ISSUE 45, SPRING 2003 XCELL JOURNAL XILINX, INC.

Xcell journal

THE AUTHORITATIVE JOURNAL FOR PROGRAMMABLE LOGIC USERS

Issue 45
Spring 2003

Evolving Low-Cost Solutions

SERIAL TSUNAMI

Serial Tsunami Verification and Planning

SOFTWARE

ISE 5.2i Further Reduces Costs

RealFast RTOS

EMBEDDED PROCESSORS

New PicoBlaze RISC Reference Design

FPCs for Cost-Sensitive Applications

APPLICATIONS

Processors for Software Defined Radio

How to Make Smart Antenna Arrays

NEW PRODUCTS

MultiPRO – New Prototyping Tool

COVER STORY

Serial Tsunami High-Speed I/O Technology... Higher Performance, Lower Costs

XILINX®

Will the RealDigital™ CPLD please stand up.

CoolRunner-II™ If you are not using the latest breakthrough in CPLD technology, you are not getting the most out of your CPLD. Only the new 1.8V CoolRunner-II RealDigital CPLD from Xilinx, offering a 100% digital core, gives you the performance, low power, and features you are looking for with no price premium.

The best system performance in the industry

The CoolRunner-II RealDigital CPLD features the most I/O per macrocell, and advanced I/O interfacing including HSTL and SSTL. System performance exceeds 400 MHz with the lowest dynamic current and lowest standby current in the industry (20 times lower compared to all other 1.8V devices!) The unique Clock Divider means no more dividing the clock off-chip. And the new RealDigital CPLDs offer the smallest packaging with the highest level of design security.

HIGH CURRENT SENSE AMPS

1.8V CORE VOLTAGE CPLD COMPETITIVE CHART

Manufacturer	Xilinx	Lattice	Altera
Device Family	CoolRunner-II	ispMACH4000C	None
Standby Current	<100 µA	2 mA	N/A
Clock Divider	YES	NO	N/A
Clock Doubler	YES	NO	N/A
I/O Standards Support	LVTTL, LVCMOS, HSTL, SSTL	LVTTL, LVCMOS	N/A
I/O Banks (max)	4	2	N/A

Get the new CoolRunner-II Design Kit

Get started now with the new CoolRunner-II Design Kit, including a populated board, programming cable, design guide, and software resource CDs… all for under fifty dollars! Just contact your local Xilinx distributor or visit **www.xilinx.com/coolrunner2** to order your kit today. And don't forget the new CoolRunner-II RealDigital CPLDs are fully supported by the easy-to-use ISE WebPACK™ tool, which you can download FREE right now.

XILINX®
The Programmable Logic Company℠

www.xilinx.com/realcpld

FORTUNE® 2002
100 BEST COMPANIES TO WORK FOR

Why You Should Read *Xcell*

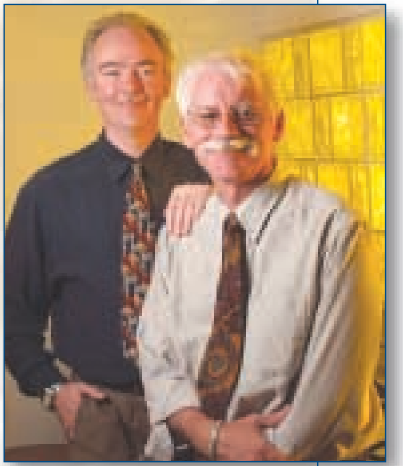

EDITOR IN CHIEF	Carlis Collins editor@xilinx.com 408-879-4519
MANAGING EDITOR	Tom Durkin tom.durkin@xilinx.com 530-271-0899
XCELL ONLINE EDITOR	Tom Pyles tom.pyles@xilinx.com 720-652-3883
CONTRIBUTING EDITORS	Jane Middleton Charmaine Hussain Peter Bergsman
DESIGN & ILLUSTRATION	Scott Blair Dan Teie

Xcell journal

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

Xcell is published quarterly. XILINX, the Xilinx logo, CoolRunner, Rocket Chips, Rocket IP, Spartan, and Virtex are registered trademarks of Xilinx, Inc. ACE Controller, ACE Flash, Alliance Series, AllianceCORE, ChipScope, CORE Generator, Fast CONNECT, Fast FLASH, Fast Zero Power, Foundation, HDL Bencher, IRL, J Drive, LogiCORE, MicroBlaze, MultiLINX, NanoBlaze, PicoBlaze, QPro, Real-PCI, RocketIO, SelectIO, SelectRAM, SelectRAM+, Smart-IP, System ACE, Virtex-II Pro, Virtex-II EasyPath, WebFITTER, WebPACK, WebPOWERED, XACT-Floorplanner, XACT-Performance, XAPP, XC designated products, Xilinx Foundation Series, Xilinx XDTV, and XtremeDSP are trademarks, and The Programmable Logic Company is a service mark of Xilinx, Inc. Other brand or product names are registered trademarks or trademarks of their respective owners.

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

Programmable logic is now the overwhelming choice for all types of designs, from low-cost consumer devices to high-performance switching systems. When you consider the many advantages of programmable logic, you'll see that it's not only the fastest, easiest, and most flexible way to develop new products, it's also the lowest cost and the most reliable solution in most applications. There simply is no better way to develop new products.

The *Xcell Journal* will help you bring your imagination into reality through programmable logic technology. *Xcell* is written by engineers who understand the challenges you face every day, and we strive to bring you the latest information about the products and services from Xilinx and our entire Partner Ecosystem, so you can make the best and most informed choices. We show you how to save time, effort, and money, while creating better, more profitable products.

Xilinx is the undisputed world leader in programmable logic. We invented the Field Programmable Gate Array (FPGA) back in 1984, and since that time we have continually created faster, cheaper, more capable products with each new generation. In 1994, a basic 25K-gate XC4025 FPGA, built on 0.6µ technology, cost you $654. This year the price of our smallest Spartan™-IIE device with twice as many system gates is just $6.00. Plus, this year we'll be the first in the industry to introduce 90 nm technology to even further reduce costs. As you can see, programmable logic has matured rapidly, and that's why its use is growing faster than any other market segment.

Our Partner Ecosystem is also an important part of what makes Xilinx technology so attractive. Our partners are the best in the business, and they provide a wide range of development tools, intellectual property (IP), design services, and support products that help make your job easier while making your products more successful. Together, we are changing the future of logic design, bringing you advanced design solutions, helping you create new realities.

The *Xcell Journal* is your best source for information about this exciting technology.

Carlis Collins

Carlis Collins
Editor-in-Chief

To subscribe to the printed *Xcell Journal*,

or to view the Web-based *Xcell Online* Journal, visit:

www.xilinx.com/publications/xcellonline/

SPRING 2003, ISSUE 45

Xcell journal

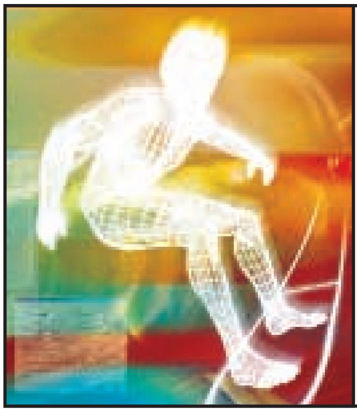

Serial Tsunami Requires Changes to SI Verification and Planning

Mentor Graphics' HyperLynx products are designed to facilitate the transition from parallel to serial connectivity.

17

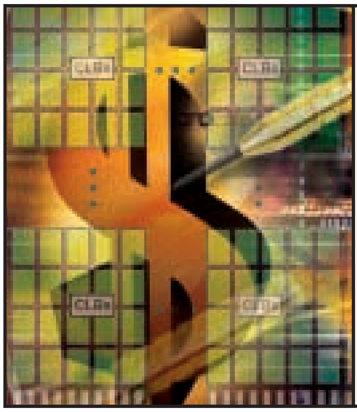

Reduce Your RISC with a PicoBlaze Reference Design

Build your own soft microcontroller in a CoolRunner-II CPLD. It's easy, fast, and free.

27

Get RealFast RTOS with Xilinx FPGAs

Real-time operating systems implemented in Xilinx FPGAs enhance performance, improve predictability, simplify design, and lower system costs.

58

How to Make Smart Antenna Arrays

The Nallatech BenADIC card combines a 20-channel data acquisition system with Xilinx XtremeDSP technology and Virtex FPGAs for high-performance digital signal processing.

76

Working with the Best

Xilinx now ranks number four on FORTUNE magazine's list of the "100 Best Companies to Work For." Here's why it matters to you.

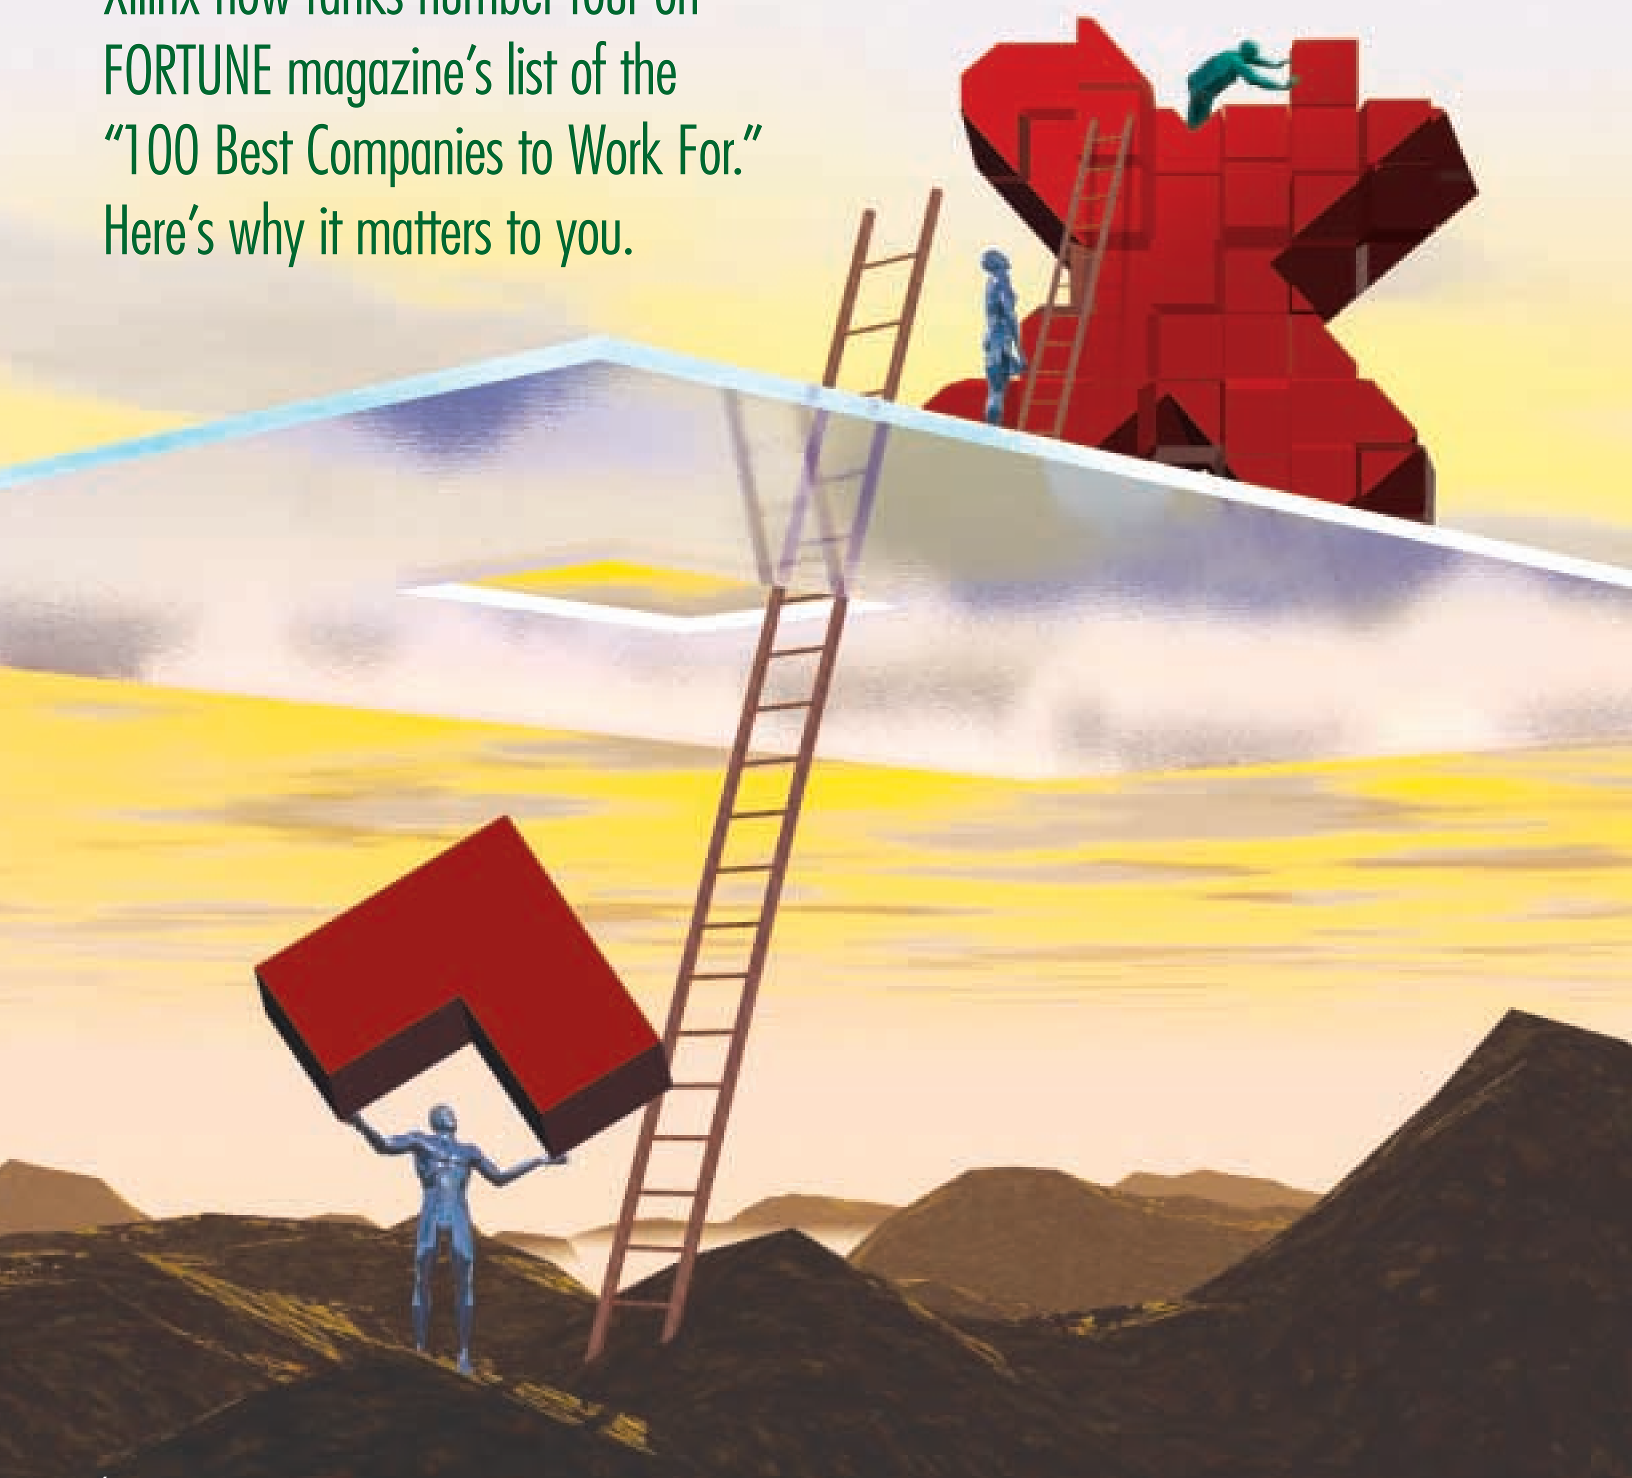

by Wim Roelandts
CEO, Xilinx Inc.

Xilinx is a certainly great company to work for, as acknowledged by *FORTUNE* magazine's ranking, and we are a great company to work with, because we have not slowed our fast pace of innovation nor decreased our customer support activities – all critical to the success of our customers. We weathered the recent economic downturn in ways that have made Xilinx stronger across the board. How did we keep our technology advancing at a high rate, continue to expand our support, and gain market share, in a downward spiraling economy? Here's how.

A large part of our business came from the telecommunications industry, which was severely hit by the downturn. Our revenues were cut almost in half, overnight. Our competitors were affected just as severely and as a result, many were forced to lay off a sizeable part of their workforce and to reduce both their product development and their customer support activities. We chose to make layoffs only as a last resort, and yet we needed to cut our expenses dramatically.

We chose to take pay cuts, on a sliding scale, instead of doing layoffs. The average pay cut was 6%, rising to 20% for myself. Everyone shared the burden according to their ability, and everyone was very happy to have some job security in a time when many of our colleagues in other companies were losing their jobs. Not only did our morale remain high, but our productivity increased as well, and we continued to produce our new technologies even faster than before.

Studies have shown that companies that can avoid layoffs rebound more quickly when the economy turns around, because they can take advantage of every new opportunity with a full staff. Companies that must layoff their workers suffer from lower morale, less productivity, and a slower return to profitability. That's what we are seeing now: Xilinx is gaining market share, while our competitors are struggling to keep up.

What We've Accomplished

The technology development that we are doing today will not show up in your hands for several years. That's why it is imperative that we keep our technology advancing, even when current revenues are slumping. Otherwise, we would find that when the economy turns around, we would not be ready to support your demand. Here are some of the important advances we are now introducing.

90 nm Low-Cost Fabrication Technology

Using IBM's most advanced, copper-based, 90 nm semiconductor manufacturing process technology, IBM and Xilinx are manufacturing a new FPGA design in IBM's new 300 mm chip fabrication facility. This technology is a major reason why our FPGAs will continue to lead the industry in cost reduction.

This new process technology has resulted in a 5% to 80% percent chip-size reduction compared to any competing FPGA. IBM plans to manufacture this new product, in high volumes, in the second half of 2003. The new IBM $2.5 billion, 300 mm chip-making facility combines – for the first time anywhere – IBM chip-making breakthroughs such as copper interconnects, silicon-on-insulator (SOI), and low-k dielectric insulation on 300 mm wafers.

Our investment in 90 nm manufacturing technology will enable us to drive pricing down to under $25 for a one-million-gate FPGA, which represents a savings of 35% to 70% compared to any competitive offering. Such a significant reduction in pricing is possible due to the remarkable economies of scale involved with moving to next-generation manufacturing processes at increasingly finer geometries. Now we can achieve greater device densities and higher yields, making our FPGAs the logical alternative to ASICs.

The rising costs of developing ASICs on more advanced processes are well known. Not only are non-recurring engineering (NREs) charges rising to over one million dollars per design, but the engineering cost of developing and testing a complex ASIC on advanced processes such as 150 nm or 130 nm technology can run up to 10 times that amount. With the deployment of our new FPGAs on 90 nm technology, Xilinx has resolved all of the deep sub-micron design challenges for you. Using these new FPGAs, you will get all of the substantial cost advantages of the 90 nm technology without being forced to worry about the detailed circuit design issues associated with ASICs. You can concentrate on getting your system designed rather than on getting the chip itself to function. With our increased density, performance, and system features, you get all the benefits of an ASIC in a flexible, programmable

FPGA, without the risk and the huge NRE costs. This is dramatically expanding the total market for FPGAs, both in new applications in existing markets and in totally new markets.

Serial Tsunami

We have pioneered the development of very high-speed serial I/O technology – we call it the Serial Tsunami Initiative. This new technology will solve many design challenges, allowing you to replace old parallel busses with a far less expensive solution. With Serial Tsunami, you can significantly reduce costs, produce faster designs, reduce your PC board area, and create products that were never possible before.

Other advantages of Serial Tsunami include reduced EMI, noise, cross talk, and skew, which makes your overall design more reliable. You can easily expand the I/O for increased bandwidth, and the physical interface can drive long signal traces on your PC boards, making them ideal for backplanes.

The move to serial I/O technology is inevitable – there is no better way of cutting costs while keeping pace with current and future bandwidth requirements. Our Virtex-II Pro™ FPGAs with embedded RocketIO™ 3.125 Gbps transceivers, and the accompanying IP cores, reference designs, and support infrastructure, provide the best possible way for you to realize all the advantages of serial I/O technology without the pitfalls.

Virtex-II Pro FPGAs support major emerging serial interfaces such as PCI Express™, Gb Ethernet PHY, 10 Gb Ethernet XAUI, Fibre Channel, OC-48, OC-192 and OC-768 SONET for backplanes, Serial RapidIO™, and InfiniBand™. Each embedded RocketIO transceiver in the Virtex-II Pro FPGAs is based on several generations of customer-proven Mindspeed SkyRail™ technology and can run from 622 Mbps to 3.125 Gbps; there are up to 24 of these transceivers available in one FPGA.

Virtex-II Pro FPGAs also support parallel interface standards such as SPI-3 (POS PHY Level 3), SPI-4.1 (Flexbus 4), SPI-4.2 (POS PHY™ Level 4), 10 Gb Ethernet Media Independent Interface (XGMII), RapidIO, PCI, PCI-X, CSIX, HyperTransport™, XSBI, and SFI-4. Therefore, Virtex-II Pro FPGAs are the only devices in production today that enable you to bridge between the parallel and serial interfaces, making them the ultimate connectivity platform.

Xilinx and partners are delivering pre-engineered IP cores for serial connectivity protocols. Cores for Gb Ethernet MAC with PHY, 10 Gb Ethernet MAC with XAUI, PCI Express, Fibre Channel, and reference designs for SONET OC-48 backplanes are available now, and more are being added. You can also build higher-level protocols using our embedded IBM PowerPC™ processors. Reference designs and evaluation/prototype boards help you verify the performance of the transceivers in real hardware.

Creating designs with speeds of 622 MHz, 3.125 GHz, 10 GHz, and beyond will present challenges with PCB design and signal integrity. Therefore, Xilinx and leading EDA partners are solving this dilemma with tools such as the Cadence SPECCTRAQuest™ and HSPICE models. We provide in-depth characterization data for the RocketIO transceivers for flawless system design. You will know exactly what to expect for your specific design situation, plus we provide best design practices and PCB layout guidelines to help you succeed.

We have also developed a new, open, scalable, lightweight serial standard called Aurora to help you transition from parallel to serial interfaces. It is a link layer protocol that can encapsulate and transport any higher-level protocol. It is very resource-efficient with low latency. A single-lane reference design and the specification are available for free download at *www.xilinx.com/aurora/*. A quad-link reference design will be available during the first half of 2003.

And Much More

I've only mentioned a few of our most important technologies. As you can see, the current downturn has not slowed our innovation, and we are fully ready to support your design requirements, both now and in the future. We are the fourth best company to work for, and the number one company to work with.

Ride the Crest of the Serial Tsunami

The Xilinx Serial Tsunami Initiative is a comprehensive set of programmable serial I/O solutions and system strategies that will help you simplify designs, increase performance, and lower system costs.

