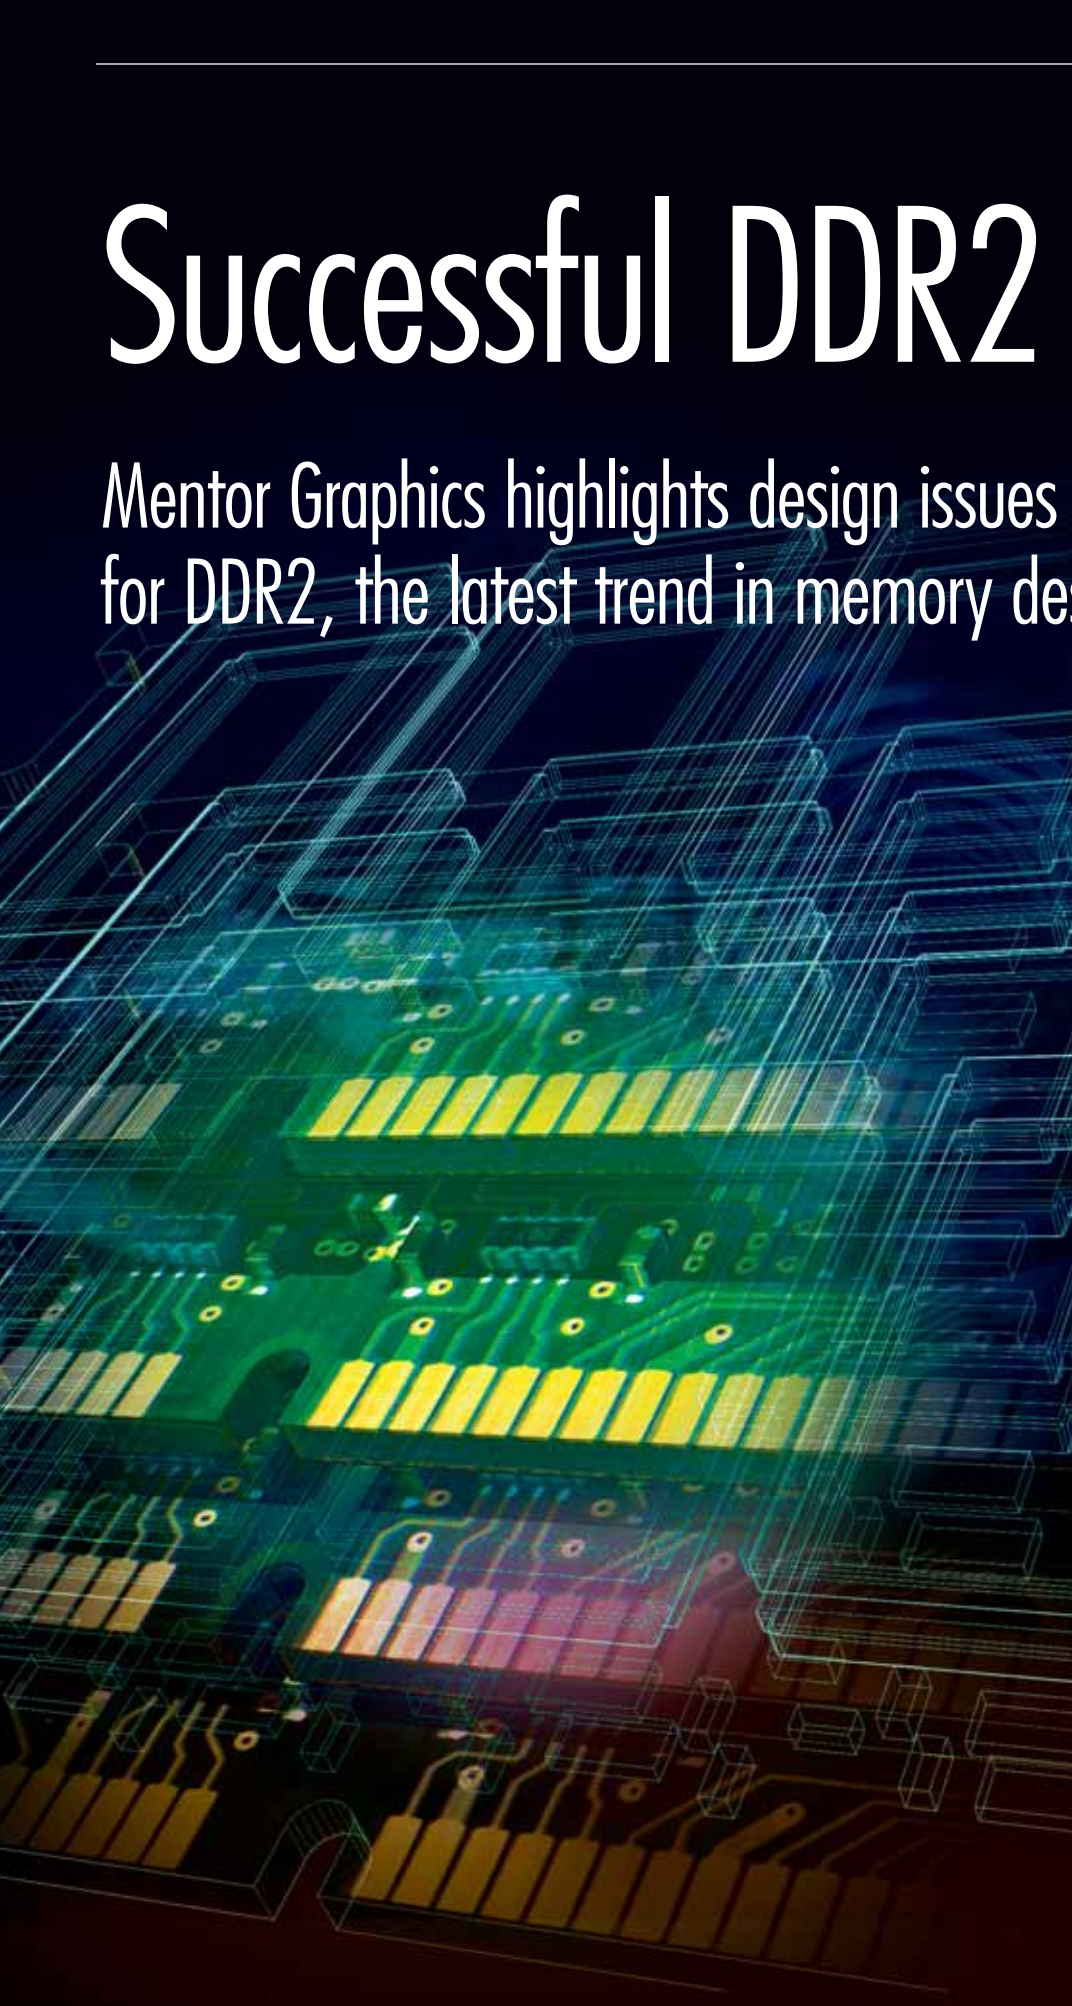

by Steve McKinney
HyperLynx Technical Marketing Engineer
Mentor Graphics
steven_mckinney@mentor.com

The introduction of the first SDRAM interface, in 1997, marked the dawn of the high-speed memory interface age. Since then, designs have migrated through SDR (single data rate), DDR (double data rate), and now DDR2 memory interfaces to sustain increasing bandwidth needs in products such as graphics accelerators and high-speed routers. As a result of its high-bandwidth capabilities, DDR and DDR2 technology is used in nearly every sector of the electronics design industry – from computers and networking to consumer electronics and military applications.

DDR technology introduced the concept of "clocking" data in on both a rising and falling edge of a strobe signal in a memory interface. This provided a 2x bandwidth improvement over an SDR interface with the same clock speed. This, in addition to faster clock frequencies, allowed a single-channel DDR400 interface with a 200 MHz clock to support up to 3.2 GB/s, a 3x improvement over the fastest SDR interface. DDR2 also provided an additional 2x improvement in bandwidth over its DDR predecessor by doubling the maximum clock frequency to 400 MHz. Table 1 shows how the progression from SDR to DDR and DDR2 has allowed today's systems to maintain their upward growth path.

SDR		DDR				DDR2			
PC100	PC133	DDR - 200	DDR - 266	DDR - 333	DDR - 400	DDR2 - 400	DDR2 - 533	DDR2 - 667	DDR2 - 800
0.8	1.1	1.6	2.1	2.7	3.2	3.2	4.266	5.33	6.4
Single Channel Bandwidth (GB/s)									

Table 1 – The progression from SDR to DDR and DDR2 has allowed today's systems to maintain their upward growth path. Speed grades and bit rates are shown for each memory interface.

With any high-speed interface, as supported operating frequencies increase it becomes progressively more difficult to meet signal integrity and timing requirements at the receivers. Clock periods become shorter, reducing timing budgets to a point where you are designing systems with only picoseconds of setup or hold margins. In addition to these tighter timing budgets, signals tend to deteriorate because faster edge rates are needed to meet these tight timing parameters. As edge rates get faster, effects like overshoot, reflections, and crosstalk become more significant problems on the interface, which results in a negative impact on your timing budget. DDR2 is no exception, though the JEDEC standards committee has created several new features to aid in dealing with the adverse effects that reduce system reliability.

Some of the most significant changes incorporated into DDR2 include on-die termination for data nets, differential strobe signals, and signal slew rate derating for both data and address/command signals. Taking full advantage of these new features will help enable you to design a robust memory interface that will meet both your signal integrity and timing goals.

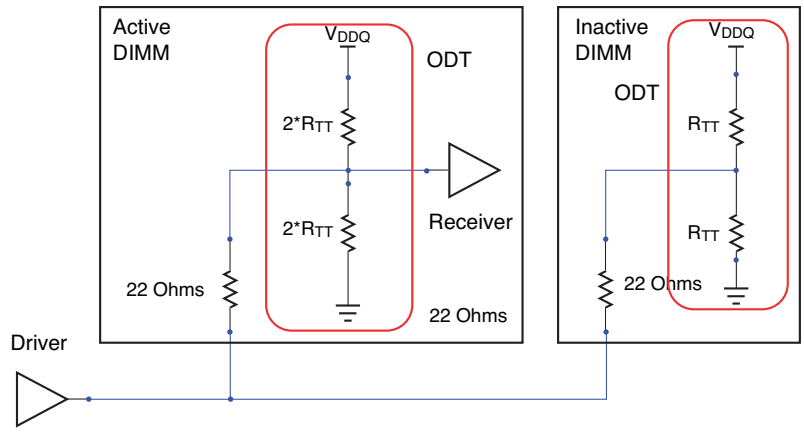

Figure 1 – An example of ODT settings for a write operation in a 2 DIMM module system where R_{TT} = 150 Ohms.

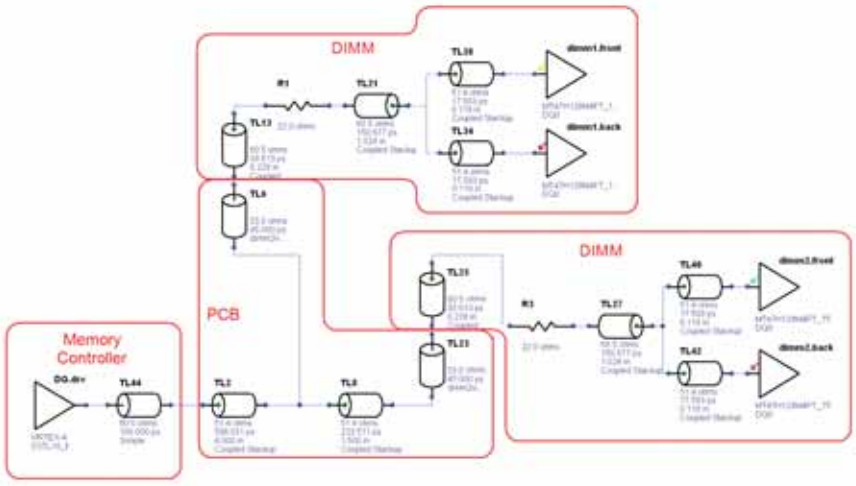

Figure 2 – The HyperLynx free-form schematic editor shows a pre-layout topology of an unbuffered 2 DIMM module system. Transmission line lengths on the DIMM are from the JEDEC DDR2 unbuffered DIMM specification.

On-Die Termination

The addition of on-die termination (ODT) has provided an extra knob with which to dial in and improve signal integrity on the DDR2 interface. ODT is a dynamic termination built into the SDRAM chip and memory controller. It can be enabled or disabled depending on addressing conditions and whether a read or write operation is being performed, as shown in Figure 1. In addition to being able to turn termination off or on, ODT also offers the flexibility of different termination values, allowing you to choose an optimal solution for your specific design.

It is important to investigate the effects of ODT on your received signals, and you can easily do this by using a signal integrity software tool like Mentor Graphics' HyperLynx product. Consider the example design shown in Figure 2, which shows a DDR2-533 interface (266 MHz) with two unbuffered DIMM modules and ODT settings of 150 Ohms at each DIMM. You can simulate the effects of using different ODT settings and determine which settings would work best for this DDR2 design before committing to a specific board layout or creating a prototype.

With the 150 Ohm ODT settings, Figure 3 shows significant signal degrada-

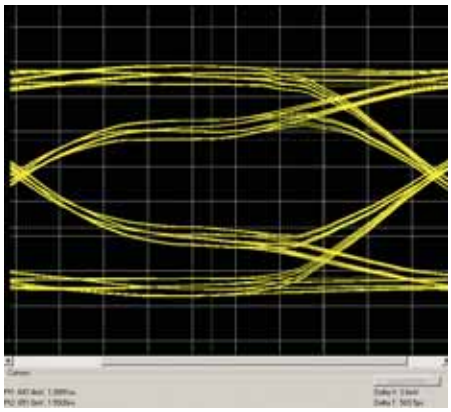

Figure 3 – The results of a received signal at the first DIMM in eye diagram form. Here, ODT settings of 150 Ohms are being used at both DIMM modules during a write operation. The results show there is an eye opening of approximately 450 ps outside of the VinAC switching thresholds.

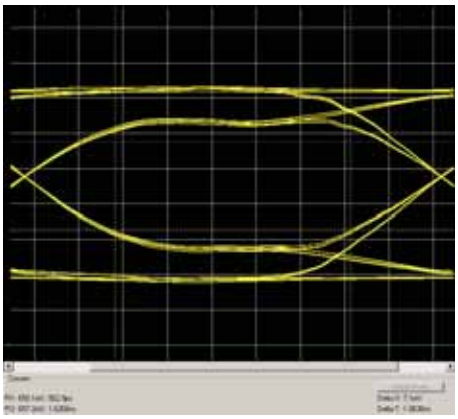

Figure 4 – This waveform shows a significant improvement in the eye aperture with a new ODT setting. Here, the ODT setting is 150 Ohms at the first DIMM and 75 Ohms at the second DIMM. The signal is valid for 1.064 ns with the new settings, which is an increase of 614 ps from the previous ODT settings.

tion at the receiver, resulting in eye closure. The eye shows what the signal looks like for all bit transitions of a pseudo-random (PRBS) bitstream, which resembles the data that you might see in a DDR2 write transaction. Making some simple measurements of the eye where it is valid outside the VinhAC and VinlAC thresholds, you can see that there is roughly a 450 ps window of valid signal at the first DIMM module.

It is appropriate to try to improve this eye aperture (opening) at the first DIMM if possible, and changing the ODT setting is one of the options available for this. To improve the signal quality at the first DIMM, you must change the ODT value at the second DIMM. Setting the ODT at the second DIMM to 75 Ohms and re-running the simulation, Figure 4 shows more than a 100 percent increase in the eye aperture at the first DIMM, resulting in a 1.06 ns eye opening. As you can see, being able to dynamically change ODT is a powerful capability to improve signal quality on the DDR2 interface.

With respect to a DDR interface, ODT allows you to remove the source termination, normally placed at the memory controller, from the board. In addition, the pull-up termination to VTT at the end of the data bus is no longer necessary. This reduces component cost and significantly improves the layout of the board. By removing these terminations, you may be able to reduce layer count and remove unwanted vias on the signals used for layer transitions at the terminations.

Signal Slew Rate Derating

A challenging aspect of any DDR2 design is meeting the setup and hold time requirements of the receivers. This is especially true for the address bus, which tends to have significantly heavier loading conditions than the data bus, resulting in fairly slow edge rates. These slower edge rates can consume a fairly large portion of your timing budget, preventing you from meeting your setup and hold time requirements.

To enable you to meet the setup and hold requirements on address and data buses, DDR2's developers implemented a fairly advanced and relatively new timing concept to improve timing on the interface: "signal slew rate derating." Slew rate derating provides you with a more accurate picture of system-level timing on the DDR2 interface by taking into account the basic physics of the transistors at the receiver.

For DDR2, when any memory vendor defines the setup and hold times for their component, they use an input signal that has a 1.0V/ns input slew rate. What if the signals in your design have faster or slower slew rates than 1.0V/ns? Does it make sense to still meet that same setup and hold requirement defined at 1.0V/ns? Not really. This disparity drove the need for slew rate derating on the signals specific to your design.

To clearly understand slew rate derating, let's consider how a transistor works. It takes a certain amount of charge to build up at the gate of the transistor before it switches high or low. Consider the 1.0V/ns slew rate input waveform between the switching region, Vref to Vin(h/l)AC, used to define the setup and hold times. You can define a charge area under this 1.0V/ns curve that would be equivalent to the charge it takes to cause the transistor to switch. If you have a signal that has a slew rate faster than 1.0V/ns, say 2.0V/ns, it transitions through the switching region much faster and effectively improves your timing margin. You've added some amount of timing margin into your system, but that was with the assumption of using the stan-

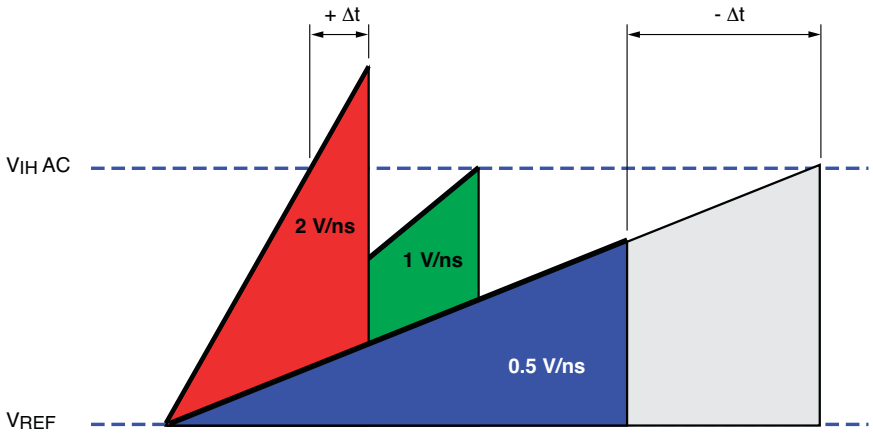

Figure 5 – A 1V/ns signal has a defined charge area under the signal between Vref and VinhAC. A 2V/ns signal would require a + Δt change in time to achieve the same charge area as the 1V/ns signal. A 0.5V/ns signal would require a - Δt change in time to achieve the same charge area as the 1V/ns signal. This change in time provides a clearer picture of the timing requirements needed for the receiver to switch.

dard setup and hold times defined at 1.0V/ns. In reality, you haven't allowed enough time for the transistor to reach the charge potential necessary to switch, so there is some uncertainty that is not being accounted for in your system timing budget. To guarantee that your receiver has enough charge built up to switch, you have to allow more time to pass so that sufficient charge can accumulate at the gate.

Once the signal has reached a charge area equivalent to the 1.0V/ns curve between the switching regions, you can safely say that you have a valid received signal. You must now look at the time difference between reaching the VinAC switching threshold and the amount of time it took for the 2.0V/ns to reach an equivalent charge area, and then add that time difference into your timing budget, as shown in Figure 5.

Conversely, if you consider a much slower slew rate, such as 0.1V/ns, it would take a very long time to reach the switching threshold. You may never meet the setup and hold requirements in your timing budget with that slow of a slew rate through the transition region. This could cause you to overly constrain the design of your system, or potentially limit the configuration and operating speed that you can reliably support. But again, if you consider the charge potential at the gate with this slow slew rate, you would be able to subtract some time out of your budget (as much as 1.42 ns under certain conditions) because the signal reached an equivalent charge area earlier than when it crossed the VinAC threshold.

To assist you in meeting these timing goals, the memory vendors took this slew rate information into account and have constructed a derating table included in the DDR2 JEDEC specification (JESD79-2B on *www.jedec.com*). By using signal derating, you are now considering how the transistors at the receiver respond to charge building at their gates in your timing budgets. Although this adds a level of complexity to your analysis, it gives you more flexibility in meeting your timing goals, while also providing you with higher visibility into the actual timing of your system.

Determining Slew Rate

To properly use the derating tables, it is important to know how to measure the slew rate on a signal. Let's look at an example of a slew rate measurement for the rising edge of a signal under a setup condition.

The first step in performing signal derating is to find a nominal slew rate of the signal in the transition region between the Vref and Vin(h/l)AC threshold. That nominal slew rate line is defined in the JEDEC specification as the points of the received waveform and Vref and VinhAC for a rising edge, as shown in Figure 6.

It would be a daunting task to manually measure each one of your signal edges to determine a nominal slew rate for use in the derating tables toward derating each signal. To assist with this process, HyperLynx simulation software includes built-in measurement capabilities designed specifically for DDR2 slew rate measurements. This can reduce your development cycle and take the guesswork out of trying to perform signal derating. The HyperLynx oscilloscope will automatically measure each of the edge transitions on the received waveform, reporting back the minimum and maximum slew rate values, which can then be used in the JEDEC derating tables. The scope also displays the nominal slew rate for each edge transition, providing confidence that the correct measurements are being made (see Figure 7).

The nominal slew rate is acceptable for use in the derating tables as long as the

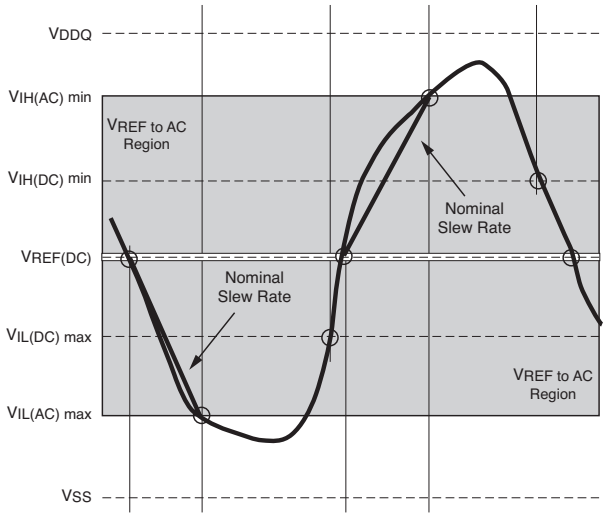

Figure 6 – The waveform illustrates how a nominal slew rate is defined for a signal when performing a derating in a setup condition. The waveform is taken from the DDR2 JEDEC specification (JESD79-2B).

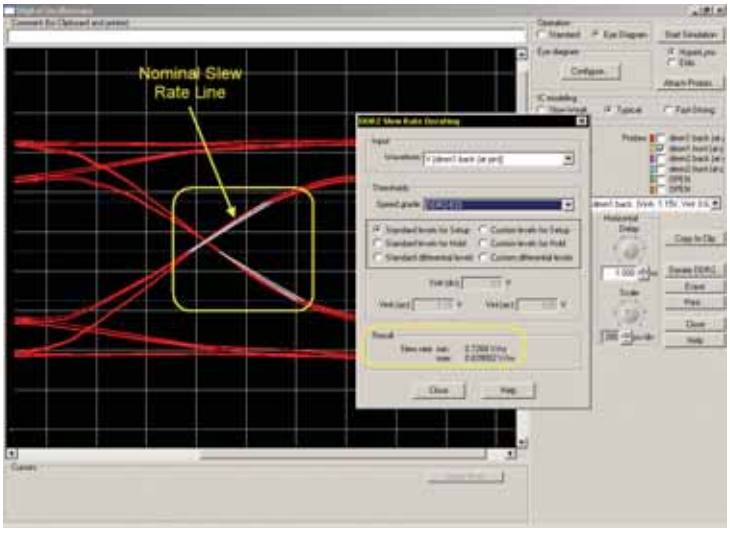

Figure 7 – The HyperLynx oscilloscope shows an automated measurement of the nominal slew rate for every edge in an eye diagram with the DDR2 slew rate derating feature. The measurement provides the minimum and maximum slew rates that can then be used in the DDR2 derating tables in the JEDEC specification.

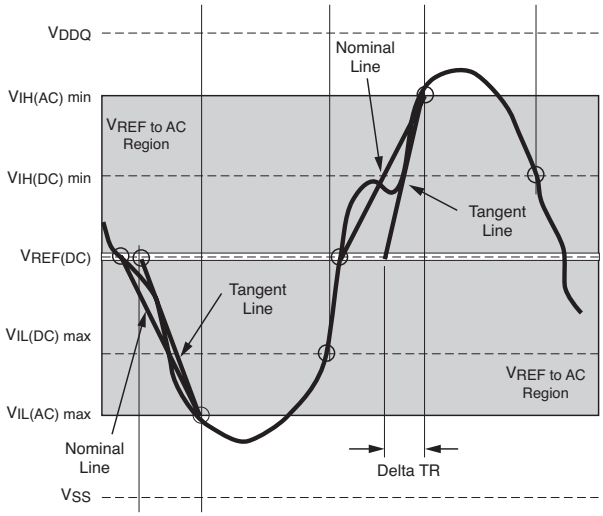

Figure 8 – This waveform, taken from the DDR2 JEDEC specification, shows how a tangent line must be found if any of the signal crosses the nominal slew rate line. The slew rate of this tangent line would then be used in the DDR2 derating tables.

received signal meets the condition of always being above (for the rising edge) or below (for the falling edge) the nominal slew rate line for a setup condition. If the signal does not have clean edges – possibly having some non-monotonicity or "shelf"-type effect that crosses the nominal slew rate line – you must define a new slew rate. This new slew rate is a tangent line on the received waveform that intersects with VinhAC and the received waveform, as shown in Figure 8. The slew rate of this new tangent line now becomes your slew rate for signal derating.

You can see in the example that if there is an aberration on the signal edge that would require you to find this new tangent line slew rate, HyperLynx automatically performs this check for you. If necessary, the oscilloscope creates the tangent line, which becomes part of the minimum and maximum slew rate results. As Figure 9 shows, the HyperLynx oscilloscope also displays all of the tangent lines, making it easier to identify whether this condition is occurring.

For a hold condition, you perform a slightly different measurement for the slew rate. Instead of measuring from Vref to the VinAC threshold, you measure from VinDC to Vref to determine the nominal slew rate (shown in Figure 10). The same conditions regarding the nominal slew rate line and the inspection of the signal to determine the necessity for a tangent line for a new slew rate hold true here as well.

Conclusion

With the new addition of ODT, you've seen how dynamic on-chip termination can vastly improve signal quality. Performing signal derating per the DDR2 SDRAM specification has also shown that you can add as much as 1.42 ns back into your timing budget, giving you more flexibility in your PCB design and providing you with a better understanding of system timing.

Equipped with the right tools and an understanding of underlying technology, you will be able to move your designs from DDR to DDR2 in a reasonably pain-free process – realizing the added performance benefits and component-count reductions promised by DDR2.

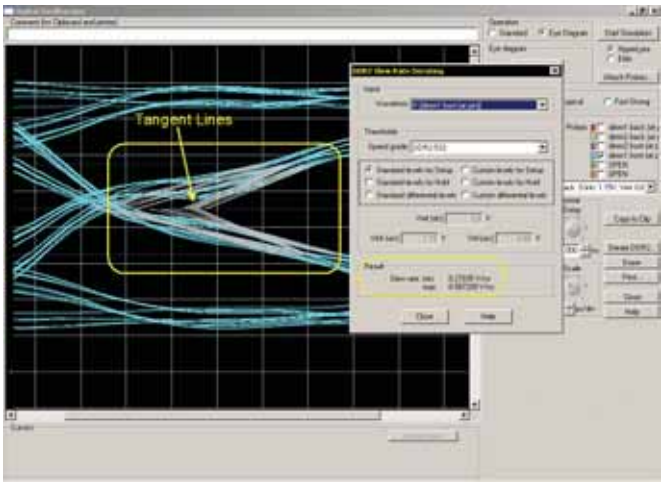

Figure 9 – The HyperLynx oscilloscope shows how the tangent line is automatically determined for you in the DDR2 slew rate derating feature. The slew rate lines in the display indicate that they are tangent lines because they no longer intersect with the received signal and Vref intersection. The oscilloscope determines the slew rate of these new tangent lines for you and reports the minimum and maximum slew rates to be used in the derating tables.

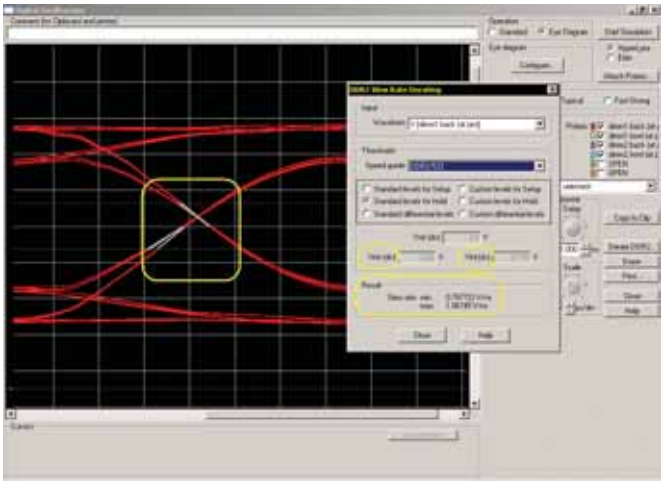

Figure 10 – The oscilloscope shows how a derating for a hold condition is being performed on the received signal. The DC thresholds are used in place of the AC switching thresholds, which are noted in the DDR2 derating dialog.

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

FREE on-line training *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®

www.xilinx.com/dod

Pb-free devices available now

PERFORMANCE

Tired of spinning in circles to meet timing on your FPGA designs?

Synplicity's **Synplify® Premier** solution solves FPGA timing closure through a unique graph-based physical synthesis technology providing highly accurate correlation to final timing and 5 to 20% better performance in a single-pass flow.

FPGA Solution

For more information on the Synplify Premier product and how it can help you quickly reach aggressive timing goals for your FPGAs, visit:
www.synplicity.com/products

Synplicity®

Simply Better Results

Implementing High-Performance Memory Interfaces with Virtex-4 FPGAs

You can center-align clock-to-read data at "run time" with ChipSync technology.

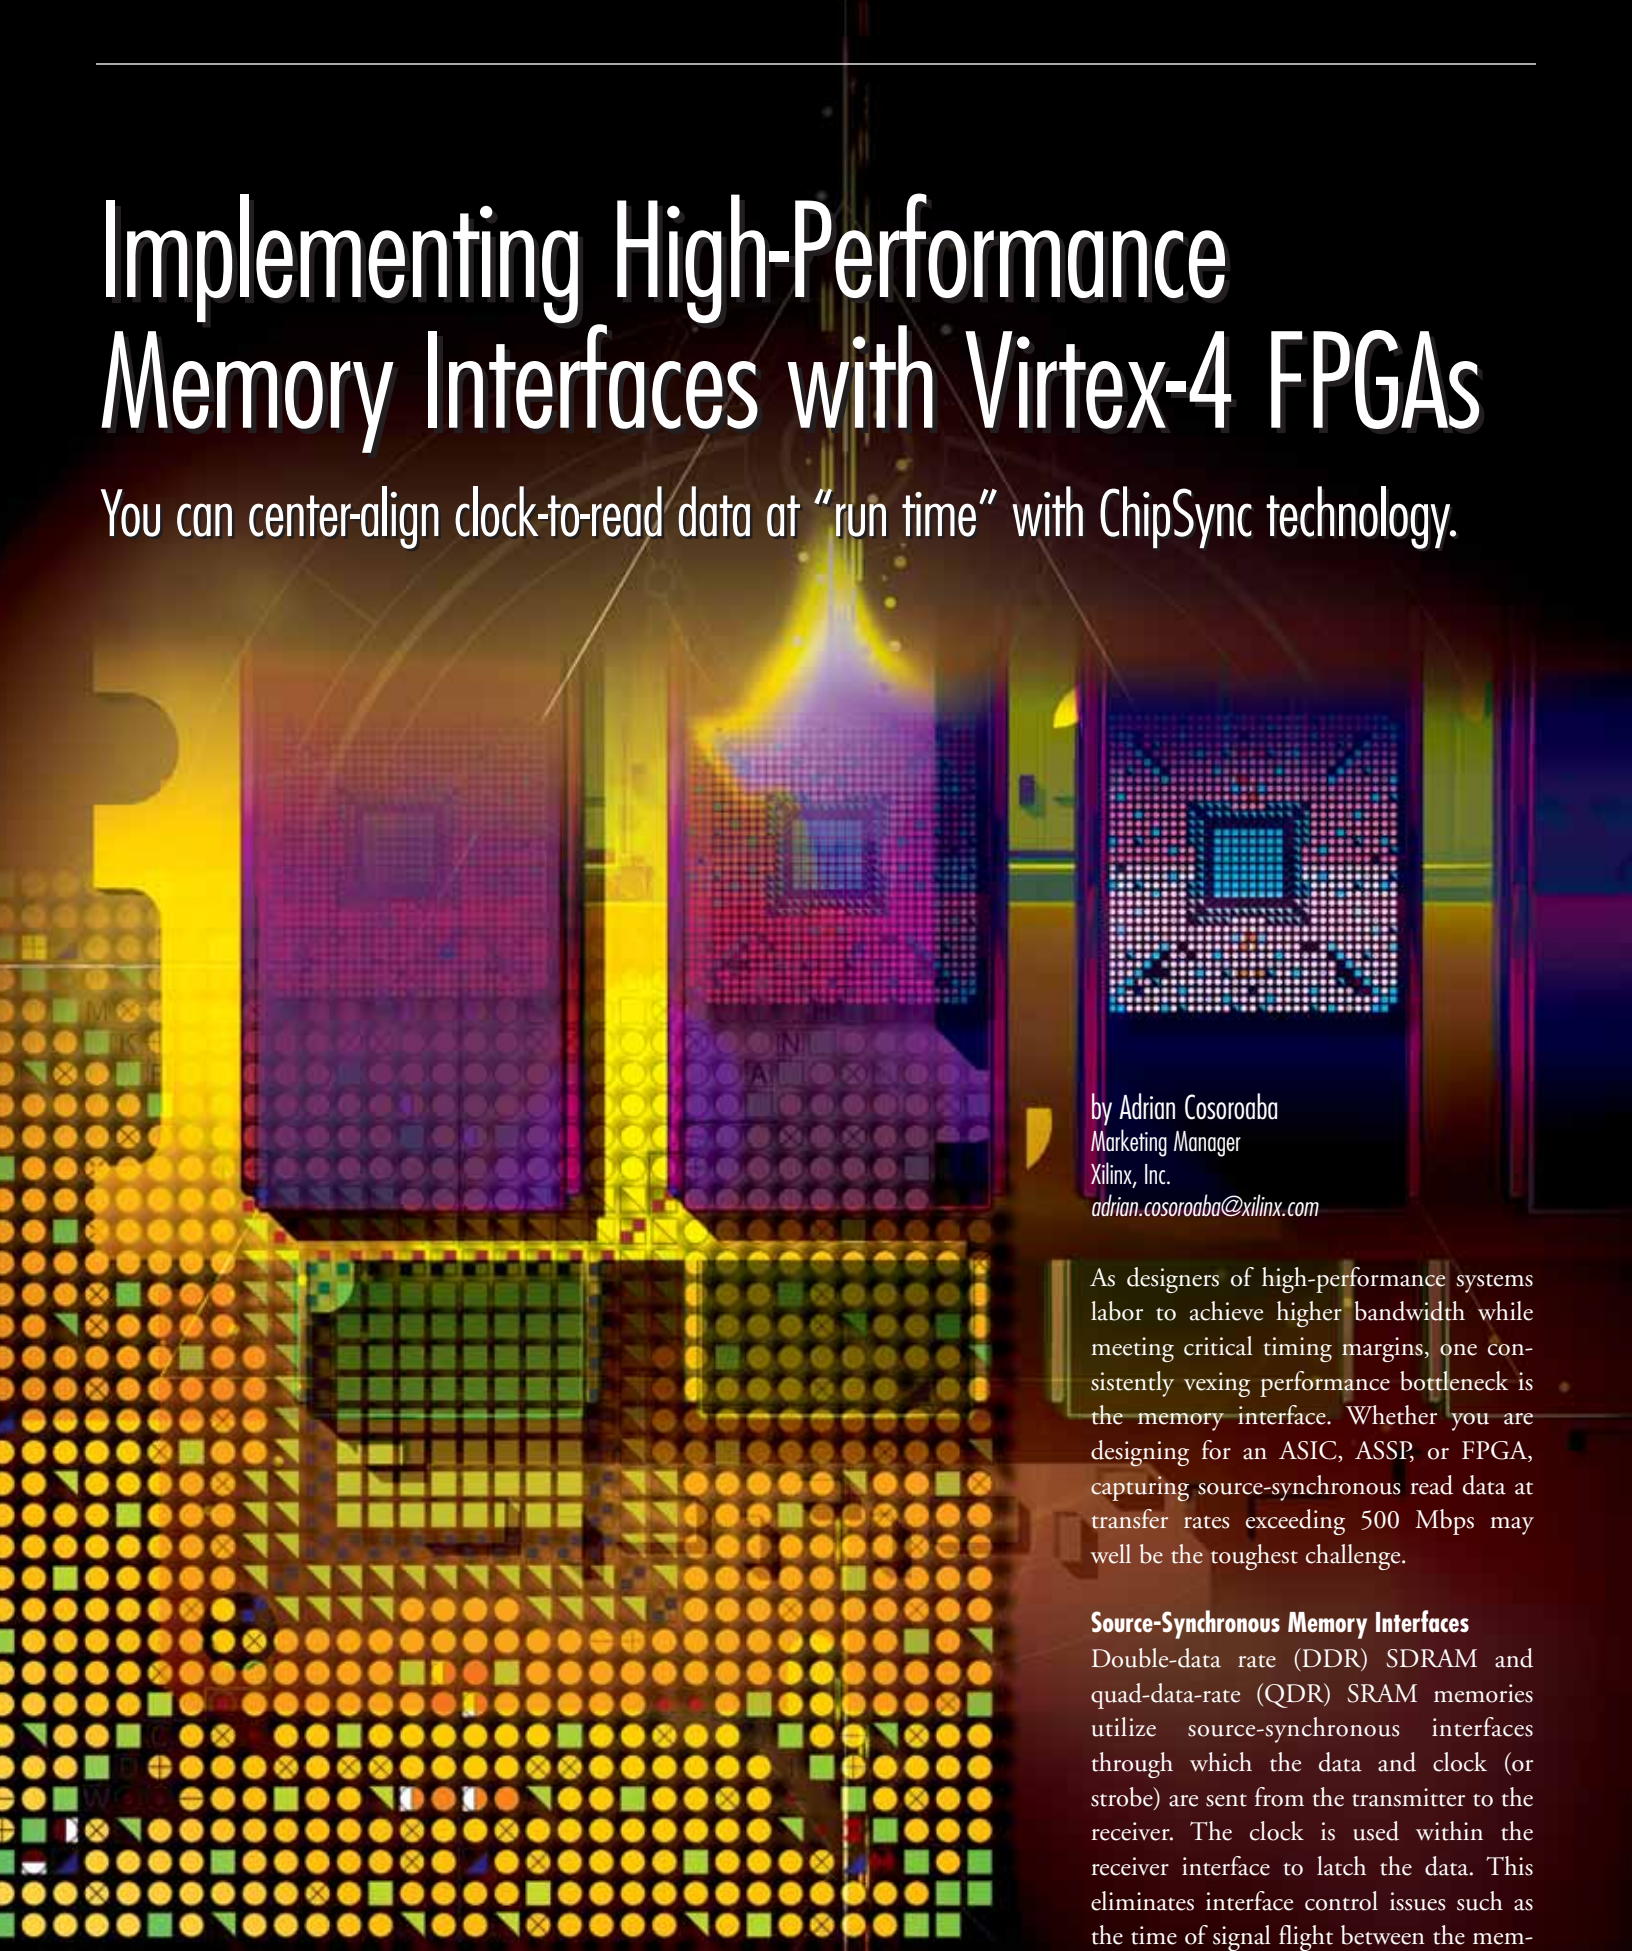

by Adrian Cosoroaba
Marketing Manager
Xilinx, Inc.
adrian.cosoroaba@xilinx.com

As designers of high-performance systems labor to achieve higher bandwidth while meeting critical timing margins, one consistently vexing performance bottleneck is the memory interface. Whether you are designing for an ASIC, ASSP, or FPGA, capturing source-synchronous read data at transfer rates exceeding 500 Mbps may well be the toughest challenge.

Source-Synchronous Memory Interfaces

Double-data rate (DDR) SDRAM and quad-data-rate (QDR) SRAM memories utilize source-synchronous interfaces through which the data and clock (or strobe) are sent from the transmitter to the receiver. The clock is used within the receiver interface to latch the data. This eliminates interface control issues such as the time of signal flight between the memory and the FPGA, but raises new challenges that you must address.

One of these issues is how to meet the various read data capture requirements to implement a high-speed source-synchronous interface. For instance, the receiver must ensure that the clock or strobe is routed to all data loads while meeting the required input setup and hold timing. But source-synchronous devices often limit the loading of the forwarded clock. Also, as the data-valid window becomes smaller at higher frequencies, it becomes more important (and simultaneously more challenging) to align the received clock with the center of the data.

Traditional Read Data Capture Method

Source-synchronous clocking requirements are typically more difficult to meet when reading from memory compared with writing to memory. This is because the DDR and DDR2 SDRAM devices send the data edge aligned with a non-continuous strobe signal instead of a continuous clock. For low-frequency interfaces up to 100 MHz, DCM phase-shifted outputs can be used to capture read data.

Capturing read data becomes more challenging at higher frequencies. Read data can be captured into configurable logic blocks (CLBs) using the memory read strobe, but the strobe must first be delayed so that its edge coincides with the center of the data valid window. Finding the correct phase-shift value is further complicated by process, voltage, and temperature (PVT) variations. The delayed strobe must also be routed onto low-skew FPGA clock resources to maintain the accuracy of the delay.

The traditional method used by FPGA, ASIC, and ASSP controller-based designs employs a phase-locked loop (PLL) or delay-locked loop (DLL) circuit that guarantees a fixed phase shift or delay between the source clock and the clock used for capturing data (Figure 1). You can insert this phase shift to accommodate estimated process, voltage, and temperature variations. The obvious drawback with this method is that it fixes the delay to a single value predetermined during the design phase. Thus, hard-to-predict variations within the system itself – caused by different routing to different memory devices, variations between FPGA or ASIC devices, and ambient system conditions (voltage, temperature) – can easily create skew whereby the predetermined phase shift is ineffectual.

These techniques have allowed FPGA designers to implement DDR SDRAM memory interfaces. But very high-speed 267 MHz DDR2 SDRAM and 300 MHz QDR II SRAM interfaces demand much tighter control over the clock or strobe delay.

System timing issues associated with setup (leading edge) and hold (trailing edge) uncertainties further minimize the valid window available for reliable read data capture. For example, 267 MHz (533 Mbps) DDR2 read interface timings require FPGA clock alignment within a .33 ns window.

Other issues also demand your attention, including chip-to-chip signal integrity, simultaneous switching constraints, and board layout constraints. Pulse-width distortion and jitter on clock or data strobe signals also cause data and address timing problems at the input to the RAM and the FPGA's I/O blocks (IOB) flip-flop. Furthermore, as a bidirectional and non-free-running signal, the data strobe has an increased jitter component, unlike the clock signal.

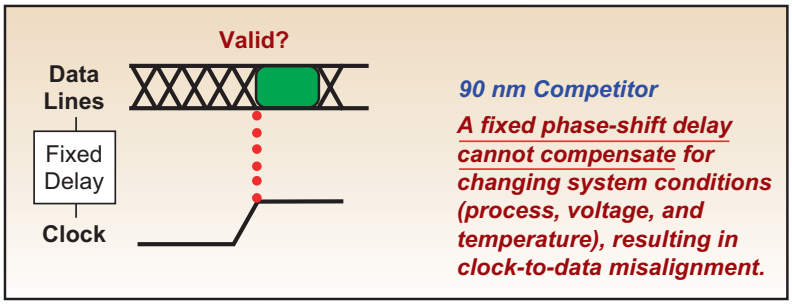

Figure 1 – Traditional fixed-delay read data capture method

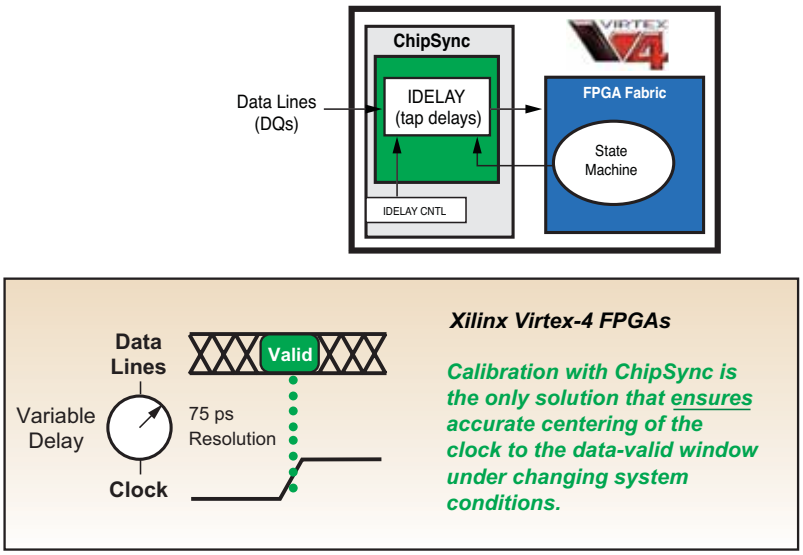

Figure 2 – Clock-to-data centering using ChipSync tap delays

Clock-to-Data Centering Built into Every I/O

Xilinx® Virtex™-4 FPGAs with dedicated delay and clocking resources in the I/O blocks – called ChipSync™ technology – answer these challenges. These devices make memory interface design significantly easier and free up the FPGA fabric for other purposes. Moreover, Xilinx offers a reference design for memory interface solutions that center-aligns the clock to the read data at "run time" upon system initialization. This proven methodology ensures optimum performance, reduces engineering costs, and increases design reliability.

ChipSync features are built into every I/O. This capability provides additional flexibility if you are looking to alleviate board layout constraints and improve signal integrity.

ChipSync technology enables clock-to-data centering without consuming CLB resources. Designers can use the memory read strobe purely to determine the phase relationship between the FPGA's own DCM clock output and the read data. The read data is then delayed to center-align the FPGA clock in the read data window for data capture. In the Virtex-4 FPGA architecture, the ChipSync I/O block includes a precision delay block known as IDELAY that can be used to generate the tap delays necessary to align the FPGA clock to the center of the read data (Figure 2).

Memory read strobe edge-detection logic uses this precision delay to detect the edges of the memory read strobe from which the pulse center can be calculated in terms of the number of delay taps counted between the first and second edges. Delaying the data by this number of taps aligns the center of the data window with the edge of the FPGA DCM output. The tap delays generated by this precision delay block allow alignment of the data and clock to within 75 ps resolution.

The first step in this technique is to determine the phase relationship between the FPGA clock and the read data received at the FPGA. This is done using the memory read strobe. Based on this phase relationship, the next step is to delay read data to center it with respect to the FPGA clock. The delayed read data is then captured directly in input DDR flip-flops in the FPGA clock domain.

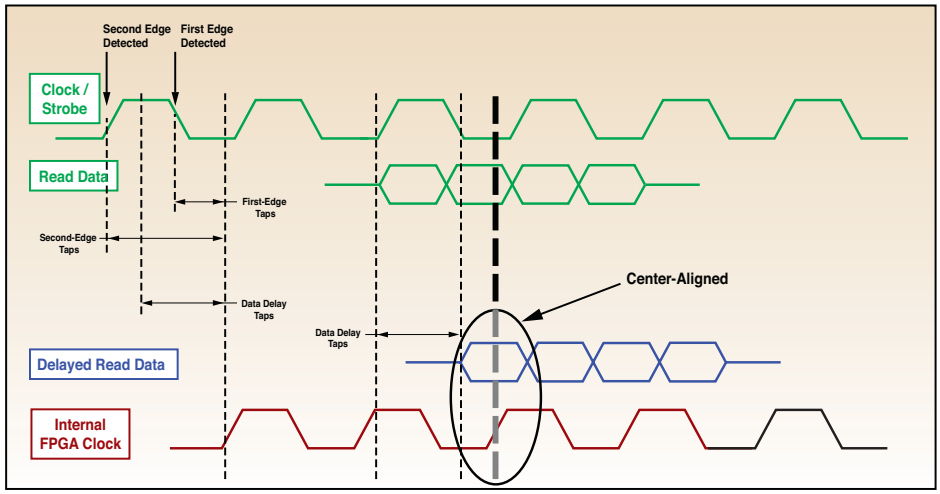

Figure 3 – Clock-to-data centering at "run time"

The phase detection is performed at run time by issuing dummy read commands after memory initialization. This is done to receive an uninterrupted strobe from the memory (Figure 3).

The goal is to detect two edges or transitions of the memory read strobe in the FPGA clock domain. To do this, you must input the strobe to the 64-tap IDELAY block that has a resolution of 75 ps. Then, starting at the 0-tap setting, IDELAY is incremented one tap at a time until it detects the first transition in the FPGA clock domain. After recording the number of taps it took to detect the first edge (first-edge taps), the state machine logic continues incrementing the taps one tap at a time until it detects the second transition (second-edge taps) in the FPGA clock domain.

Having determined the values for first-edge taps and second-edge taps, the state machine logic can compute the required data delay. The pulse center is computed with these recorded values as (second-edge taps – first-edge taps)/2. The required data delay is the sum of the first-edge taps and the pulse center. Using this delay value, the data-valid window is centered with respect to the FPGA clock.

ChipSync features are built into every I/O. This capability provides additional flexibility if you are looking to alleviate board layout constraints and improve signal integrity.

Each I/O also has input DDR flip-flops required for read data capture either in the delayed memory read strobe domain or in the system (FPGA) clock domain. With these modes you can achieve higher design performance by avoiding half-clock-cycle data paths in the FPGA fabric.

Instead of capturing the data into a CLB-configured FIFO, the architecture provides dedicated 500 MHz block RAM with built-in FIFO functionality. These enable a reduction in design size, while leaving the CLB resources free for other functions.

Clock-to-Data Phase Alignment for Writes

Although the read operations are the most challenging part of memory interface design, the same level of precision is required in write interface implementation. During a write to the external memory device, the clock/strobe must be transmitted center-aligned with respect to data. In the Virtex-4 FPGA I/O, the clock/strobe is generated using the output DDR registers clocked by a DCM clock output (CLK0) on the global clock network. The write data is transmitted using the output DDR registers clocked by a

DCM clock output that is phase-offset 90 degrees (CLK270) with respect to the clock used to generate clock/strobe. This phase shift meets the memory vendor specification of centering the clock/strobe in the data window.

Another innovative feature of the output DDR registers is the SAME_EDGE mode of operation. In this mode, a third register clocked by a rising edge is placed on the input of the falling-edge register. Using this mode, both rising-edge and falling-edge data can be presented to the output DDR registers on the same clock edge (CLK270), thereby allowing higher DDR performance with minimal register-to-register delay.

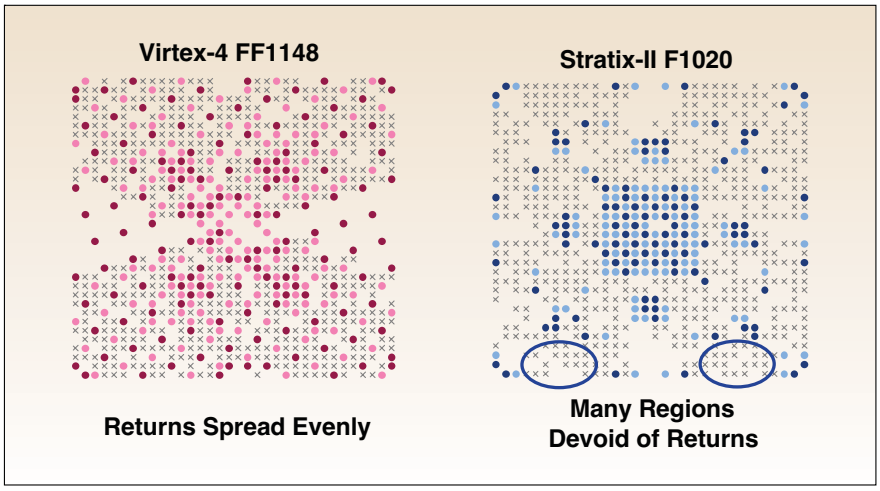

Figure 4 – Pin-out comparison between Virtex-4 and Stratix-II FPGAs

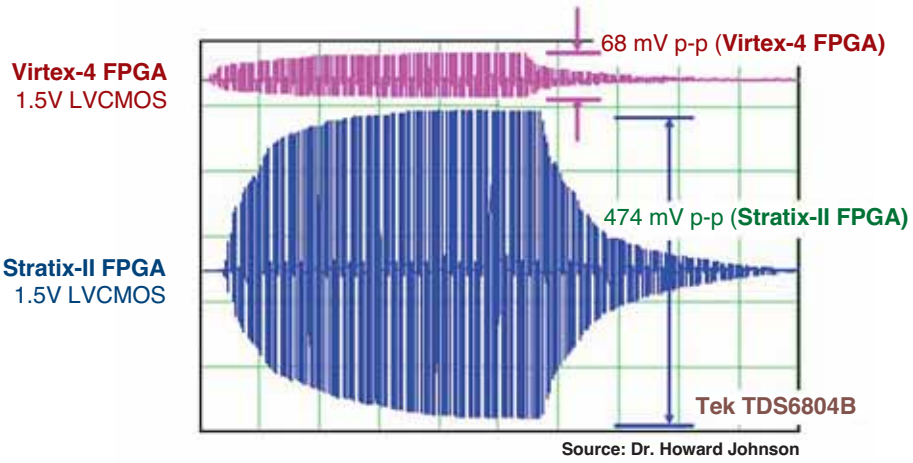

Figure 5 – Signal integrity comparison using the accumulated test pattern

Signal Integrity Challenge

One challenge that all chip-to-chip, high-speed interfaces need to overcome is signal integrity. Having control of cross-talk, ground bounce, ringing, noise margins, impedance matching, and decoupling is now critical to any successful design.

The Xilinx column-based ASMBL architecture enables I/O, clock, and power and ground pins to be located anywhere on the silicon chip, not just along the periphery. This architecture alleviates the problems associated with I/O and array dependency, power and ground distribution, and hard-IP scaling. Special FPGA packaging technology known as SparseChevron enables distribution of power and ground pins evenly across the package. The benefit to board designers is improved signal integrity.

The pin-out diagram in Figure 4 shows how Virtex-4 FPGAs compare with a competing Altera Stratix-II device that has many regions devoid of returns.

The SparseChevron layout is a major reason why Virtex-4 FPGAs exhibit unmatched simultaneous switching output (SSO) performance. As demonstrated by signal integrity expert Howard Johnson, Ph.D., these domain-optimized FPGA devices have seven times less SSO noise and crosstalk when compared to alternative FPGA devices (Figure 5).

Meeting I/O placement requirements and enabling better routing on a board requires unrestricted I/O placements for an FPGA design. Unlike competing solutions that restrict I/O placements to the top and bottom banks of the FPGA and functionally designate I/Os with respect to address, data and clock, Virtex-4 FPGAs provide unrestricted I/O bank placements.

Finally, Virtex-4 devices offer a differential DCM clock output that delivers the extremely low jitter performance necessary for very small data-valid windows and diminishing timing margins, ensuring a robust memory interface design.

These built-in silicon features enable high-performance synchronous interfaces for both memory and data communications in single or differential mode. The ChipSync technology enables data rates greater than 1 Gbps for differential I/O and more than 600 Mbps for single-ended I/O.

Conclusion

As with most FPGA designs, having the right silicon features solves only part of the challenge. Xilinx also provides complete memory interface reference designs that are hardware-verified and highly customizable. The Memory Interface Generator, a free tool offered by Xilinx, can generate all of the FPGA design files (.rtl, .ucf) required for a memory interface through an interactive GUI and a library of hardware-verified designs.

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

SIGNAL INTEGRITY

Xilinx Low-Noise FPGAs Meet SI Challenges

Unique chip package and I/Os accelerate system development

The good news is plentiful. Advances in silicon technology are enabling higher system performance to satisfy the requirements of networking, wireless, video, and other demanding applications. At the same time, an abundance of I/O pins with faster data edge-rates enables higher interface speeds.

However, every advance creates new design challenges. Wide buses with hundreds of simultaneously switching outputs (SSOs) create crosstalk. The undesirable effects of this electrical noise include jitter that threatens the stability of high-bandwidth interfaces. Sharp edge-rates further exacerbate noise problems.

> "High signal integrity demands a low-noise chip."

Howard Johnson
The world's foremost authority on signal integrity

Essential noise control

To address these issues, careful printed circuit-board (PCB) design and layout are critical for controlling system-level noise and crosstalk. Another important consideration is the electrical characteristics of the components mounted on the PCB. With its Virtex™-4 FPGAs, Xilinx, Inc., uses special chip I/O and packaging technologies that significantly improve signal-integrity not only at the chip level, but also at the system-level.

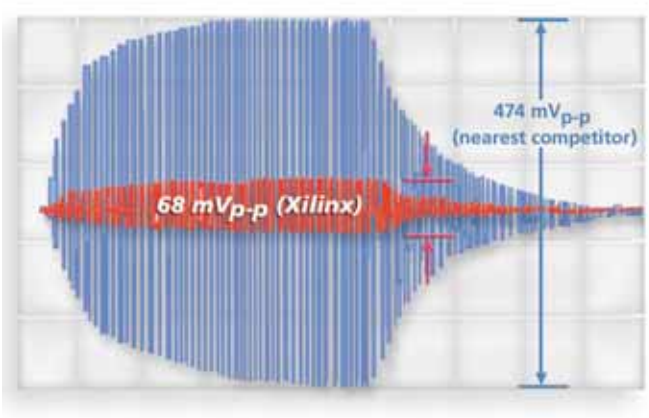

Design Example: 1.5 volt LVCMOS 4mA, I/O, 100 aggressors shown

"Xilinx preempts signal-integrity issues at the chip level," says Xilinx Senior Director of Systems and Applications Engineering Andy DeBaets. "This reduces board development and debug effort, and may even make the difference between a successful and scrapped board design."

In high-speed systems, a significant source of crosstalk is inductive coupling within the PCB via field under the BGA package. In systems with sub-optimal design, noise from simultaneously switching outputs can reach levels that severely degrade performance. In extreme cases, this noise can even lead to system failure. A properly designed BGA package minimizes inductive coupling between I/Os by placing a power/ground pin pair next to every signal pin.

"Xilinx's innovative SparseChevron™ package design minimizes crosstalk problems that can degrade system performance. This is particularly important for wide, high-speed DDR2 SDRAM or QDR II SRAM memory designs," DeBaets says.

Seven times less crosstalk

Analysis by independent signal-integrity expert Dr. Howard Johnson verifies the ability of SparseChevron technology to control SSO noise/crosstalk at the chip level. "High signal integrity demands a low-noise chip" states Johnson. "Compared to competing 90nm FPGAs, the Virtex-4 FPGAs in SparseChevron packaging demonstrate seven times less SSO noise," Johnson stresses

Xilinx Virtex-4 devices include several other features to help designers improve system-level signal integrity prior to board layout. Programmable edge rates and drive strength minimize noise while meeting other design objectives. Xilinx Digitally Controlled Impedance (DCI) technology enables designers to implement on-chip line termination for single-ended and differential I/Os. By eliminating the need for external termination resistors, DCI technology enables designers to minimize system component count and significantly simplify board layout and manufacturing.

To learn more about how Virtex-4 FPGAs can help control noise in your system, visit www.xilinx.com/virtex4.

For more information on signal integrity, go to
www.xilinx.com/platformsi

XILINX®

FPGA-Based Solutions for Storage-Area Networks

Xilinx FPGAs provide a solid foundation for value-add software.

by Sriram R. Chelluri
Marketing Manager
Xilinx, Inc.
sriram.chelluri@xilinx.com

Bryce Mackin
Manager, Vertical Marketing Storage & Servers
Xilinx, Inc.
bryce.mackin@xilinx.com

David Gamba
Senior Marketing Manager,
Vertical Marketing & Partnerships
Xilinx, Inc.
david.gamba@xilinx.com

FPGAs have played a critical role in the development of emerging technologies for SAN infrastructure products such as hubs, switches, RAID arrays, and storage appliances. In these products, FPGAs have been used in data compression, encryption, protocol translation, RAID striping, and virtualization.

In today's SAN market, a company's main value-add is in the software. Most emerging and established companies turn to software instead of hardware to give them cost and feature advantages. FPGAs allow companies to keep their intellectual value-add in software and still offer the hardware benefits of increased performance and scalability to meet requirements for current data-center products. An FPGA-based design allows you to create value-add in software that enables using license keys for adding product features and functionality – without having to rip and replace products to support additional functionality.

Other benefits of FPGA-based product designs are reduced time to market, field upgradability to support enhanced protocols, and configurable scalability to meet various protocol and performance requirements often encountered by product manufacturers. FPGAs also offer a lower total cost of development (TCD) in the long term, providing a flexible design option as product requirements change to meet business requirements. Compared to ASIC- or CPU-based designs that might be "too little too late," FPGAs allow you to add functionality to meet ever-changing market demands.

User demand for features and functionality in storage products has outpaced any CPU- or ASIC-based solutions that are bound by performance and scalability limitations. In this article, we'll explore some design solutions based on Xilinx® FPGAs for a subset of SAN ecosystem products found in typical data centers, such as security appliances, storage virtualization devices, and multi-protocol switches.

Data Security Solutions

Securing data in transit from the fabric to the disk is of critical importance for enterprise IT managers. System designers are coming up with new products for securing data at the server level, at the core of the fabric (like a switch), at storage/tape devices, or as an appliance that can be placed in the data path and integrated into the current infrastructure.

Keeping enterprise data secure and meeting the performance and scalability requirements of applications has become a challenge – and no vendor wants to be the performance bottleneck in a storage infrastructure. Therefore, product designers use FPGAs as the "glue" chip in interconnecting encryption ASICs that provide standard algorithms and key management functionality for performing the wire-speed packet inspection and processing required for application up-time.

Figure 1 shows how a security appliance can use a combination of FPGAs and off-the-shelf standard algorithm encryption chips to protect data at wire speed. Additional value-add features such as custom encryption, threat monitoring, and virus checking can be added in the FPGAs and enabled as licensed features without having to rip and replace product designs to enable these new features.

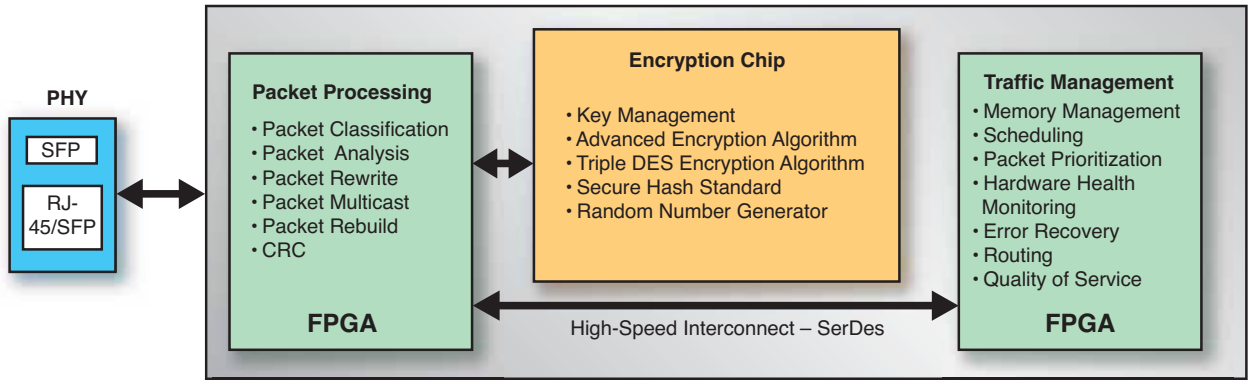

Figure 1 – FPGA-based security appliance

Storage Virtualization

Fibre Channel switching architectures allowed companies to scale servers and storage independently. Server and storage virtualization is the next evolution of the data center. A number of vendors offer server virtualization based on the high-speed InfiniBand protocol. Server virtualization allows IT administrators to allocate compute time to applications based on available resources.

Similarly, by virtualizing storage in a data center, system administrators can allocate storage based on performance, cost, and application policies. Storage virtualization can be at the file level (as in a network attached storage [NAS] box) or at the block level (as in IBM's SAN Volume Controller [SVC]). These classes of devices provide storage services such as management of available disks, data mirroring, and replication.

To achieve performance and scalability requirements, vendors often use FPGAs to provide the necessary hardware acceleration and meet the latency requirements of the SCSI/iSCSI protocol. Some of the FPGA-based hardware assistance functions implemented in storage virtualization products are:

- TCP offload
- Packet classification
- Ethernet: TCP or UDP
- Fibre Channel
- InfiniBand
- Traffic manager for queuing and scheduling
- Cyclic redundancy check (CRC) processing
- Protocol rate matching
- High-speed interconnect

Figure 2 shows how FPGAs fit into the design of a storage virtualization product.

FPGAs are primarily considered to be the glue chip in product design, providing hardware assistance for compute- or I/O-intensive applications. Additional functionalities such as packet classification and TCP offload can be added to FPGAs that have powerfully built CPU cores.

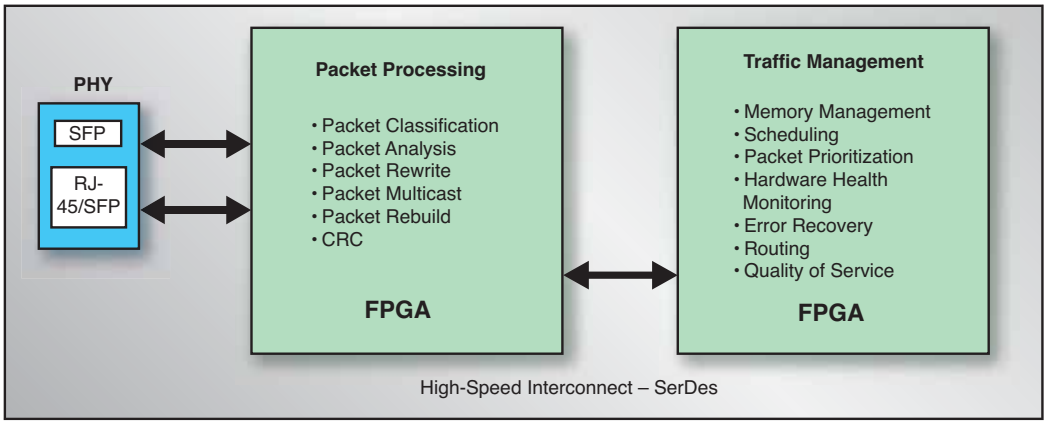

Figure 2 – FPGA-based storage virtualization product

Besides OSI layer 2-5 functionality, features that provide a competitive edge – encryption, mirroring and replication, fast lookup algorithms, XML processing, and anti-virus scanning – can all be built into the product and enabled by software license keys, without having to build ASICs or undertake complicated board designs.

Benefits

As costs continuously decrease for raw disk storage, vendors are looking to software services to provide value-add features such as volume management, mirroring, and replication services. To meet the performance and scalability requirements and yet keep costs down, you can use FPGAs to provide an opportunity to maintain margins in software without having to develop ASICs or add additional CPUs to meet application requirements.

Multi-Protocol Switch

Over the years, Ethernet and Fibre Channel switches have increased in port density, performance, and scalability. Enterprise customers are looking to multi-protocol switches to provide Ethernet, Fibre Channel, and gateway functionality. Some vendors already ship multi-protocol switches with dedicated or selectable Fibre Channel and Ethernet ports for SANs. As the media interconnect, link speeds, and protocol functionality become common across protocols, FPGAs with CPU cores can accommodate flexible design. Figure 3 shows how modern FPGAs can be used to support a variety of functions, including:

- Multi-protocol packet processing
- Protocol conversion
- Multicast and broadcast
- Routing
- Block and file search
- Buffer management

Conclusion

Established companies and venture-funded startups see software services for enterprise SAN products as the key enabler to valuations and margins. Solutions that require deep data inspection, such as security and storage appliances, firewalls, XML processors, and web-load balancers are an ideal choice for FPGAs or FPGA-based co-processing solutions.

In order to "beat the cost" out of the product, systems-level designers are looking to FPGAs instead of ASICs or general CPUs as the enabling technology to help them meet performance and scalability requirements, while at the same time keeping TCD within budget.