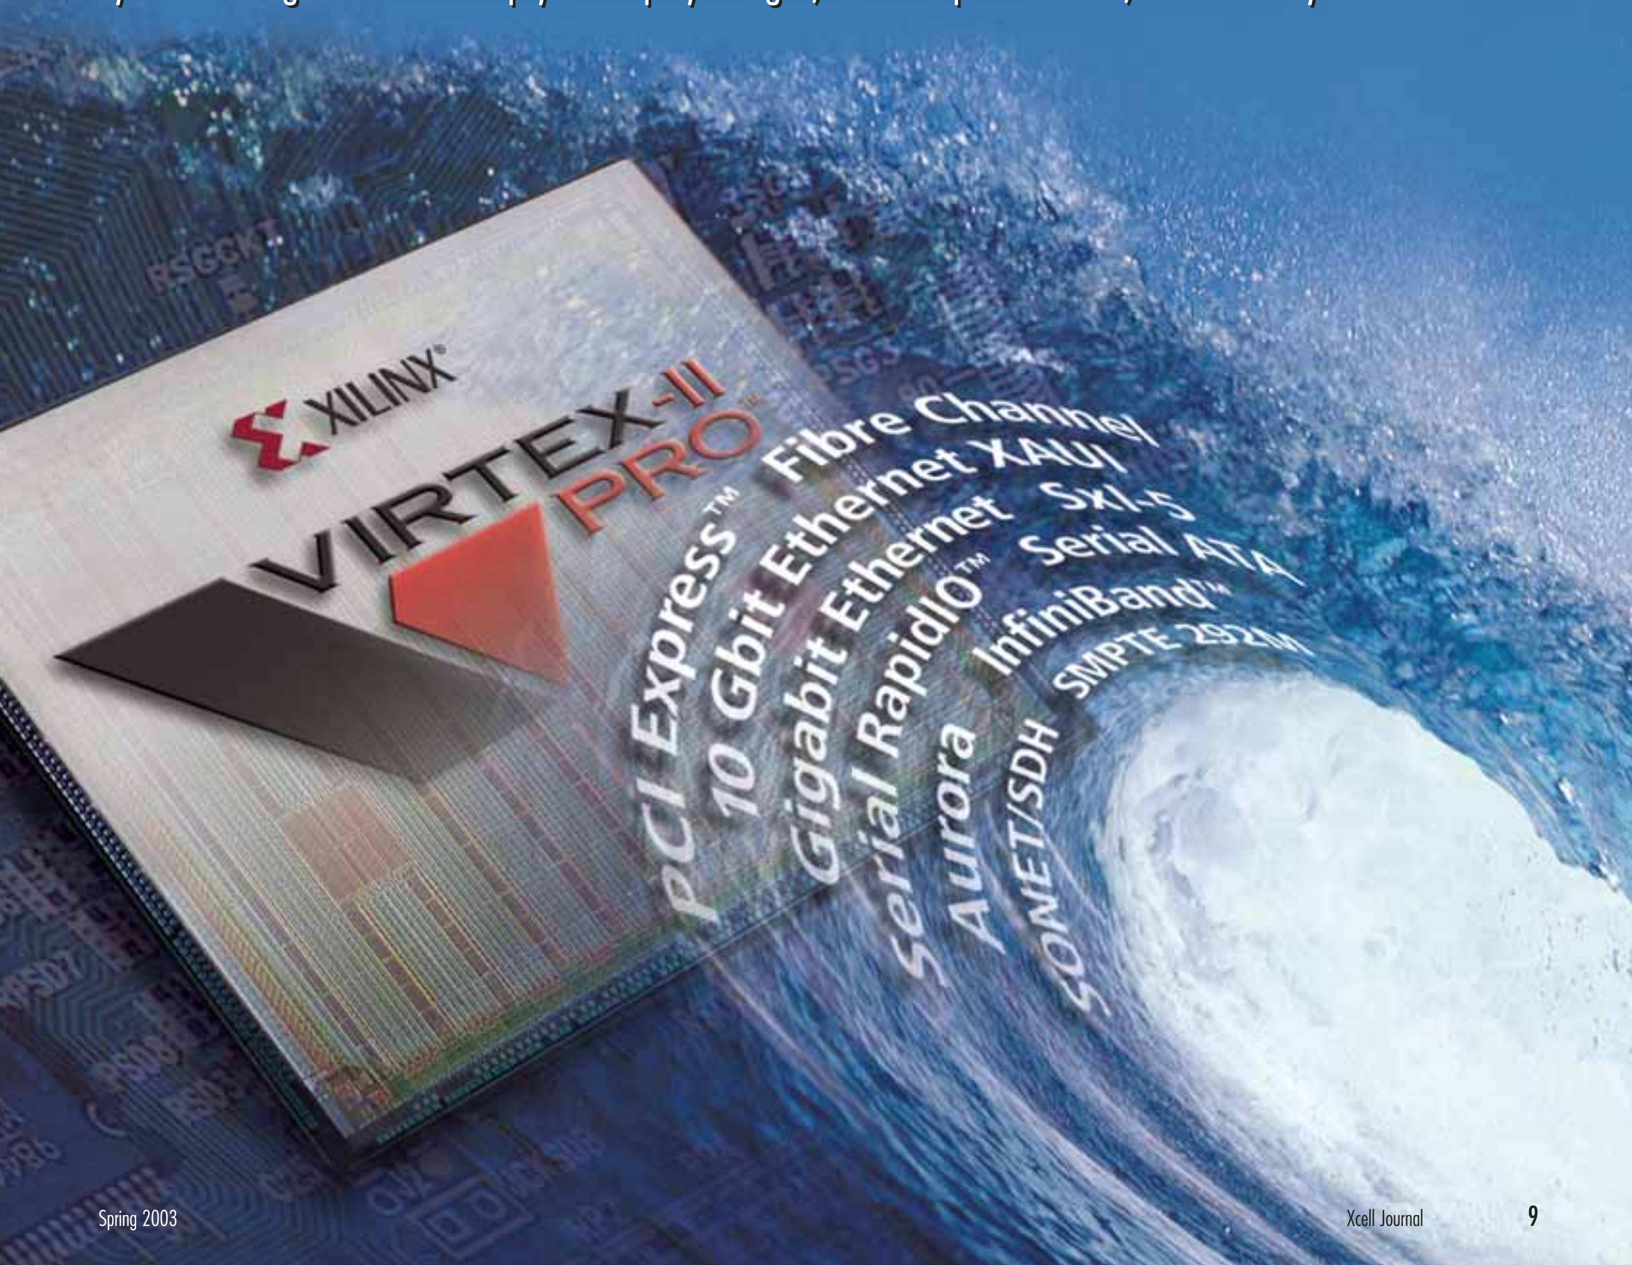

by Anil Telikepalli
Marketing Manager, Virtex Solutions
Xilinx, Inc.
anil.telikepalli@xilinx.com

Remember the days when serial I/O brought either USB or IEEE 1394 to mind? Not any more. A veritable tsunami of new and evolving serial I/O standards is washing over the technological landscape, delivering promises of higher performance, lower costs, and simpler designs. Remarkable advances in semiconductor technology and the availability of low-power CMOS serial transceivers are driving a migration of tidal wave proportions from parallel to serial interfaces.

Xilinx is at the forefront of this movement. We have been shipping our flagship Virtex-II Pro™ FPGAs (*www.xilinx.com/virtex2pro*) since the beginning of 2002. The Virtex-II Pro devices are the only Platform FPGAs with embedded 3.125 Gbps RocketIO™ CMOS serial transceivers. Designing cutting edge serial I/O technologies is a challenging endeavor, but using serial I/Os to build your systems need not be.

Broad Trend Toward Serial Connectivity

Experts agree that both single-ended and differential parallel I/Os have reached their physical limitations and cannot provide a reliable and cost-effective means for data rates greater than 1 Gbps. Serial I/O provides performance benefits to high-speed systems and cost benefits to low-speed systems. This double benefit has propagated waves of new serial interface standards development.

The inevitable result is the current widespread migration toward serial I/O across many segments of the industry, including PC and consumer, storage and servers, communications networking, industrial computing and control, and test equipment.

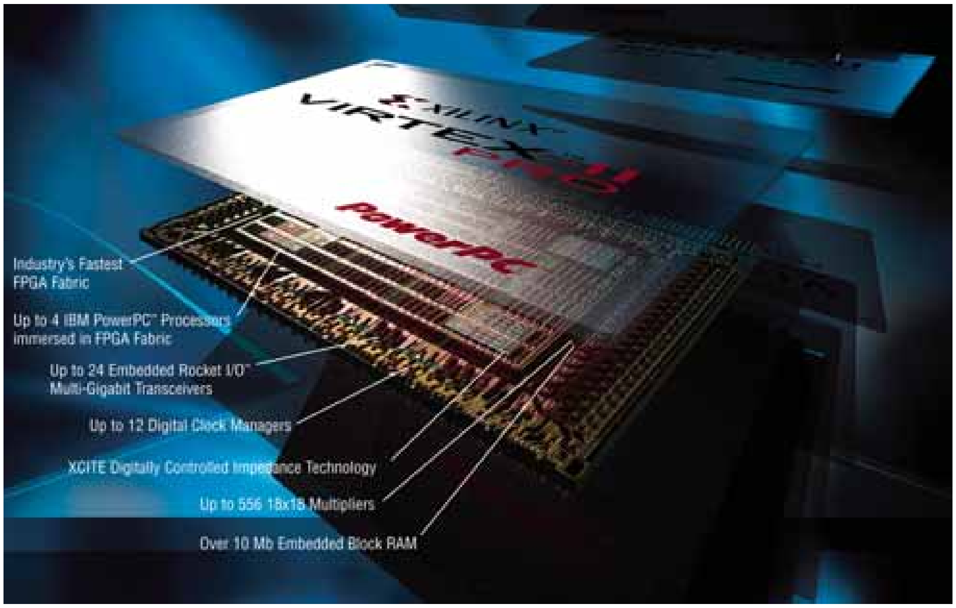
XILINX
VIRTEX-II PRO
PowerPC
Industry's Fastest FPGA Fabric
Up to 4 IBM PowerPC™ Processors immersed in FPGA Fabric
Up to 24 Embedded Rocket I/O™ Multi-Gigabit Transceivers
Up to 12 Digital Clock Managers
XCITE Digitally Controlled Impedance Technology
Up to 556 18x18 Multipliers
Over 10 Mb Embedded Block RAM

Figure 1 - Virtex-II Pro Platform FPGA

Serial system interfaces such as PCI Express™, Serial RapidIO™, InfiniBand™, 1 Gb Ethernet, 10 Gb Ethernet XAUI (10 gigabit attachment unit interface), Fibre Channel, Serial ATA, SxI-5, and TFI-5 are all available today, with many more coming to address specific needs.

Serial Tsunami Initiative

Xilinx launched the Serial Tsunami Initiative (*www.xilinx.com/connectivity*) to sail the crest of the industry move from parallel to serial interfaces, to reduce costs, and to keep pace with current and future bandwidth requirements. The initiative is a vision for delivering a complete suite of serial connectivity solutions – including Platform FPGAs, IP cores, design software and methodologies, reference designs, solution boards, extensive characterization data, and training classes – to enable you to design your next-generation products.

Unlike point solutions, the Xilinx Serial Tsunami Initiative help you surf the surges of multiple, evolving serial standards all the way across devices, IP, and software as you build your systems.

The foundation of the Serial Tsunami Initiative was our aqusition of RocketChips Inc. two years ago. Today, a dedicated R&D team in the Communications Technology Division (CTD) is wholly focused on improving, developing, and delivering the capabilities that make up the Serial Tsunami vision. Commenting on this strategy, Wim Roelandts, president and CEO of Xilinx, said: “The underlying RocketIO technology that’s making it possible for Xilinx to support multi-gigabit systems is truly ‘rocket science’ – not something you’d want to simply license from an external IP provider. Having the internal expertise is a critical element of our strategy to make serial a mainstream technology by providing designers with a comprehensive, scalable, and cost-effective solution.”

> “The underlying RocketIO technology that’s making it possible for Xilinx to support multi-gigabit systems is truly ‘rocket science’ – not something you’d want to simply license from an external IP provider”

“As the only FPGA vendor shipping plat-

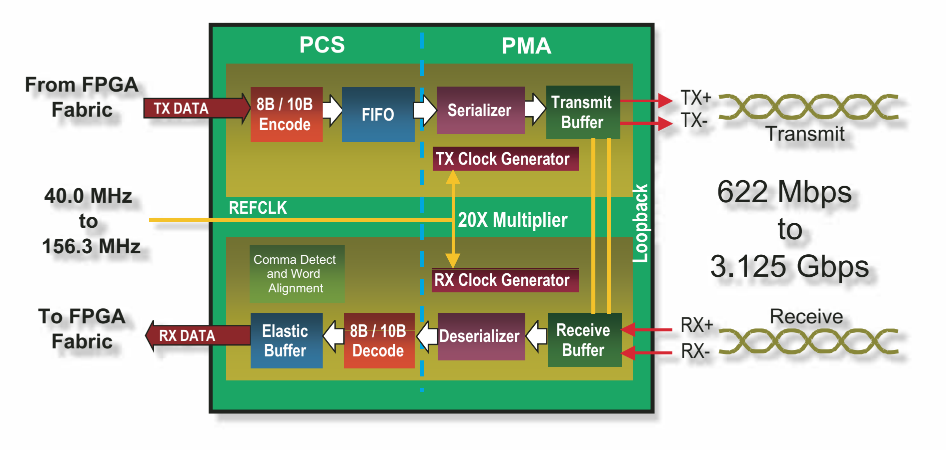

Figure 2 - RocketIO serial transceiver in Virtex-II Pro FPGAs

form devices with programmable 3.125 Gbps transceivers and IP cores for key serial interface standards, it's evident that our approach is paying off," he added. "Watch for many more exciting developments in the coming months as we continue to extend our market lead in serial connectivity."

Steve Berry, principal analyst for Electronic Trend Publications, evaluated the benefits of the serial I/O trend and the role Xilinx plays in the movement: "Through its leading technology and well-established strategic partnerships, Xilinx is poised to lead the industry in the transition to serial interfaces. System architects will experience dramatic improvements in bandwidth, pin count, power, and signal integrity."

Virtex-II Pro Platform FPGAs

Although it is widely accepted that serial I/O delivers significant advantages over parallel I/O methods, until now there was no flexible, cost-effective, general-purpose silicon support. The standards wars do not have clear winners, and the transition path is not obvious.

Virtex-II Pro Platform FPGAs (Figure 1) deliver state-of-the-art serial I/O with as many as 24 RocketIO transceivers embedded in the highest performance FPGA fabric in production. The RocketIO transceivers (Figure 2) operate at speeds from 622 Mbps to 3.125 Gbps. The product is available in a wide range of programmable logic densities in 10 devices and several packages. Virtex-II Pro FPGAs support all major emerging serial interfaces such as PCI Express, 1 Gb Ethernet PHY, 10 Gb Ethernet XAUI, Fibre Channel, OC-48, OC-192, and OC-768 SONET for backplanes, and Serial RapidIO.

Virtex-II Pro FPGAs also support parallel interface standards such as SPI-3 (POS PHY Level 3), SPI-4.1 (Flexbus 4), SPI-4.2 (POS PHY™ Level 4), 10 Gb Ethernet Media Independent Interface (XGMII), RapidIO, PCI, PCI-X, CSIX, HyperTransport™, XSBI, and SFI-4.

The Virtex-II Pro FPGA is the only device available today that enables bridging across all these interface classes and generations, making it the ultimate connectivity platform.

Xilinx Tools and Solutions for Serial Connectivity

Using the 3.125 Gbps RocketIO integrated transceivers in Virtex-II Pro FPGAs, Xilinx and its partners are delivering pre-engineered IP cores for serial connectivity protocols. Cores for 1 Gb Ethernet MAC with PHY, 10 Gb Ethernet MAC with XAUI, PCI Express, Fibre Channel, and reference designs for SONET OC-48 backplanes are available now, and more are being added. You can also build higher level protocols using the embedded IBM PowerPC™ processors. Reference designs and evaluation/prototype boards help you verify the performance of transceivers in real hardware.

Designing with parallel I/O meant you were limited to speeds of 33 MHz to 133 MHz. In the serial world, these speeds leap to 622 MHz, 3.125 GHz, 10 GHz, and beyond. This raises challenges with PCB design and signal integrity. Xilinx and leading EDA partners such as Cadence are solving this dilemma with tools such as SPECCTRAQuest™ transmission media and HSPICE™ (highly accurate simulation program with integrated circuit emphasis) models.

Xilinx provides in-depth characterization data (Figure 3) for the RocketIO transceivers for flawless system design using our Virtex-II Pro FPGAs. You will know exactly how your design is expected to operate, and you will be supported every step of the way with best design practices and PCB layout guidelines.

Aurora Reference Design

Aurora is a new, open, lightweight, scalable serial interface provided by Xilinx to help you transition from parallel to serial interfaces. It supports any transport protocol, has a compact architecture, and delivers low latency. A single-lane reference design is available for free download at: *www.xilinx.com/aurora*. A quad link reference design will be available during the first half of 2003.

Xilinx takes an active role in industry standards organizations – including PICMG, RapidIO Trade Association, NPF, OIF, PCI-SIG, XFP, SMPTE, and

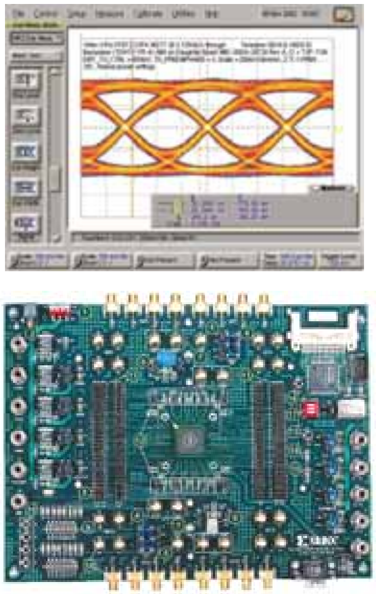

Figure 3 - RocketIO characterization using ML320 hardware platform: eye diagram at receiver shows 3.125 Gbps on 44-inch FR4, two connectors with 33% pre-emphasis.

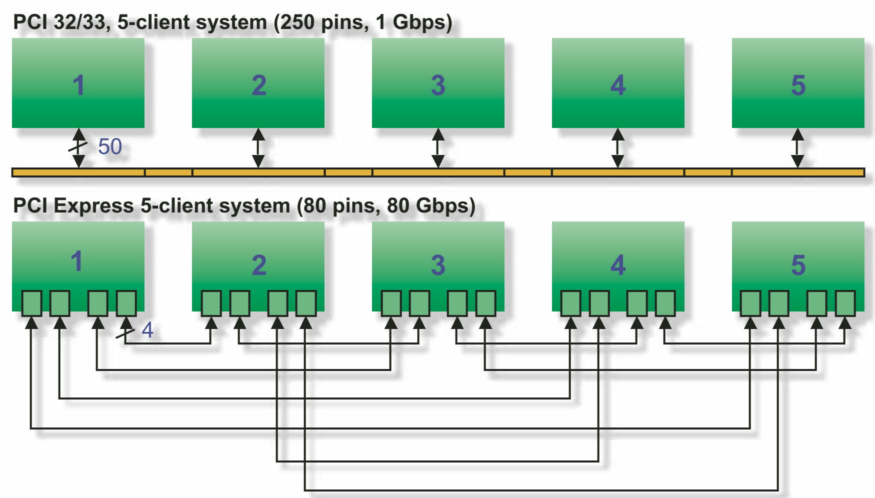
PCI 32/33, 5-client system (250 pins, 1 Gbps)
1
2
3
4
5
50
PCI Express 5-client system (80 pins, 80 Gbps)
1
2
3
4
5
4

Figure 4 - The PCI Express serial standard delivers 80X bandwidth with less than one-third of the pins required by the PCI parallel protocol.

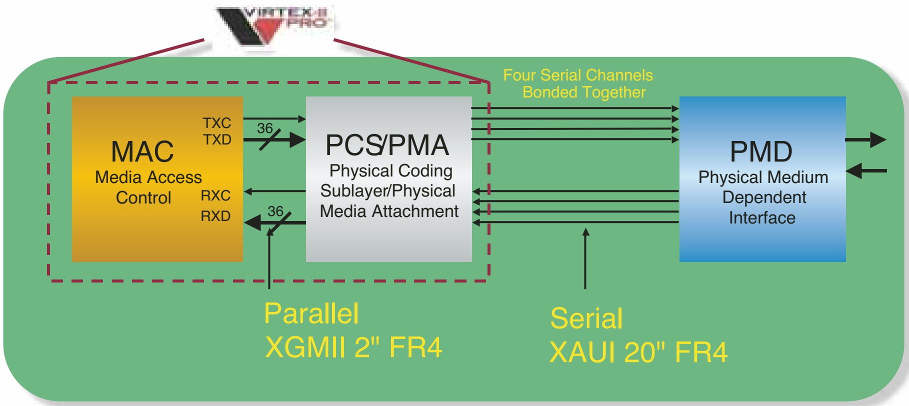
VIRTEX-II PRO
Four Serial Channels Bonded Together
MAC
Media Access Control
TXC
TXD
36
RXC
RXD
36
PCS/PMA
Physical Coding Sublayer/Physical Media Attachment
PMD
Physical Medium Dependent Interface
Parallel
XGMII 2" FR4
Serial
XAUI 20" FR4

Figure 5 - Serial XAUI allows 10 times longer PCB trace, lower pin count, and lower jitter/skew than XGMII.

others – with the goal of providing complete solutions synchronized with the availability of new standards. An example is the industry's first PCI Express core, which Xilinx released on the same date that the PCI-SIG ratified the specification. [*Ed. note: See "Xilinx Technology Enabled Instant Deployment of Real-PCI Express" in this issue.*]

The Serial Tsunami Initiative also offers several levels of training about working with serial technologies. These range from an online introduction to in-depth face-to-face classes. In addition, design support and services are available from experienced designers within Xilinx Design Services.

Manage Interoperability

The multiplicity of available and emerging serial standards is reflected in the range of interfaces embraced by available ASSPs. Your optimum design could easily incorporate multiple ASSPs equipped with a variety of interfaces. Xilinx and ASSP vendors work closely together to ensure interoperability with Virtex-II Pro features and electrical I/O standards as well as interface IP cores jointly verified in hardware. [*Ed. note: For a discussion of ASSP and Virtex-II Pro FPGA interoperability, see "Virtex-II Pro Platform FPGAs Deliver Proven Interoperability" in this issue.*]

Bridge Across Any Standard

The leading parallel standards are not going to vanish overnight, and you might find it necessary to accommodate an older parallel standard in a design that is intended to promote a newer serial standard. So how do you manage the transition? Virtex-II Pro FPGAs have solved this for you by supporting both serial and parallel interfaces within the same FPGA. Xilinx delivers IP cores and reference designs to interface with parallel standards such as 10 Gb Ethernet XGMII, RapidIO, SPI-3, SPI-4.1, SPI-4.2, HyperTransport, PCI, PCI-X, CSIX, XSBI, SFI4, and many others.

With Virtex-II Pro FPGAs you can bridge across parallel and serial interfaces to make a seamless transition. Plus, you can continue to interface with the ASSPs that best suit your needs based on their function, not on their interface support.

Case Studies – Serial vs. Parallel

Serial helps you break the bandwidth bottleneck with higher data rates using fewer pins. This lowered pin count delivers many advantages over traditional parallel implementations.

- Low device pin counts: With fewer pins per connection, you save costs from small board real estate, smaller packages, fewer PCB traces and layers, and even smaller connectors.
- Expansion and scalability: High-speed serial pipes scale to support higher data rates as needs change.
- Improved EMI and noise immunity: With embedded clock-data mechanism that is serial vs. wide parallel data and clock, you get higher clock rates at lower EMI, noise, cross talk, and skew.
- Physical interfaces: Serial I/O physical interfaces can also drive long PCB traces on backplanes or even external cable (copper or optical).

One obvious way to analyze the costs and benefits of I/O interfaces is to compare performance per pin. Let us take a look at two examples – PCI Express vs. PCI, and 10 Gb Ethernet XAUI vs. XGMII.

Case 1 – PCI and PCI-Express

PCI is a higher bandwidth, parallel, shared-bus standard, while the PCI Express protocol is a new serial version. A 5-client communication example (Figure 4) contrasts the two standards:

- A 32 bit/33 MHz PCI interface requires 50 pins, delivering about 1 Gbps

(32 bits x 33 MHz = ~1 Gbps) total bandwidth that would be shared among all 5 clients (5 x 50 pins = 250 pins).

- The PCI-Express interface operates at 2 Gbps data rate in each direction, delivering 4 Gbps full-duplex bandwidth over just four pins (that is, two differential pairs) for each point-to-point connection. The total aggregate bandwidth provided in the same 5-client configuration is 80 Gbps (4 Gbps/connection x 4 connections/-client x 5 clients = 80 Gbps).

Case 2 – 10 Gigabit Ethernet XGMII and XAUI

XGMII is a full-duplex interface between the MAC and PHY layers. XAUI is a serialized version of this interface (Figure 5).

- XGMII is a full-duplex, parallel interface operating at 312.5 Mbps per wire using 74 pins for data, clock, and control. Its data rate is 10 Gbps each way. Due to signal count, skew, and other problems, XGMII fails to support multiple interfaces in a single chip or routes longer than a few centimeters. Hence, it is restricted to be only a chip-to-chip interface and has a maximum FR4 trace limitation of 2 inches.

- XAUI is a 4-lane, full-duplex, serial interface, with each lane running at 2.5 Gbps data rate (3.125 Gbps baud rate). It requires 4 differential signal pairs in each direction and hence, 16 pins in total, delivering 10 Gbps aggregate data bandwidth. Automatic de-skew and pre-emphasis allows XAUI to route as much as 20 inches FR4 on PCBs, backplanes, and even cable. Its low pin count makes it highly scalable.

Fewer pins, higher bandwidth, scalability, and significant cost savings are all driving the move to serial.