For additional information, visit *www.xilinx.com/products/silicon_solutions/fpgas/easypath/index.htm.*

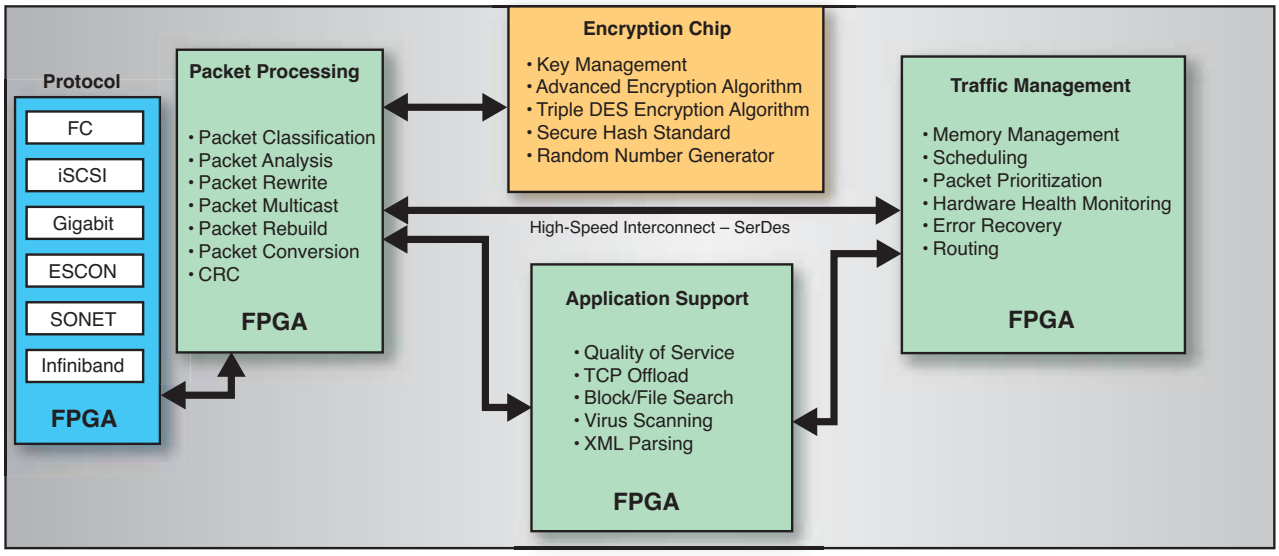

Figure 3 – FPGA-based multi-protocol switches

Generating Efficient Board Support Packages

The Xilinx Platform Studio toolset enables quick and easy BSP generation for Virtex FPGAs with immersed PowerPC processors.

by Rick Moleres
Manager of Software IP
Xilinx, Inc.
rick.moleres@xilinx.com

Milan Saini
Technical Marketing Manager
Xilinx, Inc.
milan.saini@xilinx.com

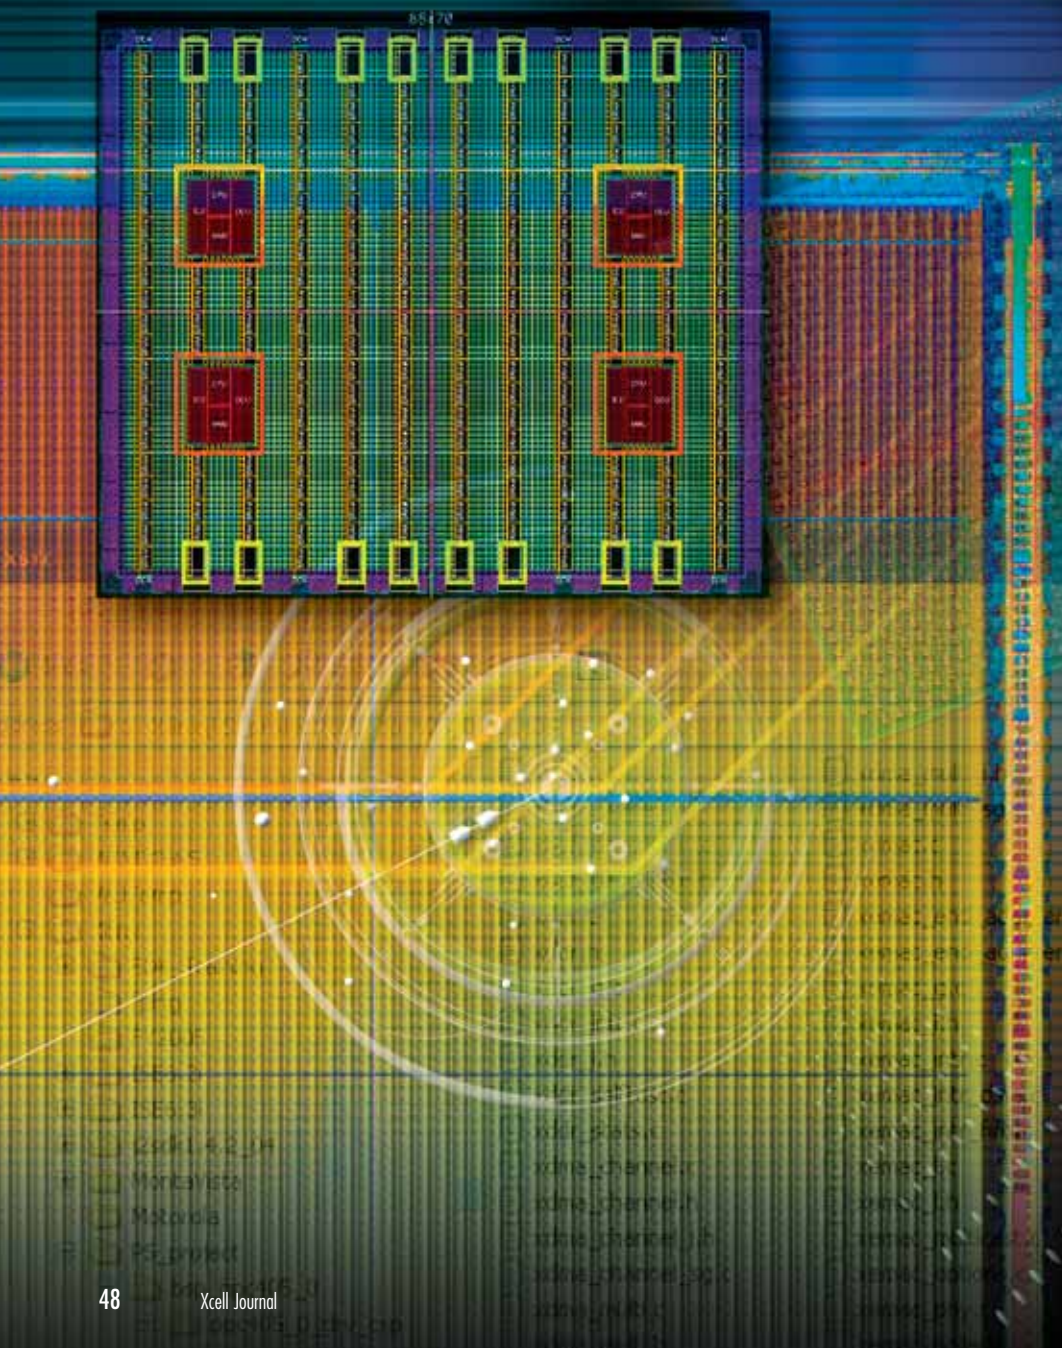

Platform FPGAs with embedded processors offer you unprecedented levels of flexibility, integration, and performance. It is now possible to develop extremely sophisticated and highly customized embedded systems inside a single programmable logic device.

With silicon capabilities advancing, the challenge centers on keeping design methods efficient and productive. In embedded systems development, one of the key activities is the development of the board support package (BSP). The BSP allows an embedded software application to successfully initialize and communicate with the hardware resources connected to the processor. Typical BSP components include boot code, device driver code, and initialization code.

Creating a BSP can be a lengthy and tedious process that must be incurred every time the microprocessor complex (processor plus associated peripherals) changes. With FPGAs, fast design iterations combined with the inherent flexibility of the platform can make the task of managing the BSP even more challenging (Figure 1). This situation clearly underscores the need for and importance of providing an efficient process for managing BSPs.

In this article, we'll describe an innovative solution from Xilinx that simplifies the creation and management of RTOS BSPs. We chose the WindRiver VxWorks flow to illustrate the concept; however, the underlying technology is generic and equally applicable to all other OS solutions that support Xilinx® processors.

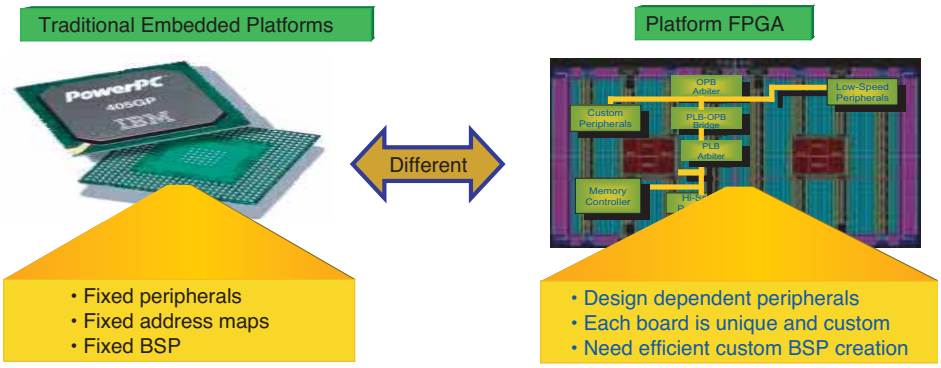

Figure 1 – Platform FPGA flexibility requires the software BSP generation process to be efficient.

Xilinx Design Flow and Software BSP Generation

Designing for a Xilinx processor involves a hardware platform assembly flow and an embedded software development flow. Both of these flows are managed within the Xilinx Platform Studio (XPS) tool, which is part of the Xilinx Embedded Development Kit (EDK).

You would typically begin a design by assembling and configuring the processor and its connected components in XPS. Once the hardware platform has been defined, you can then configure the software parameters of the system. A key feature of Platform Studio is its ability to produce a BSP that is customized based on your selection and configuration of processor, peripherals, and embedded OS. As the system evolves through iterative changes to the hardware design, the BSP evolves with the platform.

An automatically generated BSP enables embedded system designers to:

- Automatically create a BSP that completely matches the hardware design
- Eliminate BSP design bugs by using pre-certified components
- Increase designer productivity by jump-starting application software development

Creating BSPs for WindRiver VxWorks

Platform Studio can generate a customized Tornado 2.0.x (VxWorks 5.4) or Tornado 2.2.x (VxWorks 5.5) BSP for the PowerPC™ 405 processor and its peripherals in Xilinx Virtex™-II Pro and Virtex-4 FPGAs. The generated BSP contains all of the necessary support software for a system, including boot code, device drivers, and VxWorks initialization.

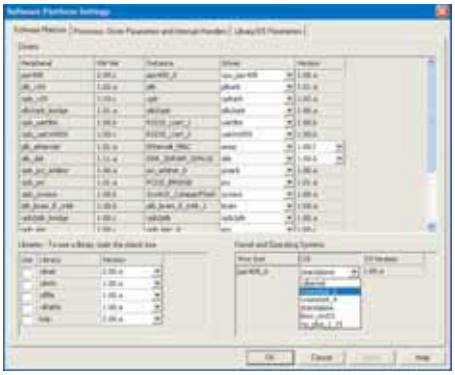

Figure 2 – Setting the embedded OS selection

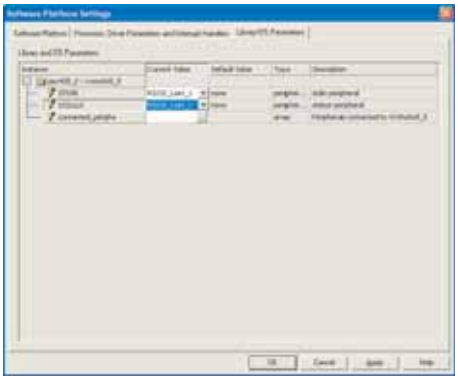

Figure 3 – Configuring the OS-specific parameters

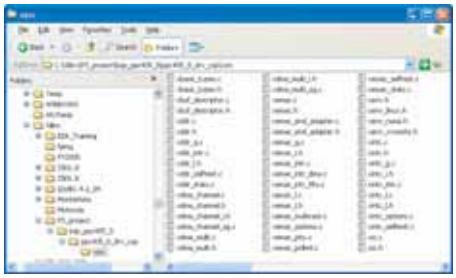

Figure 4 – Generated BSP directory structure

Once a hardware system with the PowerPC 405 processor is defined in Platform Studio, you need only follow these three steps to generate a BSP for VxWorks:

1. Use the Software Settings dialog box (see Figure 2) to select the OS you plan to use for the system. Platform Studio users can select vxworks5_4 or vxworks5_5 as their target operating system.

2. Once you have selected the OS, you can go to the Library/OS Parameters tab, as shown in Figure 3, to tailor the Tornado BSP to the custom hardware. You have the option of selecting any UART device in the system as the standard I/O device (stdin and stdout). This results in the device being used as the VxWorks console device.

   You can also choose which peripherals are connected peripherals, or which devices will be tightly integrated into the VxWorks OS. For example, the Xilinx 10/100 Ethernet MAC can be integrated into the VxWorks Enhanced Network Driver (END) interface. Alternately, the Ethernet device need not be connected to the END interface and can instead be accessed directly from the VxWorks application.

3. Generate the Tornado BSP by selecting the Tools > Generate Libraries and BSP menu option. The resulting BSP resembles a traditional Tornado BSP and is located in the Platform Studio project directory under ppc405_0/bsp_ppc405_0 (see Figure 4).

Note that ppc405_0 refers to the instance name of the PowerPC 405 processor in the hardware design. Platform Studio users can specify a different instance name, in which case the subdirectory names for the BSP will match the processor instance name.

The Tornado BSP is completely self-contained and transportable to other directory locations, such as the standard Tornado installation directory for BSPs at target/config.

Customized BSP Details

The XPS-generated BSP for VxWorks resembles most other Tornado BSPs except for the placement of Xilinx device driver code. Off-the-shelf device driver code distributed with Tornado typically resides in the target/src/drv directory in the Tornado distribution directory. Device driver code for a BSP that is automatically generated by Platform Studio resides in the BSP directory itself.

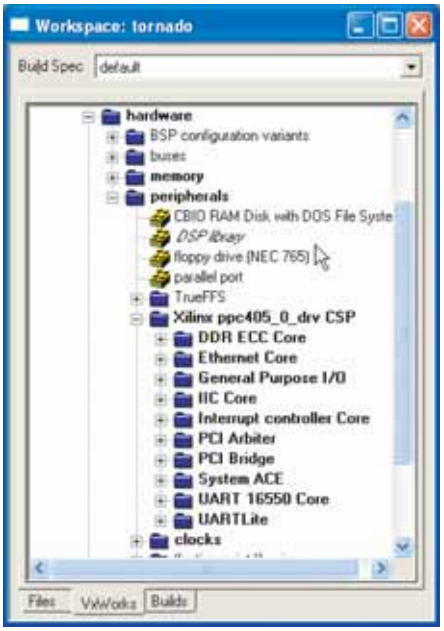

Figure 5 – Tornado 2.x Project: VxWorks tab

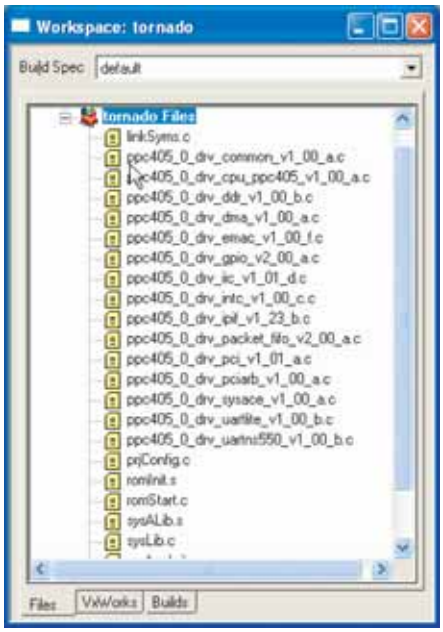

Figure 6 – Tornado 2.x Project: Files tab

This minor deviation is due to the dynamic nature of FPGA-based embedded systems. Because an FPGA-based embedded system can be reprogrammed with new or changed IP, the device driver configuration can change, calling for a more dynamic placement of device driver source files. The directory tree for the automatically generated BSP is shown in Figure 4. The Xilinx device drivers are placed in the ppc405_0_drv_csp/xsrc subdirectory of the BSP.

The Tornado BSP created by Platform Studio has a makefile that you can modify at the command line if you would rather use the diab compiler instead of the gnu compiler.

Xilinx device drivers are implemented in C and are distributed among several source files, unlike traditional VxWorks drivers, which typically consist of single C header and implementation files. In addition, there is an OS-independent implementation and an optional OS-dependent implementation for device drivers.

The OS-independent part of the driver is designed for use with any OS or any processor. It provides an application program interface (API) that abstracts the functionality of the underlying hardware. The OS-dependent part of the driver adapts the driver for use with an OS such as VxWorks. Examples are Serial IO drivers for serial ports and END drivers for Ethernet controllers. Only drivers that can be tightly integrated into a standard OS interface require an OS-dependent driver.

Xilinx driver source files are included in the build of a VxWorks image in the same way that other BSP files are included in the build. For every driver, a file exists named ppc405_0_drv_<driver_version>.c in the BSP directory. This file includes the driver source files (*.c) for the given device and is automatically compiled by the BSP makefile.

This process is analogous to how VxWorks' sysLib.c includes source for Wind River-supplied drivers. The reason why Xilinx driver files are not simply included in sysLib.c like the rest of the drivers is because of namespace conflicts and maintainability issues. If all Xilinx driver files are part of a single compilation unit, static functions and data are no longer private. This places restrictions on the device drivers and would negate their OS independence.

Integration with the Tornado IDE

The automatically generated BSP is integrated into the Tornado IDE (Project Facility). The BSP is compilable from the command line using the Tornado make tools or from the Tornado Project. Once the BSP is generated, you can simply type make vxWorks from the command line to compile a bootable RAM image. This assumes that the Tornado environment has been previously set up, which you can do through the command line using the host/x86-win32/bin/torVars.bat script (on a Windows platform). If you are using the Tornado Project facility, you can create a project based on the newly generated BSP, then use the build environment provided through the IDE to compile the BSP.

In Tornado 2.2.x, the diab compiler is supported in addition to the gnu compiler. The Tornado BSP created by Platform Studio has a makefile that you can modify at the command line if you would rather use the diab compiler instead of the gnu compiler. Look for the make variable named TOOLS and set the value to "diab" instead of "gnu." If using the Tornado Project facility, you can select the desired compiler when the project is first created.

The file 50ppc405_0.cdf resides in the BSP directory and is tailored during creation of

the BSP. This file integrates the device drivers into the Tornado IDE menu system. The drivers are hooked into the BSP at the Hardware > Peripherals subfolder. Below this are individual device driver folders. Figure 5 shows a menu with Xilinx device drivers.

The Files tab of the Tornado Project Facility will also show the number of files used to integrate the Xilinx device drivers into the Tornado build process. These files are automatically created by Platform Studio and you need only be aware that the files exist. Figure 6 shows an example of the driver build files.

Some of the commonly used devices are tightly integrated with the OS, while other devices are accessible from the application by directly using the device drivers. The device drivers that have been tightly integrated into VxWorks include:

- 10/100 Ethernet MAC
- 10/100 Ethernet Lite MAC
- 1 Gigabit Ethernet MAC
- 16550/16450 UART
- UART Lite
- Interrupt Controller
- System ACE™ technology
- PCI

All other devices and associated device drivers are not tightly integrated into a VxWorks interface; instead, they are loosely integrated. Access to these devices is available by directly accessing the associated device drivers from the user application.

Conclusion

With the popularity and usage of embedded processor-based FPGAs continuing to grow, tool solutions that effectively synchronize and tie the hardware and software flows together are key to helping designer productivity keep pace with advances in silicon.

Xilinx users have been very positive about Platform Studio and its integration with VxWorks 5.4 and 5.5. Xilinx fully intends to continue its development support for the Wind River flow that will soon include support for VxWorks 6.0 and Workbench IDE.

Microprocessor Library Definition (MLD)

The technology that enables dynamic and custom BSP generation is based on a Xilinx proprietary format known as Microprocessor Library Definition (MLD). This format provides third-party vendors with a plug-in interface to Xilinx Platform Studio to enable custom library and OS-specific BSP generation (see Figure 7). The MLD interface is typically written by third-party companies for their specific flows. It enables the following add-on functionality:

- Enables custom design rule checks
- Provides the ability to customize device drivers for the target OS environment
- Provides the ability to custom-produce the BSP in a format and folder structure tailored to the OS tool chain
- Provides the ability to customize an OS/kernel based on the hardware system under consideration

The MLD interface is an ASCII-based open and published standard. Each RTOS flow will have its own set of unique MLD files. An MLD file set comprises the following two files:

- A data definition (.mld) file. This file defines the library or operating system through a set of parameters set by the Platform Studio. The values of these parameters are stored in an internal Platform Studio database and intended for use by the script file during the output generation.
- A .tcl script file. This is the file that is called by XPS to create the custom BSP. The file contains a set of procedures that have access to the complete design database and hence can write a custom output format based on the requirements of the flow.

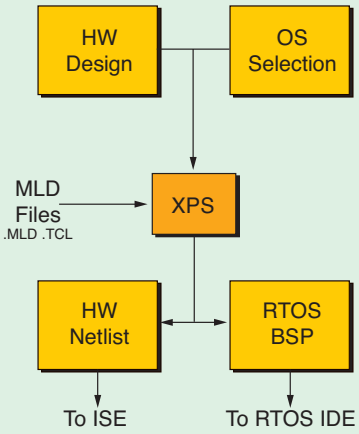

Figure 7 – Structure of an MLD flow

The MLD syntax is described in detail in the EDK documentation (see "Platform Specification Format Reference Manual" at *www.xilinx.com/ise/embedded/psf_rm.pdf*). You can also find MLD examples in the EDK installation directory under sw/lib/bsp.

Once MLD files for a specific RTOS flow have been created, they need to be installed in a specific path for Xilinx Platform Studio to pick up on its next invocation. The specific RTOS menu selection now becomes active in the XPS dialog box (Project > SW Platform Settings > Software Platform > OS).

Currently, the following partners' MLD files are available for use within XPS:

- Wind River (VxWorks 5.4, 5.5) (included in Xilinx Platform Studio)
- MontaVista (Linux) (included in Xilinx Platform Studio)
- Mentor Accelerated Technologies (Nucleus) (download from *www.xilinx.com/ise/embedded/mld/*)
- GreenHills Software (Integrity) (download from *www.xilinx.com/ise/embedded/mld/*)
- Micrium (µc/OS-II) (download from *www.xilinx.com/ise/embedded/mld/*)
- µcLinux (download from *www.xilinx.com/ise/embedded/mld/*).

"High Signal Integrity Demands a Low Noise Chip".

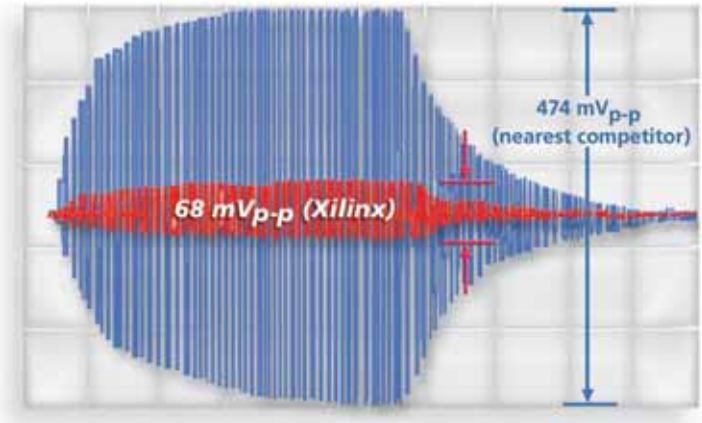

Design Example: 1.5 volt LVCMOS 4mA, I/O, 100 aggressors shown.

Best Signal Integrity: 7x Less SSO Noise

Virtex-4 FPGAs deliver the industry's best signal integrity, allowing you to pre-empt board issues at the chip level, for high-speed designs such as memory interfaces. Featuring a unique SparseChevron™ pin out pattern, the Virtex-4 family provides the highest ratio of VCCO/GND pin pairs to user I/O pins available in any FPGA. By strategically positioning one hard power pin and one hard ground pin adjacent to every user I/O on the device, we've reduced signal path inductance and SSO noise to levels far below what you can attain with a virtual ground or soft ground architecture.

The Industry's Highest Signal Integrity, Proven By Industry Experts

Incorporating continuous power and ground planes, plus integrated bypass capacitors, we're eliminating power-supply noise at its source. In addition, we provide on-chip termination resistors to control signal ringing. The lab tests speak for themselves. As measured by signal integrity expert Dr. Howard Johnson, no competing FPGA comes close to achieving the low-noise benchmarks of Virtex-4 devices.

Visit ***www.xilinx.com/virtex4/si*** today, and choose the right high-performance FPGA before things get noisy.

Dr. Howard Johnson, author of *High-Speed Digital Design,* frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. Visit www.sigcon.com to register.

FPGA and Structured ASIC JOURNAL Reader's Choice Awards

Highest customer satisfaction:
- Virtex & Spartan FPGAs
- ISE design tools
- Support

XILINX®
The Programmable Logic Company℠
www.xilinx.com/virtex4/si

View The TechOnLine Seminar Today

Breakthrough Performance At The Lowest Cost

Optimizing Embedded Systems for Broadband 10 Gigabit Ethernet Connectivity

FPGA technology is revolutionizing broadband systems development.

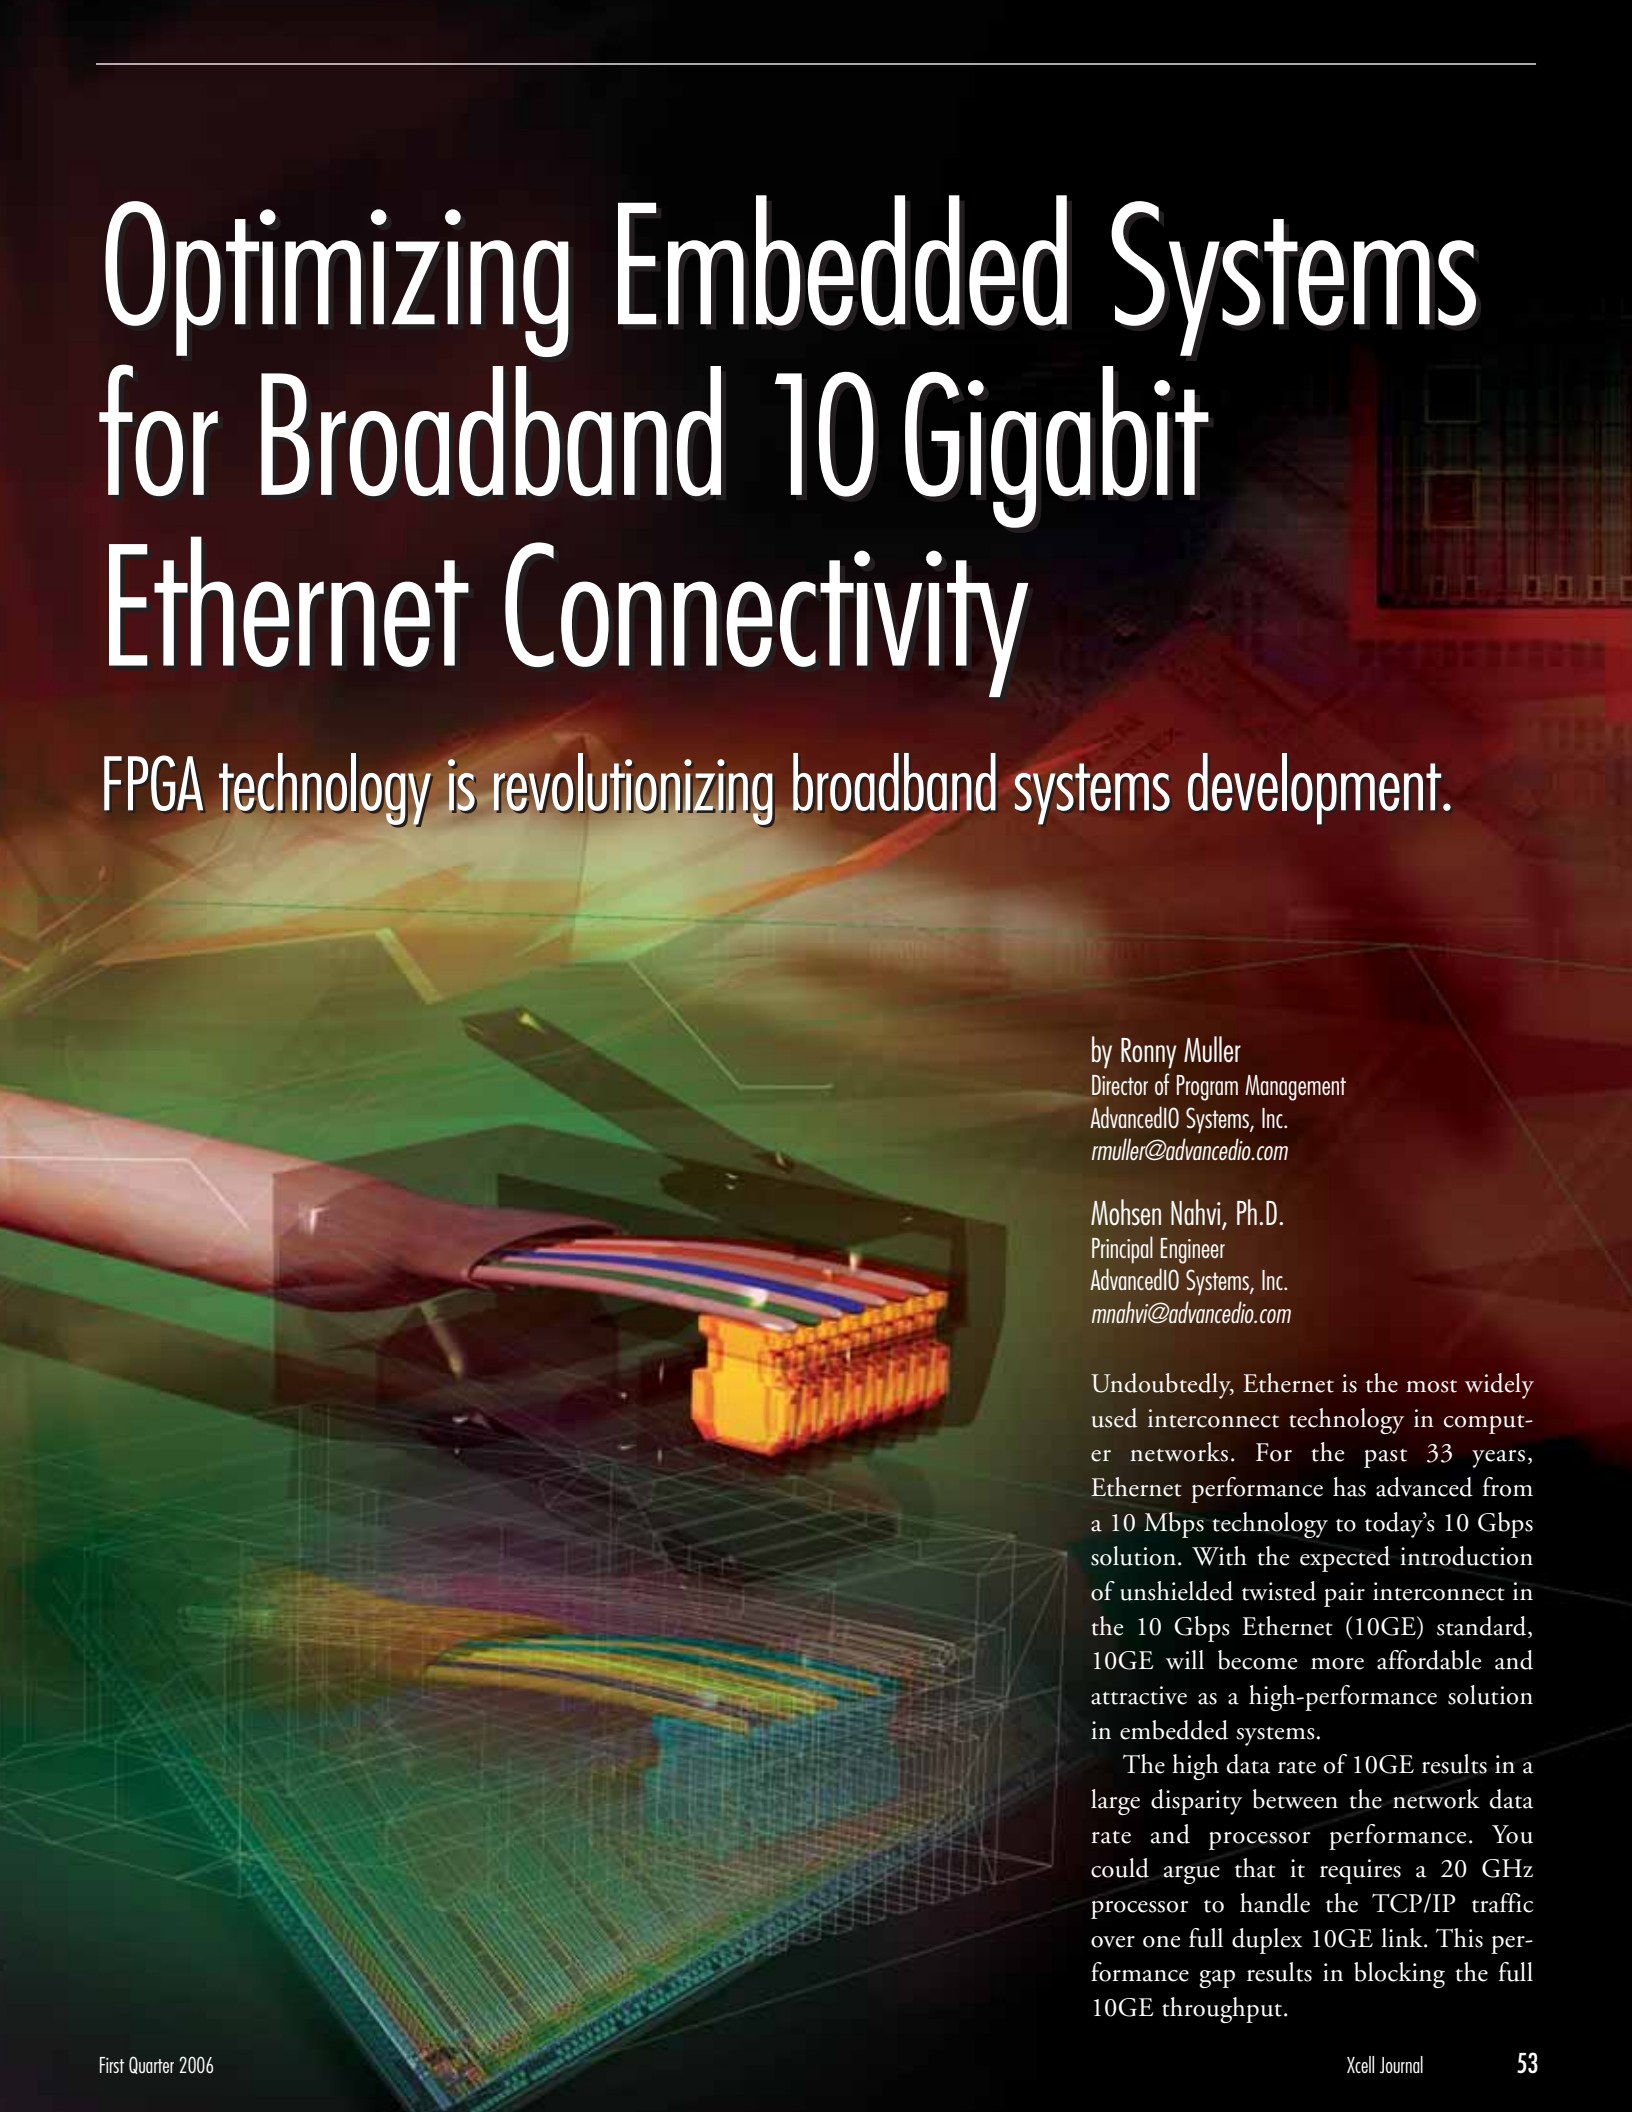

by Ronny Muller
Director of Program Management
AdvancedIO Systems, Inc.
rmuller@advancedio.com

Mohsen Nahvi, Ph.D.
Principal Engineer
AdvancedIO Systems, Inc.
mnahvi@advancedio.com

Undoubtedly, Ethernet is the most widely used interconnect technology in computer networks. For the past 33 years, Ethernet performance has advanced from a 10 Mbps technology to today's 10 Gbps solution. With the expected introduction of unshielded twisted pair interconnect in the 10 Gbps Ethernet (10GE) standard, 10GE will become more affordable and attractive as a high-performance solution in embedded systems.