Conclusion

The Xilinx Serial Tsunami Initiative is revolutionizing system architectures. State-of-the-art serial I/O delivers scalable, high-bandwidth at low cost. As the industry rapidly moves to serial connectivity, and away from current parallel interface schemes, you must either sink or swim. Xilinx and Virtex-II Pro Platform FPGAs offer a whole boatload of serial solutions from the initial concept to the finished product. Check it out at: *www.xilinx.com/serialsolution/*.

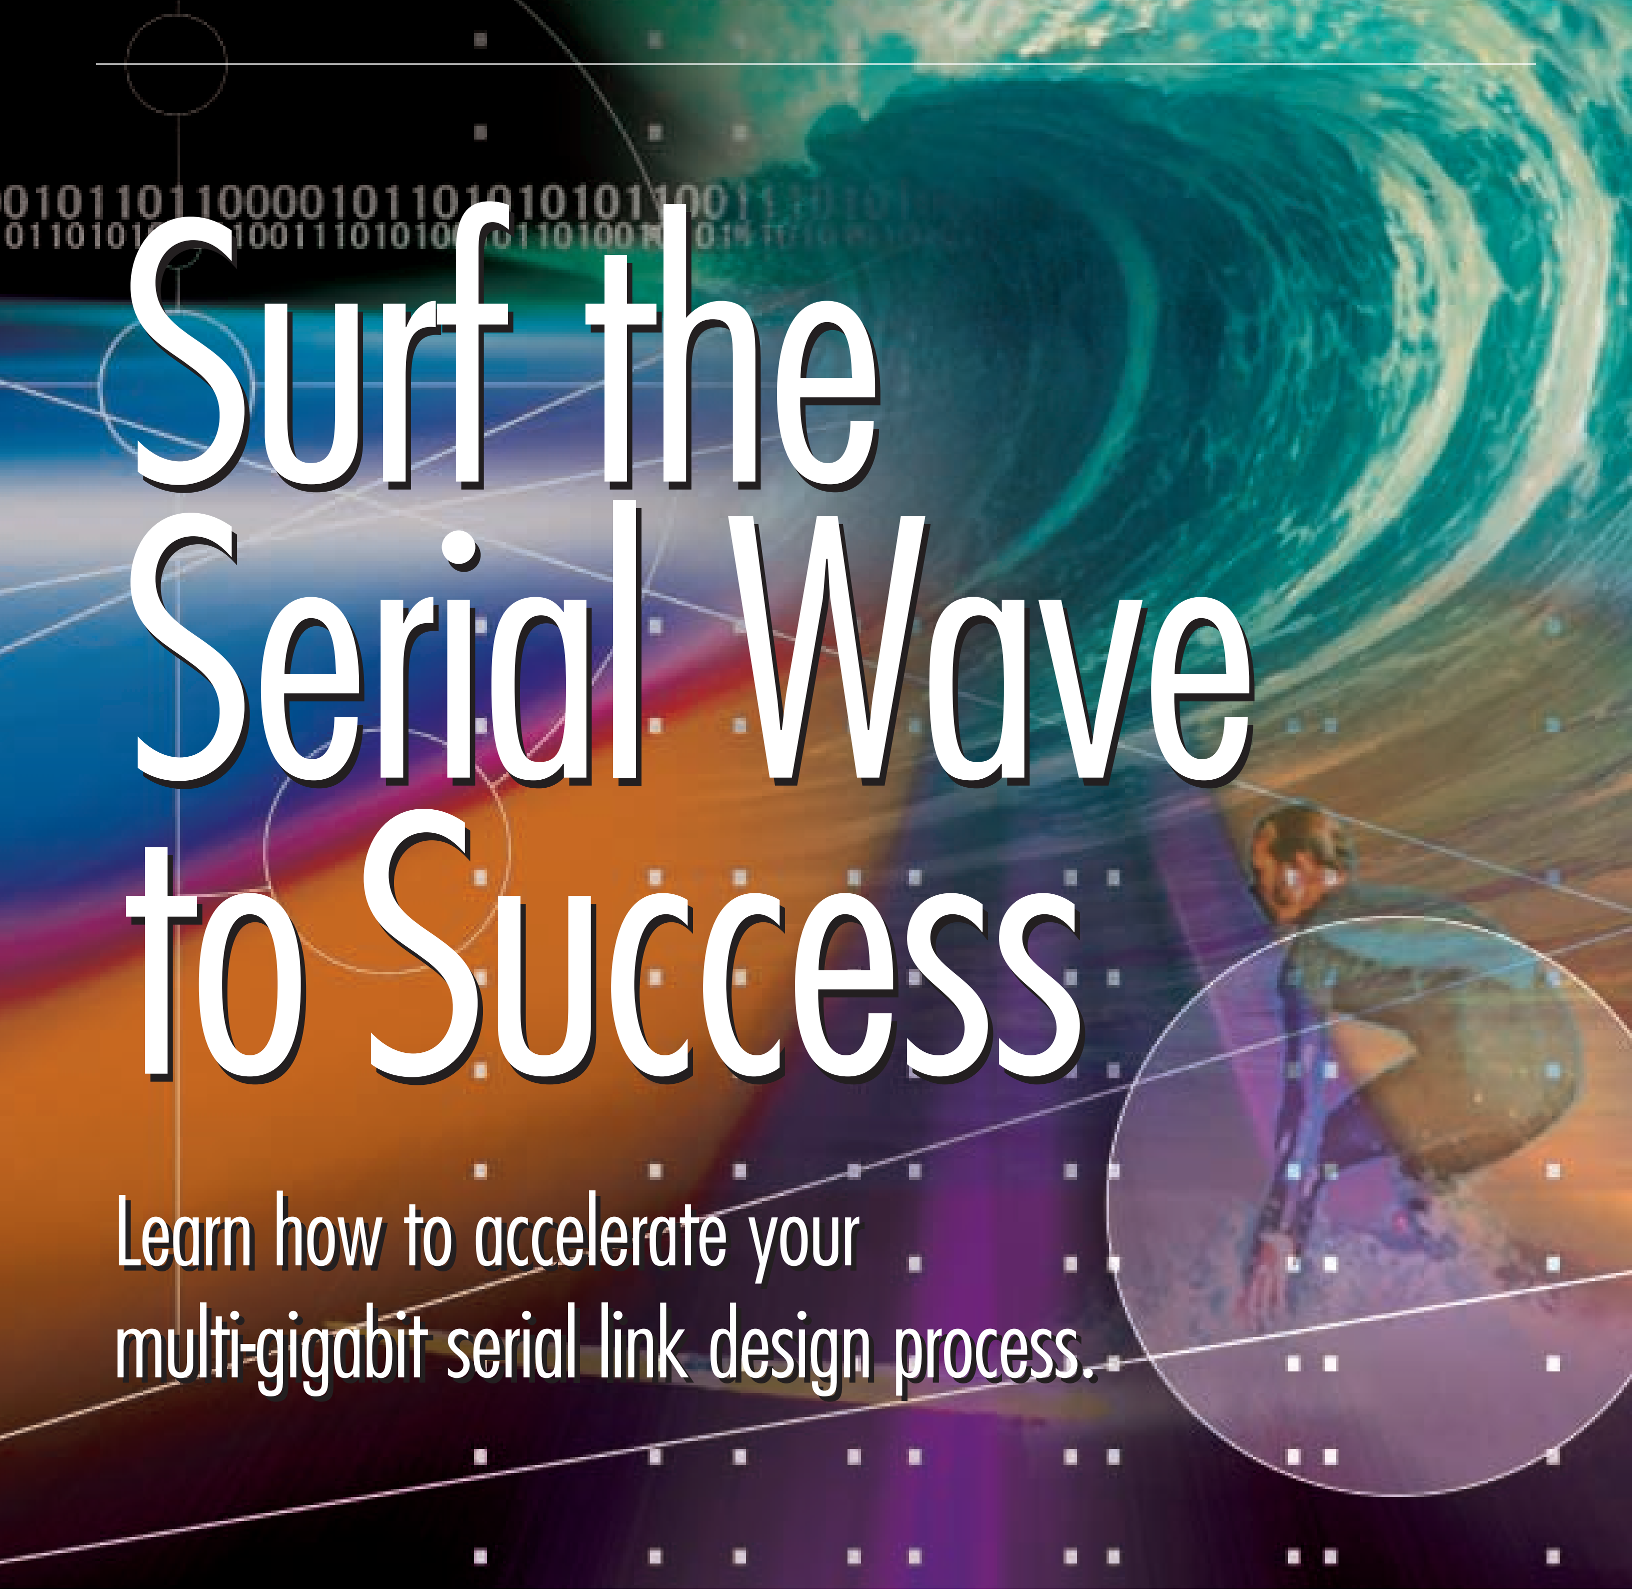

Surf the Serial Wave to Success

Learn how to accelerate your multi-gigabit serial link design process.

by Donald Telian
Technologist
Cadence Design Systems, Inc
donaldt@cadence.com

The move to multi-gigahertz (MGHz) serial technology represents a sea change of tsunami proportions. A variety of industry forces are revolutionizing product design to accommodate the speed and throughput of serial data transfer.

This article will show you how to harness the power of the Xilinx Serial Tsunami Initiative. We'll look at helpful tools and effective techniques already being used by engineers today.

Effective Data Transfers

Whether you're moving bits down an MGHz serial link or moving money into your bank account, it's important to make sure all the data is transferred correctly. You simply cannot afford to lose data.

Figure 1 illustrates two types of data transfers relevant in MGHz design. The first row shows the serial link itself. Data is sourced by the transmitter (Tx), transferred through the differential interconnect, and latched onto by the receiver (Rx). If all elements are not tuned to each other, the data is not transferred effectively. The transmission medium must be designed carefully – and all three elements (Tx, transmission medium, Rx) must be well-matched.

Similarly, the second row of Figure 1 shows the design chain between Xilinx serial technology and your design process. Just as in the case of the serial link, the Xilinx technology must be delivered to you in a medium that matches your design process to ensure a clean data transfer.

In cooperation with Cadence, Xilinx has developed the SPECCTRAQuest Design Kit as a way to effectively communicate the operation of the RocketIO™ MGHz transceivers found in the Xilinx Virtex-II Pro™ FPGAs. Later, we'll examine the Xilinx-Cadence partnership and how you can use it to accelerate your design process and improve your products. But first, let's take a closer look at the MGHz serial link itself.

How Serial Links Work

Measuring the "opening" on an eye diagram is a common way to judge the effectiveness of serial transmission. Figure 2 superimposes eye diagrams of received signals in three slightly different test cases. In the green signal's circuit, the transmitter, interconnect, and receiver impedances are well-matched. Here, all three elements are working together to produce an acceptably wide eye opening.

The other two waveforms in Figure 2 show what happens when only one of the three elements becomes imbalanced. In the blue signal's circuit, a mismatch in the impedance of the transmission line causes erratic signal behavior and a collapse of the eye opening. Changing the transmitter's impedance, however, causes an even further collapse in the red signal's circuit behavior.

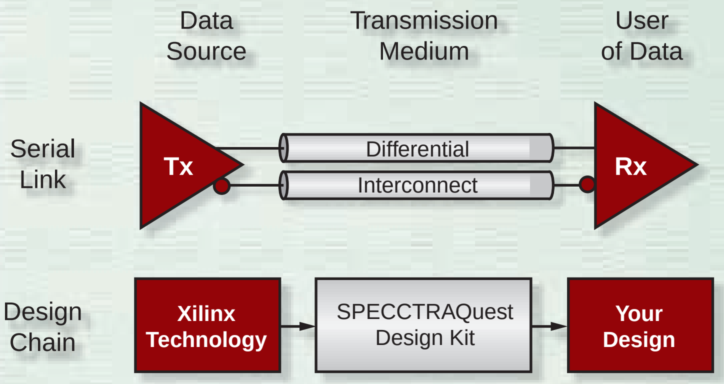

Figure 1 - How data is effectively transferred in a serial link, and its design chain.

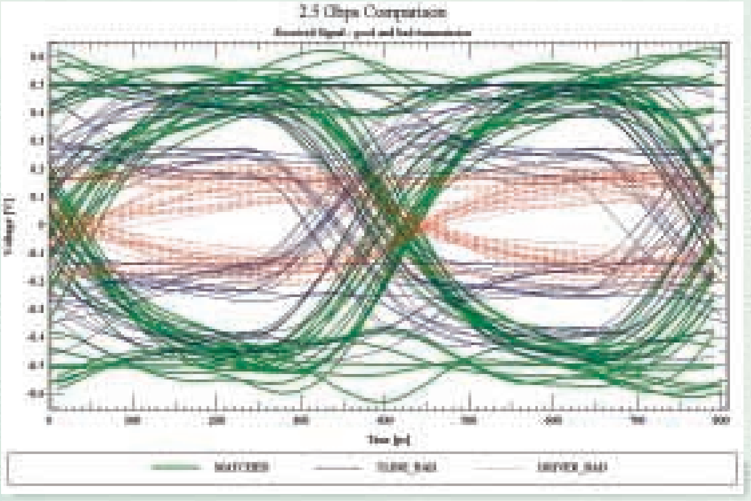

Figure 2 - Mismatch in any one of the serial link components can adversely affect performance.

Although the red signal appears more deterministic than the blue case, the transmitter in this case is not delivering enough voltage swing to the circuit to meet the thresholds in the receiver to extract the serial data.

Items that make an MGHz serial link work right include:

- Proper sizing of the transmitter for the required voltage swing
- An understanding of the differential impedance of the transmission medium (Z_differential is typically 2*[Z_uncoupled – Z_coupled])
- Matching that impedance with a termination resistor between the two nets at the receiver's inputs
- Thorough characterization and accounting for the interconnect's discontinuities and behaviors (such as vias, connectors, dielectric loss).

The SPECCTRAQuest RocketIO Design Kit

Recognizing that the MGHz design process has discontinuities too, Xilinx proactively developed the RocketIO Design Kit for Cadence's SPECCTRAQuest high-speed PCB design tool. This kit was first introduced with the Virtex-II Pro FPGA in March 2002, and was described in an *Xcell Journal* article at that time (see *support.xilinx.com/publications/xcellonline/partners/xc_speckit42.htm*). The kit helps you implement the RocketIO technology by providing the electronic files and models that match and can be inserted directly into your design process. Multimedia tutorials within the kit help you quickly understand the steps involved.

Mohammad W. Ali, Ph.D., a technologist at Tellabs, found the kit to offer significant improvements in both the throughput and quality of his design process. He states, "The new silicon package board solutions in the design kit save me a lot of time, particularly for my multiboard simulations that involve different styles of routed 2.5 GHz differential pairs. With the new interfaces in this SPECCTRAQuest Kit, I can accomplish my simulation task 10 to 20 times faster."

With RocketIO transceivers, signaling throughput has increased an order of magnitude. And with the accompanying design kit, the throughput of the design process has increased similarly as well – even with the challenges of MGHz design.

Design Chain Optimization

Great technology that is hard to use isn't really all that great. New technologies have failed because they were just too hard to access or too complex to work with. That's why Figure 1 shows the two parallel challenges that must be solved for high-speed serial communication to succeed:

1. Proper transmission of serial data from transmitter to receiver, and
2. Proper transfer of serial technology from Xilinx to you.

Focusing on the second challenge is what "design chain optimization" is all about. Design chain optimization is the only way to achieve the 10X to 20X design task improvement that the RocketIO kit has to offer.

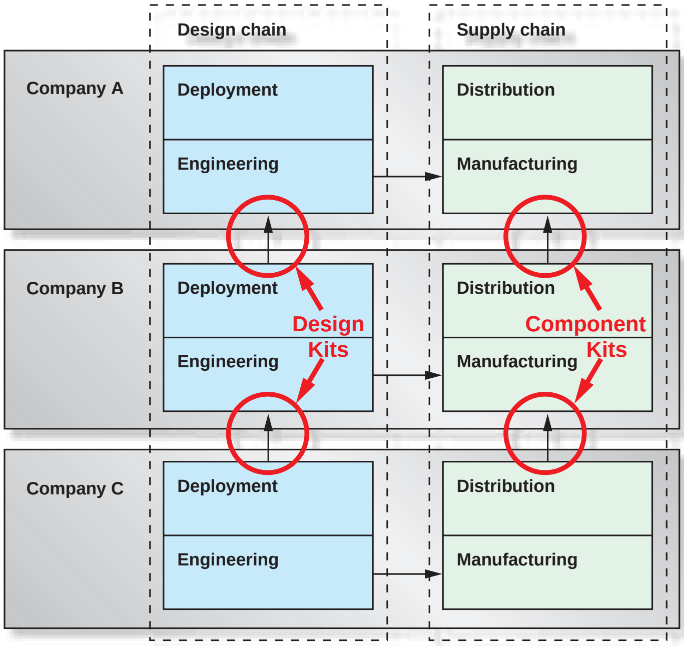

Figure 3 - Design and supply chains and the role of design kits.

Figure 3 illustrates the design chain. Because the term "design chain" is not as common as "supply chain," both are shown to help you understand their function and relationship to each other. Within the design chain, design kits of "virtual components" (in the form of models, EDA files, and databases) are transferred from the technology deployment group at one company to the engineering group of another.

In our example, the RocketIO kit effectively communicates the nuances of MGHz technology to Xilinx customers. This is done by avoiding the vagaries of textual datasheets, instead providing electronic files that can be easily inserted into your design process. These files are "executable specifications" that can quickly be understood by engineers all over the world, because the tool shows the RocketIO serial transceiver in a context with which they are familiar.

Bridging IC to PCB

Just as all elements in a serial link must be matched, so must the elements in the serial design chain. But here Xilinx had a challenge: the model formats commonly used by PCB designers would not work with this new technology. In fact, the only accurate representation of the RocketIO transceiver was the model used to design the silicon – an IC-level model that only worked in IC design tools.

As Figure 1 shows, the SPECCTRAQuest Design Kit answered this challenge and became an efficient "transmission medium" to bridge the worlds of IC and PCB modeling. New technology in the SPECCTRAQuest kit allows you to simulate arbitrary PCB layouts with complex IC models – all from the SPECCTRAQuest user interfaces commonly found in the high-speed PCB design process. If Xilinx had required PCB engineers to learn new IC simulation tools, it would have caused a mismatch in the design chain and hindered the adoption of RocketIO transceivers.

Wenwei Qiao, an engineer at Applied Materials, prefers using the Xilinx and Cadence kit's pre-packaged complex silicon models within the SPECCTRAQuest environment because they can be manipulated much like simpler IBIS-style models. "In only 10 minutes after installation, I was able to begin simulating my multigigabit solution," he reports. The user interface helps him focus on the design task and improve his product's quality instead of wading through thousands of lines of text-based models and netlists.

Stéphane Tessier, a hardware engineer at Radical Horizon, a Montreal-based software-defined radio (SDR) solution provider, agrees that the kits are a "must-have" for developing multi-gigabit links. He found that the tutorial information in the kits shortened his learning curve, and he believes use of the kits will "reduce the number of board iterations."

Conclusion

A survey of engineers currently using the kits revealed that they unanimously find them valuable for serial MGHz design. Already, 75% of the engineers believe that using the combined Xilinx/Cadence kit has helped them improve their product's quality.

During 2002, the integration of the SPECCTRAQuest and RocketIO design kits have become an integral part of the Xilinx Serial Tsunami Initiative – listed among *EDN* magazine's top 100 products for 2002.

The serial tsunami is here and growing. As you join fellow engineers in riding the serial wave, be sure to download your free copy of the SPECCTRAQuest Design Kit. It will help you put the power of MGHz signaling into your next design.

For More Information

The SPECCTRAQuest RocketIO Design Kit can be downloaded free of charge at: *support.xilinx.com/support/software/spice/spice-request.htm.* Registration and click-license NDA are required.

Information about Cadence SPECCTRAQuest (SQ) and other free SQ design kits is available at *www.specctraquest.com.*

An executive white paper on design chain optimization is available at *http://register.cadence.com/register.nsf/designChain/.*

Serial Tsunami Requires Changes to SI Verification and Planning

Mentor Graphics' HyperLynx products are designed to facilitate the transition from parallel to serial connectivity.

by Dave Kohlmeier
Director of Engineering, Simulation and Analysis Products, System Design Division
Mentor Graphics, Inc.
dave_kohlmeier@mentor.com

We see third-generation I/O (3GIO) serial connectivity standards showing up everywhere: HyperTransport™, InfiniBand™, and PCI Express™. Why are all the major platform vendors moving to these new interconnect schemes? Basically, the speed of this "Serial Tsunami" is being driven by the inability to cost-effectively resolve signal and power integrity issues prevalent in standard synchronous, parallel, multidrop bus designs. Power and ground noise due to large high-frequency current demands in large voltage, swing-single-ended designs; reflections from stubs due to multidrop connections; impossible delay/skew constraints for data versus clock nets; and other limitations have made cost-effective high-speed parallel design unattainable.

At the crest of the serial tsunami are differential signaling, embedded clocks, pre-emphasis, and point-to-point interconnect; all these elements give us a way around the electromagnetic barriers of parallel multidrop design. These changes require us to change our verification methodologies to assure successful designs. Let's take a look at some of the changes in detail, focusing on how they affect the need for up-front planning and verification.

Differential Signaling

Low voltage differential signaling (LVDS) is the real basis for the transition to 3GIO. LVDS, a two-wire system where return currents are expressly dealt with by the second closely coupled wire, is used to minimize ground-return loops and the cumulative effect produced by wide parallel buses on that return loop (otherwise known as simultaneous switching noise, or SSN). At these high frequencies, though, it is still very important to maintain a complete ground return path (no cuts or voids in the ground plane), especially for common mode currents.

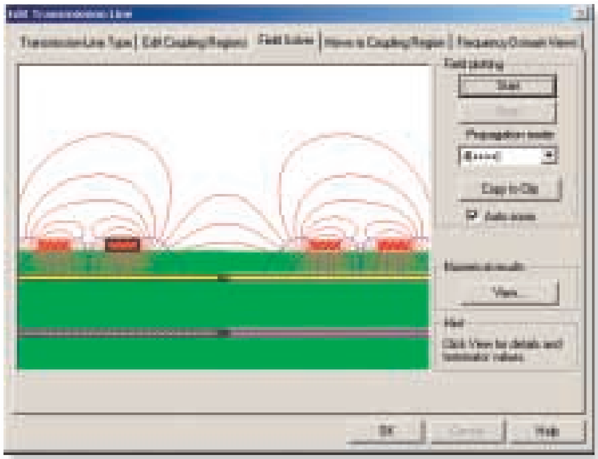

Figure 1 - Plotted EM field lines for two adjacent differential pairs

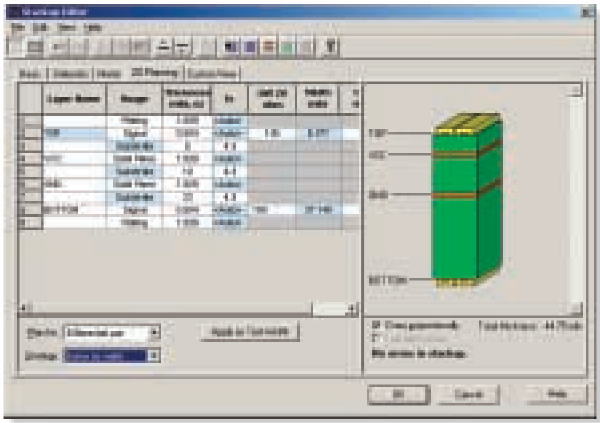

Figure 2 - Differential impedance stackup planning using the stackup editor

New 3GIO constraints for differential (diff) pair topologies include:

- Minimize coupling from adjacent pairs
- Match length of traces in a pair
- Minimize the number of vias used, but match use and location within pairs when they are used.

The idea here is that the more the diff pair is coupled, the more external noise will be rejected. An aggressor pair will induce a signal in the victim pair, but as such, it would induce the same signal in both wires (if they are closely coupled). Any noise agent that affects both wires in a diff pair will have no effect on the resulting differential waveform.

So, it is important that a verification tool understands the self-coupling of each diff pair and the coupling between pairs, including broadside (up and down a layer) and edge-coupled (side to side on same layer) pairs. Figure 1 shows the EM field lines calculated by the Mentor Graphics HyperLynx™ 2D field solver for two diff pairs. Notice the difference in the number of lines (which indicate field strength) between members of a pair versus those between the pairs themselves.

Impedance planning is an important step with diff pairs where internal terminations might require you to maintain specific differential impedance. For this reason, we have included an impedance-planning dialog in the HyperLynx stackup editor, as shown in Figure 2. In the editor, you can simply set a priority parameter, such as trace width, and request a spacing value that allows you to reach your goal (say, a differential impedance of 100 ohms).