The high data rate of 10GE results in a large disparity between the network data rate and processor performance. You could argue that it requires a 20 GHz processor to handle the TCP/IP traffic over one full duplex 10GE link. This performance gap results in blocking the full 10GE throughput.

. . . AdvancedIO provides an FPGA framework in which these high-speed interfaces are mapped, placed and routed, and guaranteed to meet timing.

The problem is even worse in high-performance embedded systems. These embedded systems typically rely on high-density distributed processing elements (PE), where each element is optimized to perform specific DSP functions and has no free cycles to perform the complex task of TCP/IP traffic processing.

To address the huge performance gap, you need TCP/IP offload engines (TOE) and protocol acceleration schemes at the network interface; in effect, smart 10GE solutions that go beyond the simple task of a network adaptor. A few ASIC-based solutions on the market do offload part of the TCP/IP stack. However, because of the huge gap between the network and the processor performance, a robust solution is essential that will not only help with offloading the TCP/IP protocol, but also enable you to address other challenges specific to applications like real-time data encryption/decryption and packet classification.

In this context, FPGA-based solutions offer an ideal solution because of their flexibility, configurability, and shorter time-to-market capabilities. An FPGA-based network adaptor allows you to offload the TCP/IP protocol processing from the host system, enabling you to implement other functionality into these intelligent adaptors to improve overall embedded system performance.

In this article, we'll discuss a configurable 10GE solution that enables you to address your high-performance network needs and solve the 10GE challenges in your system.

The AdvancedIO Architecture

AdvancedIO Systems is focused on solving problems of high-speed packet switching in broadband applications using programmable logic. The company offers a line of I/O products and system solutions based on a common architecture.

At the heart of AdvancedIO's architecture is a Xilinx® FPGA, which implements the 10GE physical coding sub-layer (PCS), the 10GE media access controller (MAC), and the RapidIO (or PCIe) switched fabric interface. These standard interfaces are available as optimized logic cores from a variety of vendors.

As shown in Figure 1, AdvancedIO provides an FPGA framework in which these high-speed interfaces are mapped, placed and routed, and guaranteed to meet timing. Consequently, AdvancedIO greatly reduces the engineering effort required for designing new applications. All interfaces are fully controllable from the on-chip PowerPC™ 405 processors, allowing low-data-rate verification of broadband applications.

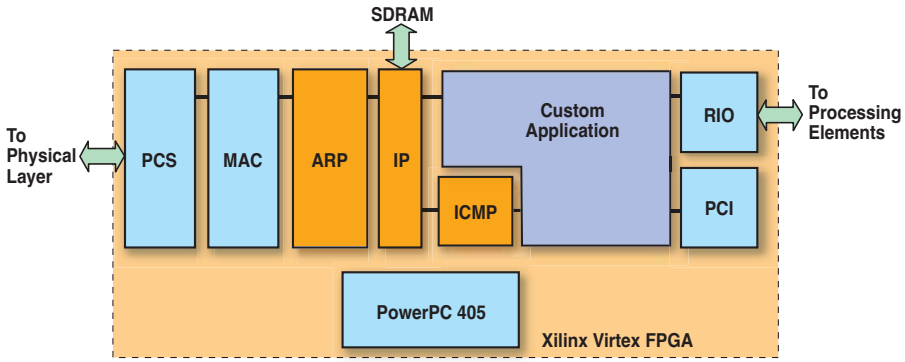

Figure 1 – AdvancedIO 10GE architecture

The V1010 is the industry's first configurable 10GE module in an XMC form factor (VITA 42), with the design based on AdvancedIO's architecture for embedded applications. One particular configuration of the module bridges between RapidIO and 10GE provides the ability to perform 10GE layer 2 and 3 packet processing at wire speed.

Three additional logic blocks facilitate the implementation of a 10GE network endpoint within the V1010 module:

- ARP. This block takes incoming IP frames and converts them into Ethernet frames by appending the Ethernet destination and source MAC addresses. It implements the network address to hardware address request and response protocol and maintains an ARP table in hardware.

Figure 2 – AdvancedIO V1010 configurable 10GE XMC module

- IP. This block terminates IP. It implements IP fragmentation and de-fragmentation at wire speed by buffering fragmented datagrams in SDRAM until the full datagram has been received. Moreover, it checks and generates IP checksums and also performs IP routing. The IP routing tables are configured by the on-chip PowerPC 405 processor.
- ICMP. This block implements the required ICMP protocol. It responds to ping/trace-route commands, for example, and reports and counts errors.

Figure 2 is a photograph of the module currently shipping.

Distributed Processing Made Easier

As shown in Figure 3, the AdvancedIO architecture allows full-speed access to switched fabrics such as RapidIO or PCIe, and can send and receive full IP datagrams to and from the 10GE IP network. Using this architecture, you can easily implement any protocol processing, from the very simple to the very complicated.

As an example, multiprocessor embedded systems with RapidIO connectivity can now seamlessly communicate with each other over an IP 10GE network, as shown in Figure 4. To achieve this, the V1010 module encapsulates RapidIO packets into UDP packets, keeping track of lost and out-of-order packets and reporting these errors to the processing elements.

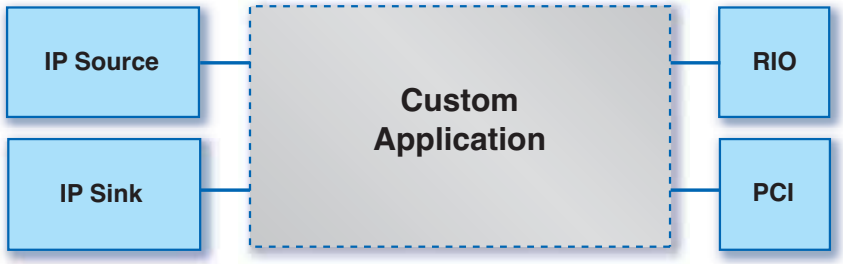

Figure 3 – AdvancedIO processing model for broadband applications

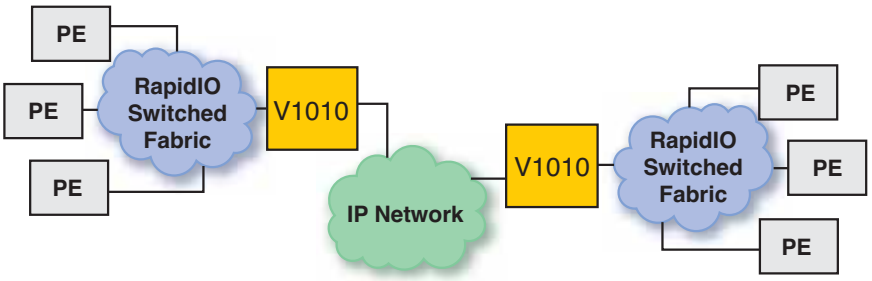

Figure 4 – Distributed remote processing over RapidIO and 10 Gigabit Ethernet

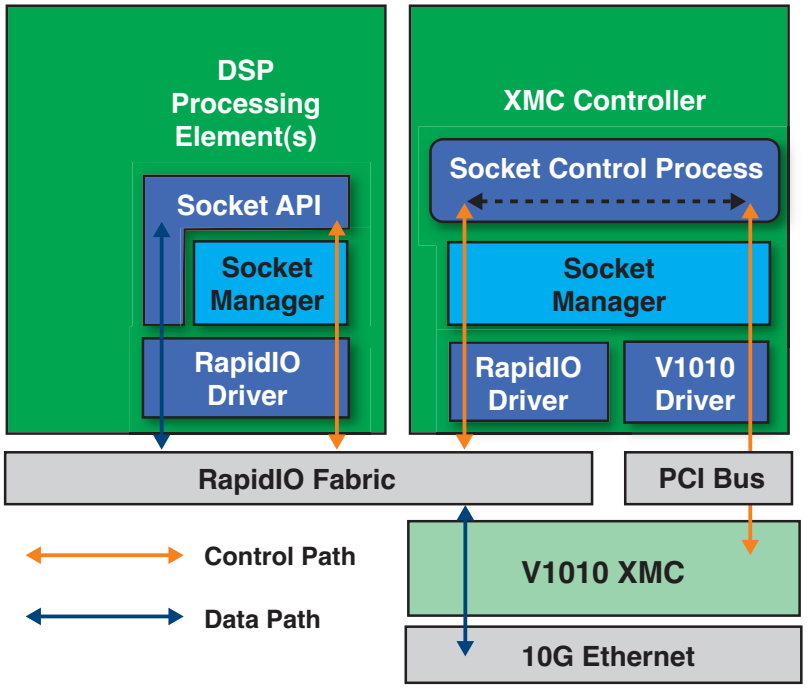

Figure 5 – AdvancedIO embedded software framework

The above tunneling system allows for inter-chassis communication over an IP network and is completely transparent to the processing elements. As all legal RapidIO packets in both message and I/O specifications are transferred over the network, you now have the capability of utilizing remote direct memory access (RDMA) over an IP network, which enables and greatly simplifies the development of complex distributed embedded applications.

The programming interface is a Unix-like socket application programming interface (API), so an application can set up and control the connectivity to the 10GE gateway from any processor element in a RapidIO network. As shown in Figure 5, the AdvancedIO software framework abstracts the complexity of the underlying functions and the 10GE data path. In a typical application, the socket manager is called, which in turn sets up the control path between a processing element and one or more 10GE gateways. Once the control path has been established, an application can fully utilize the 10GE data path.

Conclusion

With the advent of 10 Gigabit Ethernet and its foray into multiprocessor embedded systems, you need intelligent network solutions that offer more than mere TCP/IP offloading. The programmability of the AdvancedIO architecture provides the right ingredient for rapid system development and boosts the performance of embedded systems to a higher level.

The AdvancedIO architecture offers a new and unique solution to embedded computing developers, providing an efficient way for applications to take advantage of all that the 10 Gigabit Ethernet technology has to offer.

For more information on AdvancedIO's 10 Gigabit Ethernet products and services, visit *www.advancedio.com* or e-mail *contactus@advancedio.com*.

Versatile 10A Power Supply Mounts Like an IC

By Afshin Odabaee
Product Marketing Engineer, Power Management Products
Linear Technology Corp.

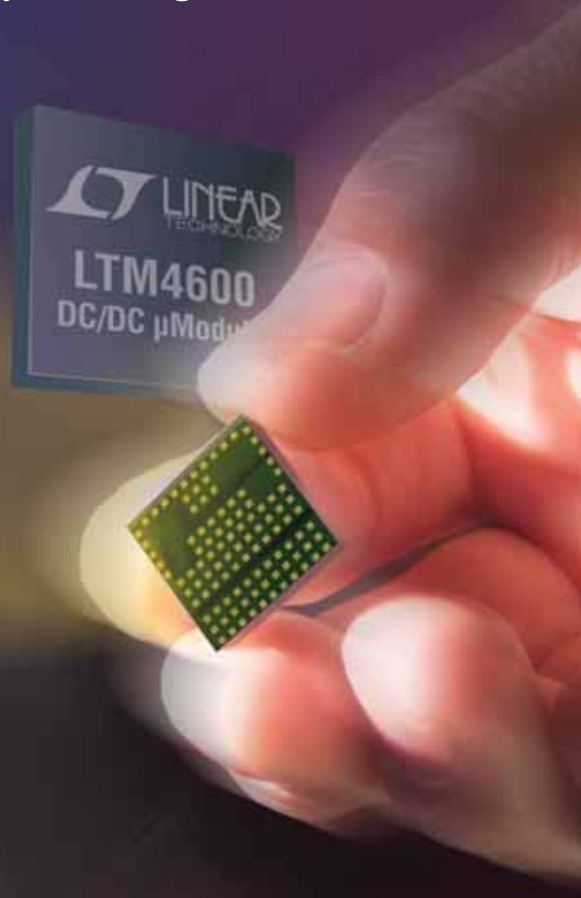

Pre-manufactured DC/DC power supplies, also referred to as point-of-load modules, have been promising simpler, smaller and quicker solutions but they fall short in meeting the demands for system assembly of densely populated embedded boards. Some solutions, labeled as complete power supplies, still require an external inductor, many additional input and output capacitors and compensation circuitry. Most are assembled on a small printed circuit board (PCB) and require hand inspection for reliability purposes since the circuit components are exposed and not encapsulated.

For heat dissipation and safe component spacing, many embedded systems specify a maximum thickness for both the top and bottom of the board. Unfortunately, high power density DC/DC modules must implement tall inductors and rely on thick PCBs to alleviate heat dissipation. The large size and tall profile inhibits their use. Therefore, a designer must either design a discrete power supply where it can be optimized to meet the profile requirement or rely on lower power DC/DC modules where thinner inductors are used.

System designers are compelled to make compromises in the selection, performance and definition of the optimum power supply. The best solution is a complete power supply with no external power components, no mathematical analysis, easy layout and a product that fully meets the surface mount assembly requirements, just as the other digital ICs on the board. The result is a solution that is easy to select, design and assemble.

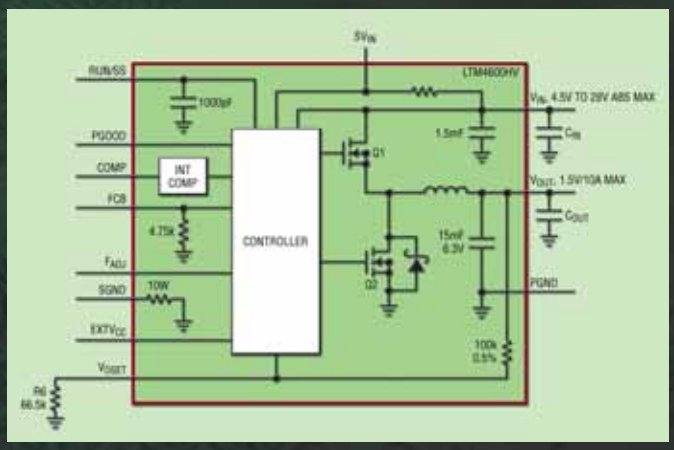

Figure 1- The LTM4600 requires minimum external components. The difficult compensation circuitry, inductor, MOSFET, DC/DC controller and input/output capacitors are onboard. 1.5V, 10A application is shown here.

Complete Power Supply in an IC Form-Factor

This IC-like solution is a 10A synchronous switchmode regulator with built-in inductor, supporting power components and compensation circuitry (Figure 1). The LTM®4600 is the DC/DC module that meets the spacing and assembly requirements of the densely populated and advanced embedded systems. This encapsulated µModule™ DC/DC power supply is housed in a 15mm x 15mm x 2.8mm LGA package. Its size is smaller than most FPGAs and processors. With only a 2.8mm profile, the LTM4600 readily can be placed on the backside of a board without adding significant thickness to a board.

This µModule is rated for 20V and 28V input operation (two versions). The output voltage is adjustable with a single resistor from 0.6V to 5V. The LTM4600 can deliver up to 10A of output current and offers excellent transient response to fast changing load current transients.

Easy Design Replication by Copying and Pasting the Layout

A common complaint among system designers is the assembly-house's alteration of a specified layout. The result is many rounds of debugging between the design- and assembly-house. This problem is alleviated with the LTM4600's fixed and simple layout.

LINEAR TECHNOLOGY

NU HORIZONS ELECTRONICS CORP.

Instant 10A Power Supply

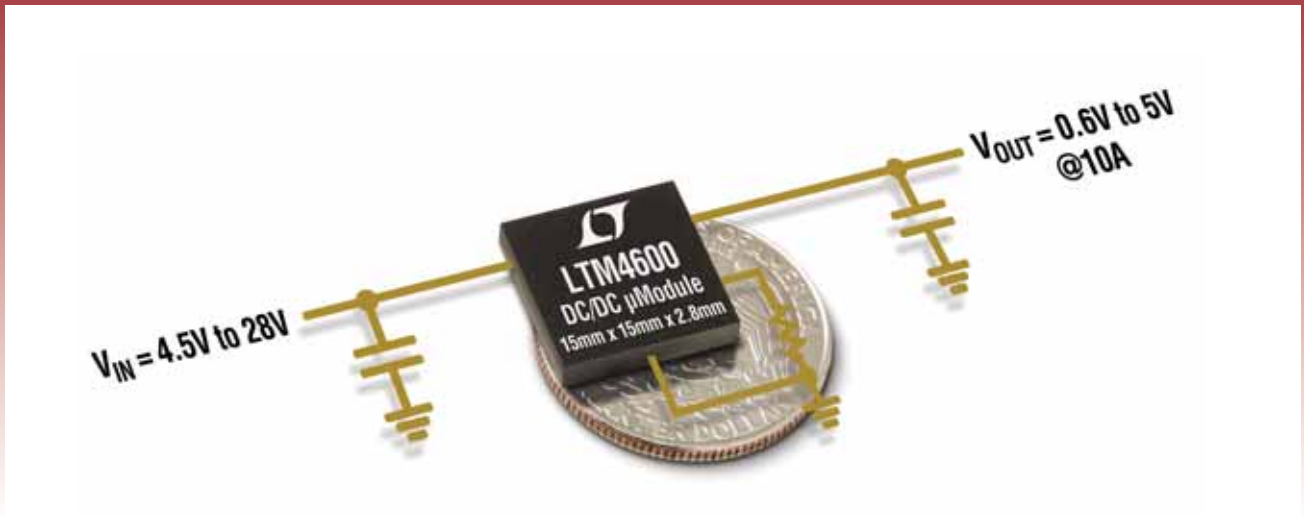

Complete, Quick & Ready.

The LTM4600 is a complete 10A switchmode step-down power supply with a built-in inductor, supporting power components and compensation circuitry. With high integration and synchronous current mode operation, this DC/DC µModule™ delivers high power at high efficiency in a tiny, low profile surface mount package. Supported by Linear Technology's rigorous testing and high reliability processes, the LTM4600 simplifies the design and layout of your next power supply.

Features

- 15mm x 15mm x 2.8mm LGA with 15°C/W JA
- Pb-Free (e4), RoHS Compliant
- Only CBULK Required
- Standard and High Voltage:
  LTM4600EV: 4.5V ≤ V_{IN} ≤ 20V
  LTM4600HVEV: 4.5V ≤ V_{IN} ≤ 28V
- 0.6V≤ V_{OUT} ≤ 5V
- I_{OUT}: 10A DC, 14A Peak
- Parallel Two µModules for 20A Output

Ultrafast Transient Response

2% ΔV_{OUT} with a 5A Step

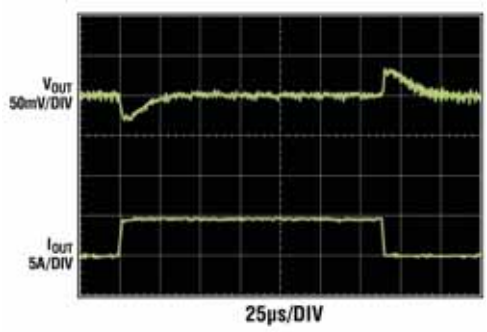

V_{IN} = 12V, V_{OUT} = 1.5V, 0A to 5A Load Step
(C_{OUT} = 3 x 22µF CERAMICS, 470µF POS CAP)

Contact Info

Nu Horizons Electronics Corp.

www.nuhorizons.com/linear

Tel: 1-888-747-NUHO

LTM4600 Application Notes
Available for Download

LINEAR TECHNOLOGY

NU HORIZONS ELECTRONICS CORP.

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!

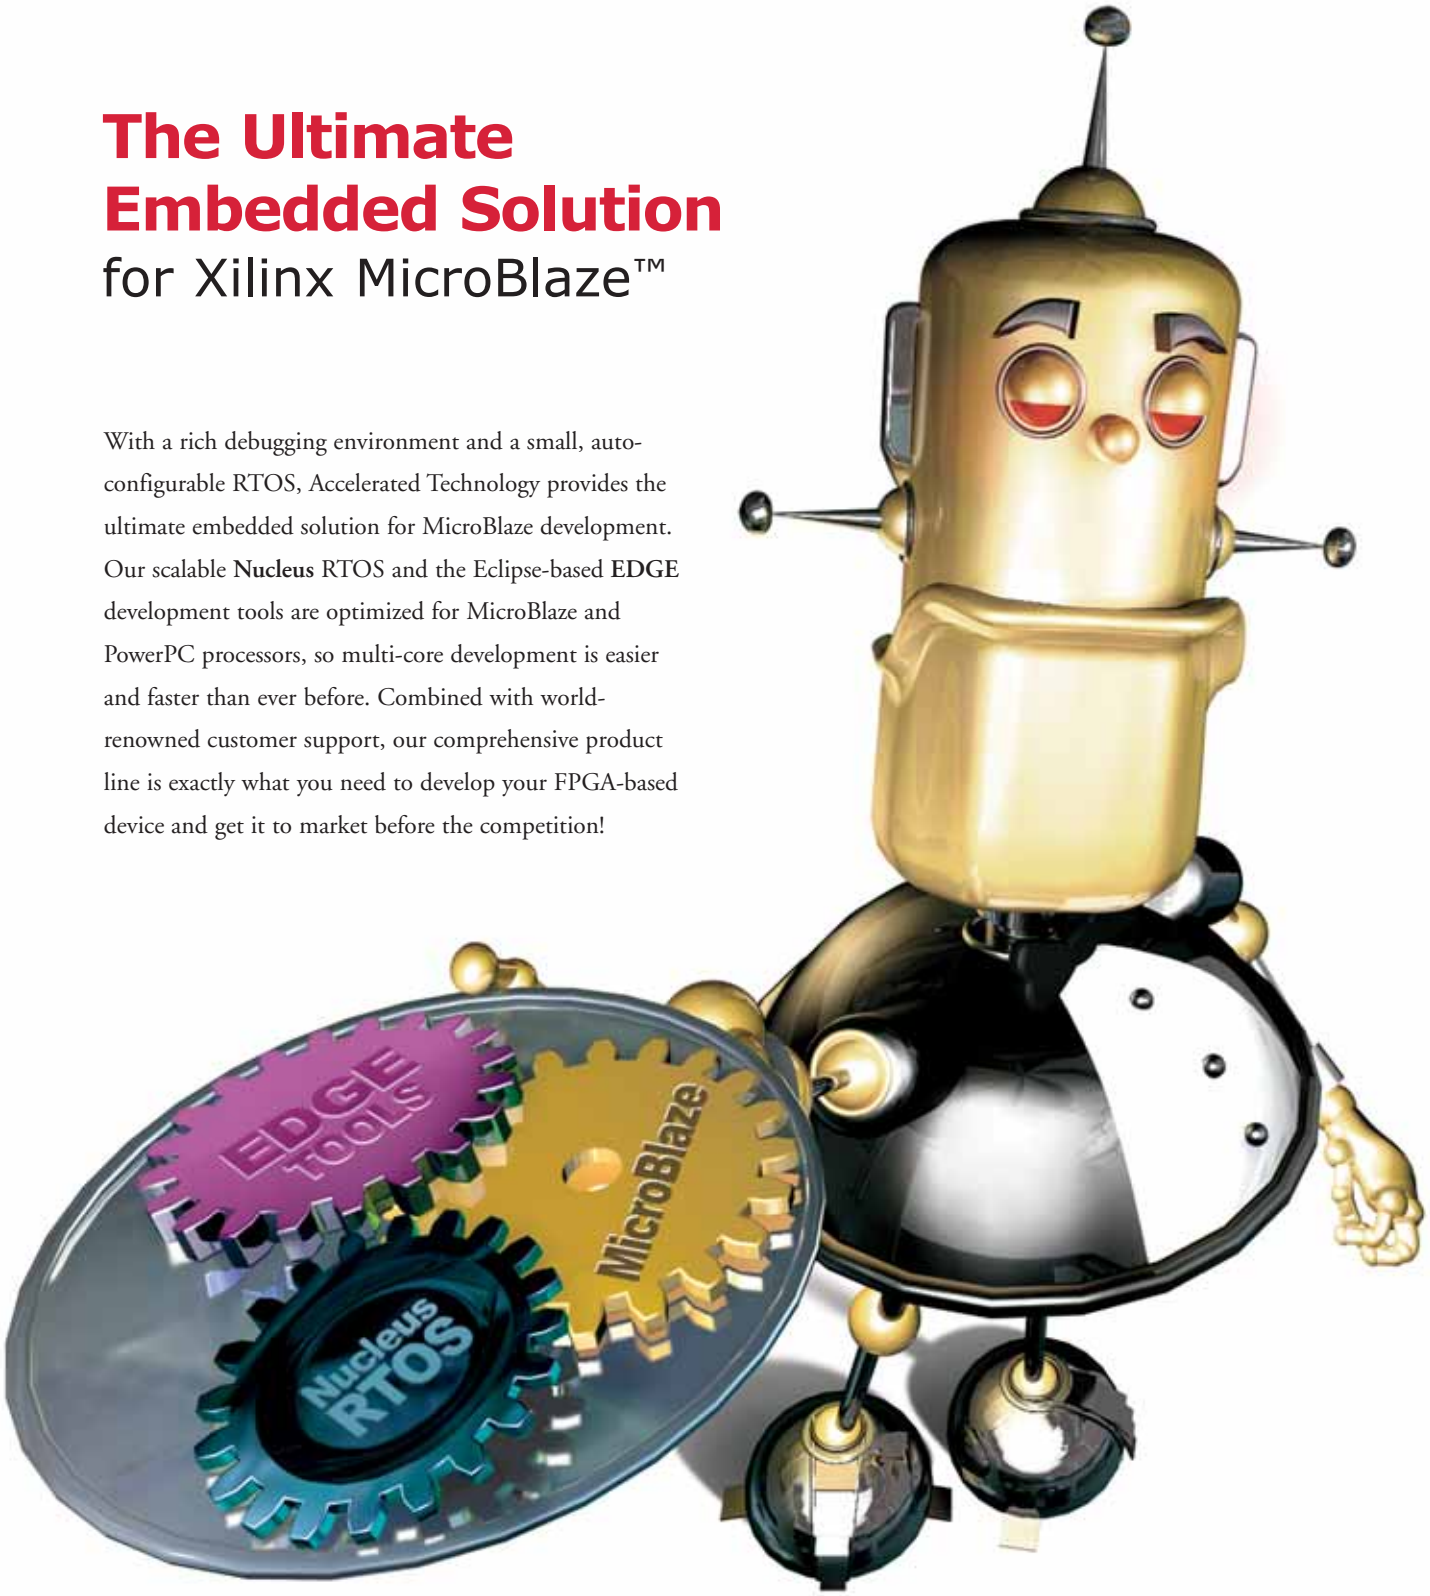

For more information on the Nucleus complete solution, go to
www.acceleratedtechnology.com/xilinx

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

Accelerated Technology®
A Mentor Graphics Division

Board Design Panacea

The 7Circuits tool algorithmically solves FPGA pinout problems and synthesizes PC board schematics.

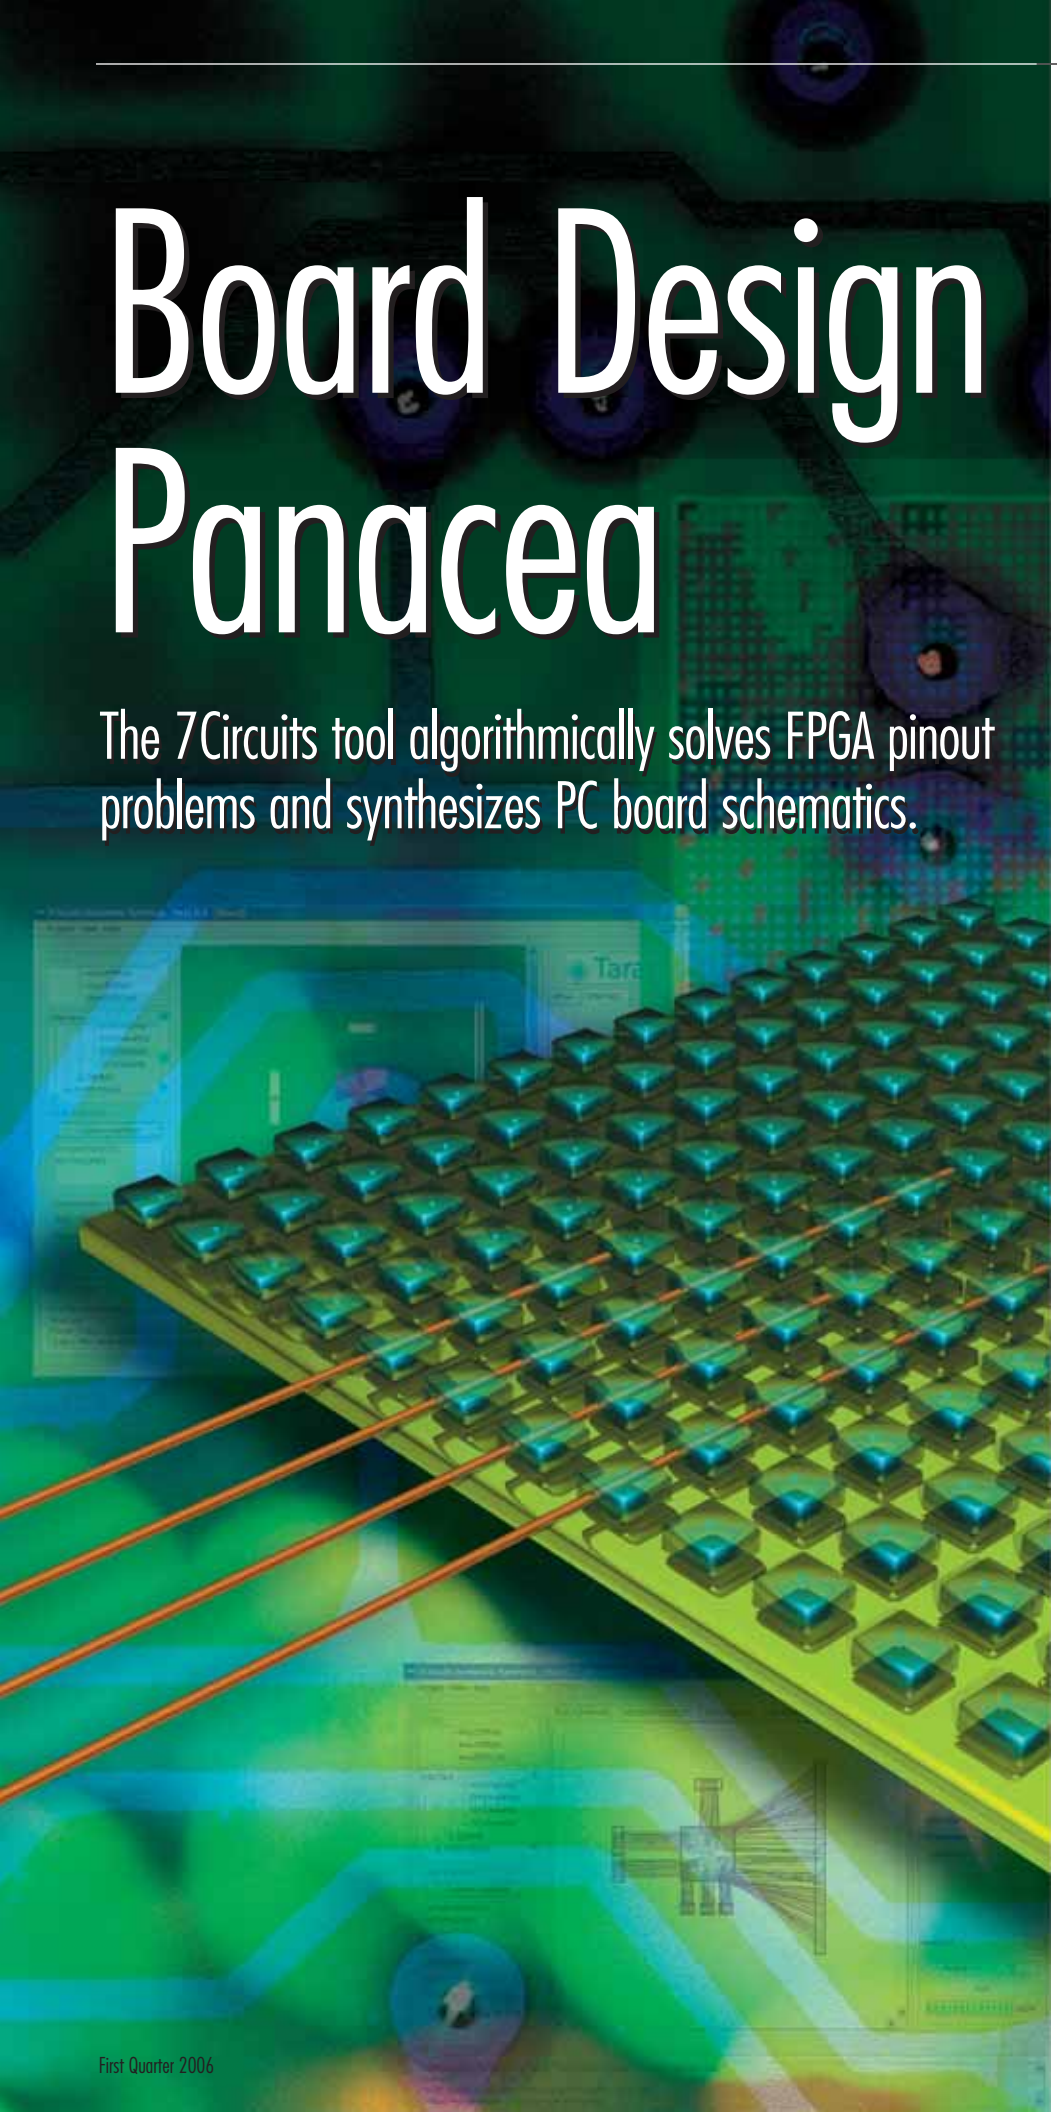

by Nagesh Gupta
Founder/CEO
Taray, Inc.
nagesh@tarayinc.com

PC board design is a cumbersome and time-consuming task. Although some of the steps require knowledge and intelligence to complete, most of the process is mundane and routine. Add FPGAs to the mix, and the complexity of the board grows significantly. FPGAs have a myriad of complex I/O rules that are multi-dimensional and can present difficult problems:

1. In most cases with large and complex designs, FPGA pinouts are hardly optimal, and non-optimal pinouts result in lower design performance. The cost of the PC board also increases because of the higher number of layers.
2. Today, pins for FPGAs are mostly selected manually. The pin selection is aided by large spreadsheets with signal names, I/O standards, clocking types, interface, and so on.
3. Drawing schematics is a fully manual process. The FPGA symbol has to be created, and then the FPGA pins have to be connected up to the interface pins. To avoid expensive mistakes, all of the pins have to be correctly connected. The configuration and power supply pins have to be connected as well.

Taray, which brought you the Xilinx® Memory Interface Generator, has developed a new tool called 7Circuits. 7Circuits solves these problems in an innovative way.

7Circuits

7Circuits is a highly intuitive tool that not only selects all of the FPGA pins but also generates PC board schematics for the FPGA and its interfaces.

7Circuits solves FPGA pin allocation problems algorithmically after considering the different constraints. At a higher level, the constraints that the tool considers are:

- Physical constraints. An example of a physical constraint is the physical placement of the FPGA and the interfaces on the PC board.

- Electrical constraints. I/O voltage levels, use of DCI termination, and I/O signaling standards form the electrical constraints.
- Logical constraints. The logical constraints are derived from the interface protocol. For example, if the FPGA is interfacing to a DDR2 memory, the DDR2 protocol will dictate the logical constraints of the interface.
- User preferences. You can tune the performance features of 7Circuits to achieve optimal results.
- FPGA. The location, type, and number of I/Os are among some of the parameters considered.

7Circuits comes with a board view on startup. You begin by placing the FPGA on the board. Next, you place the different components with which the FPGA interfaces. The FPGA and all of the components are shown to scale. The components should be located correctly with respect to the FPGA and the placement should be identical to the actual board placement. An example of the component and FPGA placement is shown in Figure 1.

7Circuits supports a large blend of standard components that you can select and place on the board. If a particular component is not already supported, 7Circuits provides a simple user interface to create the custom interface (alternately, Taray can help you create the interface). Defining the interface component correctly is key to the generation of correct outputs.

7Circuits can block off the pins selected outside the tool. Reading a UCF file with the pin location constraints supports this functionality. 7Circuits can also generate interfaces incrementally. In other words, you can open a saved project and add more interfaces to it without disturbing the existing connections.

If you want to use specific banks for certain interfaces, you can make 7Circuits do it. You can also specify the percentage of pins to be used within each bank. This enables 7Circuits to be customized for any requirement.

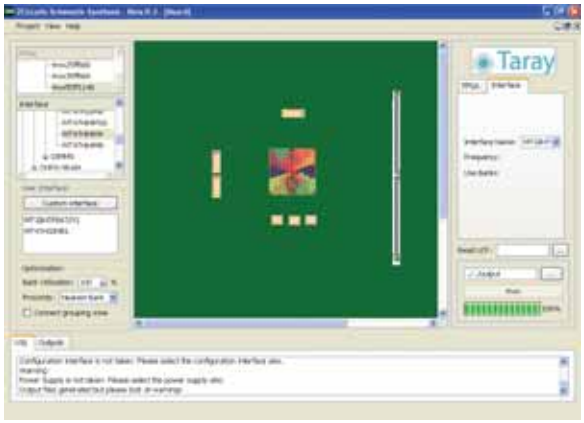

Figure 1 – Placement of the FPGA and interface components on the board

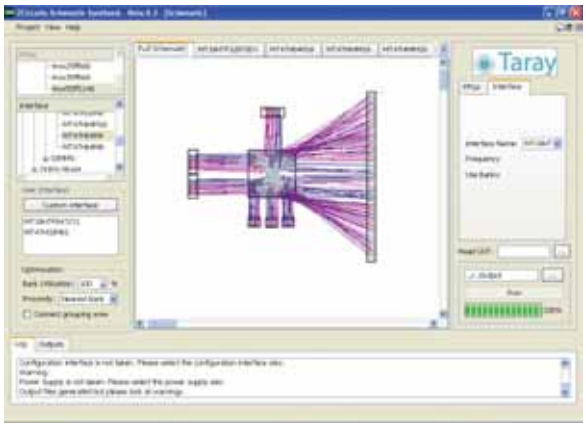

Figure 2 – A ratsnest view of the connections determined by 7Circuits

7Circuits goes through multiple optimization phases to select the pins optimally. After running through different optimization phases, 7Circuits displays the ratsnest connections to enable you to view any bowtie effects. Such interactive output at this stage is a key enabler to optimal results. You can try out different placements or different optimization options within 7Circuits to improve the bowtie effects. An example of the ratsnest is shown in Figure 2.

7Circuits produces a UCF file for pin locations; an EDIF schematics file for the FPGA, interface symbols, and schematics; and a top-level RTL file with all interface port declarations.

Key Advantages

7Circuits produces results with a holistic understanding of the problem space. This makes 7Circuits the first tool to bring system-level understanding into the FPGA solution. By doing so, 7Circuits comes up with the most optimal solution for pinout.

7Circuits reduces the time it takes to create an FPGA-based board from weeks to hours. The pinouts are very dependant on placement. In the current mode of operation, you do not have the luxury of trying out different placements to optimize results. Each placement and generation of the corresponding pinouts is at least a three-man-week task. This makes it impossible for you to try out various placements. With 7Circuits, you can try out four to five different placements and decide on the best placement within a few hours.

7Circuits offers you the added benefit of generating schematics for all of the mundane connections automatically. This task not only saves time, but also ensures correctness.

Here are some of the key advantages of using 7Circuits:

- 7Circuits connects all of the interface pins correctly. In addition, it connects up the power supplies to the right voltage levels.
- It connects Vref pins to the correct voltage levels depending on the I/O standard used.
- It reserves Vrp/Vrn pins when DCI is used. If DCI is used, the Vrp/Vrn pins are connected to the appropriate voltage levels.
- All configuration modes such as JTAG, slave serial, and master serial are supported. The connections are made automatically.

Because most of the mistakes are made in the unexciting and routine connections, the schematics are of a great benefit. They save greater than three man weeks of time and, more importantly, ensure correctness.

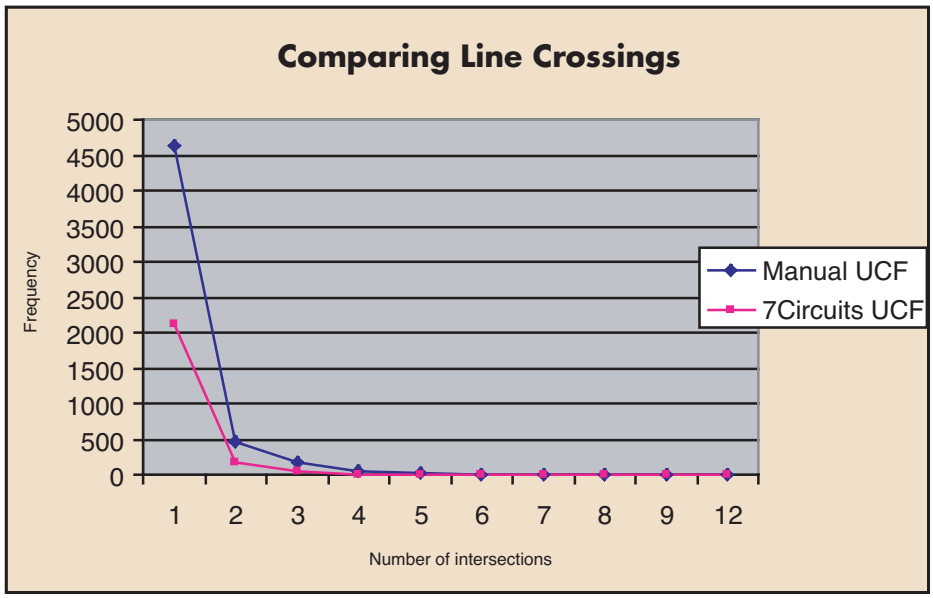

Figure 3 – Bowtie effects are significantly reduced, thus simplifying layout and reducing PCB layers.

Technology

The key to producing effective results is in the algorithms and the technology behind the tool. 7Circuits uses patent-pending technology to solve the issues identified in this article. Here are some of the key innovations in 7Circuits:

- Identifying and representing information. 7Circuits requires physical as well as architectural information on every interface and protocol. All of this information has been precisely identified for the components already supported. For new components, the tool provides a simple and intuitive GUI for you to give this information.
- Special signals are correctly identified and represented so that these signals can be associated to special pins. One example is the Xilinx RocketIO™ pins.
- 7Circuits also considers the logical and architectural aspects. Pins that are logically related will be placed together. This ensures quicker design convergence through the synthesis and PAR phases.
- 7Circuits constantly monitors the number of wire crossings and minimizes them, minimizing the number of board layers. This is key to reducing manufacturing costs.
- Length matching. Various heuristic algorithms are applied to reduce the delta length of signals that are to be length-matched. Applying these algorithms early on avoids long traces on the board. This improves signal quality and enables the PC board router to converge faster.

Results

7Circuits has been going through beta trials since Fourth Quarter 2005. Some of our customers have successfully laid out the board using our outputs.

Additionally, we have tested our results with many Xilinx reference designs. Our test process is as follows:

1. Generate a design for the same interfaces as the standard Xilinx reference board using 7Circuits.
2. Compare the ratsnest of the reference design against the ratsnest from the tool. In all cases, we found that 7Circuits produced a lower bowtie than the reference design.
3. Use the UCF generated by the tool and go through synthesis, build, map, PAR, and bitgen. Ensure that timing results from 7Circuits' UCF meet the reference design requirements.

Figure 3 shows an analytical comparison of the results for a memory reference board. The board has a Xilinx FPGA and interfaces with two DDR2 SDRAM DIMMs. This makes a 144-bit-wide interface. It also interfaces with DDR2 components to make a 24-bit-wide interface. The figure charts the frequency of line crossings against the number of line crossings. These comparisons clearly show the efficiency of the tool:

1. The original number of line crossings was 5,337. The line crossings with 7Circuits were reduced to 2,339 – a reduction of more than 50%.
2. There are 4,600 lines that cross each other manually. With 7Circuits, only 2,050 lines cross each other (1 point crossing each other).

Conclusion

Taray is committed to ensuring your success through the use of 7Circuits. Having created the Memory Interface Generator for Xilinx FPGAs, Taray's engineers have the depth of experience required to understand the issues facing you.

We are planning rich feature sets for future releases of 7Circuits, including:

- Schematics. 7Circuits will generate Orcad and DxDesigner schematics natively.
- Symbols. 7Circuits will be able to use symbols from your symbol library. Additionally, 7Circuits will also be able to use fractured (split) symbols to ensure that the schematics are consistent with your company standards.
- Parts. 7Circuits will support other Xilinx FPGA families and support more interface components.
- 7Circuits will offer a verification mode. This will be a great feature for you to check that your files are consistent and that your choices are optimal. You will be able to make incremental changes to improve your results.

A demo version of the 7Circuits tool is available at *www.tarayinc.com*. Revision 1.0 will be released in Second Quarter 2006.

Operating a NAND Flash Device Through an FPGA

You can use FPGAs to access mass storage with NAND Flash.

by Ryan Fisher
Applications Engineer
Micron Technology, Inc.
rfisher@micron.com

As product capabilities continue to expand, so does the demand for high-density static memory storage. NAND Flash is being used for media storage in a large number of systems, including digital cameras, USB stick drives, and portable music players. NAND Flash memory is prominently positioned to address these and other device needs and is evolving rapidly to meet the ever-growing demand.

The most direct approach for a host or system to use a NAND device is by using a NAND controller. The NAND controller can be internal; built into the application processor or host; or incorporated in designs as an external, standalone NAND controller chip.

An alternative method involves utilizing FPGA resources that already exist in many systems. With FPGA resources, you can create a state machine to act as a NAND controller.

To demonstrate this principle, Micron Technology has developed a NAND controller using a Xilinx® Spartan™-3 FPGA. In this article, I'll focus on the high-level principles of how NAND Flash devices operate.

NAND Flash – The Basics

Discrete programmable non-volatile memory has traditionally been reserved for operating systems (and a few embedded applications of these systems). NAND Flash devices have changed that by providing memory densities that are orders of magnitude greater than older discrete non-volatile memory. As such, NAND has one of the lowest prices per bit for discrete programmable non-volatile memory. This makes its storage resources available to other applications and functionalities that may have been otherwise too costly to implement using lower density non-volatile memory.

NAND Flash devices employ a state machine control structure to carry out various operations (READ, PROGRAM, ERASE, RESET, READ ID, READ STATUS) common throughout NAND Flash manufacturers (although some specialty commands exist). Because the base set of

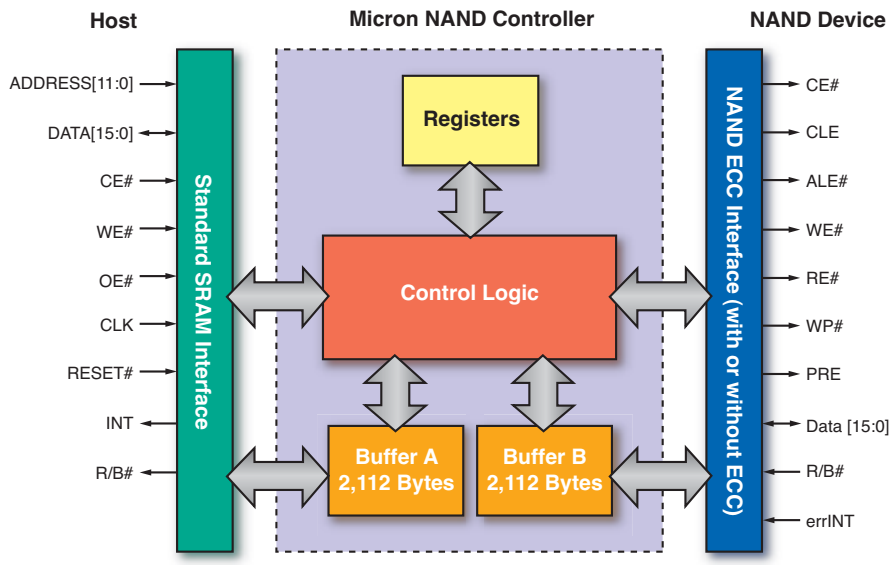

Figure 1 – Block diagram of FPGA NAND controller between host and NAND devices

commands are identical, you have the flexibility to use NAND devices from different manufacturers in the same product.

NAND Flash devices also have a highly multiplexed data/address structure that allows for a low active pin count. NAND devices have not changed their TSOP package pinouts since the introduction of the 64 Mb density, and with NAND devices now at the 8 Gb density range, this allows easy footprint migration from one memory density to another.

The combination of low price per bit, shared command structure, and signal pinouts makes NAND a powerful solution for mass storage use.

Implementing a NAND Controller with an FPGA

Micron has developed a NAND controller built with a Spartan-3 device that sits between the NAND Flash and the system host, which is illustrated in Figure 1. Most systems have an SRAM-type interface, which makes this setup of the FPGA between the NAND device and the host device easy to implement.

The host device, FPGA, and NAND device may be connected point-to-point. A pull-up resistor is needed on the R/B# line on the NAND side to provide the proper functionality for that signal. (See the 2 Gb NAND Flash datasheet at *www.micron.com/products/nand/massstorage/partlist.aspx* for more details on this and other NAND Flash signals.)

Using the FPGA, we are able to perform READ, PROGRAM, ERASE, RESET, READ ID, and READ STATUS operations to the NAND Flash device. Although we used a Spartan-3 FPGA for the implementation of the NAND controller, you could use any number of FPGAs in this process.

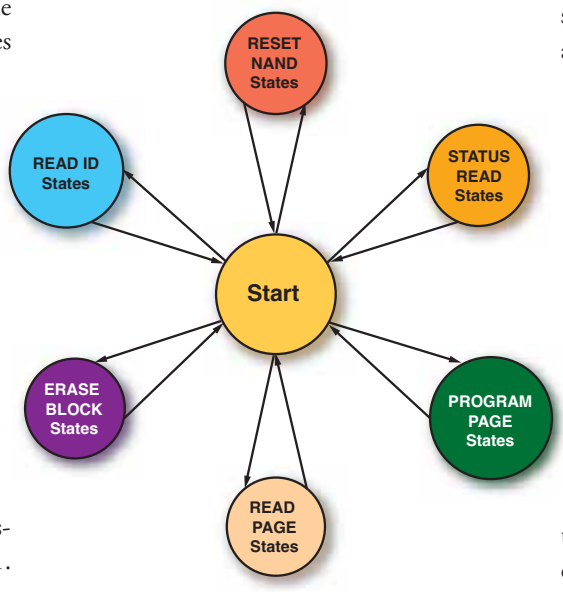

Figure 2 – FPGA NAND controller state machine diagram

FPGA Controller State Machine

NAND Flash devices operate through an internal state machine that controls the various commands of the NAND device (READ, PROGRAM, ERASE, RESET, READ ID, READ STATUS). Many of these commands are common between NAND Flash devices from different manufacturers, allowing for part-to-part compatibility. This allowed us to set up a controller that had the flexibility to communicate with the internal state machine of the NAND Flash device, which could also be used across many NAND device densities independent of the manufacturer.

Figure 2 illustrates the state machine we created on the Xilinx Spartan-3 device to communicate with NAND Flash devices. Beginning with the "start" state, we branched off to the different functions of the NAND controller state machine to operate the NAND Flash. Using these functions, we operated the NAND Flash device as a mass storage device.

Conclusion

Mass storage capabilities are no longer considered "nice to have" in end products; they are becoming a necessary feature. NAND Flash devices are a useful solution to meet the demand for high-density static mass storage. Designs that do not have an embedded or external NAND Flash controller have the disadvantage of not being able to use this resource.

By using FPGA resources already available to many platforms that do not have NAND Flash controllers, you can gain access to this valuable resource by implementing an FPGA NAND Flash controller.

You can find more information about the FPGA NAND Flash controller discussed in this article, including VHDL code, in the technical note "Micron NAND Flash Controller via Xilinx Spartan-3 FPGA (TN-29-06)," available on the Micron website at *www.mircon.com/products/nand/massstorage/technote.html*.

For more information about Micron's NAND Flash products, visit *www.micron.com/products/nand/massstorage/partlist.aspx*.

X-ray vision for your designs

Agilent 16900 Series logic analysis system with FPGA dynamic probe

- *Increased visibility with FPGA dynamic probe*
- *Intuitive Windows® XP Pro user interface*
- *Accurate and reliable probing with soft touch connectorless probes*
- *16900 Series logic analysis system prices starting at $21,000*

Agilent Direct

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

U.S. 1-800-829-4444, Ad# 7909
Canada 1-877-894-4414, Ad# 7910
www.agilent.com/find/new16900
www.agilent.com/find/new16903quickquote

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

The FPGA dynamic probe, when combined with an Agilent 16900 Series logic analysis system, allows you to access different groups of signals to debug inside your FPGA—without requiring design changes. You'll increase visibility into internal FPGA activity by gaining access up to 64 internal signals with each debug pin.

You'll also be able to speed up system analysis with the 16900's hosted power mode—which enables you and your team to remotely access and operate the 16900 over the network from your fastest PCs.

The intuitive user interface makes the 16900 easy to get up and running. The touch-screen or mouse makes it simple to use, with prices to fit your budget. Optional soft touch connectorless probing solutions provide unprecedented reliability, convenience and the smallest probing footprint available. Contact Agilent Direct today to learn more.

Agilent Technologies
dreams made real

Unveiling Nova

Using a single Virtex-II Pro FPGA to create a reprogrammable video switching system.

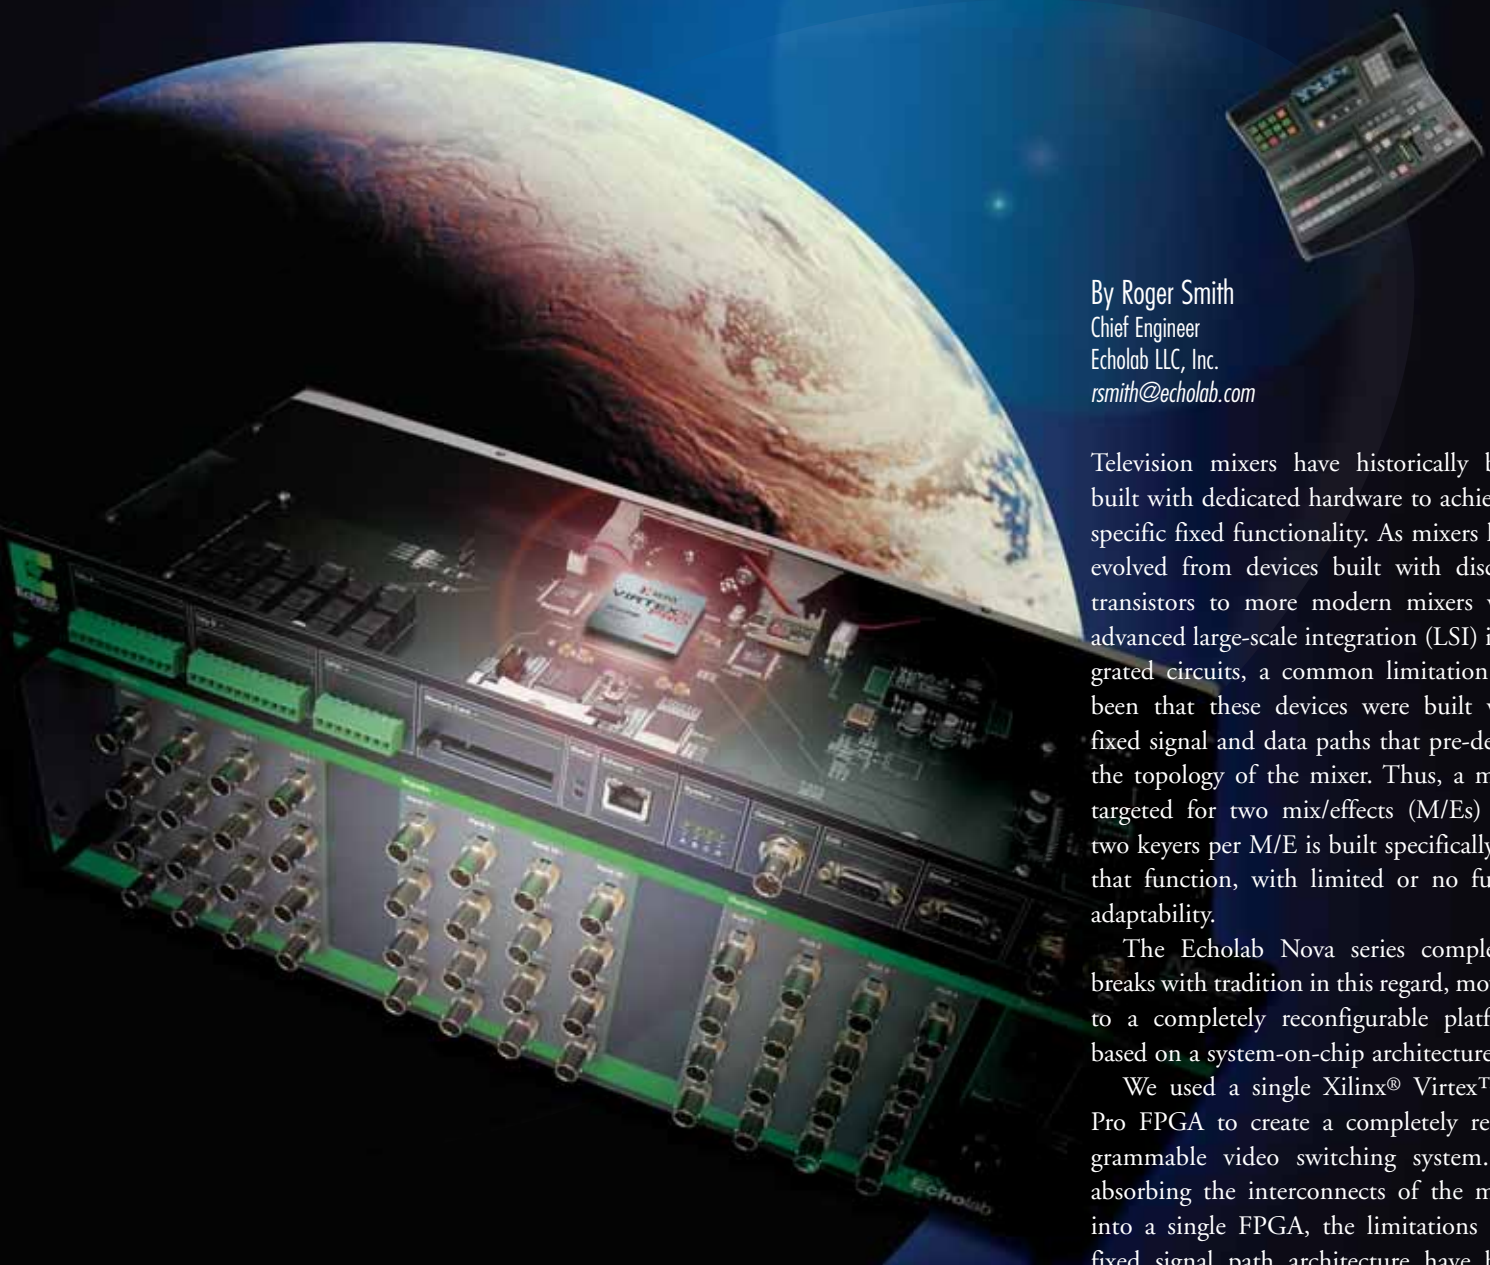

By Roger Smith
Chief Engineer
Echolab LLC, Inc.
rsmith@echolab.com

Television mixers have historically been built with dedicated hardware to achieve a specific fixed functionality. As mixers have evolved from devices built with discrete transistors to more modern mixers with advanced large-scale integration (LSI) integrated circuits, a common limitation has been that these devices were built with fixed signal and data paths that pre-define the topology of the mixer. Thus, a mixer targeted for two mix/effects (M/Es) and two keyers per M/E is built specifically for that function, with limited or no future adaptability.

The Echolab Nova series completely breaks with tradition in this regard, moving to a completely reconfigurable platform based on a system-on-chip architecture.

We used a single Xilinx® Virtex™-II Pro FPGA to create a completely reprogrammable video switching system. By absorbing the interconnects of the mixer into a single FPGA, the limitations of a fixed signal path architecture have been removed, and the topology of the mixer can be redefined again and again throughout the life of the product.

Nova Family

In 2004, Echolab launched the first member of the Nova series, the Nova 1716. This product is a full program/preset mixer with 16 inputs and 16 outputs, two downstream keyers, and a full M/E upstream complete with two effects keyers. In 2005, we introduced several new members of the family. The Nova 1932 is a 32-input program/preset mixer, with two upstream M/Es and a full complement of six keyers. The identity4 is a 1 M/E 16-input look-ahead preview mixer. The identity4 brings tremendous advances in video layering, with four upstream and two downstream keyers and five internal pattern generators, all in a flight-pack-size panel and frame.

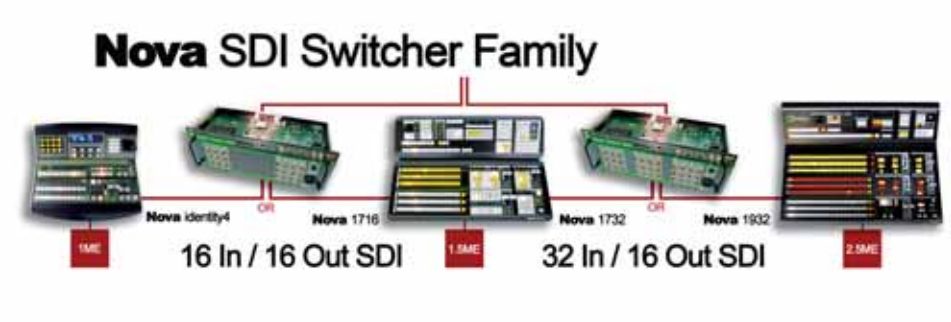

Figure 1 – The Nova switcher family

As shown in Figure 1, what is unique about these mixers is that they share a common frame and electronics, which are simply re-programmed to fit the target mixer design.

I/O

Major advances in FPGA design were necessary to undertake such a dramatic shift in mixer architectures. One of the first challenges that we had to overcome was how to get the video bandwidth into and out of the FPGA.

Even the smallest Nova family member (shown in Figure 2) has 16 SDI inputs and 16 SDI outputs. At 270 Mbps, this is an aggregate video data bandwidth of more than 8.5 Gbps. The Nova 1732 and 1932 family members have 32 inputs and 16 outputs, approaching approximately 13 Gbps.

One of the key aspects of the Virtex-II Pro device is its ability to support I/O bandwidths in excess of 400 MHz on all FPGA I/O pins. Another key feature of Xilinx parts is support for internal termination of high-speed differential inputs. Because of the high-speed nature of this video I/O, the wires must be treated as transmission lines, with great care paid to electrical termination. Historically, these termination resistors would be placed directly outside the chip, as close as possible to the end of the line. The need for a substantial amount of terminations so close to the chip would be a layout complication. Because Xilinx can support internal termination of these high-speed video signals, this has greatly simplified the architecture, making the support of large blocks of high-speed video I/O possible.

The next challenge was deciphering the SMPTE 259M stream. Because the SDI streams are coming from all over the studio, even when genlocked (synchronizing a video source with other television signals to allow the signals to be mixed), there can be large phase shifts of +/- half a line between these signals, and therefore, a common clock cannot be initially used to sample all of the incoming signals.

Figure 2 – Nova SDI I/O

Typically, most SDI mixers use individual clock and data separator circuits on each input so that the hardware can recover the individual bits from each stream. After the data streams are separated and decoded, sync detector circuits are used to write these streams into FIFO memories. A common genlock clock and reference is then used to read out the video streams from the memories for effects processing downstream.

This topology is not viable for Echolab's system-on-chip architecture. The large number of clocks created by these front-end clock and data separators would overwhelm the FPGA's support for the total number of clocks, as well as the inherent need to crowd these parts near the FPGA. The video interface to the FPGA must be simpler, and require few or no parts near the FPGA. Thus, traditional clock and data separation techniques do not work here.

Echolab applied an asynchronous data-recovery technique from Xilinx application note XAPP224 ("Data Recovery") originally developed for the networking market. The technique uses precise low-skew clocks to sample the inputs in excess of a gigahertz. The samples are examined to determine the location of the data bit cell transitions. The encoded data is extracted from the stream and can cross into the genlock clock domain without ever extracting the clock from the input stream. For more information about this technique, see *www.xilinx.com/bvdocs/appnotes/xapp224.pdf.*

The FPGA's ability to route nets and guarantee skew performances on the order of picoseconds has enabled development and implementation of this unique SDI input.

Another design challenge in digital mixers has been the implementation of video line delays and FIFOs, which have historically been used in large numbers to time internal video paths and outputs. It is often necessary to add delay lines to AUX bus outputs to keep them in time with the primary mixer outputs. These delay lines and FIFOs have typically been implemented with discrete memory devices.

The Xilinx FPGA solution contains large numbers of video line-length memo-

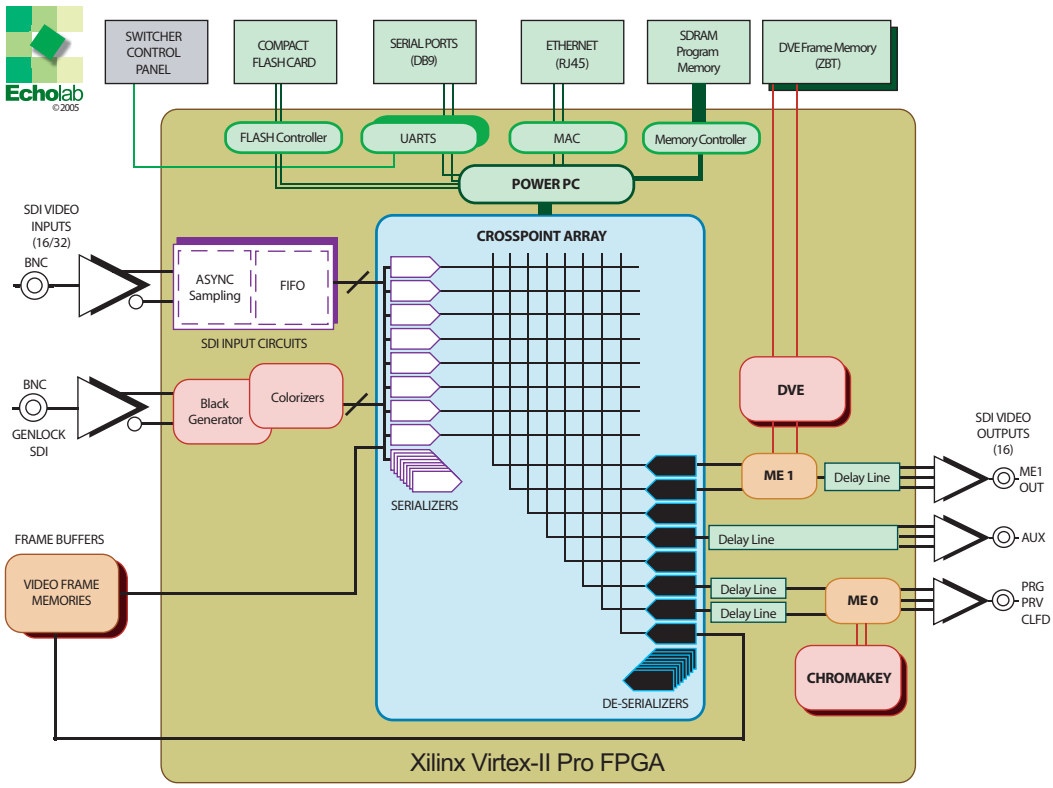

Figure 3 – Virtex-II Pro switcher implementation

ries on chip, which lend themselves naturally to delay lines. Xilinx has enough of these memories on-chip to allow all AUX bus outputs on the Nova series to be timed.

Crosspoint Array

The next challenge was to implement a crosspoint array of sufficient size to support these large mixers without overwhelming the resources of the chip. Although the crosspoint array appears to be a reasonable size from the front panel, the array is often much larger because it must support all of the internal sources and functions inside the mixer.

The Nova 1716 has 16 external inputs, but it also has 3 internal colorizers, black, three internal frame buffers, and several intermediate sources generated by the upstream M/E. For outputs from the array, each M/E requires seven video buses (A/B, key 1 cut and fill, key 2 cut and fill, as well as video for borders). There are also 12 AUX bus feeds and dedicated buses to support capture on the internal frame buffers.

By the time you add it all up, the required internal crosspoint array is easily 30 x 30 for the 16-input mixer. To develop this crosspoint array as a 10-bit wide parallel implementation would consume a large amount of resources within the FPGA.

A more effective use requires the design of high-speed serializers and de-serializers within the part. The Virtex-II Pro FPGA's ability to generate multiples of the genlock clock within the part with low skew allows you to create whole sections of the chip that can run at the SMPTE 259M bit-serial rate of 270 MHz. This implementation of an SDI serial rate crosspoint array (Figure 3) effectively limits the use of valuable FPGA fabric to less than 10 percent of the capacity of the target chip.

Figure 4 – Effects

Effects Generation

Once the streams have gone through the crosspoint array and have been de-serialized, the video buses, now "timed," can be routed to the appropriate processing blocks within the FPGA to perform various video effects (see Figure 4).

The creation of video effects within the FPGA such as wipes, mixes, and keys is easily performed with the basic building blocks of the FPGA. Most basic video effects can be performed with nothing more than simple combinations of addition, subtraction, and multiplication. Hundreds of embedded high-speed multipliers within the FPGA fabric allow a variety of video effects to be performed effortlessly with very high precision.

Embedded memory can also be used for LUTs and filter coefficient storage. An example of an often-used filter would be an interpolating filter for 4:4:4 up-sampling before a DVE or Chromakey.