Vias

At extremely high frequencies, using vias can increasingly introduce signal integrity (SI) problems. Why? Vias are a discontinuity in the transmission line.

Just visualize the nicely controlled impedance of a trace over a ground plane and then compare that to a vertical tube of copper with no corresponding ground return path – the impedance is clearly different. In these systems, it's important to keep diff pair vias close together and to add ground stitching vias in the vicinity for shielding and common mode return.

However, it's also important to simulate the impedance discontinuity. HyperLynx Version 7.0 has given you the ability to have via L and C calculated automatically, to specify L and C values for each via type, or to set a default via value for the entire board.

Transmission Line Losses

Without going into a lengthy discussion on losses in transmission lines, let's just stress the importance of your verification and planning tools in recognizing and supporting loss. As frequencies increase, AC losses are significant. Skin effect (where more current is forced to the surface of a conductor) and dielectric loss (thermal heating of the dielectric as the EM waves travel through the dielectric) are the culprits.

Especially in backplanes where the trace length on FR4 can be substantial, losses will reduce and "smear" voltages at the receiver. This reduces your ability to have a clean "eye" for the receiving IC to sync to, and extract, the correct data. HyperLynx Version 7.0 includes the trusted "W" element in its simulator for robust loss support.

High-Speed Models

Modeling and simulation go hand-in-hand. Whether you are doing analog or digital simulation, you can't get very far without models. In the SI business, the I/O Buffer Information Specification (IBIS) standard has been a huge benefit to systems designers. Virtually all IC vendors are now making I/O models available for their devices.

As we enter this next decade with 3GIO, IBIS is moving to support subcircuit models in SPICE (Simulation Program with Integrated Circuit Emphasis) or VHDL_AMS in the proposed IBIS 4.1 (now Bird 75) specification. In the meantime, simulation environments must support the models that are currently available, and those are predominately HSPICE (from MetaSoft).

HyperLynx Version 7.0 offers support for HSPICE models. These models are generally encrypted and require a license for the HSPICE simulator; Version 7.0 calls the HSPICE simulator from the HyperLynx environment. Model assignment and waveform analysis are the same as always, but behind the scene HyperLynx prepares a SPICE netlist of the entire electrical topology (including all coupling and loss elements) and then invokes HSPICE, extracts the results, and presents them in the native HyperLynx environment. This includes any multibit stimulus patterns you have defined.

For example, if you want to simulate a Xilinx RocketIO™ multi-gigabit transceiver implementation, simply assign the RocketIO model just as you would an IBIS model and use either your true topology extracted from your layout system (HyperLynx supports trace model extraction from all major PCB vendors) or a "what-if" topology defined in HyperLynx. Then hit the simulate button and view the results that are displayed in an eye diagram based on HSPICE results.

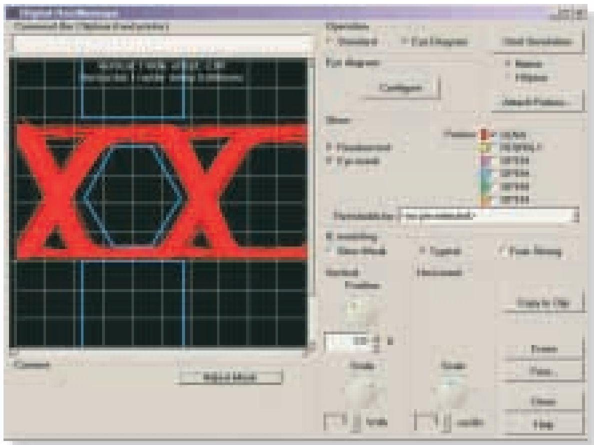

Figure 3 - Oscilloscope view of an eye diagram including eye mask

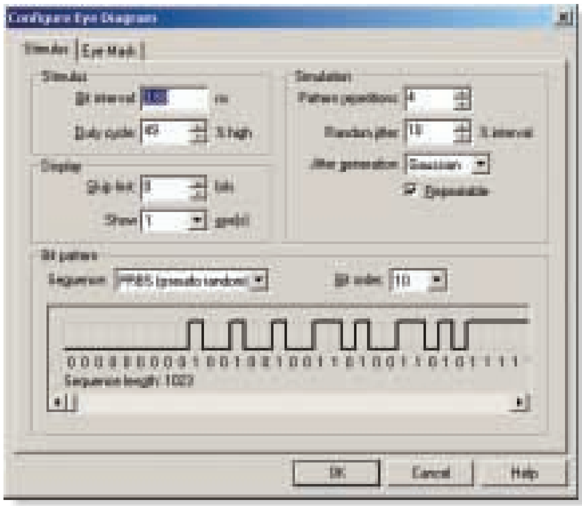

Sequence: PRBS (pseudo random)

0 0 0 0 0 0 0 0 0 1 0 0 1 0 0 1 0 0 1 1 0 1 0 0 1 1 0 1 0 1 1 1 1

Sequence length: 1023

Figure 4 - Multibit stimulus generation dialog

Eye Diagrams

For us digital designers, an eye diagram – as shown in Figure 3 – is nothing more than what we have been used to looking at in a logic analyzer for years – data bits and high and low voltage levels. The difference between a logic analyzer and an eye diagram is that in the eye, hundreds of bits are overlaid on each other so that we can see if there is a large enough window of time where the signals are reliably in one state or the other (necessary for the receiver clock recovery circuit to dependably extract the data).

As in digital simulation, we are required to create stimuli that can affect the signals in a realistic way, which in turn we see as a changing shape in the "eye." Because the bit history will affect the analog result, an-easy-to-use multibit stimulus editor in HyperLynx Version 7.0 has been added that allows you to create bit streams of ones and zeros for the simulator to drive the diff pair (Figure 4). Pseudo-random, 8-bit/10-bit encoded and customer defined patterns are supported, as well as random (uniform or Gaussian) jitter. This functionality, combined with easy-to-use differential probes, provides you with robust eye diagram support.

3GIO systems typically specify what a sufficient eye should look like for a receiver to extract the clock and data. This capability is referred to as an "eye mask." The ability to specify an eye mask is included in HyperLynx Version 7.0, giving you an easy way to "see" in the oscilloscope view if your resulting eye pattern meets the manufacturer's specification.

Conclusion

High-speed, low-voltage differential serial interconnect is the wave of the future for both interboard and intraboard interconnects. The features you need in a verification and analysis tool are changing to support this transition. Because tolerances are extremely tight in all aspects of these interconnect systems, the entire system must be simulated to assure first-pass success of the PCB design.

In summary, consider a quote from an Intel® white paper on PCI Express, "...detailed simulation and validation are necessary to guarantee a successful design." We at Mentor Graphics believe our latest release of the HyperLynx products will make this task easier and more productive. We wish you luck as you move into the gigahertz world of serial connectivity.

Could Microprocessor Obsolescence Be History?

Embedding soft processors in FPGA fabric offers a radical but robust new solution that effectively eradicates the problem of processor obsolescence.

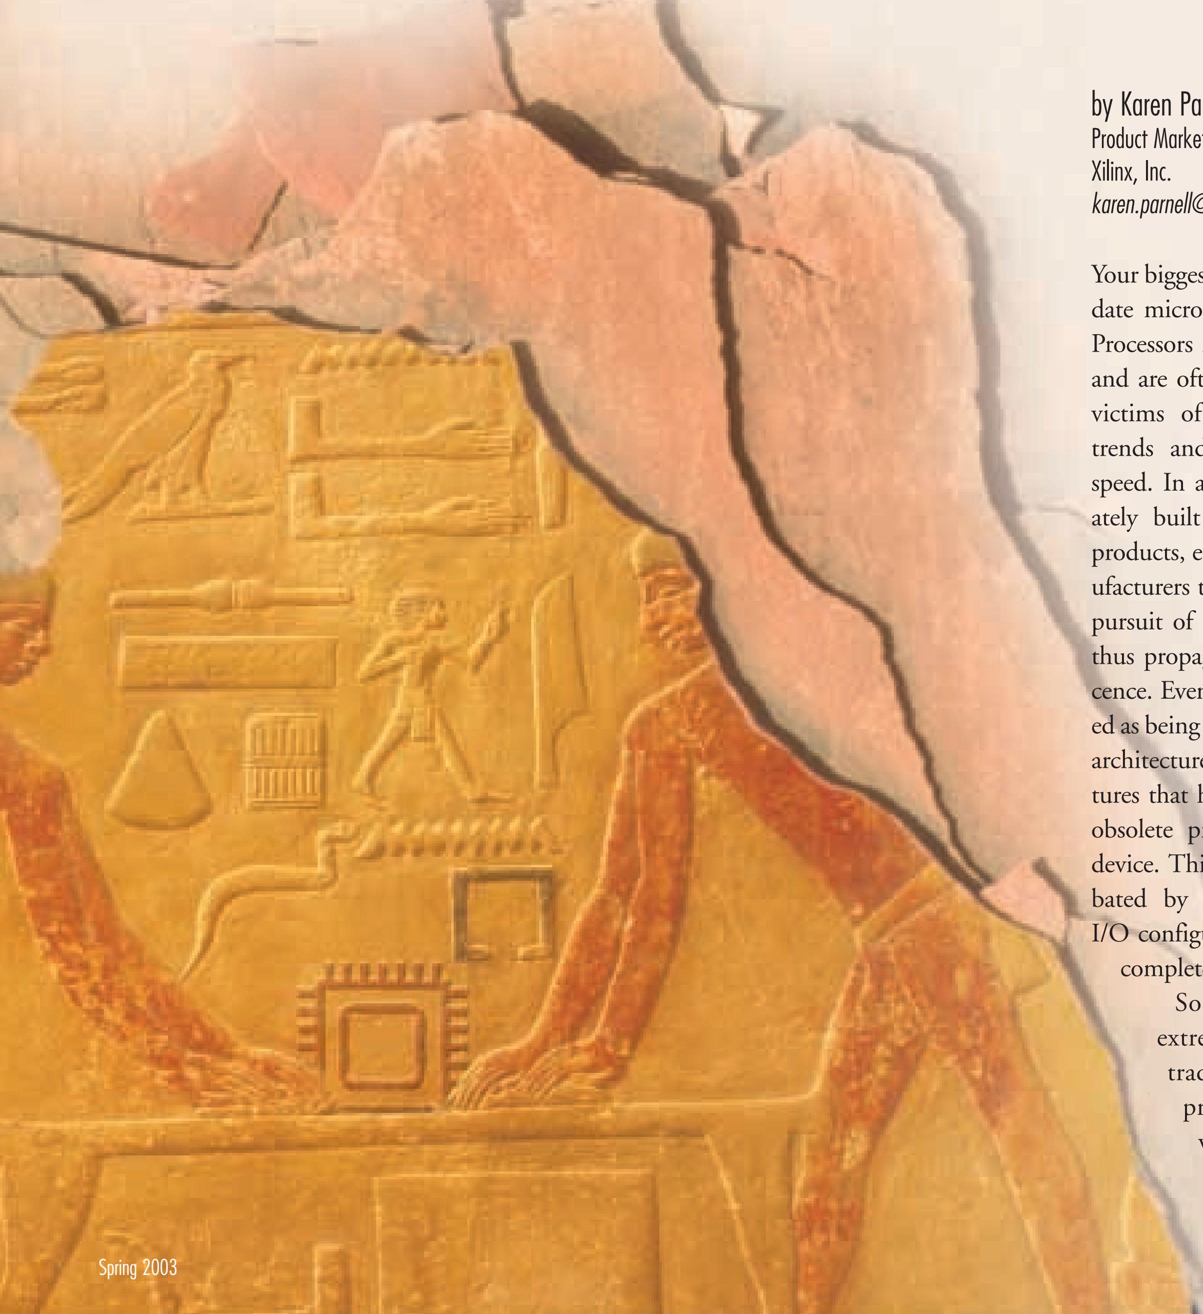

by Karen Parnell
Product Marketing Manager, Automotive
Xilinx, Inc.
karen.parnell@xilinx.com

Your biggest obsolescence problem is out-of-date microprocessors and microcontrollers. Processors have shorter life spans than ever and are often discontinued at short notice, victims of fluctuating consumer market trends and the demand for ever-greater speed. In addition, obsolescence is deliberately built into such items as consumer products, encouraging microprocessor manufacturers to abandon existing processors in pursuit of planned platform introductions, thus propagating a ripple effect of obsolescence. Even for designs coded in "C" (touted as being "portable code"), there are always architecture-specific instructions and features that hamper the changeover from the obsolete processor to the next-generation device. This changeover problem is exacerbated by different package options and I/O configurations, which can necessitate a complete board re-spin.

Soft processor cores offer an extremely attractive alternative to traditional approaches toward the problem of obsolescence, all of which are time-, labor-, and cost-intensive, and waste legacy development efforts.

A Look at Traditional Approaches

The predicament of automotive telematics equipment designers, for example, is representative of the problem. Although design and development time scales have shrunk from five years down to two, production is measured in years and the products remain active in the field for another 10 years plus. If we imagine a scenario in which every Electronic Control Unit (ECU) in a car contains at least one processor, and that every car contains up to 60 ECUs, this translates to a major problem every time a processor is obsoleted at relatively short notice.

Designers rely on several solutions to deal with processor obsolescence. The applicability of any given solution depends upon a number of variables: the value of the application software, the projected life of the system, and the amount of time and money available to solve the problem.

The most expensive solution is to redesign the system around a new processor. Depending on the volume of the code, a redesign can cost hundreds of worker-years, much of it devoted to validation and testing. Not only is the huge investment in debugging and refining the existing software lost – a clear case of throwing the baby out with the bath water – but the solution is temporary at best. If the system has a long projected life, the same problem will recur every few years.

Another solution is the Last Time Buy (LTB), which, on the surface, appears to be the most cost-effective option. The problem with this approach is that the designer must guess how much product to buy for the life of his program. If he guesses wrong, he is faced with an even more difficult problem: a larger legacy investment that must somehow be upgraded.

Inserting a new processor along with software written to emulate the old one is, at present, better in theory than in reality – although it's an appealing concept that does have some operational history. This option preserves the legacy software, so the process is relatively cheap and fast. But once again, the solution is not permanent. If the system has a long projected life, a new solution might have to be repeated every few years.

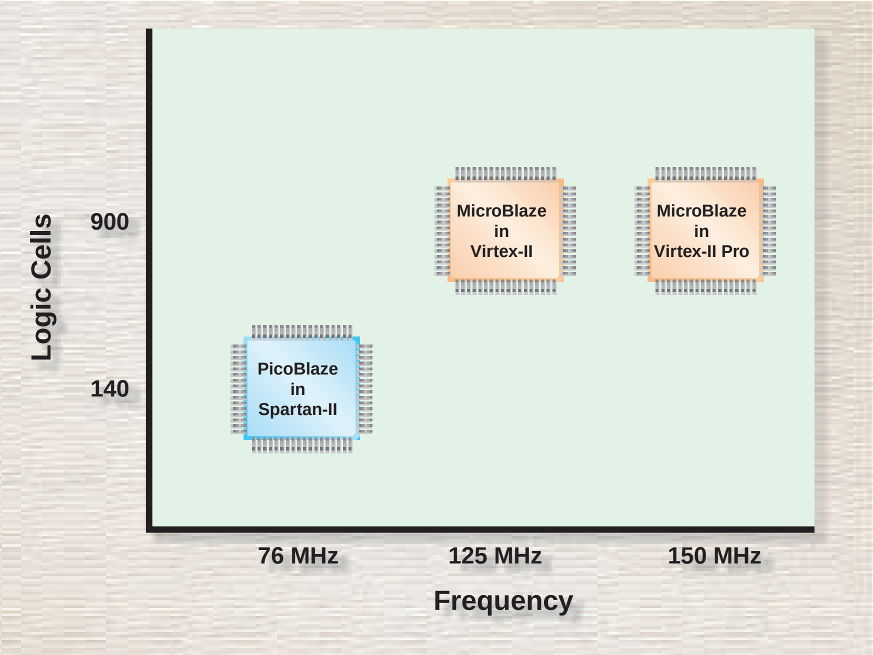

Figure 1 - Xilinx PicoBlaze and MicroBlaze soft processors

More important, software emulation is inherently a serial process. Because so much performance is consumed running the emulation rather than the application, it is slow. It has been shown empirically that emulation requires, on average, about 20 clock cycles of the new processor for every legacy instruction it executes.

In addition, emulation breeds further obsolescence, because the processor used as the emulation engine will itself become obsolete and may force an entire rewrite of the emulator.

Off-the-shelf solutions often merely substitute new headaches for old ones. Let's say you need a processor with 10 UARTs, an interrupt controller, and access to a block of external flash memory. Although there are many off-the-shelf processors offering multiple UARTs and other desired peripherals, they typically have numerous other peripherals that would go unused in your system. So, not only are you paying for the additional peripherals, but you often have to place unused peripherals into a safe mode or otherwise disable them via software.

Decommissioning unused peripherals places an additional burden on your software design team, who not only have to make the processor peripherals operate correctly, but now have to write code for the parts of the processor that are not used. Clearly, purchasing an off-the-shelf solution in this scenario would be highly wasteful, not only in terms of initial cost, but also in wasted engineering time during the design process.

The Soft Processor Solution

The soft processor solution eradicates processor obsolescence and preserves many

MicroBlaze Development Board Part Numbers	
Avnet (Silica)	
ADS-SP2-MB-EVL	MicroBlaze/Spartan-II Evaluation Kit
ADS-VE-MB-EVL	MicroBlaze/Virtex-E Evaluation Kit
ADS-VE-MB-DEV	MicroBlaze/Virtex-E Development Kit
Memec	
DS-KIT-MBLAZE-V2	Virtex-II Embedded Development Kit
Nohau	
EMUL-MicroBlaze-PC	High-Speed Debugger for MicroBlaze Soft Processor

Table 1 - EDK Development boards

years of legacy code and development. The new approach is to own the soft processor core and embed it in an FPGA host. Not only can you port the core to multiple FPGA platforms, but you can also design the peripheral set to meet your exact design requirements, thus eradicating architecture compromises and wasted peripherals.

The Xilinx MicroBlaze™ soft processor gives you the luxury of a different approach. Now, you can start with a processor core and build the peripheral set to meet your exact requirements. You've eliminated silicon wastage because you implement only what you need. You've reduced software design complexity because no code need ever be written to disable unwanted processor functionality. Creating unusual processor configurations – which can be changed at any time to suit changes in the specification – is now a simple task.

Even if after five or six years of field use, when the FPGA hardware may itself be nearing the end of its life, the soft processor core can simply be dropped into its new FPGA host, using the same "C" code. The hardware platform may need some PCB modifications but the legacy code remains usable and intact.

MicroBlaze and PicoBlaze Soft Processors

Xilinx offers both a 32-bit MicroBlaze soft processor core and an 8-bit PicoBlaze™ soft core. The PicoBlaze processor runs at speeds of 116 MHz, yet occupies a tiny footprint of just 35 configurable logic blocks (CLBs). (See Figure 1.)

The MicroBlaze 32-bit soft processor core is the industry's fastest soft processing solution. It runs at 150 MHz and delivers 100 D-MIPS. It features a RISC architecture with Harvard-style, separate 32-bit instruction and data buses running at full speed to execute programs and access data from both on-chip and external memory. A standard set of peripherals is IBM CoreConnect™ bus- enabled to offer MicroBlaze designers compatibility and reuse.

Xilinx EDK Solutions

Xilinx Embedded Development Kits (EDKs), including the soft processor core and a standard set of peripherals, are available from Xilinx and its distribution partners (see Table 1). The kits include a complete set of GNU-based software tools, including compiler, assembler, debugger, and linker.

MicroBlaze EDKs bought from Xilinx and its distribution partners also include development boards that support the Virtex™-E, Virtex-II, Spartan™-II, and Spartan-IIE series of FPGAs. Table 2 summarizes the specifications for the two processor cores.

Selected Xilinx FPGAs in the new IQ Solutions range have been qualified to operate over the -40ºC/-40ºF to 125ºC/257ºF temperature range and are designed for use in automotive applications, such as telematics systems.

Conclusion

Embedded in FPGA fabric, Xilinx soft processor cores such as MicroBlaze and PicoBlaze can eradicate processor obsolescence issues by providing a stable platform that is owned and configured by the designer. Used in combination with the new IQ Solutions extended temperature range FPGAs, they are ideal for such applications as automotive telematics.

Not only will you benefit from the flexibility, integration, and upgradeability offered by programmable logic, you can now take advantage of a processor tailored to your design needs that will not go obsolete, and will ultimately save the time and money associated with costly redesign.

For more information, visit these websites:

MicroBlaze Information
www.xilinx.com/microblaze/

PicoBlaze Information
www.xilinx.com/picoblaze/

Automotive IQ Solutions
www.xilinx.com/automotive/.

Soft Processor	Architecture	Bus	MIPS/Speed	Size	FPGA Support	Support
MicroBlaze	32-bit RISC	Harvard-style buses 32-bit instruction and data buses	100 D-MIPs 150 MHz	225 CLBs	Virtex Virtex-E Virtex-II Virtex-II Pro Spartan-II Spartan-IIE	Embedded Development Kit (EDK) – soft processor core, peripherals, GNU-based software tools (compiler, assembler, debugger, and linker)
PicoBlaze	8-bit	8-bit address and data buses	35 MIPS 116 MHz	35 CLBs	Virtex Spartan-II	Free-of-charge reference design and application note, assembler

Table 2 - Xilinx soft processors

Reduce Your RISC with a PicoBlaze Reference Design

Build your own soft microcontroller in a CoolRunner-II CPLD. It's easy, fast, and free.

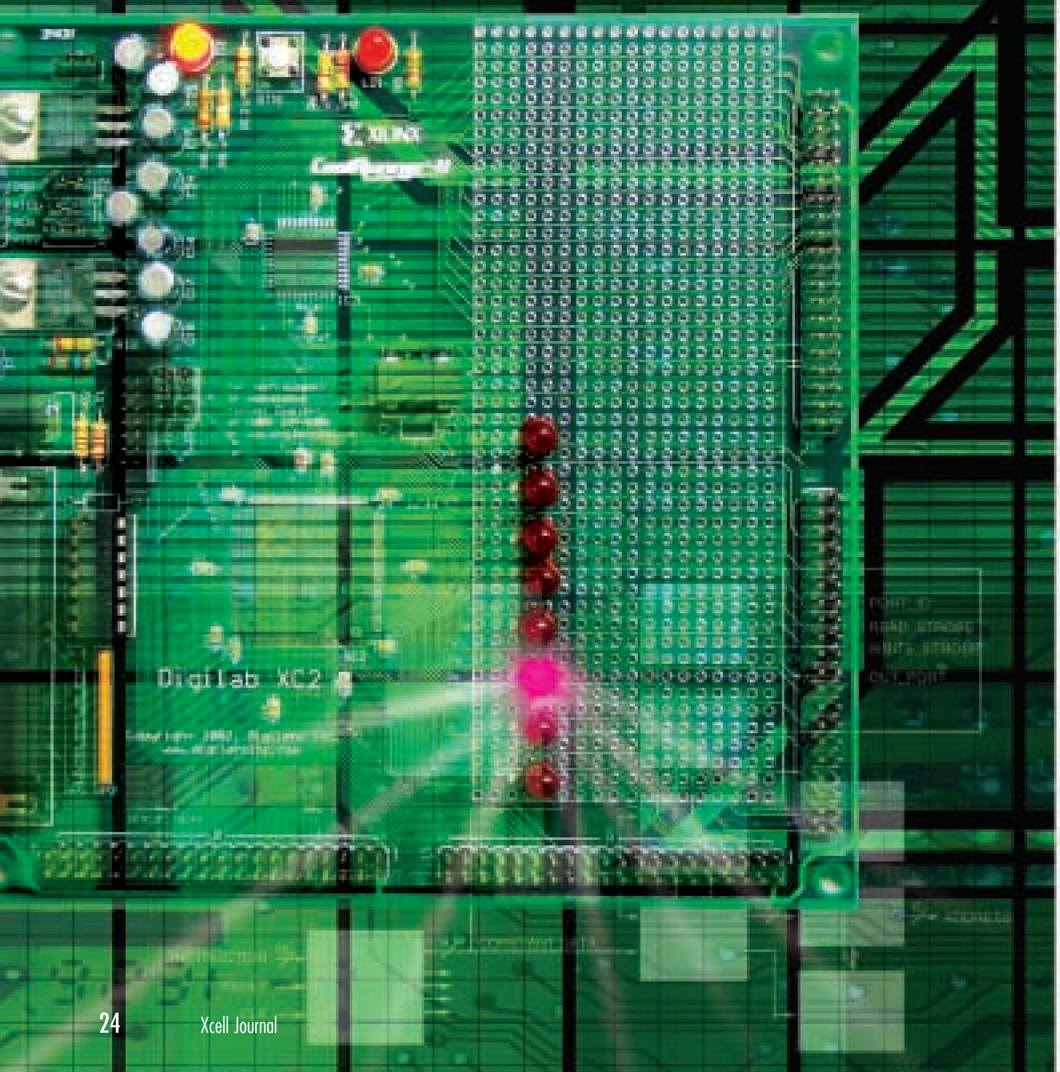

by Jesse Jenkins, Applications Manager
CPLD Business Unit
Xilinx, Inc.
jesse.jenkins@xilinx.com

Would you like to design your own microcontroller, but don't want all the hassle of starting from scratch with either the architecture or the support software?