A high-precision circle wipe requires a square-root function. This complicated mathematical function was implemented with a CORDIC core provided by Xilinx. These blocks of pre-built and tested IP from Xilinx and other third-party developers allow you to rapidly deploy designs without having to develop and test every building block from scratch. Development with these IP cores can remain at a very high level, providing quick time to market. Also, the optimized

size of these cores allows the design of complex mixers within a single chip.

Embedded PowerPC

The computer horsepower necessary to run today's large vision mixers has grown immensely over the last dozen years. The need to accurately support field-rate effects on multiple M/E banks while at the same time communicating with complex control panels and many external devices has tested the limits of earlier 8- and 16-bit microprocessors.

Most manufacturers have either used several smaller distributed processors or a larger, faster 32-bit microprocessor. Echolab's system-on-chip architecture (see Figure 5 – PowerPC™ implementation) takes advantage of two 32-bit PowerPCs (running at 270 MHz) embedded directly in the fabric of the Xilinx FPGA. This tight coupling between the processor and the mixer hardware leads to substantial performance improvements.

A vast library of pre-built processor peripherals from Xilinx and third parties makes it easy for functions like serial ports, memory controllers, and even Ethernet peripherals to be dropped right into the design. Custom peripherals are also easy to design.

Several large switcher manufacturers have gone to commercial operating systems like Windows or Linux to improve their software productivity. Although the early benefits can be appealing, the downside to this transition is a loss of control over the reliability of the switcher's code base, making a device that is not as robust and bulletproof as previous generations of mixers.

Echolab has chosen Micrium's µC/OS-II real-time operating system (RTOS) for the Nova series (see *www.micrium.com*). This OS is a priority-based, pre-emptive multitasking kernel that has been certified for use in safety-critical applications in medical and aviation instruments. Time-critical video processing is assigned to the highest priority task. Management of the file system, console I/O, and network stack are allocated to lower priority tasks, allowing the processor to utilize spare processor cycles in the background without ever interfering with the video hardware.

The tasks communicate with each other – and synchronize their activity – with thread-safe semaphores and message queues provided by the OS.

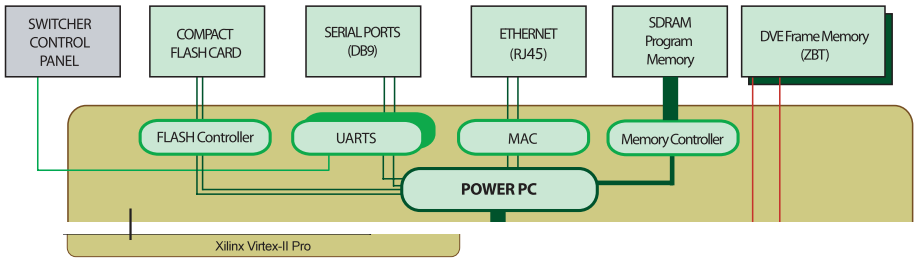

Figure 5 – PowerPC implementation

System Connectivity

All Nova series switchers have multiple ports for broadcast studio interconnectivity. An industry-standard RS-422 port allows for the implementation of industry-standard editing protocols. A standard RS-232 port is available for PC connectivity. An Ethernet port allows the switcher to be directly connected to a network.

Under control of Nova's µC/OS-II, several servers are running concurrently that provide an integrated Web server for remote status and display, an embedded XML-RPC server for remote control, as well as a full TFTP server for remote upload and download of graphics and stills.

Compact Flash and Re-Configurability

At the heart of the Nova system is an industry-standard Compact Flash card (Figure 6), designed to hold all of the firmware and software to configure and boot the Nova. With a user-accessible mode switch, you can load as many as eight different on-line configurations of Nova firmware and software from a single flash card.

The remaining storage on the card is available to store user data such as sequences and panel saves, as well as key memories and other user settings. Also, the Compact Flash is designed to hold all of the graphics and stills online. Support for cards up to 2 GB or more provides an immense storage capacity, well beyond the archaic floppy disks found in competitive products.

Given Nova's unique architecture, it is easy for the product to be extended through downloadable firmware updates. These updates can be as simple as a routine software patch, as complex as adding a keyer to an AUX bus output, or restructuring the internal video flow within an ME for a specialized application. Architectures as different as the 2 ME Nova 1716 (with its program/preset architecture) and the Nova identity4 (with a six-keyer look-ahead preview structure) can be loaded into the same hardware.

Figure 6 – Compact Flash

Roadmap to the Future

Next-generation Virtex-4 FPGA technology from Xilinx will allow Echolab to move its system-on-chip architecture to support high-definition products. Virtex-4 FPGAs bring a

higher level of performance to the embedded logic, as well as the embedded peripherals. Some of these enhancements include:

Fabric enhancements

- Larger arrays
- Faster
- Lower power

I/O enhancements

- General-purpose I/O speeds to 1 GHz
- Dedicated Rocket IO™ transceiver speeds beyond 10 GHz

Dedicated hardware resources

- Multiple tri-mode Ethernet MACs
- 500 MHz multipliers with integrated 48-bit accumulators for DSP functions
- Block memories now have dedicated address generators for FIFO support

Conclusion

Television mixers have grown larger and more complex in the last dozen years. More and more, they are the focal point for the interconnection of a wide range of studio equipment.

One of the primary benefits of the system-on-chip architecture is reduced parts count. This reduction in parts count contributes directly to lower power, reduced PCB complexity, higher reliability, and reduced cost.

Another major benefit of the system-on-chip architecture is that it is almost entirely reconfigurable. This has allowed multiple products with different video architectures to be built on a common platform. This also lends itself to easy customization for specialized applications or specific vertical markets.

As the computer network continues to play more of a role in today's modern television studio, the Nova series will be ready with support for streaming video over Gigabit Ethernet. H.264 and WMV9 codecs will drop right into the Nova's system-on-chip architecture, providing future features on today's hardware.

For more information, please feel free to see Echolab's complete line of television mixers at *www.echolab.com.*

High Definition, Low Bandwidth

Implementing a high-definition H.264 codec solution with a single Xilinx FPGA.

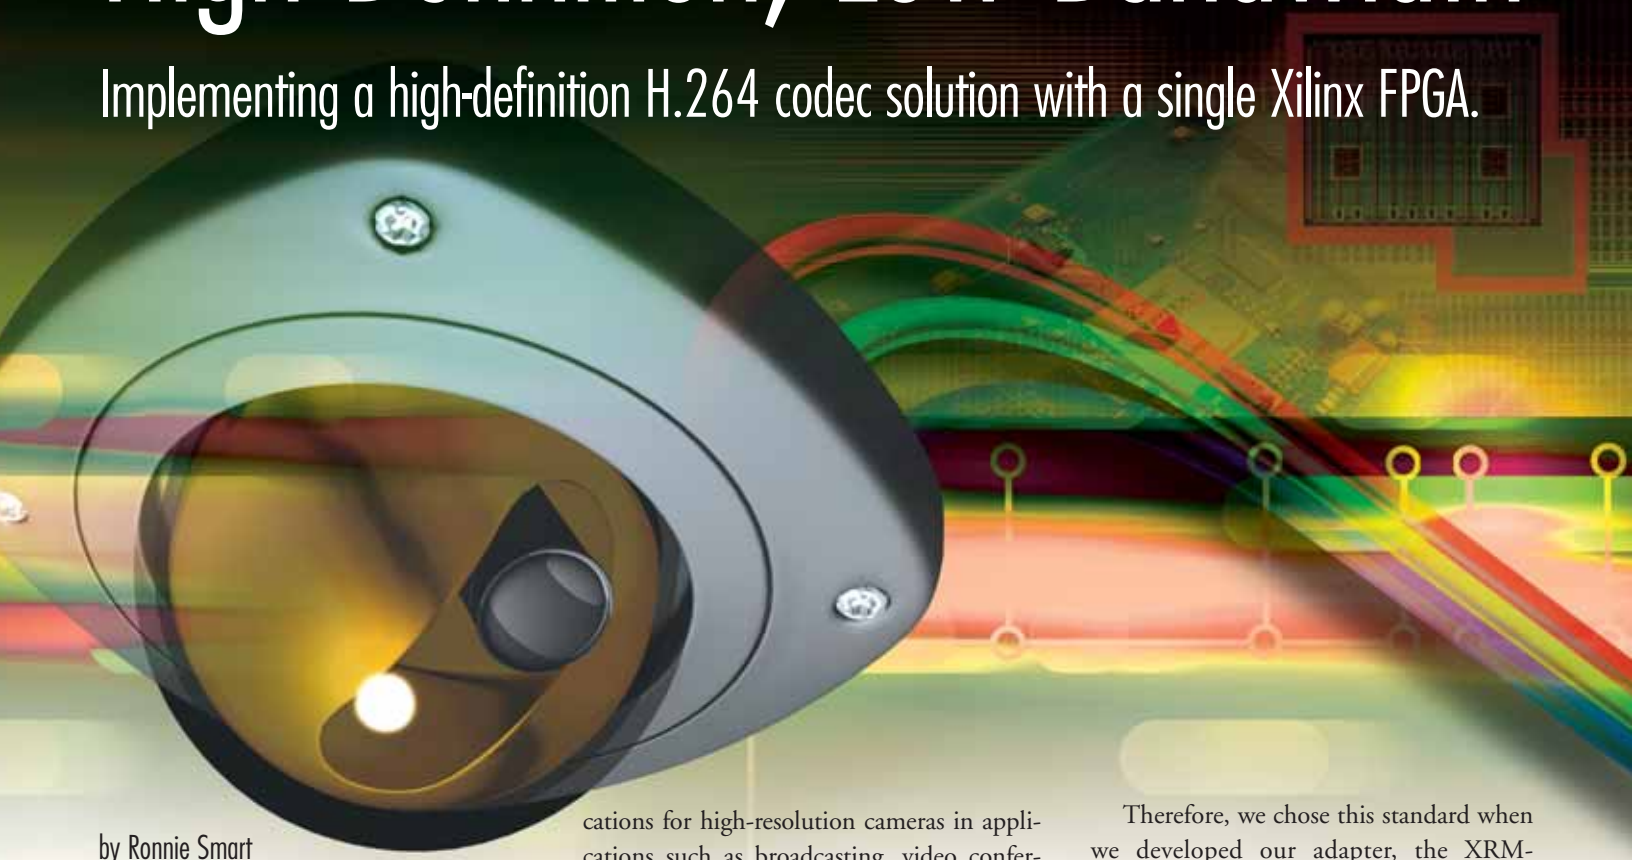

by Ronnie Smart
Software Engineer
Alpha Data Ltd.
rsjs@alpha-data.com

Rob Beattie, Ph.D.
Design Engineer
4i2i Communications Limited
rb@4i2i.com

Ryan Dalzell, Ph.D.
Design Engineer
4i2i Communications Limited
rd@4i2i.com

Andy Ray
Design Engineer
4i2i Communications Limited
ar@4i2i.com

Alpha Data, in conjunction with our partners 4i2i Communications, has implemented one of the world's first high-definition H.264 codec solutions in a single FPGA platform on a single PCI board. This solution gives you the means to develop applications for high-resolution cameras in applications such as broadcasting, video conferencing, video surveillance, and aerospace and defense. The solution uses a Xilinx® Virtex™-4 FPGA board (Alpha Data part number ADM-XRC-4) with an adapter for communication with the camera.

Camera Link Standard

Historically, the image processing and digital video market has suffered from a lack of interconnect standards for cameras and frame grabbers. Various manufacturers developed different connectors and protocols, each requiring different cables and interface logic. However, the Camera Link standard is becoming increasingly widespread, taking away much of the pain of connecting together digital video hardware. Its greatest strengths are:

- It is camera- and frame grabber-independent
- It has high data throughput, using at most two Mini D Ribbon (MDR) connectors

Therefore, we chose this standard when we developed our adapter, the XRM-CAMERALINK. The Automated Imaging Association (AIA) controls the Camera Link standard.

Channel Link

The key component of the Camera Link standard is the Channel Link chip set, a parallel-to-serial transmitter and serial-to-parallel receiver developed by National Semiconductor. Channel Link transmits data bits using the latest high-speed LVDS (low-voltage differential signaling) technology.

A conventional RS-644 LVDS requires at least 56 conductors; a single Channel Link transmits signals using 11 conductors, reducing cable forms and shielding requirements.

The Channel Link transmitter achieves this high transmission rate by serializing – for each clock cycle – 28 bits into 4 LVDS data streams. A fifth LVDS stream transmits the clock signal. The Channel Link receiver de-serializes the data streams back into 28 bits.

Configurations

The standard defines three configurations: base, medium, and full (Figure 1) using one, two, and three Channel Links, respectively. Base configuration is limited to transmitting 28 bits of video data per clock cycle, as it uses only one Channel Link over one MDR cable. You can realize higher transmission rates with the medium configuration (56 bits per clock cycle) and full configuration (84 bits per clock cycle), but they require two MDR cables.

The Camera Link standard states that 24 of the 28 bits transmitted by a Channel Link are image data bits. In base configuration, these bits are split across three 8-bit ports. The standard defines how the image data bits are assigned to the ports (the number and size of pixels transmitted on each clock cycle). Medium and full configurations have six and eight ports, respectively.

The other four bits transmitted by a Channel Link are pixel qualifier signals: frame valid (FVAL), line valid (LVAL), data valid (DVAL), and spare (SPARE). The last is reserved for future use.

Camera Control

All of the configurations include four RS-644 LVDS pairs for camera control (CC1-CC4). The Camera Link standard does not describe the signals used by camera control; therefore, camera manufacturers define their own signals.

Asynchronous Communication

Two RS-644 LVDS pairs (Tx and Rx) are used for asynchronous serial communication between the camera and frame grabber. Frame-grabber manufacturers supply an API implementation for serial communication mandated by the standard.

XRM-CAMERALINK

We designed the XRM-CAMERALINK frame grabber to work with the Virtex-II, Virtex-II Pro, and Virtex-4 series of reconfigurable computing cards. In keeping with the Camera Link standard, it is camera-independent.

The XRM-CAMERALINK has two subsystems: Channel Link receiver and port-to-pixel mappings.

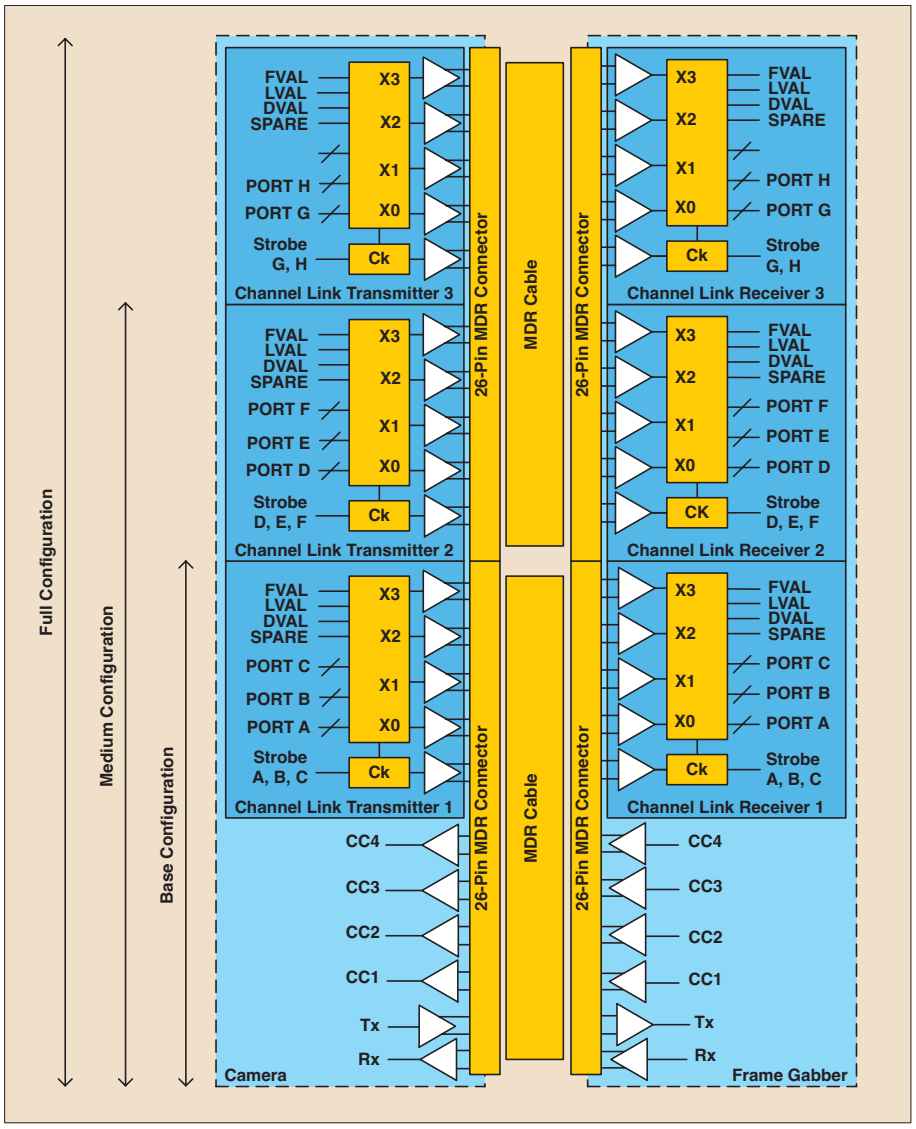

Figure 1 – Camera Link block diagram with the three Channel Links, camera control, and asynchronous serial communication

Channel Link Receiver

The first subsystem of the frame grabber, the Channel Link receiver (Figure 2), deserializes the four incoming LVDS data streams back into the three 8-bit ports and four pixel qualifier signals.

The clock signal is phase-adjusted using the digital clock manager (DCM) so that the incoming LVDS signals are sampled correctly within a data window. The DCM multiplies the clock by 3.5 to produce the pixel clock.

It is not possible to multiply the clock rate by 7 (7 bits need to be extracted from each LVDS stream on each clock cycle), as the clock frequencies would become too great for FPGA implementation. Instead, an independent clock signal is split in the I/O block and fed through a double-data-rate register (DDR), pipeline registers, and decoder to produce even and odd syncs.

Each LVDS bit is also fed through a DDR register and a set of pipeline registers to produce even and odd streams. The streams from all of the LVDS bits are fed through decoders to produce two data streams. The two streams are interleaved through a multiplexer (MUX) with even and odd syncs to form the image data bits and pixel qualifier signals.

Port-to-Pixel Mappings

The second subsystem of the XRM-CAMERALINK translates the image data

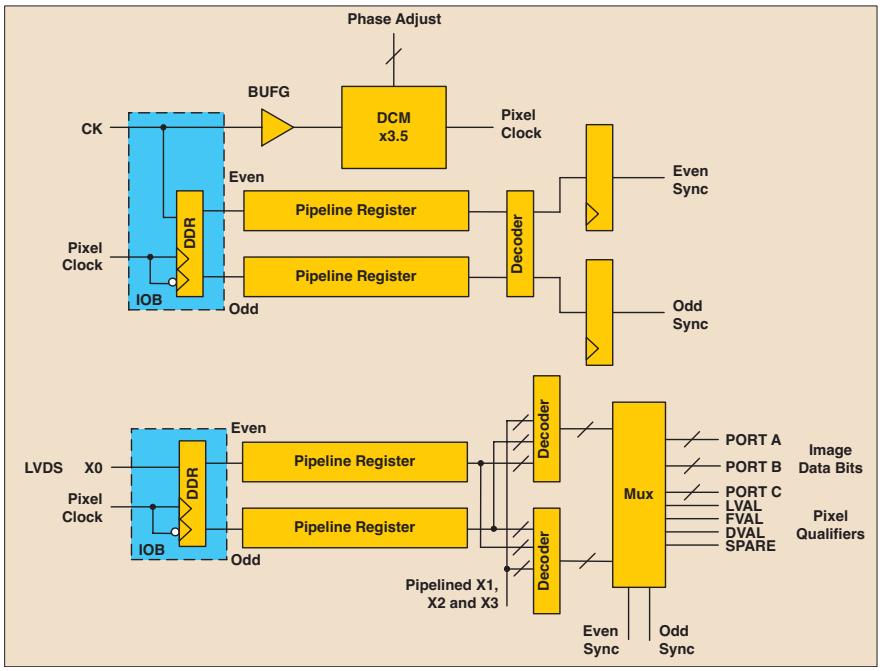

Figure 2 – XRM-CAMERALINK Channel Link receiver (base configuration) block diagram

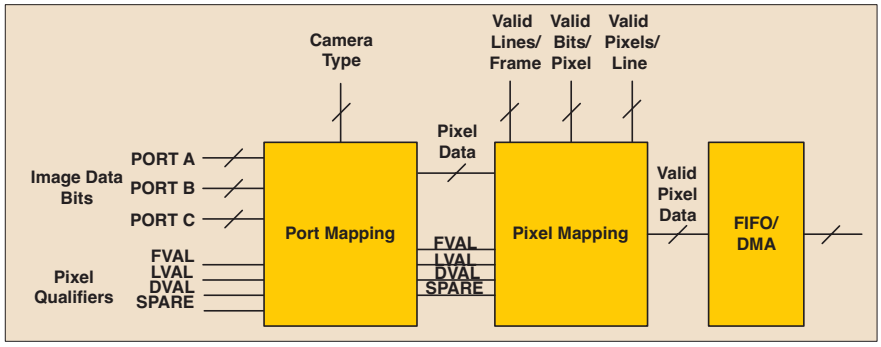

Figure 3 – XRM-CAMERALINK port-and-pixel mappings block diagram

bits and pixel qualifiers into valid pixel data (Figure 3).

The port mappings block allows you to specify (for a particular camera) the number and size of pixels transmitted by the ports for each clock cycle. The image data bits are translated into a stream of pixel data.

The timing references FVAL and LVAL indicate when frames and lines start, respectively. Invalid image bits are indicated by DVAL and are removed in the pixel mappings block.

An additional function of the pixel mappings block is to further clip the frame by defining a region of interest (ROI). Any pixel data outside the ROI is discarded. The ROI is determined by a set of FPGA registers consisting of valid lines/frame, valid bits/pixel, and valid pixels/line. By altering the ROI on successive frames, for example, you can track an object as it moves across the camera's field of view. The output from the pixel mapping block is a 32-bit word that contains valid pixel data and frame and line-count indicators. Your FPGA design will sink this data.

Camera Control and Serial Communication

Control of the camera, which is camera-specific, is achieved through the serial interface on the XRM-CAMERALINK.

Two FIFOs within the FPGA buffer data for serial communication with the camera. On the host, this communication occurs using a dedicated thread, as it is independent of the frame-grabbing logic.

H.264 Video Codec IP Core

Standard-definition (SD) 30 fps video captured using the XRM-CAMERALINK interface comprises raw video data being streamed across the interface. Cameras support a number of different formats; 4i2i Communications has interfaced cameras that provide RGB video at 8 bits per pixel and Bayer 8 bits per pixel. For the RGB video, this corresponds to a raw digital video data rate of about 248 Mbps.

Within the Virtex-4 device, this raw digital video stream is converted to YUV 4:2:0 video at a data rate of about 124 Mbps and is used as the input to the H.264 video encoder. By eliminating redundancy, this configuration reduces the bit rate to a much more manageable data rate of about 64K to 64 Mbps. The compressed video may then be efficiently stored on disk or transferred across a communication network such as the Internet.

We have also successfully used the camera link in connection with 720p high-definition (HD) cameras. Using slice-based encoding, it is possible to compress 720p 60 fps using multiple instantiations of 4i2i's H.264 IP core. About 24,000 slices are required to implement a three-slice 720p 60 fps H.264 encoder in a Virtex-4 device.

For the H.264 baseline codec IP core, 4i2i uses an architecture based on a dedicated macroblock processing pipeline. All of the video standards in common use today work on small 16 x 16 pixel blocks of video data known as macroblocks. (These are 16 x 16 luma samples, together with their corresponding chroma samples.) The algorithms generally require you to perform a sequence of operations in turn on each macroblock. Typically, these operations are prediction, transformation, quantization, and variable length encoding.

The 4i2i approach to codec implementation is to implement these operations as discrete components or processing modules, each of which process one macroblock at a time, separated by paged memory

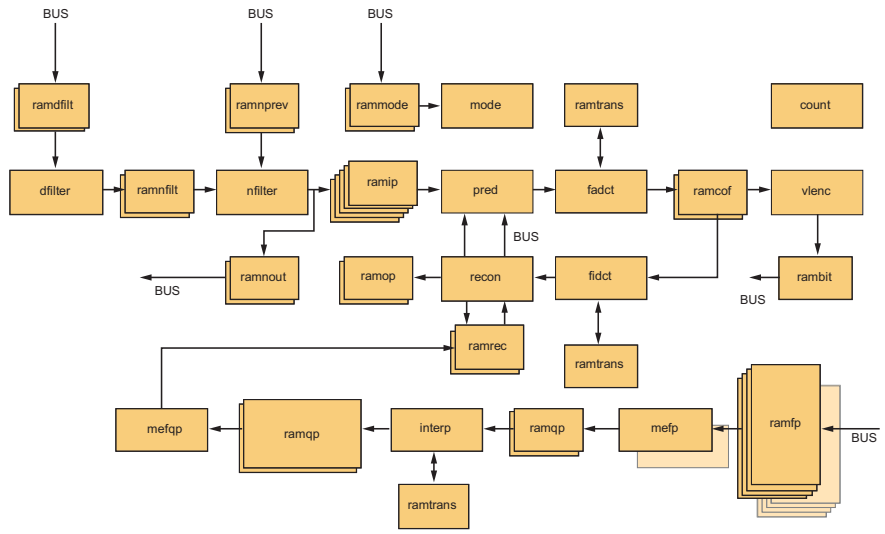

Figure 4 – H.264 encoder IP core architecture

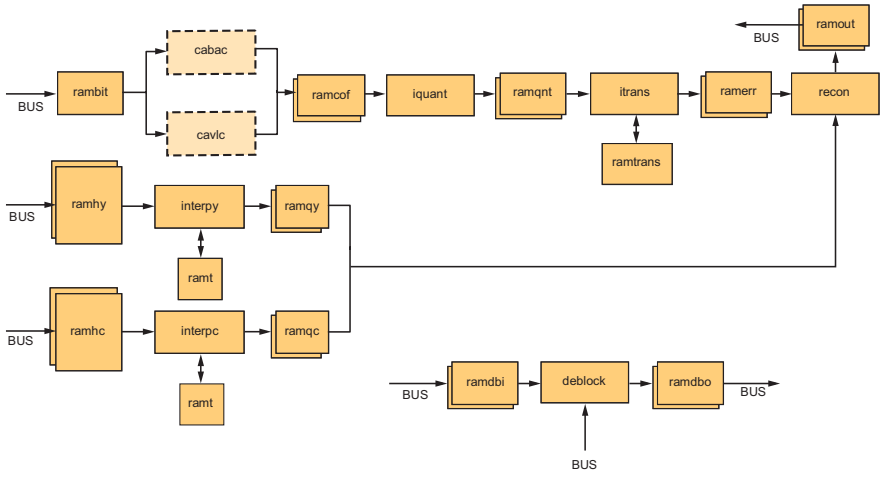

Figure 5 – H.264 decoder IP core architecture

buffers. By altering the number of pages, data scheduling may be affected. Xilinx FPGA devices are well suited to this, as they have an abundance of on-chip small block memories.

This approach has many advantages. First, it does not require any software intervention. Second, it allows all components to operate continuously at maximum throughput. Third, it results in a minimum latency implementation, because all processing operations to produce an encoded bitstream are performed on each macroblock in turn. Fourth, it allows several people to work on the design of the core independently. Finally, combining several such cores for higher throughput is a straightforward process.

An example of using this technique to implement an H.264 encoder and decoder is shown in Figures 4 and 5, respectively. 4i2i has successfully used exactly the same methodology to develop a range of codecs (SDTV and HDTV) from the ITU, ISO, and SMPTE standards at HDTV, all suitable for FPGA implementation.

Using buffer memories in this manner also has several advantages when it comes to system design. The design is independent of the type of external memory used. Several cores can share the same memory and you can use the full memory bus speed.

The compressed video data is then transferred over the PCI interface to the host PC using DMA, where it may be written to disk or streamed over the Internet.

Implementation Considerations

The complexity of modern compression algorithms makes it necessary to take care with the structures used in the RTL code to ensure an efficient use of FPGA resources. For example, writing the RTL to ensure that parametric information is stored in distributed RAM can bring considerable area savings. You can save even more resources by implementing concurrent access to more than one line of data from a macroblock buffer, using the multiple memory ports provided to on-chip RAM blocks rather than storing the data in separate register files.

Conclusion

Camera Link is fast becoming the de facto standard for high-resolution digital frame grabbers, with its high transfer rates and substantial cable reduction. A new connector called the Mini CL is one-third smaller than MDR and further reduces cabling, resulting in smaller cameras.

The XRM-CAMERALINK module provides an all-in-one interface, allowing you to develop applications for any digital camera that conforms to the Camera Link standard.

Camera Link, in connection with a high-performance state-of-the-art video codec such as 4i2i Communications's H.264 IP core, enables huge raw video bandwidth to be reduced to a much more manageable data rate that you can then store on disk or stream across the Internet.

The ADM-XRC-4 features the Virtex-4 FPGA (SX55, LX100, or LX160). Including Camera Link, it supports a range of front-panel adapters for video I/O, such as CCIR, S-Video, and HD-SDI.

For more information about the FPGA development platform or the Camera Link interface, contact Alpha Data at *info@alpha-data.com* or *www.alpha-data.com*. For more information about the H.264 IP core, contact 4i2i Communications at *contact@4i2i.com* or *www.4i2i.com*.

Priced to go.

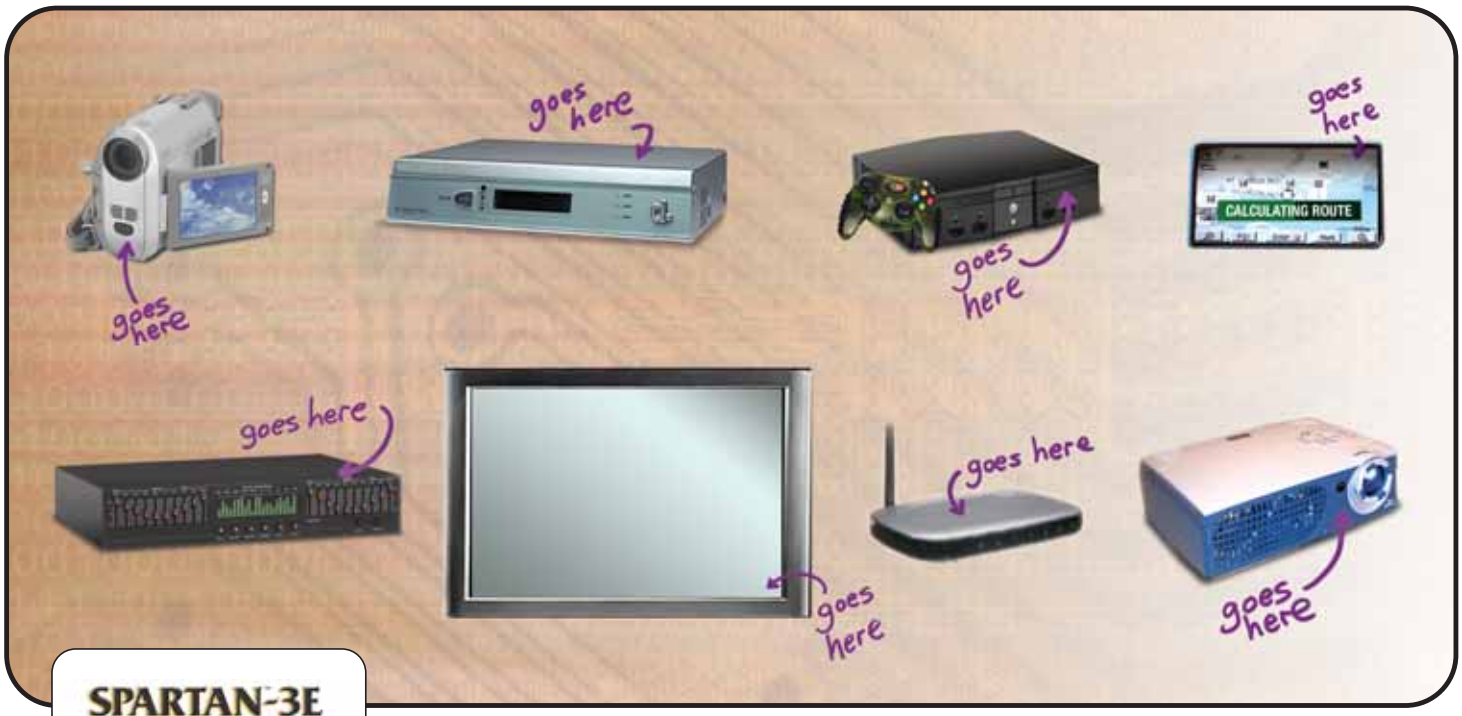

SPARTAN-3E FPGA

The industry's first 100K gate FPGA for under $2.00*

Spartan-3E Platform FPGAs offer an amazing feature set for just $2.00! You get 100K gates, embedded multipliers for high-performance/low-cost DSP, plenty of RAM, digital clock managers, and all the I/O support you need. All this in production now with a density range up to 1.6 million gates.

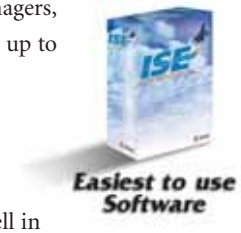

Easiest to use Software

Perfect for digital consumer apps and much more!

With the Spartan-3E series, we've reduced the previous unit cost benchmark by over 30%. Optimized for gate-centric designs, and offering the lowest cost per logic cell in the industry, Spartan-3E FPGAs make it easy to replace your ASIC with a more flexible, faster-to-market solution. Compare the value for yourself . . . and get going on your latest design!

Over 10 million Spartan-3 devices shipped!