Would you like to update older code from earlier microcontrollers to run faster and with the new memory standards?

Would you like to take the same code to substantially lower power dissipation?

Would you like to gain insight into the inner workings of a microcontroller with in-depth simulation as well as code creation support?

If the answer is yes to any of these questions, read on about the new Xilinx PicoBlaze™ microcontroller reference design. It's here now, it's fast, and maybe most important of all, it's free. This article details the CPLD version of the PicoBlaze reference design, including where to get the application code, examples, and the cross assembler. With that under your belt, you are ready to start – but first, a little more detail.

PicoBlaze Explained

The PicoBlaze soft microcontroller is an 8-bit design that supports an 8-bit data bus and 16-bit instruction bus. As you may have guessed, the PicoBlaze design is based on the RISC (reduced instruction set computer) "Harvard architecture" model with separate data and instruction ports. The PicoBlaze design is written in VHDL, and is intentionally documented so that the accompanying cross assembler directly tracks the architecture.

The PicoBlaze version currently shipping supports 49 instructions that operate within any of several Xilinx CoolRunner™-II CPLDs. The speed will vary depending on exactly which instructions you wish to support and which version of the architecture you choose.

For instance, with the full instruction

set and all instructions held outside the CPLD, you can expect to achieve about 30 MHz performance. By streamlining either the instruction set or the program, you can triple performance to 90 MHz.

Naturally, the PicoBlaze microcontroller architecture takes advantage of the two key CoolRunner-II features – high-speed execution and low power consumption.

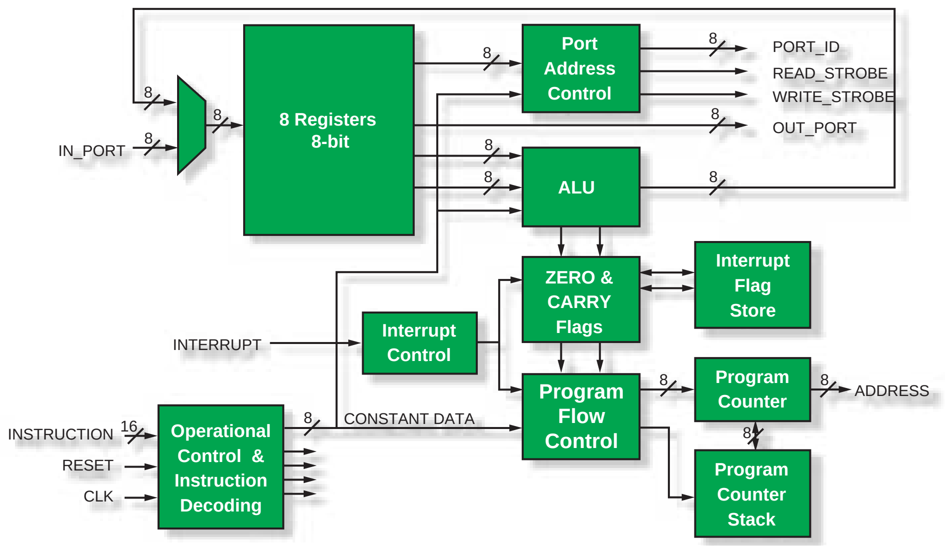

Figure 1 - PicoBlaze architecture

Add or Delete Instructions

Figure 1 shows the PicoBlaze base architecture, but don't restrict your thinking to that architecture alone. Think of it as a starting point. You are free to either add or delete capability as you see fit.

For instance, you can trim instructions from the instruction set by merely commenting them out of the VHDL. If you wish, you can also remove them from the assembler, but that is not required. You can also add instructions, if you have some application that can take advantage of essential instructions beyond those currently supplied.

It's also possible to do both – cut some instructions and add some instructions. For instance, most programmers use about 20 instructions in their day-to-day programming. Select the 20 you typically use, remove the rest, and then program. If you discover a bottlenecking "inner loop" that could benefit from a single instruction customized for that specific task, go ahead and write the VHDL that will do it at hardware speeds. Remember, the PicoBlaze microcontroller uses DualEDGE flip-flops within the processor to accomplish computation on both clock edges.

A DSP Example

To illustrate the ability of the PicoBlaze architecture to adapt, let's look at an example from DSP. The code to bit-reverse a bus is a fundamental operation used in Fast Fourier Transforms. The value is then typically driven out on the address lines as a critical step in the base algorithm. To do this in "standard" instructions would take multiple "mask and rotate" commands, creating a processing bottleneck.

Figure 2 shows the basic operations in assembler-like steps to display register contents. The algorithm starts with a byte of data labeled A-H. This byte is first internally swapped (four rotates), then successively, inner bits are picked off with Boolean AND/OR into a target register that will build up the results, two bits at a time. One pass through this results in the final register with the original contents reversed. Depending on algorithm details, it can take approximately 12 to 18 instructions. In this case, we dispense with adding the overhead of loop management with pointers and counters and unroll it.

In Figure 3, the "flip" instruction is added to the VHDL, the design is recompiled, and the processor is "rewired" to add in this key instruction. This method collapses many instructions down to some gate rewiring with the synthesis tools. Very many bit-level operations boil down to simply rewiring the CPU, and best of all, the synthesizer does the work. Many other examples exist. See the links at the end of this article.

A	B	C	D	E	F	G	H	Initial value
E	F	G	H	A	B	C	D	Swap nibbles
0	0	0	1	0	0	0	1	Mask 2 bits
0	0	0	H	0	0	0	D	Or Int. Result
D	E	F	G	H	A	B	C	Right Rot Swap
0	0	H	0	0	0	D	0	Left Rot Int. Result
0	0	0	1	0	0	0	1	Mask 2 bits
0	0	H	G	0	0	D	C	Or Int. Result
C	D	E	F	G	H	A	B	Right Rot Swap
0	H	G	0	0	D	C	0	Left Rot Int. Result

0	0	0	1	0	0	0	1	Mask 2 bits
0	H	G	F	0	D	C	B	Or Int. Result
B	C	D	E	F	G	H	A	Right Rot Swap
H	G	F	0	D	C	B	0	Left Rot Int. Result
0	0	0	1	0	0	0	1	Mask 2 bits
H	G	F	E	D	C	B	A	OR Int. Result

Final result is bit-reversed

Figure 2 - Bit-reversal code steps

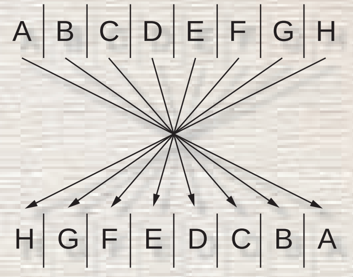

Figure 3 - Recoded instruction for bit-reversal ("flip") operation

Processor Enhancement

We just mentioned instruction set optimization, but it's also possible to add functionality. Remember that many microcontrollers include on-board function blocks that have a payoff beyond the instruction set. For instance, many 8-bit microcontrollers include internal peripheral counters or timers, interrupt handlers, and DMA circuits. With PicoBlaze, just add the right set of peripheral capability within the chip, depending on the density of the CoolRunner-II CPLD chosen. Table 1 shows the densities available in CoolRunner-II CPLDs, and Table 2 gives some estimates of macrocell usage for various add-on functions.

One very important thing to remember is that when choosing a function to add in, select just the functionality actually needed, so you will get the best usage from your choice. Bill Carter, one of the founders of Xilinx, frequently comments that most people don't really want or need a UART (universal asynchronous receiver/transmitter), but only an "RT." That is, they select a function that comes with 50 options, then only use two. They end up carrying along lots of unused circuitry that they pay for but never really use. Don't fall into that trap. Select the functions you will need and get the cheapest, fastest, lowest power solution possible. Note that most of the items listed in Table 2 exist as separate reference designs and may be found on the Xilinx website.

Performance Improvement

Getting the most out of your design will be another step. A classic way to improve the design is to "tune" it. Observe its performance behavior, identify where the processor is spending its time, discover what it is doing, and think through the best set of operations to improve. Then, implement a new version of the architecture and/or code and evaluate it again.

One easy way to do that is with the CoolRunner-II Design Kit (see Figure 4). Many target designs easily fit onto the 256-macrocell XC2C256 that resides on that board. There is also a blank pinout site for adding a 64 macrocell XC2C64, with signals already attached to the XC256. Simply construct a small hardware performance monitor that will time various code sections and report back the execution time. That way, by examining the behavior over address space and time, you can determine just how much time is spent doing the various tasks.

Figure 5 shows a simple approach to doing this operation. With care, both

CoolRunner-II	XC2C32	XC2C64	XC2C128	XC2C256	XC2C384	XC2C512
Macrocells	32	64	128	256	384	512
MAX I/O	33	64	100	184	240	270
TPD (ns)	3.5	4.0	4.5	5.0	5.5	6.0
TSU (ns)	1.7	2.0	2.1	2.2	2.3	2.4
TCO (ns)	2.8	3.0	3.4	3.8	4.2	4.6
Fsystem1 (MHz)	333	270	263	238	217	217

Table 1 - CoolRunner-II macrocell capacities and pertinent data

Function	Macrocells
IrDA and UAR/T	87
Timer/Counter	16 and up
DMA Port	16-32
Manchester Encoder/Decoder	55
Wireless XCVR	156
16b/20b Encoder/Decoder	76
Flash NAND Interface	9
SPI Interface	135
UAR/T	61
DDR SDRAM Interface	128
SMBus Controller	158

Table 2 - Common functions and approximate CoolRunner-II macrocell usage

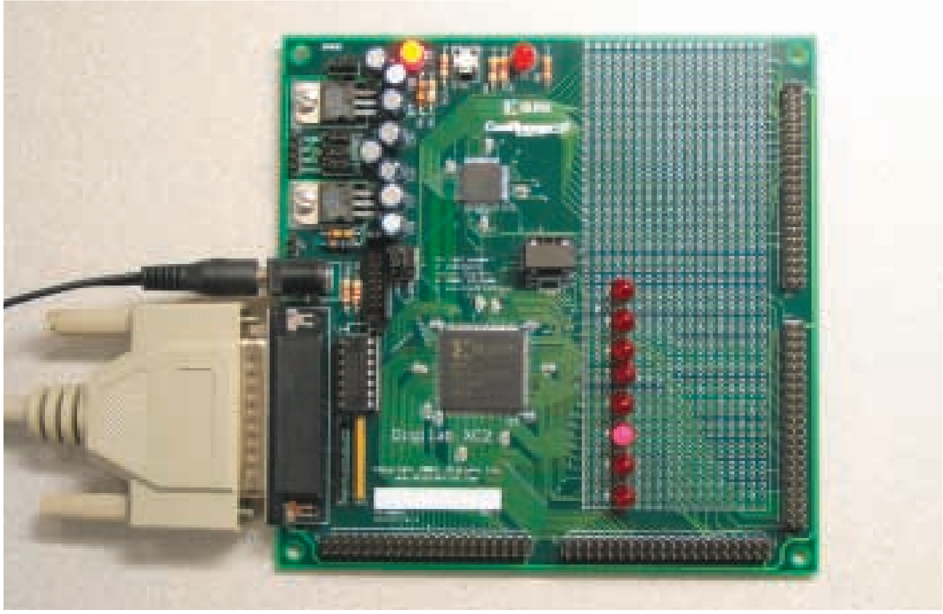

Figure 4 - CoolRunner-II Design Kit

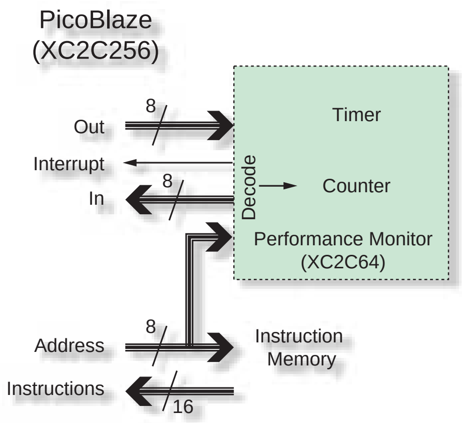

Figure 5 - PicoBlaze performance monitor

CPLDs can communicate through a PC parallel port via their JTAG boundary scan chains. Performance monitoring can help decide which aspects of the PicoBlaze design to perform in software and which to embody in hardware. One beautiful thing about building functions out of programmable logic is that different experiments to focus on specific performance targets are easily developed – giving you highly tuned, fast designs.

And you don't have to worry about power enhancement. CoolRunner-II CPLDs are already the lowest power CPLDs available today, and PicoBlaze is a very competitive low power microcontroller.

PicoBlaze Cross Assembler

As mentioned before, the PicoBlaze cross assembler is so well documented that a direct correspondence between assembly code and VHDL in the PicoBlaze design file already exists. The translator is written in ANSI-C and is assembled on Microsoft assemblers. The cross assembler is highly transportable and supports multiple output file types. For instance, it produces a binary output file ready to load into external EPROM in Intel hex format.

The cross assembler also produces the essential modeling files for the VHDL simulator. You can instantly analyze your produced code with high-speed simulations to determine the functionality and effectiveness of the code you have implemented. Then download the code into the CoolRunner-II Design Kit and see that they actually do work as expected.

Conclusions and Recommendations

This introduction to designing PicoBlaze microcontrollers was written to stimulate your imagination to discover the fascinating world of creating your own CPUs. Once you start, it can be addictive. You can easily alter the processor to be 16 or even 32 bits wide – or even non-binary powers. Then you will discover which instructions burn through the macrocells and what the new speed limits will be.

Do you need some instructions for 8 bits and others for 16? It's up to you. Application areas where these kinds of processors make sense include industrial control, low-power portable DSP (brainwaves, EKG, medical), and cryptography. Did you know that most cryptographic operations are bit-level operations and almost never do floating point arithmetic?

The PicoBlaze microcontroller reference design for CPLDs has been built, tested, and is now available over the Internet, free to the user.

Acknowledgements

Xilinx engineer Ken Chapman, who received additional support and encouragement from Henk van Kampen at Mediatronix BV, developed the original reference design. The CPLD version was created by Scott Lien, who also wrote the PicoBlaze cross assembler and the application note that is on the Xilinx website. The VHDL and cross assembler (source and executable) links are shown in xapp387 (see below).

For more information, see the following:

CoolRunner-II Design Kit:
www.xilinx.com/products/cpldsolutions/demoboard.htm (purchasing details)

CoolRunner-II Application Notes
www.xilinx.com/xapp/xapp375.pdf (timing model)
www.xilinx.com/xapp/xapp376.pdf (logic engine)
www.xilinx.com/xapp/xapp377.pdf (low-power design)
www.xilinx.com/xapp/xapp378.pdf (advanced features)
www.xilinx.com/xapp/xapp379.pdf (high-speed design)
www.xilinx.com/xapp/xapp380.pdf (cross point switch)
www.xilinx.com/xapp/xapp381.pdf (demo board)
www.xilinx.com/xapp/xapp382.pdf (I/O characteristics)
www.xilinx.com/xapp/xapp383.pdf (single error correction, double error detection)
www.xilinx.com/xapp/xapp384.pdf (DDR SDRAM interface)
www.xilinx.com/xapp/xapp387.pdf (PicoBlaze microcontroller)
www.xilinx.com/xapp/xapp388.pdf (on-the-fly reconfiguration)
www.xilinx.com/xapp/xapp389.pdf (powering CoolRunner-II CPLDs)

CoolRunner-II Data Sheets
www.xilinx.com/bvdocs/publications/ds090.pdf (CoolRunner-II CPLD family data sheet)
www.xilinx.com/bvdocs/publications/ds091.pdf (XC2C32 data sheet)
www.xilinx.com/bvdocs/publications/ds092.pdf (XC2C64 data sheet)
www.xilinx.com/bvdocs/publications/ds093.pdf (XC2C128 data sheet)
www.xilinx.com/bvdocs/publications/ds094.pdf (XC2C256 data sheet)
www.xilinx.com/bvdocs/publications/ds095.pdf (XC2C384 data sheet)
www.xilinx.com/bvdocs/publications/ds096.pdf (XC2C512 data sheet)

CoolRunner-II White Papers
www.xilinx.com/publications/products/cool2/wp_pdf/wp165.pdf (chip scale packaging)
www.xilinx.com/publications/whitepapers/wp_pdf/wp170.pdf (security)

Helping To Make Your CPLD Designs Better...

CoolRunner-II Development Board

Available from Nu Horizons

This Board is Picoblaze Ready

Being able to get your device to market in real time is the name of the game. The CoolRunner-II Evaluation Board from Nu Horizons is a very flexible testing platform that allows engineers to evaluate the XC2C256 low power CPLD from Xilinx in a typical environment.

What sets this board apart from other boards is its robust features and the XC9572XL that allows the engineer to interface 5V signals to the XC2C256. All at a cost of only $99.

For more information on the Nu Horizons CoolRunner-II 256 Macrocell Evaluation Board, including features, benefits, schematic and users manual, please visit us on-line at www.nuhorizons.com/cr2

If you are new to designing with CPLDs, and would like more information on how these devices can benefit your application, please visit our dedicated Xilinx CPLD Solution Center for a complete source of Xilinx CPLD products and support software. www.nuhorizons.com/CPLD

CoolRunner-II™

1-888-747-NUHO

XILINX

NU HORIZONS ELECTRONICS CORP.

Xilinx FPCs Target Cost-Sensitive Applications

Xilinx field programmable controllers combine the power of the 32-bit MicroBlaze soft processor cores with the versatility of Spartan-IIE FPGAs to deliver the greatest number of I/Os for low-cost processing solutions.

by Helen Yu
Processor Solutions Marketing Manager
Xilinx, Inc.
helen.yu@xilinx.com

Staying ahead of the competition is getting tougher everyday. Cost pressures, changing standards, and device obsolescence are just a few of the challenges. To maintain leadership, you need a low-cost, competitive processing solution that's customizable throughout the entire design cycle and can quickly be brought into high-volume production.

The Xilinx field programmable controller (FPC) solution allows you to create low-cost, customized processors with the peripherals, memory, and logic you want – all on a single, cost-optimized Spartan™-IIE FPGA. With the flexibility to allow integration of other intellectual property (IP) cores on the FPGA fabric, the Spartan-IIE family presents an ideal embedded solution.

Traditional Design Issues

The introduction of embedded soft processors has offered substantial benefits to the world of digital electronics. The industry's huge appetite for increasingly intelligent and sophisticated control systems has dictated the rapid advancement of processor technology. Perhaps most prevalent of all is the huge leap in demand for embedded microcontrollers and microprocessors.

If you are using a traditional microcontroller unit (MCU) or microprocessor unit (MPU) in your application, selecting the proper device is one of the most critical decisions that will ultimately determine the success or failure of your design. Typically, you must address a number of issues, including:

- Is the MCU or MPU affordable? Does it minimize the overall cost of the system while still fulfilling the design specification?
- Does the unit have the required number of I/Os? Too few can't do the job; too many can lead to excessive cost.
- Does the device have all the required peripherals? Can you add your own? Does the unit include peripherals you don't need?
- Are you paying for unneeded IP?
- Will the MCU or MPU be available over the long term? Will the processor become obsolete?

These questions are just some of the criteria to consider when choosing a traditional MCU or MPU for your application. The Xilinx FPC solution, however, makes many of these concerns irrelevant.

The FPC Solution

Using the 32-bit MicroBlaze™ soft processor core in the Spartan-IIE FPGA, the FPC offers a true low-cost solution for real-time processor control. The combination of a high-performance soft processor and a low-cost FPGA enables you to rapidly develop programmable systems for cost-sensitive applications.

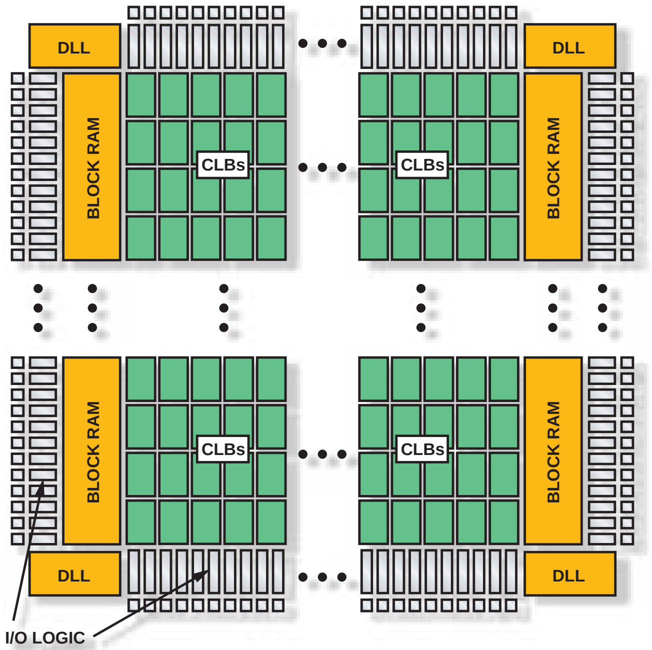

Figure 1- Spartan-IIE FPGA block diagram

The enhanced integration of embedded applications means fewer interface issues, and processor-based designs developed in the FPGA can more easily be updated without changing the PC board.

Spartan-IIE FPGAs

The Spartan-IIE 1.8V family of FPGAs achieves high-performance, low-cost operation through advanced architecture and the latest semiconductor technology. The seven-member family (see "Spartan-IIE Family Grows" in this issue) offers densities ranging from 50,000 to 600,000 system gates. Spartan-IIE devices also provide system clock rates beyond 200 MHz.

The Spartan-IIE FPGAs have a flexible, programmable architecture of configurable logic blocks (CLBs), surrounded by a perimeter of programmable input/output blocks (IOBs). There are four delay-locked loops (DLLs), one at each corner of the die. Two columns of block RAM lie on opposite sides of the die, between the CLBs and the IOB columns. The XC2S400E has four columns and the XC2S600E has six columns of block RAM. A powerful hierarchy of versatile routing channels interconnects these functional elements. Figure 1 shows a block diagram of a Spartan-IIE FPGA device.

This flexible platform is an ideal base for implementing a controller system. An embedded microcontroller takes the concept of integration one stage further by permitting you to embed the controller system into a small section of a programmable device. No longer does the microcontroller have to exist in a stand-alone package; it can now be embedded deep within custom hardware.

Spartan-IIE FPGAs are customized by loading configuration data into internal static memory cells while permitting unlimited reprogramming cycles to become a viable upgrade path for future product enhancements. Therefore, Spartan-IIE FPGAs are ideal for shortening product development cycles while offering a cost-effective solution for high-volume production.

In addition, the Spartan-IIE family delivers a cost-effective platform with high numbers of I/Os to provide excellent I/O expansion (up to 514 user I/Os). Table 1 compares the number of I/Os in two traditional microcontrollers against the number of I/Os delivered by two Spartan-IIE FPGAs.

MicroBlaze Soft Processor

The MicroBlaze 32-bit RISC soft processor is a true 32-bit processor, supporting 32-bit bus widths. The core is a RISC-based engine with a 32-bit LUT RAM-based register file with separate instructions for data and memory access.

The MicroBlaze soft processor supports both on-chip block RAM and external memory. All peripherals use the same IBM CoreConnect™ OPB bus as the IBM PowerPC™ processor – which means the processor peripherals are hardware compatible with the PowerPC processors on Virtex-II Pro™ FPGAs.

The MicroBlaze embedded system, including the MicroBlaze core and selected processor peripherals, is shown in Figure 2. Several peripherals are available to support the MicroBlaze processor, including memory controllers, UART, GPIO, I^2C, 10/100 Ethernet MAC, and many more.

The MicroBlaze core offers the flexibility and scalability of embedded processor programmable logic devices. The MicroBlaze processor requires less than half the logic resources yet offers more than twice the performance of competing soft processors, as measured by industry-standard Dhrystone-MIPS (D-MIPS) benchmarks. Delivering 49 D-MIPS of performance at 75 MHz, the MicroBlaze processor occupies only 1,050 logic cells in the Spartan-IIE FPGA.

FPC Applications

As shown in Figure 3, FPCs have significant applications in the traditional 16- and 32-bit microcontroller and microprocessor markets, which include automotive, industrial, and high-end consumer applications.