MAKE IT YOUR ASIC

XILINX®

The Programmable Logic Company℠

For more information visit
www.xilinx.com/spartan3e

* Pricing for 500K units, second half of 2006

Earth Friendly

Pb-free devices available now

Product Selection Matrix

		Virtex-4 LX (Logic)								Virtex-4 SX (Signal Processing)			Virtex-4 FX (Embedded Processing & Serial Connectivity)					
		XC4VLX15	XC4VLX25	XC4VLX40	XC4VLX60	XC4VLX80	XC4VLX100	XC4VLX160	XC4VLX200	XC4VSX25	XC4VSX35	XC4VSX55	XC4VFX12	XC4VFX20	XC4VFX40	XC4VFX60	XC4VFX100	XC4VFX140
CLB Resources	CLB Array (Row x Column)	64 x 24	96 x 28	128 x 36	128 x 52	160 x 56	192 x 64	192 x 88	192 x 116	64 x 40	96 x 40	128 x 48	64 x 24	64 x 36	96 x 52	128 x 52	160 x 68	192 x 84
	Slices	6,144	10,752	18,432	26,624	35,840	49,152	67,584	89,088	10,240	15,360	24,576	5,472	8,544	18,624	25,280	42,176	63,168
	Logic Cells	13,824	24,192	41,472	59,904	80,640	110,592	152,064	200,448	23,040	34,560	55,296	12,312	19,224	41,904	56,880	94,896	142,128
	CLB Flip Flops	12,288	21,504	36,864	53,248	71,680	98,304	135,168	178,176	20,480	30,720	49,152	10,944	17,088	37,248	50,560	84,352	126,336
Memory Resources	Max. Distributed RAM Bits	98,304	172,032	294,912	425,984	573,440	786,432	1,081,344	1,425,408	163,840	245,760	393,216	87,552	136,704	297,984	404,480	674,816	1,010,688
	Block RAM/FIFO w/ECC (18 kbits each)	48	72	96	160	200	240	288	336	128	192	320	36	68	144	232	376	552
	Total Block RAM (kbits)	864	1,296	1,728	2,880	3,600	4,320	5,184	6,048	2,304	3,456	5,760	648	1,224	2,592	4,176	6,768	9,936
Clock Resources	Digital Clock Managers (DCM)	4	8	8	8	12	12	12	12	4	8	8	4	4	8	12	12	20
	Phase-matched Clock Dividers (PMCD)	0	4	4	4	8	8	8	8	0	4	4	0	0	4	8	8	8
I/O Resources	Max Select I/O™	320	448	640	640	768	960	960	960	320	448	640	320	320	448	576	768	896
	Total I/O Banks	9	11	13	13	15	17	17	17	9	11	13	9	9	11	13	15	17
	Digitally Controlled Impedence	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
	Max Differential I/O Pairs	160	224	320	320	384	480	480	480	160	224	320	160	160	224	288	384	448
	I/O Standards	LDT-25, LVDS-25, LVDSEXT-25, BLVDS-25, ULVDS-25, LVPECL-25, LVCMOS25, LVCMOS18, LVCMOS15, PCI33, LVTTL, LVCMOS33, PCI-X, PCI66, GTL, GTL+, HSTL I (1.5V,1.8V), HSTL II (1.5V,1.8V), HSTL III (1.5V,1.8V), HSTL IV (1.5V,1.8V), SSTL2I, SSTL2II, SSTL18 I, SSTL18 II																
DSP Resources	XtremeDSP™ Slices	32	48	64	64	80	96	96	96	128	192	512	32	32	48	128	160	192
Embedded Hard IP Resources	PowerPC™ Processor Blocks	—	—	—	—	—	—	—	—	—	—	—	1	1	2	2	2	2
	10/100/1000 Ethernet MAC Blocks	—	—	—	—	—	—	—	—	—	—	—	2	2	4	4	4	4
	RocketIO™ Serial Transceivers	—	—	—	—	—	—	—	—	—	—	—	0	8	12	16	20	24
Speed Grades	Commercial (slowest to fastest)	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11, -12	-10, -11
	Industrial (slowest to fastest)	-10, -11	-10, -11	-10, -11	-10, -11	-10, -11	-10, -11	-10, -11	-10	-10, -11	-10, -11	-10, -11	-10, -11	-10, -11	-10, -11	-10, -11	-10, -11	-10
	Configuration Memory Bits	4,765,568	7,819,904	12,259,712	17,717,632	23,291,008	30,711,680	40,347,008	51,367,808	9,147,648	13,700,288	22,745,216	4,765,568	7,242,624	13,550,720	21,002,880	33,065,408	47,856,896
	EasyPath™ Cost Reduction Solutions[1]	—	XCE4VLX25	XCE4VLX40	XCE4VLX60	XCE4VLX80	XCE4VLX100	XCE4VLX160	XCE4VLX200	XCE4VSX25	XCE4VSX35	XCE4VSX55	—	XCE4VFX20	XCE4VFX40	XCE4VFX60	XCE4VFX100	XCE4VFX140

Package[1]	Area	MGT[2]	Pins	4VLX15	4VLX25	4VLX40	4VLX60	4VLX80	4VLX100	4VLX160	4VLX200	4VSX25	4VSX35	4VSX55	4VFX12	4VFX20	4VFX40	4VFX60	4VFX100	4VFX140
SF363	17 x 17 mm	—	240	240	240										240					
FF668	27 x 27 mm	—	448	320	448	448	448					320	448		320					
FF1148	35 x 35 mm	—	768			640	640	768	768	768				640						
FF1513	40 x 40 mm	—	960						960	960	960									
FF672	27 x 27 mm	12	352													320 (8)[3]	352 (12)[3]	352 (12)[3]		
FF1152	35 x 35 mm	20	576														448 (12)[3]	576 (16)[3]	576 (20)[3]	
FF1517	40 x 40 mm	24	768																768 (20)[3]	768 (24)[3]
FF1760	42.5 x 42.5 mm	24	896																	896 (24)[3]

Notes: 1. SFA Packages (SF): flip-chip fine-pitch BGA (0.80 mm ball spacing).
FFA Packages (FF): flip-chip fine-pitch BGA (1.00 mm ball spacing).
All Virtex-4 LX and Virtex-4 SX devices available in the same package are footprint-compatible.
2. MGT: RocketIO Multi-Gigabit Transceivers.
3. Number of available RocketIO Muti-Gigabit Transceivers.
4. EasyPath solutions provide conversion-free path for volume production.

Pb-free solutions are available. For more information about Pb-free solutions, visit www.xilinx.com/pbfree.

Important: Verify all data in this document with the device data sheets found at http://www.xilinx.com/partinfo/databook.htm

Product Selection Matrix

	CLB Resources					Memory Resources			DSP	CLK Resources		I/O Features				Speed		PROM	
	System Gates (see note 1)	CLB Array (Row x Col)	Number of Slices	Equivalent Logic Cells	CLB Flip-Flops	Max. Distributed RAM Bits	# Block RAM	Block RAM (bits)	Dedicated Multipliers	DCM Frequency (min/max)	# DCMs	Digitally Controlled Impedance	Number of Differential I/O Pairs	Maximum I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Configuration Memory (Bits)	EasyPath
Spartan-3E Family – 1.2 Volt																			
XC3S100E	100K	16 x 22	960	2,160	1920	15K	4	72K	4	5/326	2	NO	40	108	Single-ended LVTTL, LVCMOS3.3/2.5/1.8/1.5/1.2, PCI 3.3V – 32/64-bit 33/66MHz, PCI-X 100MHz, SSTL I 1.8/2.5, HSTL I 1.8, HSTL III 1.8 Differential LVDS2.5, Bus LVDS2.5, mini-LVDS, RSDS, LVPECL	-4 -5	-4	0.6M	
XC3S250E	250K	26 x 34	2448	5,508	4896	38K	12	216K	12	5/326	4	NO	68	172		-4 -5	-4	1.4M	
XC3S500E	500K	34 x 46	4656	10,476	9312	73K	20	360K	20	5/326	4	NO	92	232		-4 -5	-4	2.3M	
XC3S1200E	1200K	46 x 60	8672	19,512	17344	136K	28	504K	28	5/326	8	NO	124	304		-4 -5	-4	3.8M	
XC3S1600E	1600K	58 x 76	14752	33,192	29504	231K	36	648K	36	5/326	8	NO	156	376		-4 -5	-4	5.9M	
Spartan-3 and Spartan-3L Families – 1.2 Volt (see note 2)																			
XC3S50		50K	16 x 12	768	1,728	1,536	12K	4	72K	4	24/280	2	YES	56	124	Single-ended LVTTL, LVCMOS3.3/2.5/1.8/1.5/1.2, PCI 3.3V – 32/64-bit 33MHz, SSTL2 Class I & II, SSTL18 Class I, HSTL Class I, III, HSTL1.8 Class I, II & III, GTL, GTL+ Differential LVDS2.5, Bus LVDS2.5, Ultra LVDS2.5, LVDS_ext2.5, RSDS, LDT2.5, LVPECL	-4 -5	-4	.4M	
XC3S200		200K	24 x 20	1,920	4,320	3,840	30K	12	216K	12	24/280	4	YES	76	173		-4 -5	-4	1.0M	
XC3S400		400K	32 x 28	3,584	8,064	7,168	56K	16	288K	16	24/280	4	YES	116	264		-4 -5	-4	1.7M	
XC3S1000	XC3S1000L	1000K	48 x 40	7,680	17,280	15,360	120K	24	432K	24	24/280	4	YES	175	391		-4 -5	-4	3.2M	
XC3S1500	XC3S1500L	1500K	64 x 52	13,312	29,952	26,624	208K	32	576K	32	24/280	4	YES	221	487		-4 -5	-4	5.2M	✔
XC3S2000		2000K	80 x 64	20,480	46,080	40,960	320K	40	720K	40	24/280	4	YES	270	565		-4 -5	-4	7.7M	✔
XC3S4000	XC3S4000L	4000K	96 x 72	27,648	62,208	55,296	432K	96	1,728K	96	24/280	4	YES	312	712		-4 -5	-4	11.3M	✔
XC3S5000		5000K	104 x 80	33,280	74,880	66,560	520K	104	1,872K	104	24/280	4	YES	344	784		-4 -5	-4	13.3M	✔

Note: 1. System Gates include 20-30% of CLBs used as RAMs.
2. Spartan-3L devices offer reduced quiescent power consumption. Package offerings may vary slightly from those offered in the Spartan-3 family. See Package Selection Matrix for details.

Package Options and UserI/O[1]

		Spartan-3E (1.2V)							Spartan-3 (1.2V)								
			XC3S100E	XC3S250E	XC3S500E	XC3S1200E	XC3S1600E		XC3S50	XC3S200	XC3S400	XC3S1000	XC3S1500	XC3S2000	XC3S4000	XC3S5000
Pins	Area[2]	I/O's	108	172	232	304	376		I/O's	124	173	264	391	487	565	712	784
PQFP Packages (PQ) – wire-bond plastic QFP (0.5 mm lead spacing)																	
208	30.6 x 30.6 mm			158	158					124	141	141					
240	34.6 x 34.6 mm																
VQFP Packages (VQ) – very thin QFP (0.5 mm lead spacing)																	
100	16.0 x 16.0 mm		66	66						63	63						
TQFP Packages (TQ) – thin QFP (0.5 mm lead spacing)																	
144	22.0 x 22.0 mm		108	108						97	97	97					
Chip Scale Packages (CP) – wire-bond chip-scale BGA (0.5 mm ball spacing)																	
132	8 x 8 mm			92	92					89							
Chip Scale Packages (CS) – wire-bond chip-scale BGA (0.8 mm ball spacing)																	
144	12 x 12 mm																
FGA Packages (FT) – wire-bond fine-pitch thin BGA (1.0 mm ball spacing)																	
256	17 x 17 mm			172	190	190					173	173	173				
FGA Packages (FG) – wire-bond fine-pitch BGA (1.0 mm ball spacing)																	
256	17 x 17 mm																
320	19 x 19 mm				232	250	250					221	221	221			
400	21 x 21 mm					304	304										
456	23 x 23 mm											264	333	333	333		
484	23 x 23 mm						376										
676	27 x 27 mm												391	487	489	489	
900	31 x 31 mm														565	633	633
1156	35 x 35 mm															712	784
BGA Packages (BG) – wire-bond standard BGA (1.27 mm ball spacing)																	
256	27 x 27 mm																

Notes: 1. Numbers in table indicate maximum number of user I/Os.
2. Area dimensions for lead-frame products are inclusive of the leads.

Pb-free solutions are available. For more information about Pb-free solutions visit *www.xilinx.com/pbfree*.

Platform Flash Features

	XCF01S	XCF02S	XCF04S	XCF08P	XCF16P	XCF32P
Density	1 Mb	2 Mb	4 Mb	8 Mb	16 Mb	32 Mb
JTAG Prog	✔	✔	✔	✔	✔	✔
Serial Config	✔	✔	✔	✔	✔	✔
SelectMap Config				✔	✔	✔
Compression				✔	✔	✔
Design Revisions				✔	✔	✔
VCC (V)	3.3	3.3	3.3	1.8	1.8	1.8
VCCO (V)	1.8 – 3.3	1.8 – 3.3	1.8 – 3.3	1.5 – 3.3	1.5 – 3.3	1.5 – 3.3
VCCJ (V)	2.5 – 3.3	2.5 – 3.3	2.5 – 3.3	2.5 – 3.3	2.5 – 3.3	2.5 – 3.3
Clock(MHz)	33	33	33	40	40	40
Packages	VO20	VO20	VO20	FS48 VO48	FS48 VO48	FS48 VO48
Pb-Free Pkg	VOG20	VOG20	VOG20	FSG48 VOG48	FSG48 VOG48	FSG48 VOG48
Availability	Now	Now	Now	Now	Now	Now

Important: Verify all data in this document with the device data sheets found at http://www.xilinx.com/partinfo/databook.htm

For the latest information and product specifications on all Xilinx products, please visit the following links:

FPGA and CPLD Devices
www.xilinx.com/devices/

Packaging
www.xilinx.com/packaging/

Development Reference Boards
www.xilinx.com/board_search/

Platform Flash
www.xilinx.com/products/silicon_solutions/proms/pfp/

Configuration and Storage Systems
www.xilinx.com/configsolns/

Software
www.xilinx.com/ise/

IP Reference
www.xilinx.com/ipcenter/

Global Services
www.xilinx.com/support/gsd/

XILINX®
The Programmable Logic Company℠

Xilinx Spartan™-3 FPGAs and Platform Flash

www.xilinx.com/devices/

Product Selection Matrix – CoolRunner™ Series

		System Gates	Macrocells	Product Terms per Macrocell	Input Voltage Compatible	Output Voltage Compatible	I/O Features: Maximum I/O	I/O Features: I/O Banking	Speed: Min. Pin-to-pin Logic Delay (ns)	Speed: Commercial Speed Grades (fastest to slowest)	Speed: Industrial Speed Grades (fastest to slowest)	Speed: IQ Speed Grade	Clocking: Global Clocks	Clocking: Product Term Clocks per Function Block
	CoolRunner-II Family – 1.8 Volt													
CoolRunner-II	XC2C32A	750	32	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	33	2	3.8	-4 -6	-6	-6	3	17
	XC2C64A	1,500	64	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	64	2	4.6	-5 -7	-7	-7	3	17
	XC2C128	3,000	128	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	100	2	5.7	-6 -7	-7	-7	3	17
	XC2C256	6,000	256	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	184	2	5.7	-6 -7	-7	-7	3	17
	XC2C384	9,000	384	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	240	4	7.1	-7 -10	-10	-10	3	17
	XC2C512	12,000	512	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	270	4	7.1	-7 -10	-10	-10	3	17
	CoolRunner XPLA3 Family – 3.3 Volt													
CoolRunner XPLA3	XCR3032XL	750	32	48	3.3/5	3.3	36		5	-5 -7 -10	-7 -10	-10	4	16
	XCR3064XL	1,500	64	48	3.3/5	3.3	68		6	-6 -7 -10	-7 -10	-10	4	16
	XCR3128XL	3,000	128	48	3.3/5	3.3	108		6	-6 -7 -10	-7 -10	-10	4	16
	XCR3256XL	6,000	256	48	3.3/5	3.3	164		7.5	-7 -10 -12	-10 -12	-12	4	16
	XCR3384XL	9,000	384	48	3.3/5	3.3	220		7.5	-7 -10 -12	-10 -12	-12	4	16
	XCR3512XL	12,000	512	48	3.3/5	3.3	260		7.5	-7 -10 -12	-10 -12	-12	4	16

Package Options and User I/O

Pins	Area[1]	CoolRunner-II: XC2C32A	XC2C64A	XC2C128	XC2C256	XC2C384	XC2C512	CoolRunner XPLA3: XCR3032XL	XCR3064XL	XCR3128XL	XCR3256XL	XCR3384XL	XCR3512XL
QFN Packages (QFG) – quad flat no-lead (0.5 mm lead spacing)													
32	5 x 5 mm	21											
48	7 x 7 mm		37										
PLCC Packages (PC) – wire-bond plastic chip carrier (1.27 mm lead spacing)													
44	17.5 x 17.5 mm	33	33					36	36				
PQFP Packages (PQ) – wire-bond plastic QFP (0.5 mm lead spacing)													
208	30.6 x 30.6 mm				173	173	173				164	172	180
VQFP Packages (VQ) – very thin QFP (0.5 mm lead spacing)													
44	12.0 x 12.0 mm	33	33					36	36				
100	16.0 x 16.0 mm		64	80	80				68	84			
TQFP Packages (TQ) – thin QFP (0.5 mm lead spacing)													
144	22.0 x 22.0 mm			100	118	118				108	120	118*	
Chip Scale Packages (CP) – wire-bond chip-scale BGA (0.5 mm ball spacing)													
56	6 x 6 mm	33	45						48				
132	8 x 8 mm			100	106								
Chip Scale Packages (CS) – wire-bond chip-scale BGA (0.8 mm ball spacing)													
48	7 x 7 mm							36	40				
144	12 x 12 mm									108			
280	16 x 16 mm										164		
FGA Packages (FT) – wire-bond fine-pitch thin BGA (1.0 mm ball spacing)													
256	17 x 17 mm				184	212	212				164	212	212
FBGA Packages (FG) – wire-bond fine-line BGA (1.0 mm ball spacing)													
324	23 x 23 mm					240	270					220	260

* JTAG pins and port enable are not pin compatible in this package for this member of the family.

Note 1: Area dimensions for lead-frame products are inclusive of the leads.

Product Selection Matrix – 9500 Series

	System Gates	Macrocells	Product Terms per Macrocell	Input Voltage Compatible	Output Voltage Compatible	I/O Features: Maximum I/O	I/O Features: I/O Banking	Speed: Min. Pin-to-pin Logic Delay (ns)	Speed: Commercial Speed Grades (fastest to slowest)	Speed: Industrial Speed Grades (fastest to slowest)	Speed: IQ Speed Grade	Clocking: Global Clocks	Clocking: Product Term Clocks per Function Block
XC9500XV Family – 2.5 Volt													
XC9536XV	800	36	90	2.5/3.3	1.8/2.5/3.3	36	1	5	-5 -7	-7	NA	3	18
XC9572XV	1,600	72	90	2.5/3.3	1.8/2.5/3.3	72	1	5	-5 -7	-7	NA	3	18
XC95144XV	3,200	144	90	2.5/3.3	1.8/2.5/3.3	117	2	5	-5 -7	-7	NA	3	18
XC95288XV	6,400	288	90	2.5/3.3	1.8/2.5/3.3	192	4	6	-6 -7 -10	-7 -10	NA	3	18
XC9500XL Family – 3.3 Volt													
XC9536XL	800	36	90	2.5/3.3/5	2.5/3.3	36		5	-5 -7 -10	-7 -10	-10	3	18
XC9572XL	1,600	72	90	2.5/3.3/5	2.5/3.3	72		5	-5 -7 -10	-7 -10	-10	3	18
XC95144XL	3,200	144	90	2.5/3.3/5	2.5/3.3	117		5	-5 -7 -10	-7 -10	NA	3	18
XC95288XL	6,400	288	90	2.5/3.3/5	2.5/3.3	192		6	-6 -7 -10	-7 -10	NA	3	18
XC9500 Family – 5 Volt													
XC9536	800	36	90	5	5	36		10	-5 -6 -10 -15	-7 -10 -15	-15	3	18
XC9572	1,600	72	90	5	5	72		10	-7 -10 -15	-10 -15	-15	3	18
XC95108	2,400	108	90	5	5	108		10	-7 -10 -15 -20	-7 -10 -15 -20	NA	3	18
XC95144	3,200	144	90	5	5	133		10	-7 -10 -15	-10 -15	NA	3	18
XC95216	4,800	216	90	5	5	166		10	-10 -15 -20	-10 -15 -20	NA	3	18
XC95288	6,400	288	90	5	5	192		10	-10 -15 -20	-15 -20	NA	3	18

Pb-free solutions are available. For more information about Pb-free solutions visit *www.xilinx.com/pbfree*

Package Options and User I/O

		XC9500XV				XC9500XL				XC9500					
Pins	Area[1]	XC9536XV	XC9572XV	XC95144XV	XC95288XV	XC9536XL	XC9572XL	XC95144XL	XC95288XL	XC9536	XC9572	XC95108	XC95144	XC95216	XC95288
PLCC Packages (PC) – wire-bond plastic chip carrier (1.27 mm lead spacing)															
44	17.5 x 17.5 mm	34	34			34	34			34	34				
84	30.2 x 30.2 mm										69	69			
PQFP Packages (PQ) – wire-bond plastic QFP (0.5 mm lead spacing)															
100	23.3 x 17.2 mm										72	81	81		
160	31.2 x 31.2 mm											108	133	133	
208	30.6 x 30.6 mm				168				168					166	168
VQFP Packages (VQ) – very thin TQFP (0.5 mm lead spacing)															
44	12.0 x 12.0 mm	34	34			34	34			34					
64	12.0 x 12.0 mm					36	52								
TQFP Packages (TQ) – thin QFP (0.5 mm lead spacing)															
100	16.0 x 16.0 mm		72	81			72	81			72	81	81		
144	22.0 x 22.0 mm			117	117			117	117						
Chip Scale Packages (CS) – wire-bond chip-scale BGA (0.8 mm ball spacing)															
48	7 x 7 mm	36	38			36	38			34					
144	12 x 12 mm			117				117							
280	16 x 16 mm				192				192						
BGA Packages (BG) – wire-bond standard BGA (1.27 mm ball spacing)															
256	27 x 27 mm								192						
352	35.0 x 35.0 mm													166	192
FBGA Packages (FG) – wire-bond Fine-line BGA (1.0 mm ball spacing)															
256	17 x 17 mm				192				192						

Note 1: Area dimensions for lead-frame products are inclusive of the leads.

XILINX®
The Programmable Logic Company℠

Xilinx Software – ISE 8.1i

www.xilinx.com/ise

ISE™ 8.1i – Powered by ISE Fmax Technology

	Feature	*ISE WebPACK™*	*ISE Foundation™*
Platforms		Microsoft Windows 2000 / XP Red Hat Enterprise Linux 3 (32 bit)	Microsoft Windows 2000 / XP Sun Solaris 2.8 or 2.9 Red Hat Enterprise Linux 3 (32 & 64 bit)
Devices	Virtex™ Series	**Virtex:** XCV50 - XCV600 **Virtex-E:** XCV50E- XCV600E **Virtex-II:** XC2V40 - XC2V500 **Virtex-II Pro:** XC2VP2 - XC2VP7 **Virtex-4:** **LX:** XC4VLX15, XC4VLX25 **SX:** XC4VSX25 **FX:** XC4VFX12 **Virtex Q:** XQV100- XQV600 **Virtex QR:** XQVR300, XQVR600 **Virtex-E Q:** XQV600E	ALL
	Spartan™ Series	**Spartan-II/IIE:** All **Spartan-3:** XC3S50 - XC3S1500 **Spartan-3E:** All **Spartan-3L:** XC3S1000L, XC3S1500L **XA (Xilinx Automotive) Spartan-3:** All	**Spartan-II/IIE:** All **Spartan-3:** All **Spartan-3E:** All **Spartan-3L:** All **XA (Xilinx Automotive) Spartan-3:** All
	CoolRunner™ XPLA3 CoolRunner-II™ CoolRunner-IIA	All	
	XC9500™ Series	All	
Design Entry	Schematic Editor	Yes	
	HDL Editor	Yes	
	State Diagram Editor	Microsoft Windows only	
	Xilinx CORE Generator System™	Yes	
	RTL & Technology Viewers	Yes	
	PACE (Pinout & Area Constraint Editor)	Yes	
	Architecture Wizards	Yes	
	3rd Party RTL Checker Support	Yes	
	Xilinx System Generator for DSP	Sold as an Option	
Embedded System Design	Embedded Design Kit (EDK)	Sold as an Option	
Synthesis	XST - Xilinx Synthesis Technology	Yes	
	Mentor Graphics LeonardoSpectrum	Integrated Interface (EDIF Interface on Linux)	
	Mentor Graphics Precision RTL	Integrated Interface	
	Mentor Graphics Precision Physical	EDIF Interface	
	Synopsys DC-FPGA Compiler	EDIF Interface	
	Synplify/Pro/Premier	Integrated Interface	
	Synplicity Amplify Physical Synthesis	EDIF Interface	
	ABEL	CPLD (Microsoft Windows only)	

	Feature	*ISE WebPACK™*	*ISE Foundation™*
Implementation	FloorPlanner		Yes
	PlanAhead™		Sold as an Option
	Timing Driven Place & Route		Yes
	Incremental Design		Yes
	Timing Improvement Wizard		Yes
	Xplorer		Yes
Programming	IMPACT /System ACE™/ CableServer		Yes
Board Level Integration	IBIS / STAMP / HSPICE* Models		Yes
	ELDO Models* (MGT only)		Yes
Verification	ChipScope PRO™		Sold as an Option
	Graphical Testbench Editor		Microsoft Windows only
	ISE Simulator Lite		Yes
	ISE Simulator		ISE Simulator is available as an Optional Design Tool for ISE Foundation only
	ModelSim® XE III Starter		Yes
	ModelSim XE III		Sold as an option
	Static Timing Analyzer		Yes
	FPGA Editor with Probe		Yes
	ChipViewer		Yes
	XPower (Power Analysis)		Yes
	3rd Party Equivalency Checking Support		Yes
	SMARTModels for PowerPC™ and RocketIO™		Yes
	3rd Party Simulator Support		Yes

*HSPICE and ELDO Models are available at the Xilinx Design Tools Center at ***www.xilinx.com/ise***

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

Yellow Pages

Today's marketplace demands integrated world class solutions. The Xilinx Alliance Program assures you of the best available "total solutions" combining Xilinx programmable logic with key technologies from our Alliance Program members.

All Xilinx Alliance Program members contribute to a broad selection of industry-standard solutions dedicated for use with Xilinx programmable logic. Members provide at least one of the following products or services for Xilinx current device families or development flows:

- **Board-Level Products**
- **DSP Tools**
- **EDA Tools**
- **Integrated Circuits**
- **Design Services**
- **Electronic Components**
- **Embedded Processor Related Products**
- **IP Cores**

Furthermore, Xilinx Alliance Program members are a worldwide resource. Members that supply products have the staffing and infrastructure to support multiple customers from differing geographies at one time.

For a complete description of Xilinx Alliance Program members' products, services, and contact information, please visit our website at *www.xilinx.com/alliance/*.

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
3T BV *www.3t.nl* Design services: specialized in measurement instruments, fail-safe controlling, medical equipment, communication & computer peripherals		✔						
4i2i Communications Ltd. *www.4i2i.com* Video Coding, Modems (OFDM), Error Correction								✔
A2E Technologies, LLC *www.a2etechnologies.com* Verilog, VHDL, DSP, Virtex, Spartan, video, automotive, telecom		✔						✔
Accelchip Inc. *www.accelchip.com* DSP cores that directly implement matrix operations in markets such as Communications, Intelligence, Electronic Warfare, Radar, and Sonar.			✔					✔
Accelerated Technology *www.acceleratedtechnology.com* Accelerated Technology provides products and services worldwide to a wide variety of end-users, OEMs and IC vendors.						✔		

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Accent Srl *www.accent.it* Design services and platform-level intellectual property in the field of telecommunications, multimedia applications, and low-power general computing		✔						
Accord Software & Systems Pvt. Ltd. *www.accord-soft.com* Digital Base band processing		✔						
Addvalue Technologies Ltd. *www.addvaluetech.com* Digital wireless and broadband communications solution provider, focusing on wireless and digital technologies		✔						
ADI Engineering, Inc. *www.adiengineering.com* Solutions for high-speed communications, embedded computing, multimedia, and signal processing products	✔	✔						
Advanced Electronic Designs, Inc. *www.aedmt.com* Design services and manufacturing support specializing in hardware and software design for embedded systems	✔	✔						✔
Advanced Principles Group, Inc. *www.advancedprinciples.com* High-Speed, Complex FPGA designs, and Reconfigurable Computing		✔		✔				
AEE Engenharia Eletronica *www.aee.com.br* Board level design; GPS, GSM, & CDMA interface; A/D & D/A conversion; wireless, sensors and video	✔	✔						
Alatek, Inc. *www.alatek.com* Processors and peripherals					✔			
Algotronix, Ltd. *www.algotronix.com* Implementation of algorithms on FPGAs, Security and Cryptography		✔						✔
Aliathon Ltd. *www.aliathon.com* Telecom and SONET		✔						✔
Alpha Data Parallel Systems Ltd. *www.alpha-data.com* Reconfigurable computing for industry standard platforms PCI, PCI-X, CompactPCI, VME and PMC	✔	✔	✔					✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Altium Ltd. *www.altium.com* PCB/FPGA Design Flow					✔			
Amirix Systems, Inc. *www.amirix.com* Complex FPGAs, high performance boards, embedded software design	✔	✔				✔		
Amphion Semiconductor, Ltd. *www.conexant.com/products/entry.jsp?id=182* DSP, wireless communications and multimedia cores							✔	✔
Andraka Consulting Group, Inc. *www.andraka.com* Exclusively FPGAs for DSP since 1994. Applications include radar, sonar, digital communications, imaging, and video	✔	✔						✔
Apache Design Solutions Virtex-II Pro 3Gb/s RocketIO MGT models for Apache NSPICE					✔			
Apical Limited *www.apical-imaging.com* Video and image processing								✔
ARC International plc *www.arc.com* Configurable RISC/DSP 32 bit processor, peripherals (USB, Ethernet), RTOS, software and development tools						✔		✔
Array Electronics *www.array-electronics.com* DDR SDRAM memory controller, telecommunication, DSP and computational applications, high speed serial interfaces (LVDS, CDR, MGT) Automotive interfaces (MOST, I2C, Flash, SH4 bus) ASIC to FPGA migration	✔	✔						✔
Artaflex Inc. *www.artaflex.com* Design for commercial and consumer products, and technical expertise in digital video, audio, and communications		✔						
ASICS World Service, Ltd. *www.asics.ws* Networking, CPUs, DSPs, DSP functions, interface controllers, storage controllers, encryption/decryption, visualization, etc	✔	✔						✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Aspen Logic, Inc. *www.aspenlogic.com* Verification, design, and implementation of FPGA Logic for cores and systems		✔						✔
Aspire Design Ltd				✔				
Assystem Brime SAS *www.brime-sas.fr* Design services, starting from specifications up to the development, testing, and delivery of electronic and mechanic equipments		✔						
Ateme SA *www.ateme.com* DSP, video codecs	✔							✔
Atmark Techno, Inc. *www.atmark-techno.com* Linux for MicroBlaze, Embedded system designs, network		✔						
Atrenta RTL Linter					✔			
Auspy Development Inc. Design Partitioning					✔			
Avnet Design Services *www.ads.avnet.com* Xilinx FPGA Design Services specializing in the integration of customer developed and third party IP	✔	✔						✔
AVT GmbH *www.avt-ilmenau.de* Application development, FPGA-Introduction of SMEs, Development Kits, Consulting, Graphical Programming	✔	✔	✔	✔				✔
Barco-Silex *www.barco-silex.com* High Speed and High Complexity FPGA designs	✔	✔						✔
BayNet Technology, Inc. *www.baynetinc.net* Sonet, Ethernet, POS, ATM SFI-4, SPI-4, PCI Bus and Intel microprocessors		✔						✔
Bellnix Co., Ltd. *www.bellnix.com* Power supply solutions				✔				
Benchmark Electronics, Inc. *www.bench.com* EMS/JDM team; product design, support, introduction, and launch into volume production		✔						

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Birger Engineering, Inc. *www.birger.com* Image processing, compression algorithms, cameras, and high-speed ASIC/FPGA design		✔						✔
BitSim AB *www.bitsim.com* Design services; DSP, video, 3G, and WLAN designs; high-speed prototyping boards		✔						✔
Bottom Line Technologies, Inc. *www.bltinc.com* FPGA, Product, System design services firm. Networking, Datacom, Telecom, Video, PCI, Signal Processing, Military, Aerospace		✔						
BYO Solutions Design Partitioning					✔			
C & D Technologies, Inc. *www.cdtechno.com* Power supplies	✔	✔		✔				
Calyptech Pty Ltd. *www.calyptech.com* Packet over Sonet, ATM, xDSL		✔						✔
CAST, Inc. *www.cast-inc.com* General purpose IP for processors, peripherals, multimedia, networking, encryption, serial communications and bus interfaces.								✔
Celoxica Ltd. *www.celoxica.com* Development tools for immersed PowerPC processor	✔	✔	✔		✔	✔		
Centrex AS *www.centrex.no* Design services for electronic products, analog and digital hardware/software.		✔				✔		
CES Design Services *www.ces-designservices.com* Platform-oriented SoC/FPGA/PCB and firmware design in telecom, industrial and video-processing. Customized solutions: concept, verification, prototypes	✔	✔		✔		✔		✔
CG-CoreEl Programmable Solutions Pvt. Ltd. *www.cg-coreel.com* System-on-chip designs for telecommunication, networking and DSP	✔	✔		✔				✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
CHESS Embedded Technology *www.chess.nl* Embedded hardware and software design of products and services at the chip, board, and system level		✔				✔		✔
Colorado Electronic Product Design, Inc. *www.cepdinc.com* Designing embedded systems with FPGAs for DSP and interface for medical, consumer, aerospace and millitary	✔	✔				✔		
Comit Systems, Inc. *www.comit.com* SoC, embedded ARM MIPS PowerPC, IBM CoreConnect, high-speed platform FPGA design		✔						
Comm Logic Design, Inc. *www.commlogicdesign.com* Design and consulting services in hardware, FPGA, software and microcode development		✔						✔
Computex Co., Ltd Development tools for immersed PowerPC processor						✔		
Condor Engineering Inc *www.condoreng.com* Avionics network and databus; 1553 bus								✔
Controlnet India Private Limited *www.controlnetindia.com* ASIC/SoC Analog/RF FPGA Design Services, IP Development, domain: Ethernet, IEEE1394, USB, Security, PCI, WLAN, Infiniband	✔	✔						
Convergence Media Solutions, Ltd. *www.convergencelimited.com* Video broadcast and telecoms		✔						✔
Convergent Designs, LLC *www.convergent-design.com* Analog to digital video/audio conversion and processing products		✔						
Corelis, Inc. *www.corelis.com* Development tools for immersed PowerPC processor	✔	✔			✔	✔		
Coreworks, Lda *www.coreworks.pt* DSP, Audio, Video, Communications, Test & Verification, PicoBlaze, MicroBlaze, PPC and EDK		✔						✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Cottonwood Technology *www.ctwdtech.com* Military and commercial designs, complex systems, video systems and networking systems		✔						
CoWare, Inc. *www.coware.com* Hierarchical Block Diagram for DSP			✔		✔	✔		
Crossbow Technologies, Inc. *www.crossbowip.com* System Interconnect IP for on-chip and chip-to-chip multi-processing.	✔							✔
Cyntila Pty. Ltd. *www.cyntila.com* Design and consulting services in Hardware, FPGA, Software and complete turnkey product development services.		✔						
D&T Inc. *www.dntinc.com* Display (Monitor & TV, TFT LCD, PDP, and Projector).		✔						✔
DAIHEN Corporation *www.advanced-ip.jp* VxWorks for Xilinx Virtex-II Pro, MicroBlaze, 8051, and AX1610		✔						
Dark Room Technologies, Inc. *www.darkroom.com* Application development, OS device drivers, PCB design and layout, prototype assembly and test	✔	✔				✔		✔
DATA RESPONS ASA *www.datarespons.com* Embedded solutions within the industrial, military and telecom market place		✔						
Datatek Application, Inc. *www.datatekcorp.com* VHDL or Verilog, specializing in integrated hardware/software solutions		✔						
D-CLUE Technologies CO., Ltd *www.d-clue.com* designs for a wide variety of industrial and consumer electronics applications such as digital cameras, printers, data storage and communications		✔						
Delphi Engineering Group *www.delphieng.com* FPGA Design Services and Products using Xilinx FPGAs	✔	✔						

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Deltatec S.A. *www.deltatec.be* IP and services for digital imaging, DSP, multimedia, PCI and datacom.	✔	✔				✔		✔
Derivation Systems, Inc. *www.derivation.com* IP cores for high assurance hardware/software applications and embedded systems products; 32-bit Java processor core.	✔	✔						✔
Design Gateway Co., Ltd. *www.design-gateway.com* System Design / Development, LSI layout, development/Design of ASIC/FPGA and others		✔						
Design Interface Ltd *www.design-interface.com* Electronic systems design consultants	✔	✔		✔				
Digicom Answers Pty. Ltd. *www.fpga.com.au* High speed, line-rate packet processing to 40Gb/s (OC768); High speed interface bridging; Performance analysis; Test and measurement.		✔						
Digital Communications Technologies, Ltd. *www.dctl.com* Embedded JAVA solutions	✔	✔				✔		✔
Digital Core Design *www.dcd.pl* Microprocessor, microcontroller, floating point and serial communication controller cores		✔				✔		✔
Digital Design Corporation *www.digidescorp.com* Design services and products: specializing in image processing, video, communications, DSP, audio, automotive, controls, interfaces		✔						✔
Dillon Engineering, Inc. *www.dilloneng.com* FFT, DSP and Image Processing in FPGAs		✔						✔
Dolphin Integration *www.dolphin-integration.com* Bus interface, processor, peripheral and DSP cores; EDA software.					✔			✔
Drivven, Inc. *www.drivven.com* 8051, single-board computers, automotive engine control.		✔						✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
DRS Technologies *www.drs.com* Design services specializing in FPGAs for DSP, including the manufacturer of Virtex reconfigurable computing boards		✔						
DSPIA, Inc. *www.dspia.com* DSP, Communications, Networking, and Embedded Microprocessor applications		✔						
Duma Video, Inc. *www.dumavideo.com* Digital video, AES/EBU audio, MPEG and DCT cores; real-time MPEG-2 HD & SD compression	✔	✔						✔
Dynalith Systems Co. Ltd. Dynalith Systems develops and markets specialized hardware and EDA solutions that enable electronic system designers to verify their design in real hardware environments at an early design stage and, thereby significantly reduce time-to-market.	✔	✔			✔	✔		
EASii IC *www.easii-ic.com* Design services for electronic products, analog and digital hardware/software						✔		
Eden Networks, Inc. *www.edennetworks.com* Turn-key design services, specializing in telecom and networking designs (Ethernet, SONET, ATM, VoIP)	✔	✔						✔
Edgewood Technologies, LLC *klfrick@comcast.net* Telecommunications, TDM, LAN, WAN, SDRAM, EDAC, SoC, DSP		✔						
EDS Ireland Limited *www.edsire.com*			✔					
eInfochips, Ltd. *www.einfochips.com* PCI Express designs, High-speed, DSP FPGA designs, Hand held devices design	✔	✔				✔		✔
Electronic Design Associates, Inc. *edadan@comcast.net* Digital Telephony at rates from sub-DS0 through SONET		✔						
Embedded Element Science & Tech Co., Ltd. *www.eestd.com* Embedded System; FPGA; DSP; PCI/CPCI		✔						

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Embedded Systems Design, Inc. *www.esdnet.com* Embedded FPGA, DSP, and Software design		✔						
Enea Epact AB *www.epact.se* Embedded systems, RTOS, DSP, telecom/ datacom, medical, industrial,military, automotive, and consumer product designs.		✔						
Enpirion, Inc. *www.enpirion.com* Power supplies				✔				
Enterpoint Ltd. *www.enterpoint.co.uk* Design Services & Development Boards - Video, Telecom, Military and High Speed design specialities	✔	✔						✔
eSOL Co., Ltd *www.esol.co.jp/english/* RTOS for immersed PowerPC processor/ MicroBlaze SPC						✔		
Eureka Technology *www.eurekatech.com* Silicon proven IP for CPU interface, PCI, PCMCIA, memory control, and other peripheral functions for SoC designs		✔						✔
Evatronix SA *www.evatronix.pl* FPGA design services, HDL modeling (IP on demand, verificationenvironments), hardware & software development for embedded systems		✔						✔
Exalinx, Inc. *exalinx@yahoo.com* Design services; ASIC emulation for high speed communications and wireless applications; reference designs	✔	✔						✔
F.C.D. ASICS *www.fcd.co.il* Memory controllers, video algorithms, cameras, matlab conversions, utopia.		✔						
Faraday Technology Corporation *www.faraday.com.tw* CPU, DSP, MCU, Data Transmission, Peripheral, USB, SATA, Ethernet								✔
Fidus Systems, Inc. *www.fidus.ca* Design Services: FPGAs: communications, video, DSP, interface conversions, ASIC verification. Hardware, PCB layout, software, SI/EMC	✔	✔				✔		

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Fishtail Design Automation Automatic timing constraint generation.					✔			
Flexibilis Oy *www.flexibilis.com* PCI designs, Ethernet designs, Linux device drivers,etc.		✔						
Flextronics Design *www.flextronics.com* Complete product design and development services for high speed communications, reconfigurable computing, imaging, audio/video and military		✔						
Flextronics Design *www.flextronics.com* Complete product design and development services for high speed communications, reconfigurable computing, imaging, audio/video and military		✔						
FPTek, Inc. *www.fptek.com* Consulting and electronic design expertise in FPGA, H/W, DSP, S/W and embedded firmware systems		✔						
Fundacion Robotiker *www.robotiker.com* Design services for Next Generation Network technology and Software Defined Radio, from ASIC/FPGA, HW&SW to turn-key embedded systems.		✔						
GDA Technologies, Inc. *www.gdatech.com* Electronic design services	✔	✔						✔
Gemstone Communications, Inc. *www.gemstone.com.tw* Telecom/Datacom converge ASIC/FPGA Subsystem design		✔						
Global R&D Ltd. *www.global-rd.com* Electronics development for startups		✔						
GV & Associates, Inc. *www.gvassociates.com* Boards and expertise in DSP, processor, software, analog and RF design.	✔	✔						
HARDI Electronics AB *www.hardi.com* Our products are based on Xilinx devices and target ASIC Prototyping	✔				✔			

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
HCL Technologies, Ltd. *www.hcltechnologies.com* High Speed Board, ASIC, FPGA Design & Verification Services in Avionics, Medical, Networking/Telecom, Consumer Electronics		✔						
Helion Technology Limited *www.heliontech.com* Design services and IP for FPGA with a focus on Data Security, Wireless and DSP		✔						✔
iCoding Technology, Inc. *www.icoding.com* Turbo Code related error correction products including system design								✔
id3 Semiconductors *www.id3semiconductors.com* Electronics design, FPGA design	✔	✔				✔		
IDERS, Inc. *www.iders.ca* Contract electronic engineering and EMS resource, providing complete product cycle ranging from specifications to production		✔						
Iflect Technologies India Pvt Ltd *www.iflect.com* FPGA, VLSI, Embedded, Video, Interface, IP		✔						
Image Processing Techniques, Ltd. *www.imageproc.com* Broadcast video and effects	✔	✔						
Image Technology Laboratory Corp. *www.gazogiken.co.jp* Embedded system design, DSP, image-processing and advanced multimedia product design	✔	✔				✔		
Imaging Solutions Group *www.isgchips.com* Custom Camera, Image Processing and Video System Design.		✔						
Indesign, LLC *www.indesign-llc.com* Hardware/firmware design, embedded systems, DSP, FPGA, RF		✔						
Industrial Technology Research Institute *www.itri.org.tw* RF IC Technology, Mixed-Mode IC Technology, Optical IC Technology, Radio Processor Technology, IP and Platform Technology.	✔							✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Innovative Computer Technology electronic systems, board-level designs, FPGAs, and CPLDs. FPGAs have been developed for data communications, Telecommunications, laboratory instrumentation, microprocessor-based boards		✔						
Integnology Corporation *www.integnology.com* High resolution DVI, PCI-X and other high-speed serial interfaces, DDR SDRAM memory interface, and Gigabit Ethernet		✔						
Intelliga Integrated Design, Ltd. *www.intelliga.co.uk* Digital Automotive Networks, Embedded serial communications, Microcontroller design	✔	✔						✔
Intrinsix Corporation *www.intrinsix.com* IP-enabled FPGA & ASIC Design Solutions		✔						✔
Intrinsyc, Inc. *www.intrinsyc.com* High Speed FPGA Interfacing Designs and IP cores and development systems.	✔	✔				✔		✔
IO Technologies A/S *www.iotech.dk* Hard and software technology platforms, WebNet uServer and uLinux, processor systems, engineering services		✔						
iWave Systems Technologies Pvt. Ltd. *www.iwavesystems.com* Embedded hardware and software turnkey design services	✔	✔				✔		✔
JAPS Elektronik AB *www.japs.se* Embedded real-time designs from concept to serial production.		✔						
Jennic, Ltd. *www.jennic.com* ATM, SONET, wireless LAN		✔						
Knowledge Resources GmbH *www.knowres.com* Full service custom electronics designs for high-performance data and video processing devices.	✔	✔						✔
Koos Technical Services, Inc. *www.koostech.com* Communications and Digital Signal Processing.		✔						

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Lauterbach Datentechnik GmbH *www.lauterbach.com* Development tools for immersed PowerPC processor						✔		
Liewenthal Electronics Ltd. *www.liewenthal.ee* Embedded systems design services. Development of software and hardware including Virtex and Spartan FPGAs		✔						
LiveDevices Ltd. (ETAS Group) RTOS for immersed PowerPC processor						✔		
Loarant Corporation *www.loarant.com* Proprietary 16-bit RISC CPU; embedded system design.	✔	✔	✔			✔		✔
Logic Development APS				✔				
Logic Product Development *www.logicpd.com* Full service product development, including Industrial Design, Mechanical, Electrical, and Software Engineering		✔				✔		
LYRtech, Inc. *www.lyrtech.com* DSP/FPGA Development Board: Matlab® Simulink integrated and Turn-key Solutions	✔	✔	✔			✔		✔
M&M Consulting *www.mmcons.com* Design Services: Imaging, Video, Audio designs, and high-density ASIC emulation		✔						
Magma Design Automation, Inc. Electronic design automation (EDA) software that enables chip designers to meet critical time-to-market objectives, improve chip performance and handle multimillion-gate designs					✔			
MainConcept AG *www.mainconcept.com* Video codecs, H.264								✔
Martale Ltd		✔						
Mechatronics Test Equipment (I) Pvt. Ltd. *www.mte-india.com* Algorithm implementation and System design	✔	✔						✔
MEET Ltd. *www.meet-electronics.com* IP cores dedicated to industrial servo motor control.		✔						✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Memondo Graphics *www.memondo.es* Implementation of digital signal processing, communications, and image processing solutions and algorithms over Xilinx FPGAs		✔						
Mentor Graphics Corporation *www.mentorg.com/inventra* Mentor Graphics Mentor Graphics is a technology leader in electronic design automation (EDA), providing software and hardware design solutions that enable companies to develop better electronic products faster and more cost-effectively.		✔			✔	✔		✔
Mercury Computer *www.mc.com* Rapid IO								✔
Micrium RTOS for immersed PowerPC processor/ MicroBlaze SPC						✔		
Microtech International Ltd Sp.z o.o. *www.microtech.com.pl* Hardware, Software development, FPGA design services, IPCORE development	✔	✔	✔			✔		✔
Mikrokrets AS *www.mikrokrets.no* FPGA development, comprising telecom, network communication and embedded systems. PCB development		✔						
Millennium Meshwork Data Systems Co., Ltd. *www.mmds.cn* Wireless communication, RFID, and digital video.		✔						
Millogic Ltd. *www.millogic.com* Design services and synthesizable cores, expertise in PCI, video, imaging, DSP, compression, ASIC verification, communications		✔						✔
MILSTAR *www.milstar.co.il* Design services and manufacturers (industrial and military specs); DSP; communication; video; audio; high-speed boards	✔	✔						
Mind NV *www.mind.be* RTOS for immersed PowerPC processor						✔		
MindTree Consulting Pvt. Ltd. *www.mindtree.com* Wireless, Bluetooth, Storage, Communication, Automotive, Industrial automation, Consumer		✔						✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Misarc s.r.l *www.misarc.com* Architectural analysis and feasibility, hardware & software design, boards prototyping and testing, performance improvement		✔						
Mistral Software Pvt. Ltd. *www.mistralsoftware.com* Product realization and professional design services for embedded product design and development.		✔						
Model Technology™ *www.modeltechnology.com* Model Technology has built a worldwide distribution network and enjoys rapid sales growth from year to year. On December 1, 1994, the company became a wholly-owned subsidiary of Mentor Graphics. The tradition of being first, delivering high quality, easy to use, and affordable simulation tools.					✔			
ModelWare, Inc. *www.modelware.com* Datacom/telecom cores including ATM, HDLC, POS, IMA, UTOPIA, and bridges.	✔	✔						✔
MontaVista Software RTOS and development tools for immersed PowerPC processor						✔		
Multi Video Designs *www.mvd-fpga.com* Training, Expertise, Design on FPGA, PowerPC, DSP, High Speed I/O's	✔	✔						
Multiple Access Communications Ltd. *www.macltd.com* Signal processing hardware and software for wireless applications	✔	✔				✔		
Nallatech *www.nallatech.com* The High Performance FPGA Solutions Company.	✔	✔				✔		✔
Naztron Technologies, Inc. *www.naztron.com* A consultation company that provides resources specialized in ASIC and FPGA implementation and verification	✔	✔						✔
NetQuest Corporation *www.netquestcorp.com* Turn-key engineering design and production services in advanced high-speed embedded communications	✔	✔		✔				✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
NewLogic Technologies AG *www.newlogic.com* Cores for wireless applications including Bluetooth and 802.11.	✔	✔						✔
NitAl Consulting Services, Inc. *www.nital.com* PCI Express designs, High-speed, High complexity FPGA designs, Consulting services for complete system designs	✔	✔				✔		✔
North Pole Engineering, Inc. *www.npe-inc.com* Hardware and software design services for FPGA, ASIC and Embedded Systems	✔	✔				✔		✔
NorthBridge Technology Inc. *www.northbridgetech.com* All flavors of PCI designs, High-speed, High Density, High complexity FPGA designs	✔	✔						
Northwest Computer Engineering *www.nwce.com* Consulting, FPGA, and embedded microprocessor system designs	✔	✔						✔
Northwest Logic *www.nwlogic.com* High speed/complexity FPGA and board designs including PCI Express, PCI, DDR2/DDR SDRAM	✔	✔						✔
Novilit, Inc. *www.novilit.com* Teaming to bring the best communications hardware and protocol development software on the market today.					✔			
NovTech, Inc. *www.novtech.com* Security, networking, medical, and industrial controls		✔						
Nuvation Research Corporation *www.nuvation.com* Design Services; imaging and communications for defense/security, medical, consumer, and datacom/telecom markets		✔				✔		✔
Oki Information Systems Co., Ltd. *www.okijoho.co.jp* Picture transmission and Image Processing, High-speed communication, PCI, Mechanism control		✔						

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
PD3 Tecnologia *www.pd3.com.br* Digital Modems, Data Multiplexing Systems, Digital Interfaces, Ethernet Interfaces, Optical, Interfaces, Voice Interface, Different uP Platforms		✔						
Pentek, Inc. *www.pentek.com* FPGA IP Cores including FFTs, digital receivers and pulse compressions algorithms. DSP, Data Acquisition, Software Radio boards and system solutions.	✔		✔					✔
Perigee, LLC *www.perigeellc.com* DSP and image processing cores; embedded firmware, GUI and application software development		✔						
Phystream Ltd. *www.phystream.com* Integrated hardware/software data processing systems, transmission (PDH, SDH, SONET), ATM, frame relay, LAN (Ethernet, FDDI), multiservice, wireless protocols		✔				✔		✔
Pinpoint Solutions, Inc. *www.asic-design.com* Standards-based telecom and datacom IP; video	✔	✔						✔
Pixel Velocity, Inc. *www.pixel-velocity.com* Image processing, real-time machine vision, board design, FPGA design, sensor design, object oriented and embedded software development, and product development.	✔	✔						
PLC 2 *www.plc2.de* Module design and system design		✔						
PLD Applications *www.plda.com* High speed interfaces (PCI, PCI-X, PCI Express, AMBA-AHB)	✔			✔				✔
Plextek Ltd. *www.plextek.co.uk* Design services; FPGA, DSP, telematics, radio, microwave and microprocessor technologies for defence and communications	✔	✔						
Plexus Corp. *www.plexus.com* Mechanical, electrical, and software design, PCB design, prototyping services, material procurement and management		✔						

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
POLAR-Design *www.polar-design.de* High End FPGA-Design for Telecom, Networking and DSP Applications		✔						
Polybus Systems Corp. *www.polybus.com* InfiniBand based systems, DSPs, network processors, super-computers, video systems, and networking systems		✔						✔
Presco, Inc. *www.prescoinc.com* Twenty-five years of experience in custom electronics design/manufacturing for high performance Image Processing, Datacom, and Inkjet		✔						
Prevas AB *www.prevas.se* Design services for mainly Industrial and Military/Aerospace companies		✔						
Processor Systems (India) Pvt. Ltd. *www.procsys.com* Embedded (FPGA, Board & Firmware) Design in Telecom/Broadband & Consumer Areas		✔						
ProDesign Electronic & CAD-Layout GmbH *www.prodesigncad.de* Costumer-tailored solutions and experience with Xilinx products		✔			✔			
Prodrive B.V. *www.prodrive.nl* Image processing, motion control and power electronics. DSP, FPGA, PowerPC, PrPMC, VME, PCI, IEEE1394, Rapid-IO, Gigabit-Ethernet.		✔						
Product Acceleration, Inc. *www.productacceleration.com* FPGA symbol generation					✔			
Productivity Engineering GmbH *www.PE-GmbH.com* Automotive video applications, PCI-Express Designs, High-speed, High complexity FPGA Designs	✔	✔						
Programall Technologies, Inc. *www.programalltechnologies.com* Design services; embedded systems targeting to Virtex, Spartan, Virtex2 and Virtex2Pro FPGAs	✔	✔						✔
QinetiQ Limited *www.quixilica.com* Floating Point Cores, FPU for MicroBlaze		✔						✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
QualCore Logic, Inc. *www.qualcorelogic.com* Cores for microcontrollers, telecommunication, datacom, PC peripherals and bus interface standards.	✔	✔						✔
R&D Consulting *972-9-7673074* FPGA development; video; imaging; graphics; DSP; communications	✔	✔						✔
Rapid Prototypes, Inc. *www.fpga.com* Design services; high-performance reconfigurable FPGA applications requiring maximum speed or density using physical design principles	✔	✔		✔				✔
RDLABS *www.rdlabs.com* FPGA/ASIC design; wireless communications; 802.11b/a IP cores; Matlab, sysgen, ModelSim; instruments HP1661A, HP8594E, HP54520A		✔						
RealFast AB *www.realfast.se* Real-time systems and RTOS; operating system cores		✔				✔		✔
RF Engines, Ltd. *www.rfel.com* DSP; real-time, wide-band channelization/ filtering after digitization								✔
RFNet Technologies Pte Ltd *www.rfnetech.com* Wireless Communication	✔	✔						
RightHand Technologies, Inc. *www.righthandtech.com* Design services; Virtex-II Pro FPGAs, boards, and software for video gaming, storage, medical, and wireless		✔						
Rising Edge Technology, Inc. *www.risingedgetech.com* FPGA & CPLD designs, interface & control	✔	✔		✔				
Robert Bosch GmbH *www.can.bosch.com* Automotive, CAN		✔						✔
Roke Manor Research Limited *www.roke.co.uk* Design Services; IP, ATM, DSP, Imaging, Radar and Control solutions for FPGAs		✔				✔		✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Roman-Jones, Inc. *www.roman-jones.com* Design services: data acquisition, microprocessors (embedded & off chip), Automotive, PCI, analog design, DSP, Video		✔						✔
RTX Telecom A/S *www.rtx.dk* Advanced wireless communication products		✔						
Sapphire Computers, Inc. *drudolf@voyager.net* High-speed & high-utilization FPGA designs, combining expertise in the latest microprocessor & digital systems	✔	✔						✔
Secad S.A. *www.secad.fr* Image Processing, Industrial Control, Embedded Systems	✔	✔						
Shenzhen Anyware Technologies Co., Ltd. *www.anyware.com.cn* Integrated system design with high-density FPGA		✔						
Siemens AG *www.eda-services.com* Design, verification, test bench and bus model implementation, multi-level simulation, logical synthesis, static timing		✔						
Siemens AG - Medical Solutions CO MEE Medical	✔			✔				
Si-Gate GmbH *www.si-gate.com* Automotive Electronics, FPGA IP Development	✔	✔				✔		✔
Signal Integrity Software, Inc. Integrated system-level signal integrity, static timing and crosstalk analysis for high speed board design that utilize Xilinx RocketIO™ technology					✔			
Silicon & Software Systems Ltd. *www.s3group.com* World-class electronics design company uniquely combining IC, FPGA, software, and hardware design expertise	✔	✔				✔		✔
Silicon Infusion Limited *www.siliconinfusion.com* Embedded System Design, Wired & Wireless Telecoms, Multi-Processor Architectures, Modems, DSP Algorithms, Internet Protocol Processing								✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Silicon Interfaces Private Limited *www.siliconinterfaces.com* Frontend/Backend Design and Verification Services for Networking, Data Communication and Interconnect		✔						✔
Silicon Storage Technology, Inc. ATA Controller, Serial Flash, Combo-memory							✔	
Siliconexion, Inc. *www.siliconexion.com* Design services in the areas of FPGA and ASIC design, board-level design, and system hardware		✔						
Simucad *www.simucad.com* Verilog simulator support for Xilinx FPGAs					✔			
Simucad *www.simucad.com* Verilog simulator support for Xilinx FPGAs					✔			
Siscad Srl *www.siscad.it* Xilinx FPGA Design Services, IP development and integration		✔						✔
Smart Logic, Inc. *www.smart-logic.com* Design services; FPGAs for rapid prototyping including imaging, compression, and real-time control	✔	✔		✔		✔		
SoC Solutions, LLC *www.socsolutions.com* Embedded uP and software; networking, wireless, audio, GPS and handheld; peripheral cores.	✔	✔				✔		
SoleNet, Inc. *www.solenet.net* FPGA, board and system designs for wireless, telecom, consumer and reference design applications		✔						
SO-LOGIC Electronic Consulting *www.so-logic.co.at* Embedded Systems (PowerPC, MicroBlaze), System Level Design (Forge, SystemGenerator, HandleC), Xilinx training center (Austria, Eastern Europe)	✔	✔						✔
Specsoft Consulting, Inc. *www.specsoftus.com* Ethernet, PCI, USB, Wishbone bridges		✔						✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
Spectrum Signal Processing, Inc. *www.spectrumsignal.com* Designs and manufactures industry-leading high-performance signal processing engines and platforms		✔						
Starvision Technologies, Inc. *www.starvisiontech.com* MPEG2								✔
STMicroelectronics NV *www.st.com/stonline/products/families/memories/memory/mem_ctct.htm* Non-volatile memory integrated circuits							✔	
Sundance Multiprocessor Technology Ltd. *www.sundance.com* Scaleable Modules built with Virtex Pro and Virtex-4 FX family of FPGAs.	✔		✔			✔		✔
Supercomputing Systems AG *www.scs.ch* Engineering company specialized in system design and development		✔						
Symphony EDA *www.symphonyeda.com* VHDL development environment and simulator for Xilinx FPGAs					✔			
Synchronous Design, Inc. *www.synchronousdesign.com* FPGA/CPLD/board level design using synchronous digital techniques		✔						
Synective Labs AB Reconfigurable computing		✔				✔		✔
Synopsys The Xilinx Synopsys partnership is an optimum solution for today's leading high level and HDL level programmable logic designers.		✔			✔			✔
Synopsys Professional Services *www.synopsys.com* Electronic design automation (EDA) tools, intellectual property, and design services		✔						
Synplicity, Inc. The new Xilinx ISE provides flexibility and seamless integration with Synplicity tools including Synplify®, Synplify Pro™ and Amplify®, allowing engineers to achieve maximum performance in their designs.			✔		✔			
SysOnChip, Inc. *www.sysonchip.co.kr* IP and products for CDMA Cellular/PCS/WLL modem, FEC and Bluetooth								✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
SystemCrafter, Ltd. SystemC synthesis tool for Xilinx FPGAs						✔		
SysVERI Co., Ltd. RTL-level FPGA debugger		✔			✔			
Taifatech Inc. *www.taifatech.com* High performance controller targeted for Digital Home market; Switch, Wireless, PCI, VGA, and SoC.		✔						
Taray Technologies India Private Ltd. High speed memory design experts		✔						
TATA ELXSI LIMITED *www.tataelxsi.com* Networking & Communications, Wireless, Embedded software, Hardware design, DSP, Scientific Computing & Image Processing		✔						
TATA ELXSI LIMITED *www.tataelxsi.com* Networking & Communications, Wireless, Embedded software, Hardware design, DSP, Scientific Computing & Image Processing		✔						
TATA ELXSI LIMITED *www.tataelxsi.com* Networking & Communications, Wireless, Embedded software, Hardware design, DSP, Scientific Computing & Image Processing		✔						
TeamF1, Inc. Network/security stacks for immersed PowerPC processor						✔		
Technolution BV *www.technolution.nl* High-end hardware and software solutions for technical information and embedded systems	✔	✔		✔		✔		✔
Technovare Systems, Inc. *www.technovare.com* Design services for the commercial, industrial, automotive, military, and space markets		✔						
Teknosarus Embedded Systems Pvt. Ltd. *www.teknosarus.com* WiMAX PHY & MAC IP & system design	✔	✔						✔
Tensilica, Inc. *www.tensilica.com* Configurable processors cores and development tools.					✔			✔

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
TES Electronic Solutions GmbH *www.thales-ee.com* Design services for FPGA, digital IC, analog/ mixed-signal IC, RF, board, and embedded software design		✔						
TES Electronic Solutions GmbH *www.thales-ee.com* Design services for FPGA, digital IC, analog/ mixed-signal IC, RF, board, and embedded software design		✔						
The Dini Group *www.dinigroup.com* ASIC Prototyping and high-speed FPGA design.	✔	✔						
TietoEnator R&D Services AB *www.tietoenator.com* Design services for Telecom, Industrial IT and Automotive. Xilinx exclusive training provider in Scandinavia		✔						
Tokyo Electron Device Ltd. *www.teldevice.co.jp* Design expertise and custom ASIC design support		✔						
Treck Inc. Network stacks for immersed PowerPC processor/MicroBlaze SPC					✔			
TurboConcept SAS *www.turboconcept.com* Turbo code Forward Error Correction solutions								✔
Tyco Electronics Power supplies				✔				
Ultimodule, Inc. *www.ultimodule.com* MIPS processor, Memory, I/O, System control logic interfaces, O/S	✔	✔				✔		✔
Vertronix, Inc. *www.vertronix.com* ASIC, FPGA, Emulation, Verification, Verilog, System, Protocol, Communications	✔	✔						
VGVP Ltd. *www.vgvp.co.il* Design services, specializing in wireless, datacom and telecom applications	✔	✔						✔
V-Integration *www.v-integration.com* Providing first pass success in complex FPGA designs. Applications include imaging, communications and interface design		✔						

Xilinx Partner	Board-Level Products	Design Services	DSP Tools	Electronic Components	EDA Tools	Embedded Processor Products	Integrated Circuits	IP Cores
VMETRO Transtech *www.transtech-dsp.com* COTS subsystem provider of Xilinx based hardware solutions.	✔	✔	✔			✔		✔
VSL Networks, Inc. *www.gapviewcom.com* MPEG2, set top boxes, digital video broadcast (DVB) for cable, satellite and terrestrial networks; conditional access and encryption technologies.		✔						✔
Wasabi Systems, Inc. *www.wasabisystems.com* RTOS for immersed PowerPC processor						✔		
WAVELET GROUP FPGA, DSP, Hardware, Firmware, Signal and Image Processing Algorithms		✔						
Williams Consulting, Inc. *www.wciatl.com* Embedded systems hardware and software. Video, MPEG, control		✔						
Wind River Systems, Inc. *www.windriver.com* RTOS and development tools for immersed PowerPC processor						✔		
Wipro Technologies *www.wipro.com* ASIC, board design, engineering of Class 5 switches		✔						
Xelic, Inc. *www.xelic.com* SONET, SDH		✔						✔
XiChron, Inc. *www.xichron.com* Digital audio, video, Machine Vision, and DSP		✔						
Xylon d.o.o. *www.logicbricks.com* Human-machine interfaces and industrial communication controllers; video	✔	✔				✔		✔
Zuken, Inc. *www.zuken.co.jp/soc/* PCI2.2,10/100 Ethernet and Giga-bit Ethernet cores	✔	✔				✔		✔

VIRTEX-4

INCREDIBLY

INTERCONNECTED

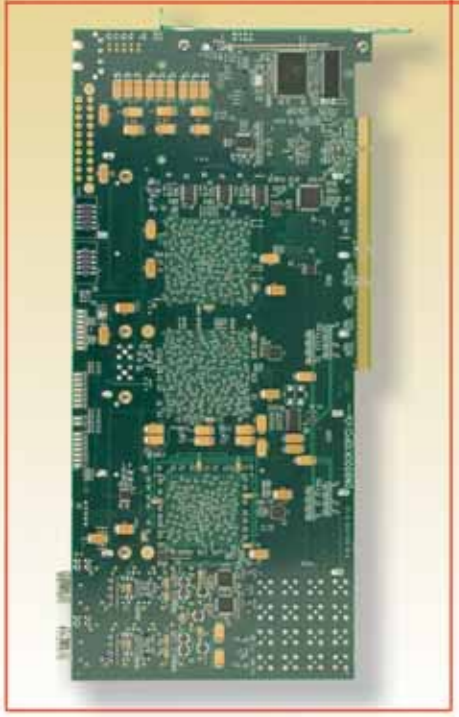

Partitioning your design onto the 3.7 million ASIC gates (LSI measure) of our new board is a lot easier. With three of the biggest, fastest, new Xilinx Virtex-4 FPGAs this PCI hosted logic prototyping system takes full advantage of the integrated ISERDES/OSERDES. 400MHz LVDS differential communication with 10X multiplexing means more than 1800 signals between FPGA A & B. Synplicity Certify™ models are provided for partitioning assistance.

A dedicated PCI Bridge (533mb/s data transfer) means that all FPGA resources are available for logic emulation. Other features are designed to ease your prototyping job:

- 5 programmable clock synthesizers
- 2 DDR2 SODIMMs (custom DIMMs for SSRAM, QDR, Flash ...)
- Two, 200 pin expansion connectors for daughter cards
- 10 GB/s serial I/O interfaces with SMA connectors and XFP or SFP Modules
- Configured by PCI, USB 2.0, or SmartMedia with partial reconfiguration support on all FPGAs

Various stuffing options and a wide selection of daughter cards let you meet your exact design requirements. Prices start at less than $10,000. Call The Dini Group today for the latest ASIC prototyping solution.

The DiNI Group

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

High VELOCITY LEARNING

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.

Don't see a seminar in a city near you? 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.

Topics Covered

- What's New in ISE 8.1i
- Faster Performance – ISE Fmax and PlanAhead
- ISE 8.1i – Focus on New Features
- Design Optimization in ISE 8.1i using XST
- PlanAhead Design Analysis Tool
- Xilinx Platform Studio Update
- Providing Power for FPGAs
- Spartan Family Update
- Virtex Family Update
- CPLD Family Update

February 2006

- February 16 - Chicago, IL
Irvine, CA
San Jose, CA
Toronto, ON
- February 21 - Dallas, TX
Ottawa, ON
- February 22 - Los Angeles, CA
Montreal, ON
- February 23 - Boston, MA
Minnesota, MN
Shenzhen, China

March 2006

- March 1 - Orlando, FL
- March 7 - Beijing, China
Dallas, TX
- March 8 - Atlanta, GA
- March 13 - Portland, OR
- March 14 - Shanghai, China
- March 16 - Irvine, CA
Philadelphia, PA
- March 20 - Baltimore, MD
- March 21 - San Jose, CA
Toronto, ON
- March 22 - Austin, TX
Los Angeles, CA
- March 23 - Boston, MA
Chicago, IL
Minnesota, MN

READY! SET! GO WITH PLANAHEAD PERFORMANCE!

LX SX FX

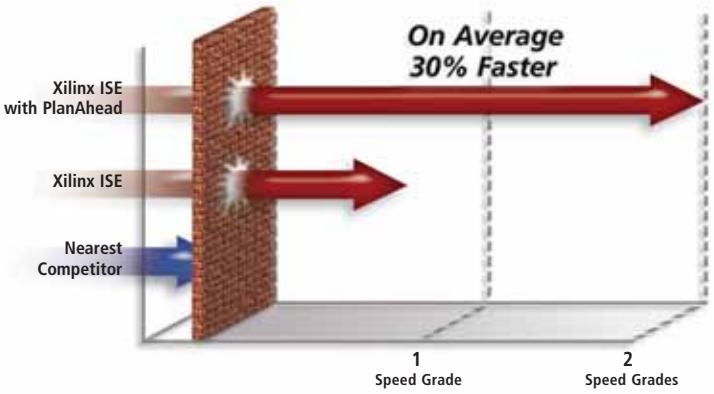

Based on benchmark data from a suite of 15 real-world customer designs targeting Xilinx and competing FPGA Solutions.

Two speed grades faster with PlanAhead software and Virtex-4

With our unique PlanAhead software tool, and our industry-leading Virtex-4 FPGAs, designers can now achieve a new level of performance. For complex, high-utilization, multi-clock designs, no other competing FPGA comes close to the Virtex-4 PlanAhead advantage:

- **30% better logic performance on average = 2 speed grade advantage**
- **Over 50% better logic performance for complex multi-clock designs**

MEET YOUR TIMING BUDGETS . . . BEAT YOUR COMPETITION TO MARKET

Meeting timing budgets is the most critical issue facing FPGA designers*. Inferior tools can hit a performance barrier, impacting your timing goals, while costing you project delays and expensive higher speed grades. To maximize the Virtex-4 performance advantage, the new PlanAhead software tool allows you to quickly analyze, floorplan, and improve placement and timing of even the most complex designs. Now, with ISE and PlanAhead you can meet your timing budgets *and* reduce design iterations, all within an easy-to-use design environment.

Download a free eval today at ***www.xilinx.com/planahead***, view the TechOnline web seminar, and prevent your next FPGA design from stalling.

* CMP: June 2005 FPGA EDA Survey

FPGA and Structured ASIC JOURNAL READER'S CHOICE AWARDS
Highest customer satisfaction:
- Virtex & Spartan FPGAs
- ISE design tools
- Support

XILINX®
The Programmable Logic Company℠
www.xilinx.com/planahead

www.techonline.com
View The TechOnLine Seminar Today

BREAKTHROUGH PERFORMANCE AT THE LOWEST COST

PN 0010916