Automotive

The modern automobile is replete with microcontroller-based systems providing automated control for just about every conceivable part of the car. Braking systems use microcontrollers to deliver advanced safety features, such as ABS and traction control. Windshield wipers are controlled to bring us timed interval wipes and even rain-sensitive automatic wiper activation. Heating controls for the vehicle interior monitor multiple zones within the passenger compartment, automatically adjusting the supply of air to maintain the desired temperature. Seats even remember the favored positions for different drivers of the car and readjust themselves automatically.

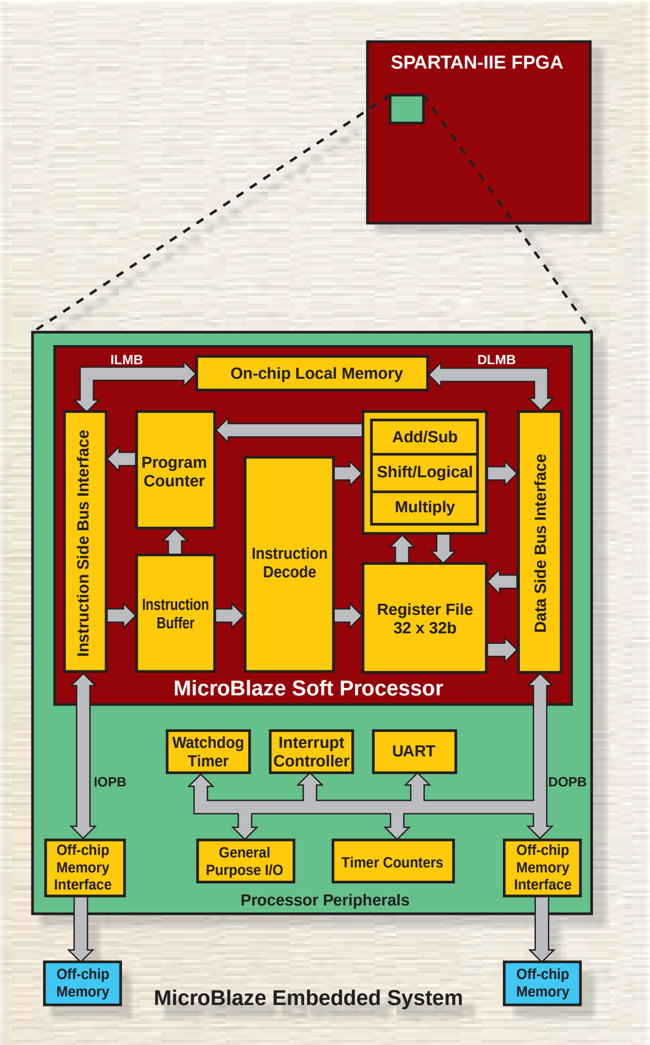

Figure 2 - MicroBlaze embedded system

Industrial

No longer must large workforces be trained to monitor a specific area of a production plant. Today, a series of microcontroller-based monitoring modules, often linked to a central station, replaces these human counterparts. Microcontrollers work tirelessly around the clock without lapses in concentration, requiring only the most minimal maintenance and supervision.

Consumer

Walk into any electronic store and you will find a host of products that use some kind of microcontroller: MP3 players, video recorders, Web tablets, televisions, plasma displays, set-top boxes, refrigerators, washing machines, telephones, answering machines, ovens, toasters, printers, and scanners all offer added functionality through the use of a microprocessor.

The Benefits of FPCs

The Xilinx FPC solution combines the low-cost Spartan-IIE FPGA family with the compact, high-performance MicroBlaze 32-bit RISC processor core. You also get complete embedded system tools (EST) support, including GNU compiler and debugger; hardware and software development tools for implementation, simulation, and verification; and more than 30 fully parameterizable processor IPs.

Overall, FPCs offer a low-cost, high-performance, and easy-to-use solution. Among the benefits of FPCs are:

- Reduced costs – By integrating your design onto a single device, you not only save time and effort, you also reduce your overall costs. Spartan-IIE FPGAs are the lowest cost programmable

Traditional Microcontrollers	# I/Os
NEC V850E/IA2 (32-bit microcontroller)	100-pin package
Motorola MC68HC912B32 (16-bit flash microcontroller)	80-pin package

Spartan-IIE FPGAs	# I/Os	Advantages
S-IIE XC2S600E	514-pin package	5x more I/O than NEC, 6x more than Motorola
S-IIE XC2S300E	329-pin package	3x more I/O than NEC, 4x more than Motorola

Table 1 - Comparison of Spartan-IIE devices and traditional microcontrollers

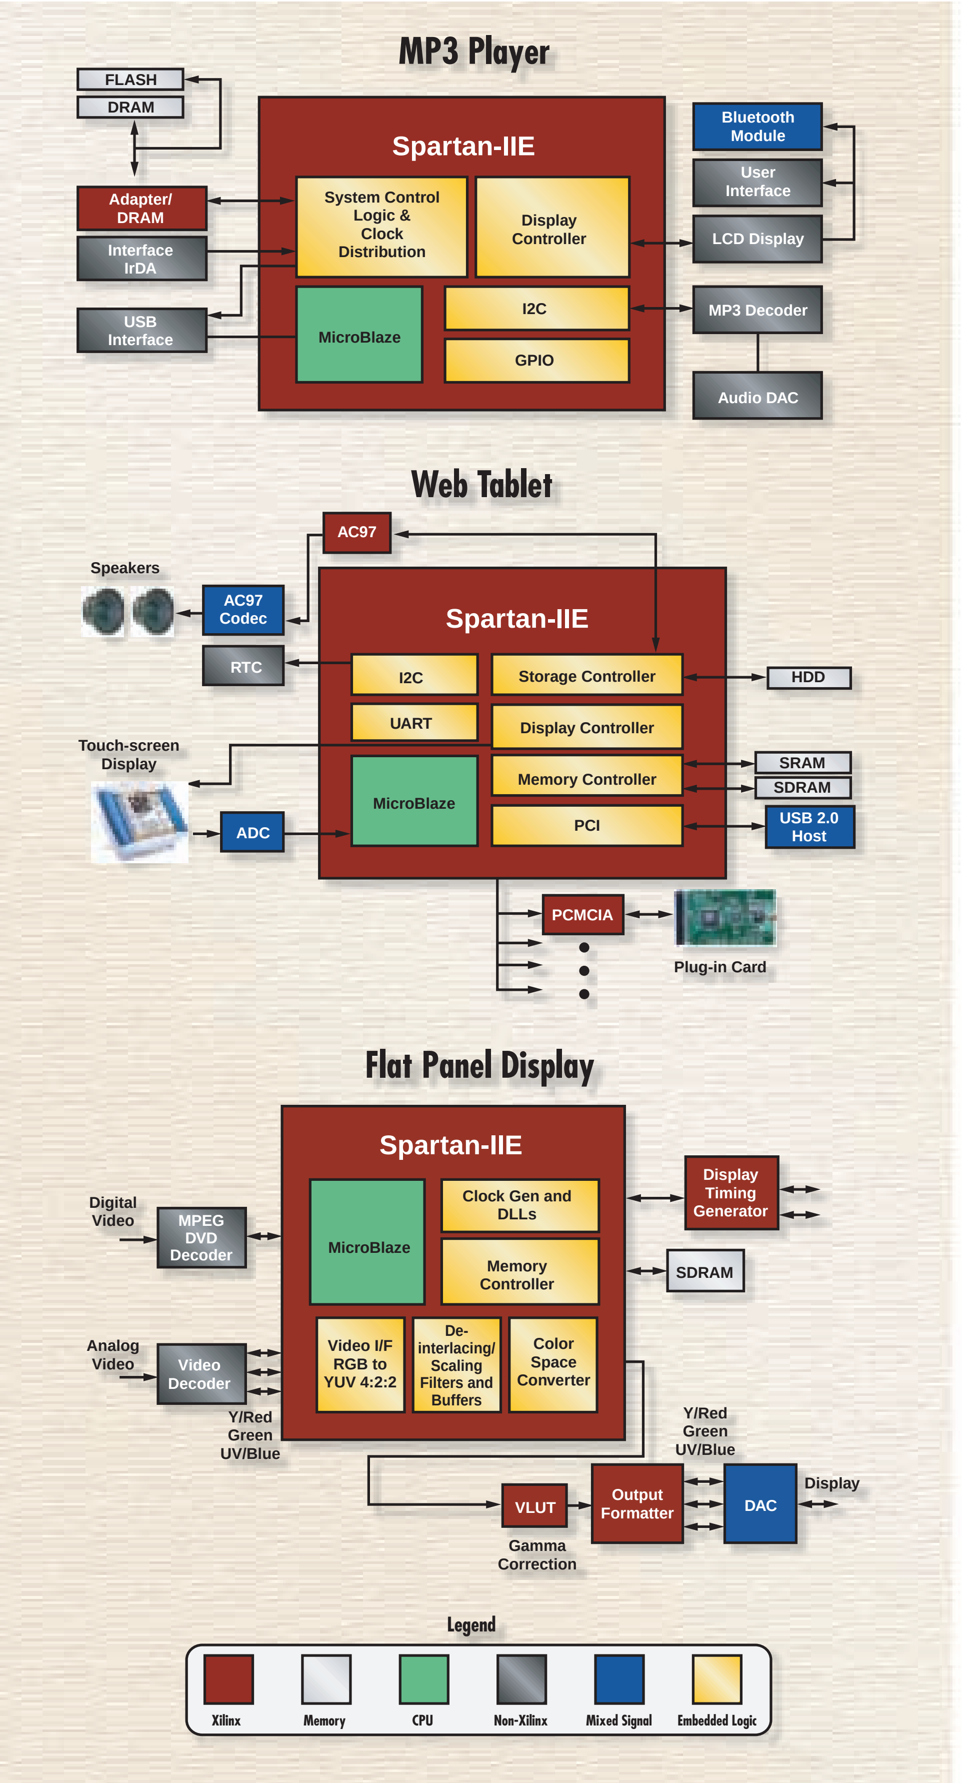

Figure 3 - FPC applications gallery

logic devices you can get. They also allow you to integrate costly board-level features such as DLLs, RAM, a variety of I/O translators – as well as MicroBlaze, the industry's fastest soft processor – into a single, compact, low-cost platform.

- More user I/Os – Spartan-IIE FPGA devices contain as many as 514 user I/Os, with more than 70% additional capacity, and the lowest cost per I/O than competing FPGAs in the same density ranges. This competitive advantage allows you to integrate more features and also shrink the form factor for each device.

- Customization – You can create a customized controller and peripheral set to meet your exact and evolving design requirements. Unlike other solutions, with FPCs you are no longer locked into a rigid, pre-selected set of peripherals – or have to pay for unused function sets. In addition, Xilinx offers more than 30 processor IPs to choose from.

- No obsolescence – Xilinx allows you to purchase the MicroBlaze soft processor core source code. This option guarantees product availability for any application you choose. You can also port the core across Xilinx product lines – even target an ASIC device.

Conclusion

FPCs from Xilinx deliver the highest I/O for low-cost processing solutions. They create a customized controller and peripheral sets to meet your exact – and evolving – design requirements. Freed from a fixed set of peripherals, you now have the power to custom tailor your peripheral set to meet your needs. In addition, the opportunity to purchase the MicroBlaze structural VHDL source code assures you of product availability well into the future. FPCs reduce your overall design cost and inventory while bringing you the highest performance from your logic devices.

For more information on the FPC solution, visit *www.xilinx.com/fpc/*.

Memec Lead by Design Training

Embedded Processor Solutions Workshop

During this hands-on training course, learn how to use the new Embedded Development Kit software to architect processor-based Xilinx solutions. Easy-to-follow lab exercises provide practical experience with:

- Defining your hardware system
- Performing simulation
- Adding peripherals
- Implementing the design
- Writing application code
- Debugging (in-system)

A MicroBlaze™ and Virtex-II Pro™ PowerPC™ version of the course focuses on the details of the processor development flow. During the MicroBlaze class, you will experiment with the new Memec Design Spartan™-IIE LC Development Kit; likewise, the Virtex-II Pro PowerPC course allows you to explore the features of the Memec Design Virtex-II Pro Development Kit.

Exclusive to Attendees: Specially Bundled Kits at Reduced Pricing!

With special kit pricing, you can continue your development efforts with the same hardware and software you used during the hands-on labs.

Virtex-II Pro PowerPC Development Kit

Explore the power of an embedded PowerPC processor, Rocket I/Os, and the most advanced FPGA architecture available.

- Virtex-II Pro Development Board
- Programming Cable
- Serial and MGT Loop-back Cables
- Power Supply
- Xilinx EDK Software
- Documentation and Reference Designs

Spartan-IIE MicroBlaze Development Kit

The ideal starter kit for developing MicroBlaze-based applications.

- Spartan-IIE LC Development Board
- Programming Cable
- Power Supply
- Xilinx EDK Software
- Xilinx ISE BaseX Software
- Documentation and Reference Designs

What Are You Waiting For?

Classes take place in March and April, in locations throughout North America.

Sign up now for the Embedded Processor Solutions Workshop nearest you.

XILINX®

Insight MEMEC

Call **800.488.4133** ext.**980** or register at www.**insight.na.memec**.com/embedded_workshops.

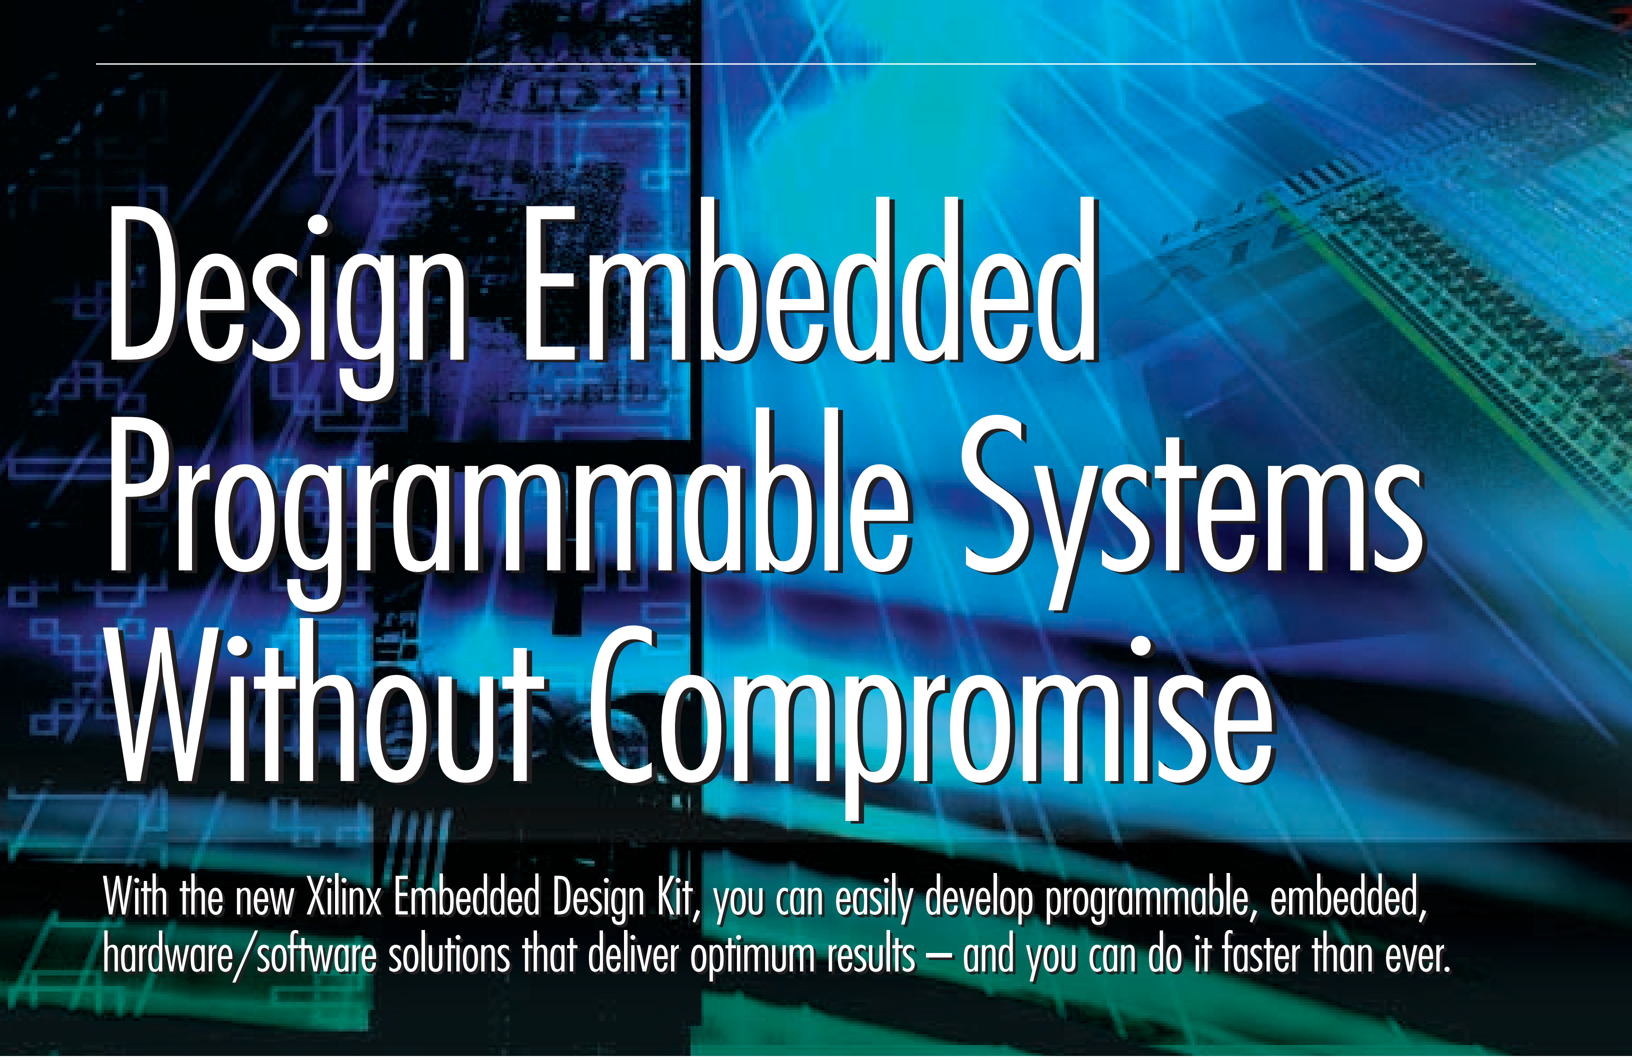

Design Embedded Programmable Systems Without Compromise

With the new Xilinx Embedded Design Kit, you can easily develop programmable, embedded, hardware/software solutions that deliver optimum results – and you can do it faster than ever.

by Ravi Pragasam
Marketing Manager, Embedded Processor Solutions
Xilinx Inc.
ravi.pragasam@xilinx.com

The trouble with embedded design up to now has been that neither ASICs nor ASSPs – the traditional choices for embedded solutions – offer an ideal fit. In an ASIC flow, hardware (HW) limitation problems discovered late in the design cycle must be dealt with in the software (SW). This drawback is becoming increasingly critical now as physical geometries shrink, and as design rules and chip fabrication processes evolve.

Likewise, ASSPs are also less than perfect because you must design your solution around standard elements, which often impose limitations. In addition, the nonstandard elements of the ASSP usually don't add enough functionality to deliver a solution that sets it apart from the competition. Typically, the resulting design has too much functionality here, and not enough there.

Long design cycles involving extensive simulation, on the one hand, and narrow market windows and changing industry standards, on the other, often necessitate compromises in performance. The result is often failure to meet original specification goals. What has been lacking is a fast, effective way to integrate HW and SW flows to produce a solution that brings out the full advantage of embedded design.

The new Xilinx Embedded Development Kit (EDK) bridges the gap between HW and SW flows by providing a single HW/SW design environment. Fusing the two early in the design flow enables you to deliver an embedded programmable system that meets your design specifications on time – without compromising performance.

A New Era of System Design

The EDK allows you to define the hardware and software platforms of your system using powerful tools based on the Platform Specification Format (PSF). PSF is an open format that provides an abstraction layer between the HW and SW sections so they can be tightly integrated during the definition and development process. This coupling ensures that the hardware platform you create will be one a software platform can match.

The EDK also enables you to rapidly match the programming environment and software capability to available hardware resources.

For the software designer, the greatest advantage afforded by PSF is access to an embedded tool that allows HW/SW codesign, and which also interfaces with industry standard software tools, such as RTOS support from Wind River Systems, Linux support from MontaVista Software, and popular open source embedded tools like GNU.

PSF Benefits

- The PSF overcomes many of the most common problems of the traditional "over-the-wall" approach to HW/SW interface definition and integration. With the PSF, you can give a hardware platform specification to a firmware or software engineer without having to wait until the hardware is available. The firmware and software developer can then configure the SW platform with full knowledge of the HW platform – with which devices are attached, memory maps, register definitions, and so on.

- Offering a single-environment approach for HW/SW platform configuration, PSF enables a rapid re-architecture of the embedded programmable systems design. This enables you to add to, or subtract from, HW resources with matching SW in the modified architecture easily.
- The unified environment, common bus structure, and common core libraries enable you to easily use a Xilinx MicroBlaze™ core, IBM PowerPC™ processor in a Virtex-II Pro™ FPGA, or both, in a single design.
- Multi-core, homogeneous (multiple MicroBlaze cores or multiple PowerPC processors), or heterogeneous (mixed PowerPC and MicroBlaze cores) in a single design (multiple bus masters), are also supported.

The PSF is a one-of-a-kind specification format for both HW and SW. Because it is an open format, it is available for adoption by customers and third parties, enabling them to integrate custom peripherals, third-party IP cores, and tools. PSF provides a common tool chain for the HW/SW platform specification, generation, development, and debug, while supporting multimaster and mixed architecture designs.

The Xilinx EDK and ISE 5.1i logic design tools provide you with the technology to help you achieve your performance requirements – faster and better than ever.

EDK Contents

The EDK includes the Xilinx Platform Studio (XPS), an all-encompassing design environment that permits you to define, configure, and generate a custom hardware and matching software/programming environment for either a stand-alone or RTOS-enabled programmable system. Hardware, software, and firmware developers who need to work in both domains can use the XPS integrated design environment.

- Embedded system tools
    - Xilinx Platform Studio
    - Tools for specifying the hardware and software platforms based on PSF
    - Xilinx microprocessor debug (XMD)
    - GNU tools for MicroBlaze and hard embedded PowerPC cores in Virtex-II Pro FPGAs (compiler and debugger)
    - Support for simulation tools
    - System generator for processors (beta version)
    - Board Support Package (BSP) generator
- Interface and infrastructure IP cores
    - Arbiters
    - Memory controllers for external memory interfaces
    - More than 40 standard IP cores (such as UART, GPIO, and timer/counter)
    - Evaluation version of high-value cores (such as 10/100 EMAC, single channel HDLC controller, and serial ATA L2)
- MicroBlaze 32-bit soft processor core
- Reference designs and examples
- Evaluation versions of Wind River and ISE 5.1i software tools.

Conclusion

With key products, such as MicroBlaze 32-bit soft processor cores and Virtex-II Pro FPGAs, Xilinx offers a complete solution that resolves many of the design challenges presented by more traditional tools and technology. With the Virtex-II Pro device, Xilinx has ushered in a new era of system design in which SW/HW flows blend to take advantage of programmability and high-performance features – such as the multi-gigabit serial transceivers available in the silicon – in a way that enables solution providers to get ahead in their markets without making performance compromises.

Working within a common HW/SW tool environment, you can rapidly construct a custom processor system consisting of processor, peripheral cores, and interconnect bus in a Xilinx FPGA. You can also integrate your own custom IP cores into the processor system.

And now, with the EDK, Xilinx has enabled the process of HW/SW co-design that has long been a dream of embedded system designers everywhere. The EDK (Part No. DO-EDK) is available now. For more information and updates to the EDK, visit *www.xilinx.com/edk/*.

For more information about Xilinx embedded processor solutions, visit Processor Central at *www.xilinx.com/processor/*.

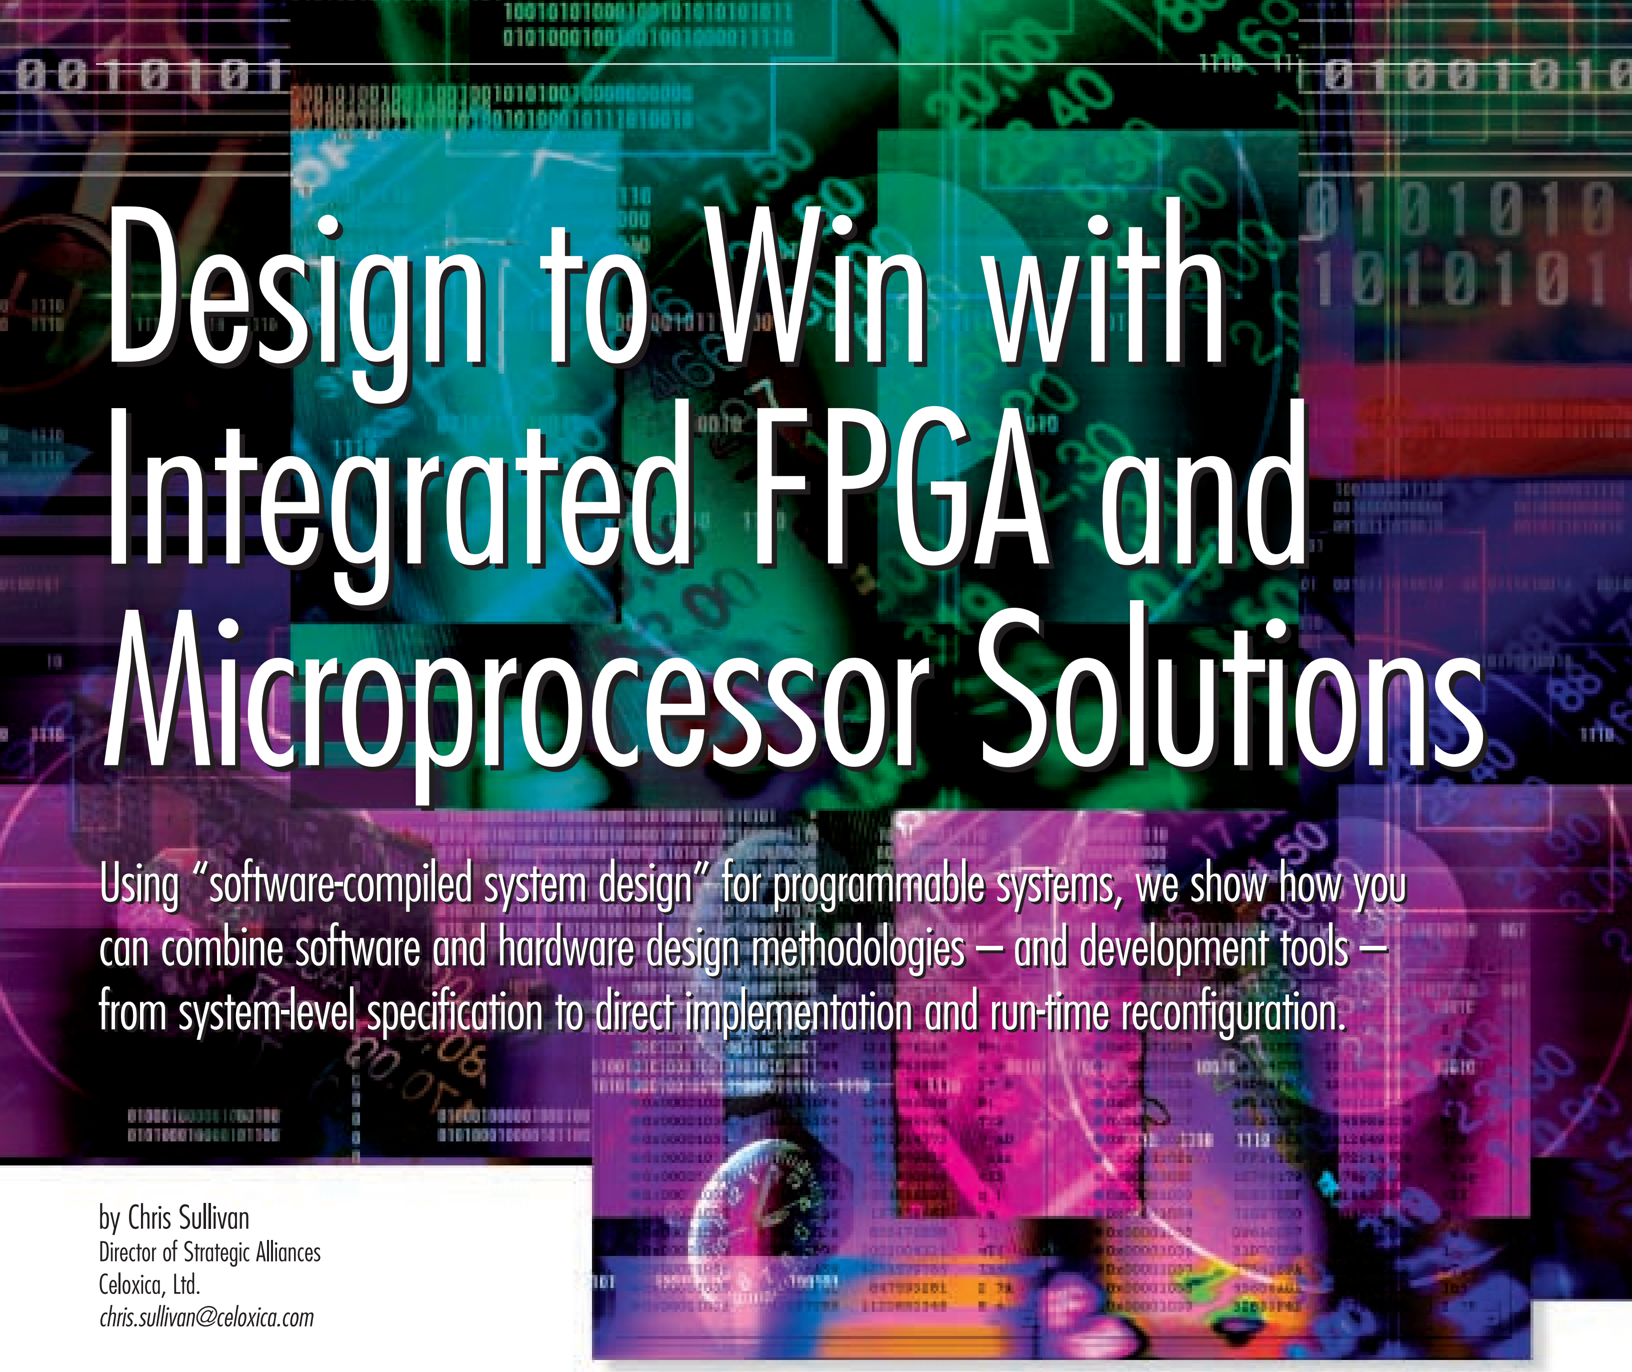

Design to Win with Integrated FPGA and Microprocessor Solutions

Using "software-compiled system design" for programmable systems, we show how you can combine software and hardware design methodologies – and development tools – from system-level specification to direct implementation and run-time reconfiguration.

by Chris Sullivan
Director of Strategic Alliances
Celoxica, Ltd.
chris.sullivan@celoxica.com

As FPGAs have developed from logic prototyping devices into fundamental system elements, there has been enthusiasm for the concept of using high-performance processors closely coupled to or immersed inside the FPGA fabric for applications that require unrivalled levels of performance and flexibility.

In this architecture, the microprocessor typically runs system applications while the FPGA manages computationally intensive tasks. Offloading processor-intensive tasks to hardware reduces the load on the processor and delivers greater bandwidth. It can also remove bottlenecks by migrating algorithms to hardware. In short, FPGAs have evolved into fully programmable systems and fast co-processors, rather than just flexible relations of the ASIC.

Existing design examples that combine Xilinx FPGA solutions with development tools from Celoxica and Wind River Systems already provide unique and tangible proof that this concept and design flow works. They form a core element of programmable system co-design, delivering a quick, efficient, and verifiable route to device-optimized implementation.

"Software-compiled system design" provides the capability to drive partitioning, co-verification, and direct implementation from the system specification. Moreover, it allows engineers to jump-start their system and software application development before actual hardware is available, thereby enabling concurrent design, saving valuable development effort and delivering the best time to market. Starting at the system level, verification becomes a whole-design life cycle activity, and by enabling system-level partitioning, you can realize a better quality of design (QoD) – right the first time, more of the time.

A Design Example

An early design example – developed by Celoxica, Wind River, and Xilinx – focused on the design methodology, tools, and run-time environments that can be applied to programmable systems. Specifically, we developed a triple-DES encryption and decryption engine to compare a programmable system solution with an alternative software implementation. A compressed video stream formed the basis of test data.

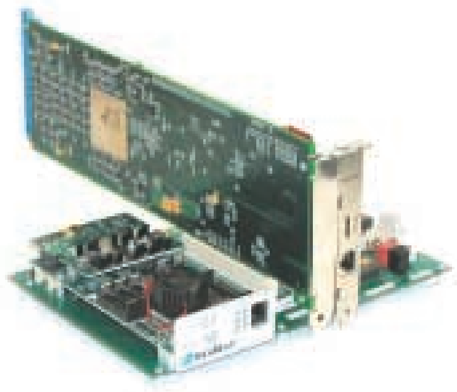

Figure 1 - PPMC750 and RC1000 connected and functioning as a prototyping platform for programmable systems

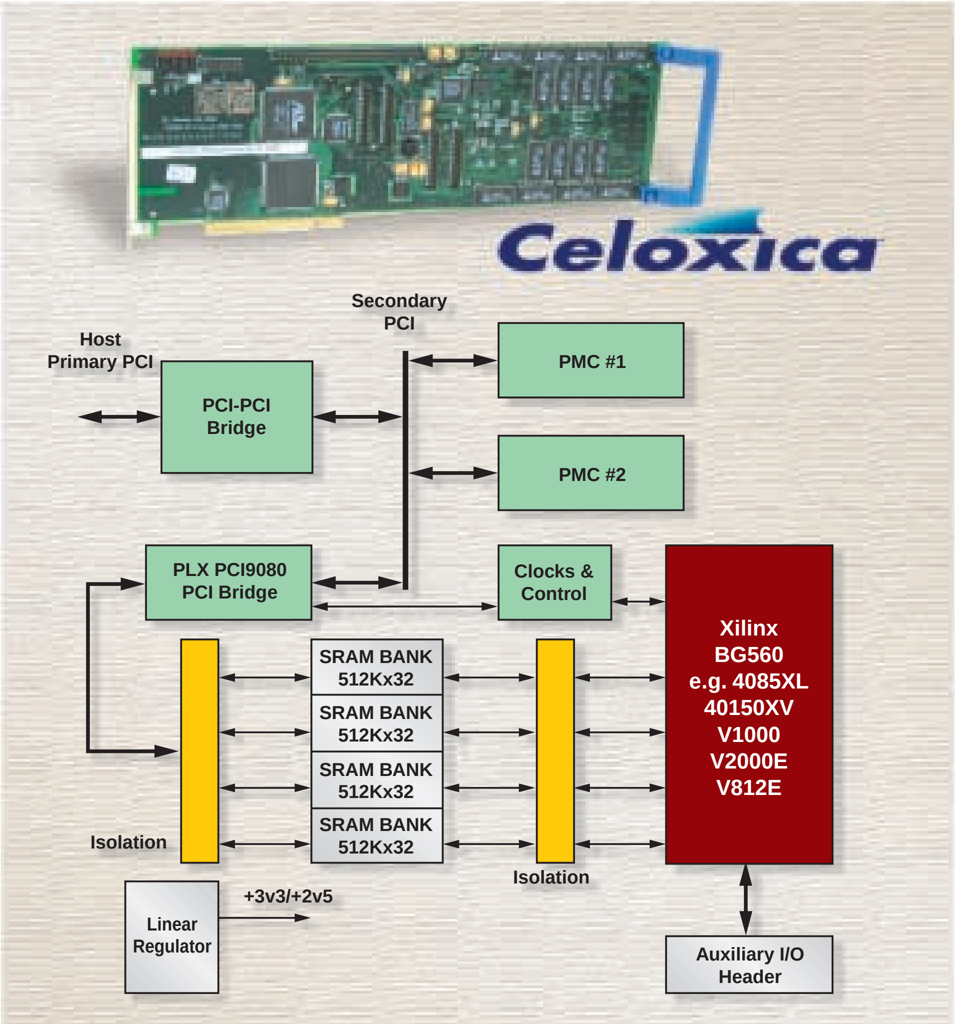

Figure 2 - Above, the RC1000 is a Celoxica PCI slot-based card, which has a secondary bus on-board for connecting other PMC based FPGA-cards as well as the existing FPGA on-board. Below the board is a schematic of the RC1000.

Hardware

The selected hardware architecture was initially based around a discrete IBM PowerPC™ processor and a Xilinx Virtex™ FPGA – effectively a first-generation Virtex-II Pro™ prototyping platform (Figure 1). We used a PPMC750 single-board computer from Wind River and Celoxica's RC1000 – a Virtex-based PCI card with a Xilinx FPGA and 8 Mb of local memory (Figure 2).

Subsequently, we deployed a newer reference platform using the PowerPC 405GP processor (Figure 3). In addition to PCI and PMC (PCI mezzanine card) connectors, this platform also featured a custom connector that allowed an FPGA daughtercard to be plugged directly onto the processor peripheral bus, thus providing even closer coupling, lower latency, and higher throughput.

Various FPGA daughtercards can be used with this reference platform, such as the ADM-XRC from Alpha Data Systems, Xilinx Durango, or the Proteus card from Wind River.

Wind River's Proteus card is equipped with a Xilinx Virtex device and memory includes 4 MB on-board SSRAM. The FPGA PMC can interface with any standard PMC slot (with an image containing a PCI soft core) or the microprocessor local bus on Wind River's SBC405GP single board computer. There is a substantial performance boost from direct processor bus connection, compared with PCI.

The design platform is completed by a simple DAC interface, enabling the FPGA card to drive a video monitor or a flat-panel LCD for standalone demonstrations.

Development Tools

The 405GP processor runs Wind River's VxWorks™ real-time operating system (RTOS), together with hardware bring-up tools that allow close control of the boot cycle for the board during the time period before control is passed to the RTOS.

The PAVE Framework API from Xilinx was used to program the FPGA with configuration files.

Determination of the system partition and application content for the FPGA were developed using Celoxica's Nexus co-design environment and DK Design Suite.

Nexus and DK

Nexus is a powerful co-design environment for programmable systems. It supports system partitioning, co-verification, and co-simulation. Nexus allows you to fully

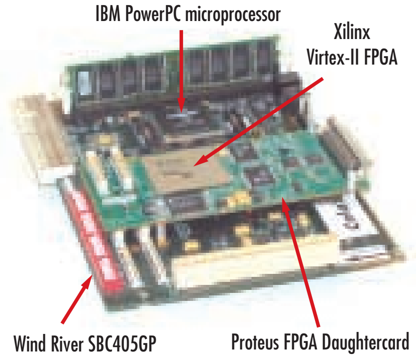

Figure 3 - Virtex-II prototyping platform

explore the design space to identify optimal system partitioning. System functionality can be simulated between hardware and software using multiple languages such as C, C++, Handel-C, SystemC, and HDLs. These models can be used throughout the design for co-verification. Nexus communicates directly during simulation with popular, third-party, hardware RTL simulators and software ISS environments.

Using DK, the resulting code may be debugged using a familiar integrated development environment (IDE), and applications are compiled direct to the FPGA fabric using device-optimized synthesis. VHDL and Verilog output is also supported for traditional RTL synthesis.

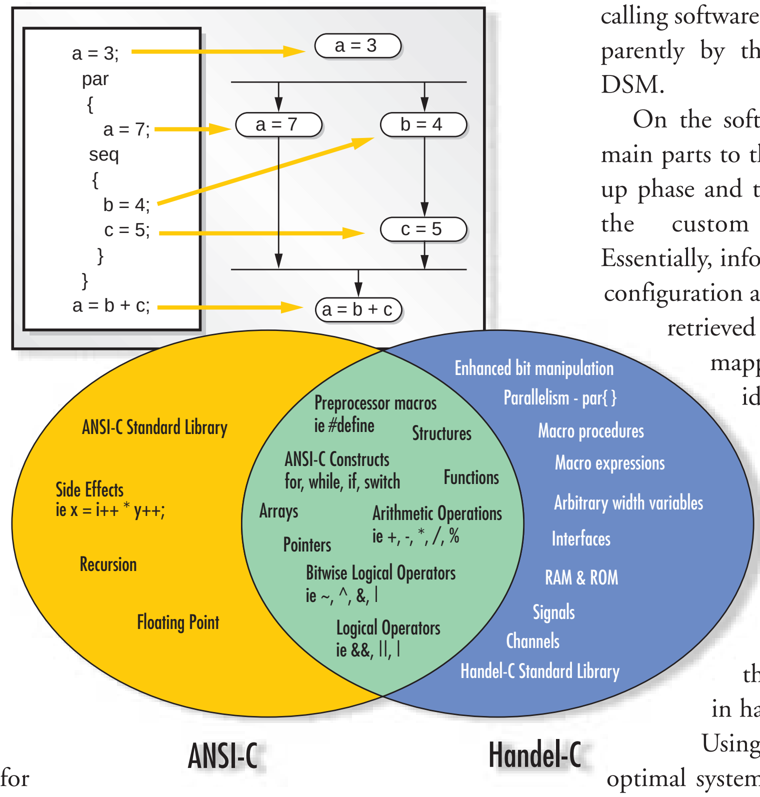

```
a = 3;
par
{
   a = 7;
  seq
  {
   b = 4;
   c = 5;
  }
}
a = b + c;
```

Figure 4 - ANSI-C/Handel-C comparison

Handel-C

We selected the Handel-C language for hardware implementation, as it provides a common level of abstraction and a common language base for both the hardware and software. The language has simple extensions to ANSI-C (Figure 4) that can be leveraged to quickly create applications that fully exploit the capabilities of a programmable system, without compromising performance or area.

As a fully synthesizable language, everything that can be described in Handel-C has translation to hardware (Figure 5). The code illustrates concepts and extensions, such as par, chan, synchronization, functions, pointers, structures, interfacing, and externing pure C functions for simulation.

With a simple timing model, each assignment in a program takes one clock cycle to execute, giving you full control over what is happening in the design at any point in time. Results are predictable and controllable, and the facility for complex sequential control flows means there are no state machines to design.

Run-Time Environment

Typically, the FPGA is connected to the microprocessor in a memory-mapped or programmed I/O fashion, but this creates the challenge of needing to develop and redevelop individual communications protocols and data-marshalling routines for each application. This problem is overcome by using DSM (Data Streaming Manager), a portable co-design API developed for hardware/software integration in programmable systems.

Data Streaming Manager

DSM is a portable hardware/software codesign API that offers a simple and transparent interface for transferring multiple independent streams of data between hardware and software. DSM supports system partitioning and final implementation; it is both bus/interconnect and OS-independent; and for the developer, it simplifies the integration between the hardware and software (Figure 6).

As an example, the hardware function reads parameters from an input port and then writes the results to an output port. All the complexities of receiving commands over the PCI or bus, routing parameters to the appropriate hardware function, and then routing the responses back to the calling software thread are handled transparently by the hardware side of the DSM.

On the software side, there are two main parts to the DSM: the control/setup phase and then the specific usage of the custom hardware function. Essentially, information about the FPGA configuration and available functions are retrieved by reading a memory-mapped register. User-defined identifiers (called function addresses) are assigned to each available hardware function, and these function addresses are later used to communicate between the application software and the functions implemented in hardware.

Using this methodology, the optimal system partition can be identified by porting blocks of software to Handel-C, for hardware prototyping, testing, and verification. DSM's portability means that multiple partitions can be rapidly evaluated, tested, and verified with the software used as a testbench throughout.

DSM also provides a functionally accurate simulation environment that allows ANSI-C programs and Handel-C applications to interact using the DSM (Figure 7). The ANSI-C program is run as a native executable on the PC. The Handel-C application is run using the simulation and debugging capability of Celoxica's co-design environment. A utility is provided through which the data passing between the applications may be monitored to assist with debugging (Figure 8). All of the API functions are provided, allowing complete system development to begin – without the development platform being available. Once working, the application can be easily transferred to the target platform for final testing.

Triple DES Encryption

Our design example was based around streaming of compressed and encrypted video data. The Autodesk FLI file format was used to compress the video, and an FLI player, developed by Celoxica, was implemented in FPGA hardware connected to the processor via the PCI bus. To benchmark the design, we loaded a cartoon animation into the memory on the processor board. A triple DES algorithm described in C ran on the PowerPC microprocessor.

The same C source code was ported to Handel-C, optimized in terms of controlling parallelism and timing, and compiled to a gate-level design that was device-optimized for the target FPGA.

A 64-bit key was used for the encryption, which subsequently allowed correct decryption of the video stream. Implementing three DES algorithms in sequence (triple DES encryption) provided further increases in this standard's security. Three 64-bit keys were used for an encrypt/decrypt/encrypt cycle in a triple DES pass, and the same keys allowed decrypt/encrypt/decrypt for decrypting the data.

This was a robust test of performance. The algorithm was inherently sequential in software, but it could be heavily pipelined for a hardware implementation.

To measure the performance improvement, we played a cartoon with each compressed frame being encrypted, decrypted, and displayed on a VGA monitor. Both hardware and software implementations were displayed together. They were triggered to start simultaneously, with the hardware version programmed to cycle continuously until the software implementation finished. Processed data from the microprocessor was fed to the FPGA, which as well as performing DES encryp-

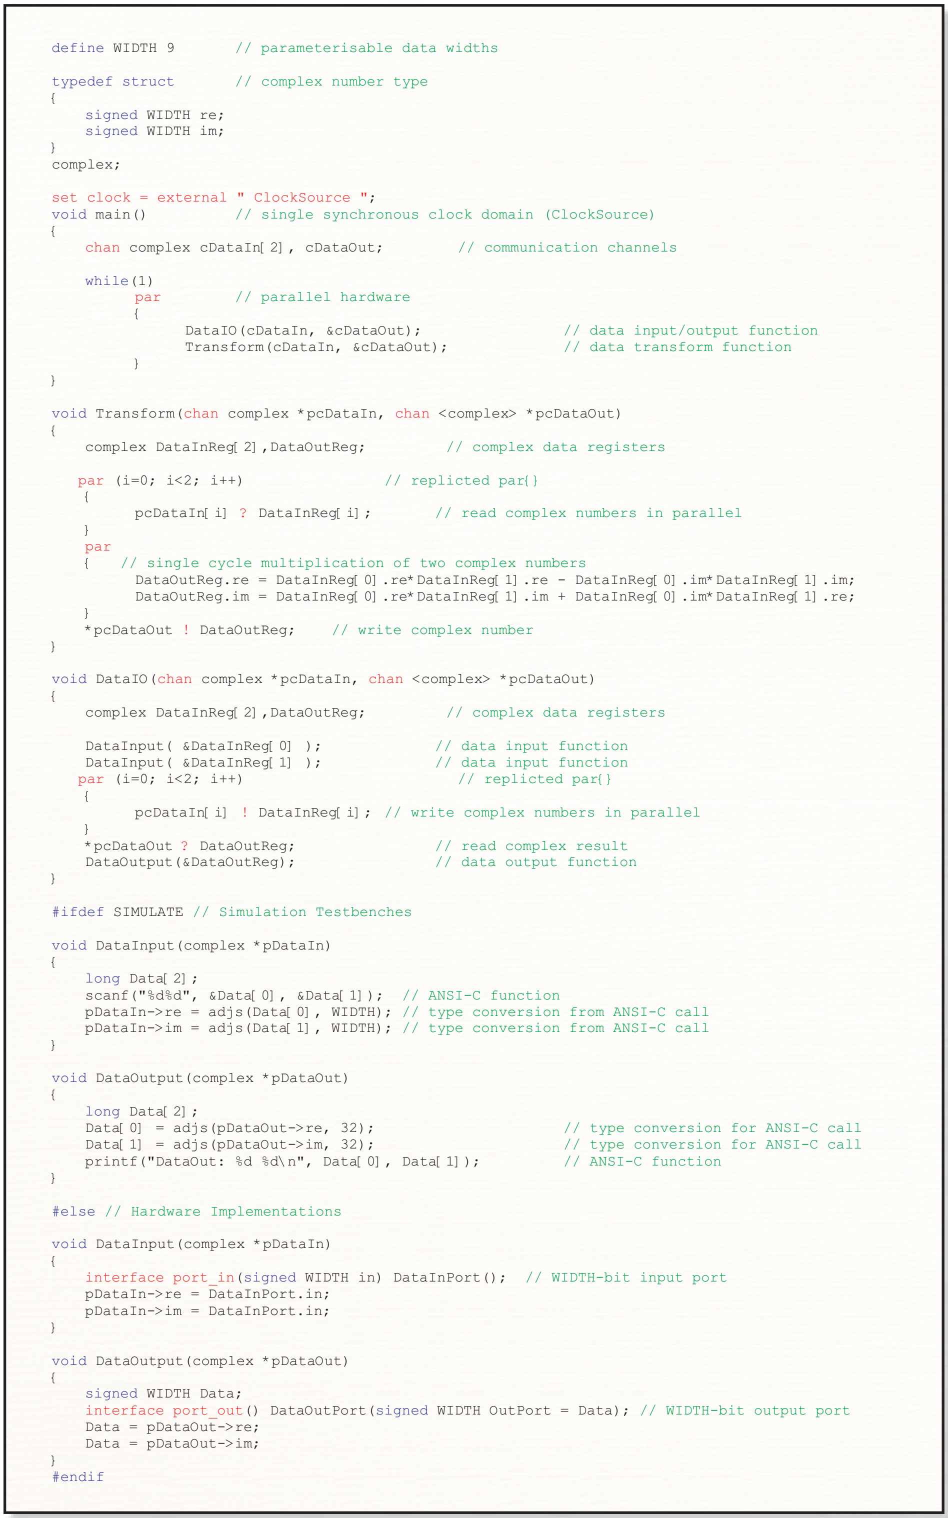

```
define WIDTH 9          // parameterisable data widths

typedef struct          // complex number type
{
    signed WIDTH re;
    signed WIDTH im;
}
complex;

set clock = external " ClockSource ";
void main()             // single synchronous clock domain (ClockSource)
{
    chan complex cDataIn[ 2] , cDataOut;          // communication channels

    while(1)
          par           // parallel hardware
          {
                DataIO(cDataIn, &cDataOut);                 // data input/output function
                Transform(cDataIn, &cDataOut);              // data transform function
          }
}

void Transform(chan complex *pcDataIn, chan <complex> *pcDataOut)
{
    complex DataInReg[ 2] ,DataOutReg;          // complex data registers

   par (i=0; i<2; i++)                 // replicated par{}
    {
          pcDataIn[ i] ? DataInReg[ i] ;        // read complex numbers in parallel
    }
    par
    {   // single cycle multiplication of two complex numbers
          DataOutReg.re = DataInReg[ 0] .re* DataInReg[ 1] .re - DataInReg[ 0] .im* DataInReg[ 1] .im;
          DataOutReg.im = DataInReg[ 0] .re* DataInReg[ 1] .im + DataInReg[ 0] .im* DataInReg[ 1] .re;
    }
    *pcDataOut ! DataOutReg;    // write complex number
}

void DataIO(chan complex *pcDataIn, chan <complex> *pcDataOut)
{
    complex DataInReg[ 2] ,DataOutReg;          // complex data registers

    DataInput( &DataInReg[ 0]  );             // data input function
    DataInput( &DataInReg[ 1]  );             // data input function
   par (i=0; i<2; i++)                           // replicated par{}
    {
          pcDataIn[ i]  ! DataInReg[ i] ; // write complex numbers in parallel
    }
    *pcDataOut ? DataOutReg;                  // read complex result
    DataOutput(&DataOutReg);                  // data output function
}

#ifdef SIMULATE // Simulation Testbenches

void DataInput(complex *pDataIn)
{
    long Data[ 2] ;
    scanf("%d%d", &Data[ 0] , &Data[ 1] );  // ANSI-C function
    pDataIn->re = adjs(Data[ 0] , WIDTH); // type conversion from ANSI-C call
    pDataIn->im = adjs(Data[ 1] , WIDTH); // type conversion from ANSI-C call
}

void DataOutput(complex *pDataOut)
{
    long Data[ 2] ;
    Data[ 0]  = adjs(pDataOut->re, 32);                     // type conversion for ANSI-C call
    Data[ 1]  = adjs(pDataOut->im, 32);                     // type conversion for ANSI-C call
    printf("DataOut: %d %d\n", Data[ 0] , Data[ 1] );          // ANSI-C function
}

#else // Hardware Implementations

void DataInput(complex *pDataIn)
{
    interface port_in(signed WIDTH in) DataInPort();  // WIDTH-bit input port
    pDataIn->re = DataInPort.in;
    pDataIn->im = DataInPort.in;
}

void DataOutput(complex *pDataOut)
{
    signed WIDTH Data;
    interface port_out() DataOutPort(signed WIDTH OutPort = Data); // WIDTH-bit output port
    Data = pDataOut->re;
    Data = pDataOut->im;
}
#endif
```

Figure 5 - Handel-C code for single cycle multiplication of two complex numbers

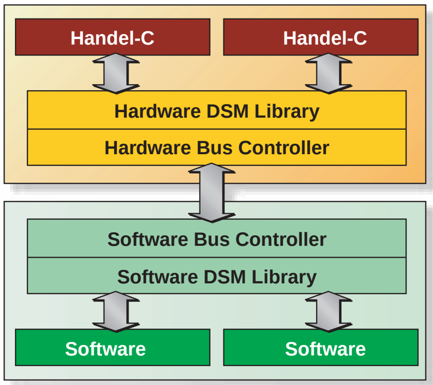

Figure 6 - DSM Data Streaming Manager

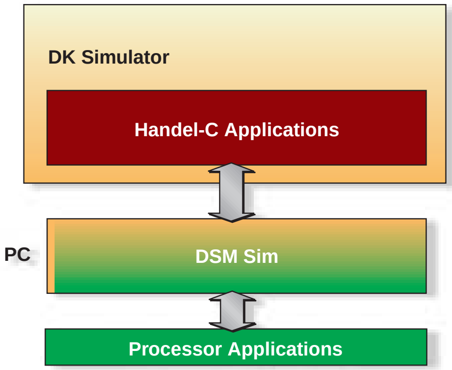

Figure 7 - DSM simulation environment

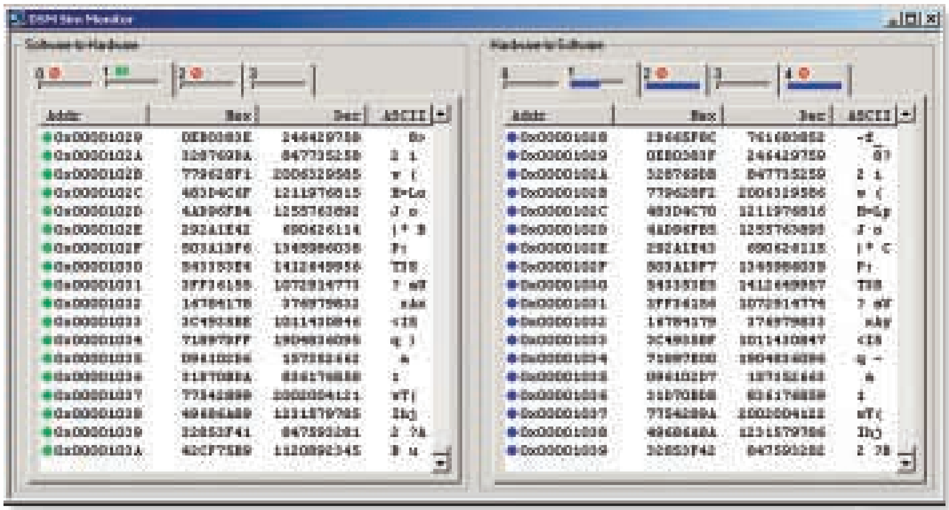

```
Software to Hardware
Addr               Hex          Dec   ASCII
0x00001029    0EB0083E    246429758   8>
0x0000102A    3287698A    847735258   2 i
0x0000102B    779620F1   2006329585   w (
0x0000102C    483D4C6F   1211976815   H=Lo
0x0000102D    4A396F84   1258762692   J o
0x0000102E    292A1E42    690626114   )* B
0x0000102F    503A13F6   1346966038   P:
0x00001030    5431532E   1412649956   T1S
0x00001031    3FF74155   1072914773   ? wU
0x00001032    14784178    376979832   xAx
0x00001033    3C49048E   1011430046   <I8
0x00001034    718979FF   1904834095   q )
0x00001035    09410236    157352462    A
0x00001036    3187088A    836176858   1
0x00001037    77342890   2002004121   wT(
0x00001038    496864A9   1231879785   Ihj
0x00001039    32853F41    847593281   2 ?A
0x0000103A    42CF75E9   1120892345   B u
```

```
Hardware to Software
Addr               Hex          Dec   ASCII
0x00001028    2D465F0C    761680652   -F_
0x00001029    0EB0383F    246429759   8?
0x0000102A    3287698B    847735259   2 i
0x0000102B    779620F2   2006329586   w (
0x0000102C    483D4C70   1211976816   H=Lp
0x0000102D    4A396F85   1258762693   J o
0x0000102E    292A1E43    690626115   )* C
0x0000102F    503A13F7   1346966039   P:
0x00001030    5431532F   1412649957   T1S
0x00001031    3FF74156   1072914774   ? wV
0x00001032    14784179    376979833   xAy
0x00001033    3C49048F   1011430047   <I8
0x00001034    71897E00   1904834096   q ~
0x00001035    09410237    157352463    A
0x00001036    3187088B    836176859   1
0x00001037    7734289A   2002004122   wT(
0x00001038    496864AA   1231879786   Ihj
0x00001039    32853F42    847593282   2 ?B
```

Figure 8 - DSM Sim Monitor for assisted debugging

	Encryption		Decryption	
	Software	FPGA	Software	FPGA
Elapsed time for 1 Mb of data	5558.8 ms	424.8 ms	5562.9 ms	424.7 ms
Cryptography rate	1.51 Mbps	19.7 Mbps	1.51 Mbps	19.8 Mbps

Table 1 - Performance statistics for triple DES encryption/decryption

tion/decryption, was programmed to generate VGA signals. Output from both the hardware and software implementations was merged to form a composite image on the monitor.

Performance Comparison

A test harness enabled triple DES performance to be benchmarked by streaming data into either the software or hardware encryption algorithm.

Theoretical performance for the FPGA was calculated as follows: The triple DES implementation produced a 64-bit word every 19 clock cycles, giving a data throughput of 85.6 Mbps for a device running at 25.4 MHz.

Actual performance was profiled using WindView, a diagnostic tool from Wind River that enables visualization and analysis of performance and timing issues in embedded systems. It allowed triggers to be set at different points in the code, and then provided accurate timing information for each trigger event. Performance statistics are detailed in Table 1.

Platform	Clock Speed	Performance
SBC405	300 MHz	1.51 Mbps
Xilinx FPGA	20 MHz	33.8 Mbps

Table 2 - Performance comparison of hardware and software encryption

This test scenario showed an FPGA throughput about 13 times faster using hardware than running software on a PowerPC microprocessor. Nevertheless, it was still about a quarter of the theoretical maximum rate. This indicated that the full benefits of placing core routines in hardware might be compromised by other system bottlenecks. Further analysis showed that there were overheads associated with offloading functionality into hardware. These overheads were associated with RAM access latency and/or bus speeds.

We also calculated the performance of hardware and software encryption in the cartoon demonstration. Results demonstrated that the hardware performed 22 times faster on a 15 times slower clock, as shown in Table 2.

Following more detailed partitioning analysis, performance closer to the theoretical limits might be realized by removing code and functionality that are not directly associated with the triple DES algorithm (for example, the FLI decoder, frame buffer, and VGA driver). Better performance would also be achievable by connecting the FPGA directly to the processor bus in a memory-mapped fashion rather than across the PCI bus.

Conclusion

The performance analysis results demonstrated significant improvements in overall system performance and quality of design. The results were achieved using a software-compiled system design methodology – specifically developed for programmable systems – that consistently delivered the fastest time to market (some 50% to 75% advantage in design time) without compromising performance or area.

For example, using the selected development tools and run-time environment, the FLI player took 10 person-days to implement, as did the triple-DES functionality. On the other hand, integrating these two blocks to produce the cartoon demonstration took just half a day. Moreover, you can very quickly explore the design space, experiment, and analyze different hardware/software trade-offs, and rapidly implement and prototype the system.

Coupling Celoxica's co-design technology with high-performance profiling tools in the development tool chain enabled further performance boosts and time-to-market efficiencies. Overall improvements in the quality of design were realized by more informed and accurate partitioning decisions, better upfront system verification, and by maximizing the speed gains of hardware implementation while minimizing the negative impact of transferring data between the FPGA and microprocessor.

The bottom line is that these system-level design qualities offer real and competitive advantages for designers of programmable systems who want to move to volume production.

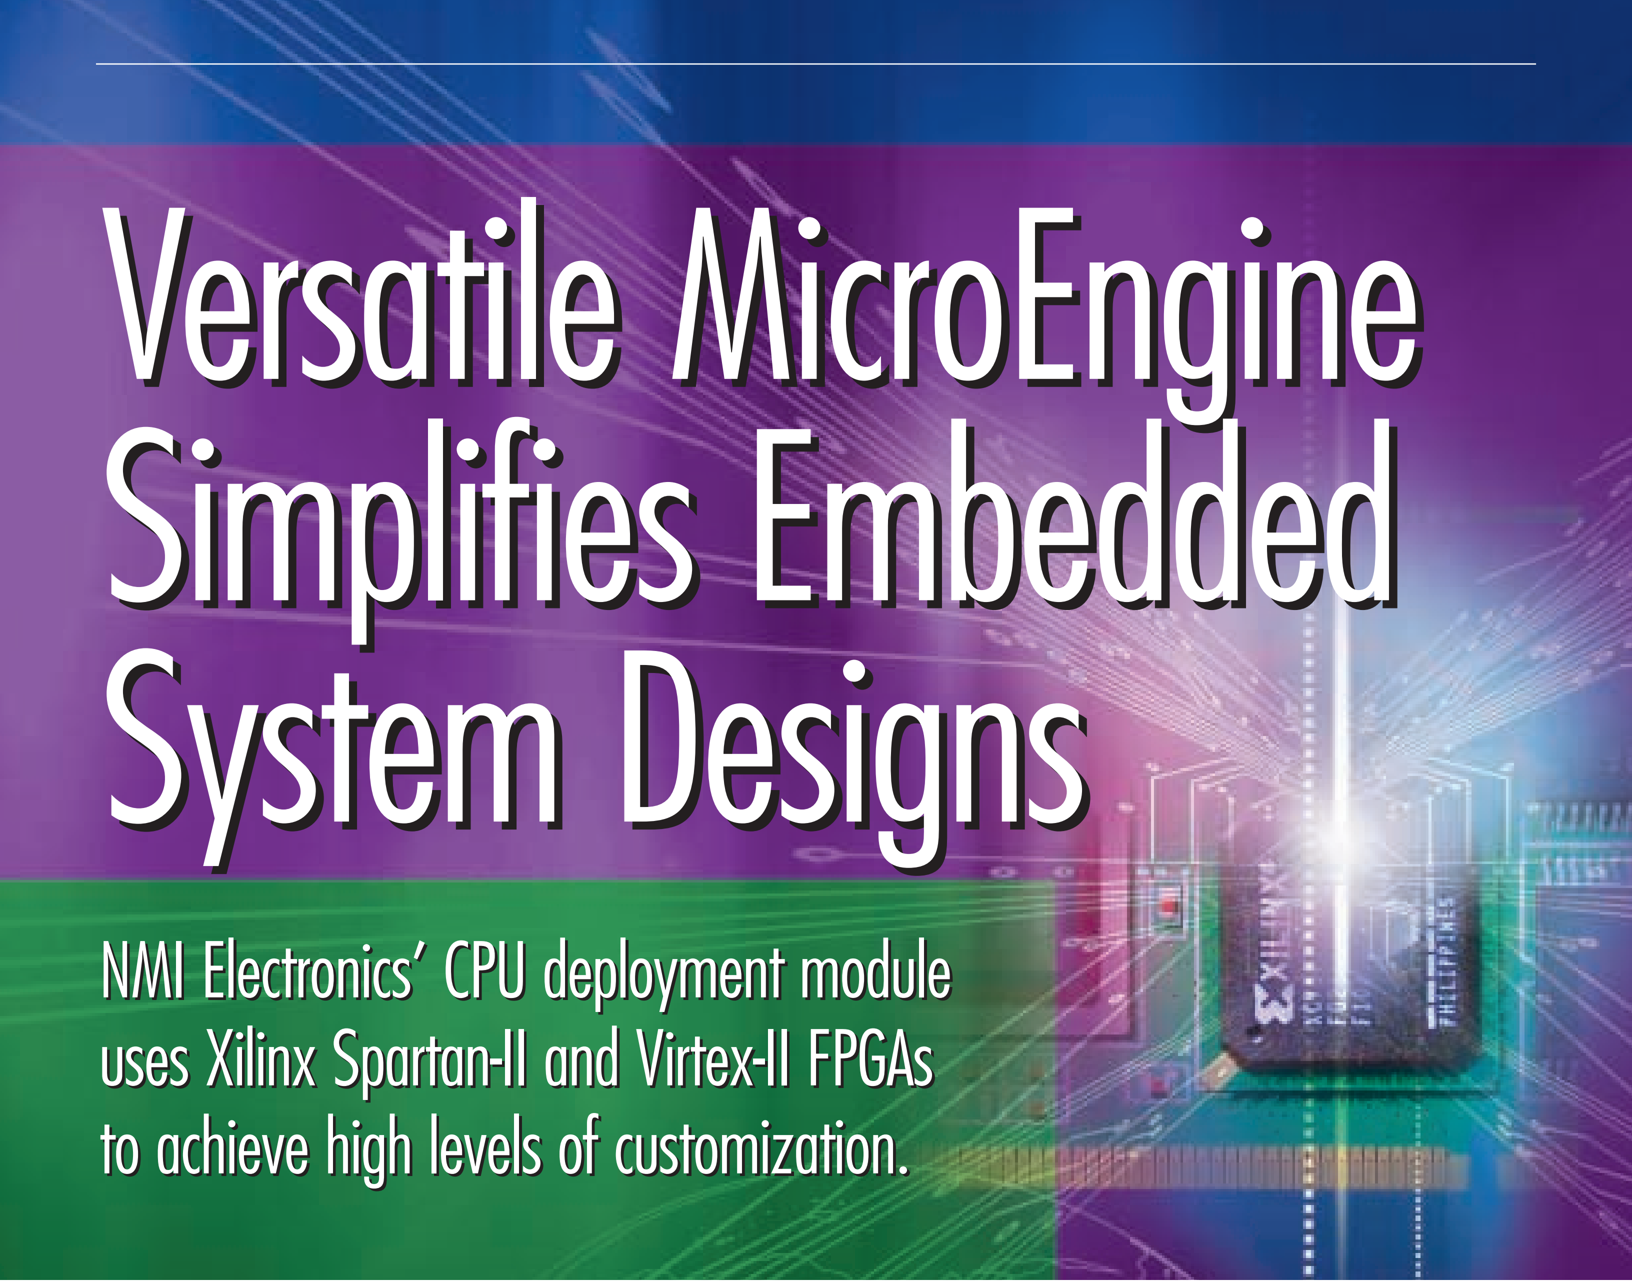

Versatile MicroEngine Simplifies Embedded System Designs

NMI Electronics' CPU deployment module uses Xilinx Spartan-II and Virtex-II FPGAs to achieve high levels of customization.

by Kevin Heawood
Vice President, Strategic Marketing
Intrinsyc Software, Inc.
kheawood@intrinsyc.com

Typically, embedded systems are designed around a specific CPU architecture that comprises a 32-bit high-performance processor unit, volatile and non-volatile memory, and a set of peripherals and interfaces specific to that system or to a particular application. With on-chip clock speeds reaching 400 MHz to 700 MHz and board-level clock speeds as fast as 133 MHz, system design has become increasingly time-consuming, complex, and risky.

Using a single-board computer (SBC) or companion chips and interface logic to reduce risk can be difficult, however, as it is not always possible to find an SBC with the correct peripheral mix and the interfaces that go with a particular CPU. Even if you do find the right SBC, the peripherals may be too expensive for your application, or otherwise inappropriate.

NMI Electronics Ltd. has developed an SBC, or more specifically, a deployment module, that greatly simplifies system design. The module contains all the main components of a 32-bit CPU system, including volatile and non-volatile memory, but it uses either a Xilinx Spartan™-II or Virtex™-II FPGA to provide a completely programmable peripheral set and an interface that can be specifically tailored to any application (Figure 1).

Programmable Interface and Companion Chip

Most SBCs have a predefined interface and set of functions. The interface is usually based on an industry standard, such as PC/104, or on a combination of a standard interface and a custom interface defined by the board manufacturer.

Peripheral functionality is normally provided by standard chipsets and is determined according to whatever the device or CPU manufacturer considers the most commonly requested functions. Beyond a few basics (for example, baud rates on UARTs or display resolutions), the mix of peripheral functions is neither flexible nor programmable, and so may be inconvenient or inappropriate for your application.

Using an FPGA effectively eliminates these restrictions. The FPGA can be placed onto the CPU local bus and closely coupled with the memory subsystem, thereby creating a highly programmable companion and interface chip.

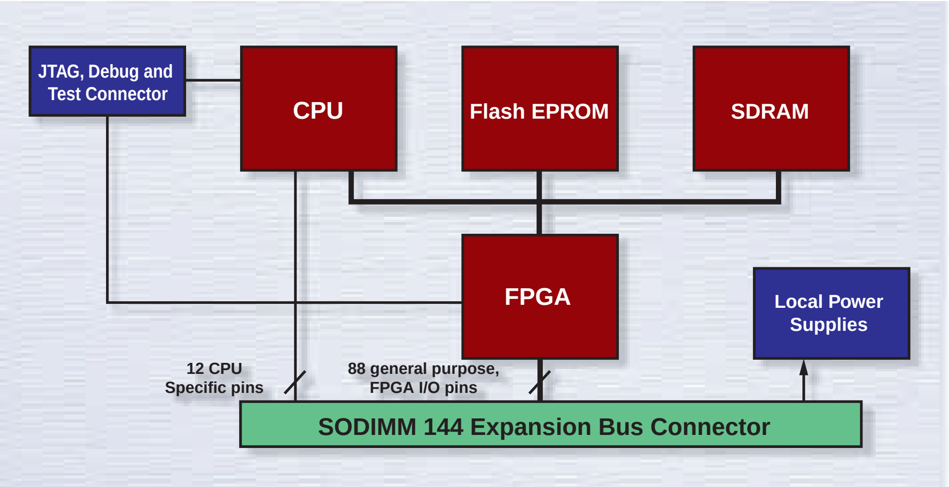
JTAG, Debug and Test Connector
CPU
Flash EPROM
SDRAM
FPGA
Local Power Supplies
12 CPU Specific pins
88 general purpose, FPGA I/O pins
SODIMM 144 Expansion Bus Connector

Figure 1 - μEngine general architecture

Interfaces

This FPGA system architecture allows you to implement the interface you require to the rest of the product hardware in a way that provides the optimum solution for your application. NMI supplies a number of standard interfaces, including the following:

- PCI
- ISA
- PCMCIA.

It is also possible to support the CPU local bus or a custom-designed interface.

Peripherals

Because almost every application has a unique set of requirements, the peripherals required for each specific system are as diverse as the applications themselves. Using the FPGA as the companion chip to fulfill these requirements, it is now relatively simple to mix and match a range of peripherals in a way that meets the exact needs of your application.

Examples of peripherals that can be included in the FPGA are:

- UARTs
- SPI, I2C, and AC97 serial interfaces
- Display controller (LCD or CRT)
- Stepper motor controller
- Camera frame grabber
- PCI slave and host interfaces
- PC/104 interface.

Accelerators

You can also place application-specific accelerators (co-processors) into the FPGA. These accelerators assist the CPU in the performance of specific functions. Examples of such accelerators include:

- 2D display assistance
- Hardware cursor support
- DSPs.

Companion Chip Application Example

Imagine you are working on an automobile navigation system that requires high-performance graphics, an interface to CAN bus, and two serial ports – one for interface to a mobile phone and one for diagnostics. It may not be possible to find a CPU and companion chip with that particular peripheral mix.

By using an FPGA-based companion chip, however, you can implement a PCI host bridge as an interface to a high-performance standard graphics chip, plus a CAN controller, and two UARTs, all within the FPGA. You could even implement a DMA controller to feed the graphics chip and service the CAN controller and UARTs, which frees up the CPU to perform more compute-intensive tasks.

NMI MicroEngine

Using the FPGA as the basis for the system interface and peripheral companion chip also makes it possible to isolate the CPU/memory/FPGA subsystem and mount this onto a small printed circuit board.

In fact, this is what we at NMI have done with our MicroEngine (μEngine). The μEngine is a small form factor deployment module that contains all the key elements of a high-performance, processor-based system (CPU, flash memory,

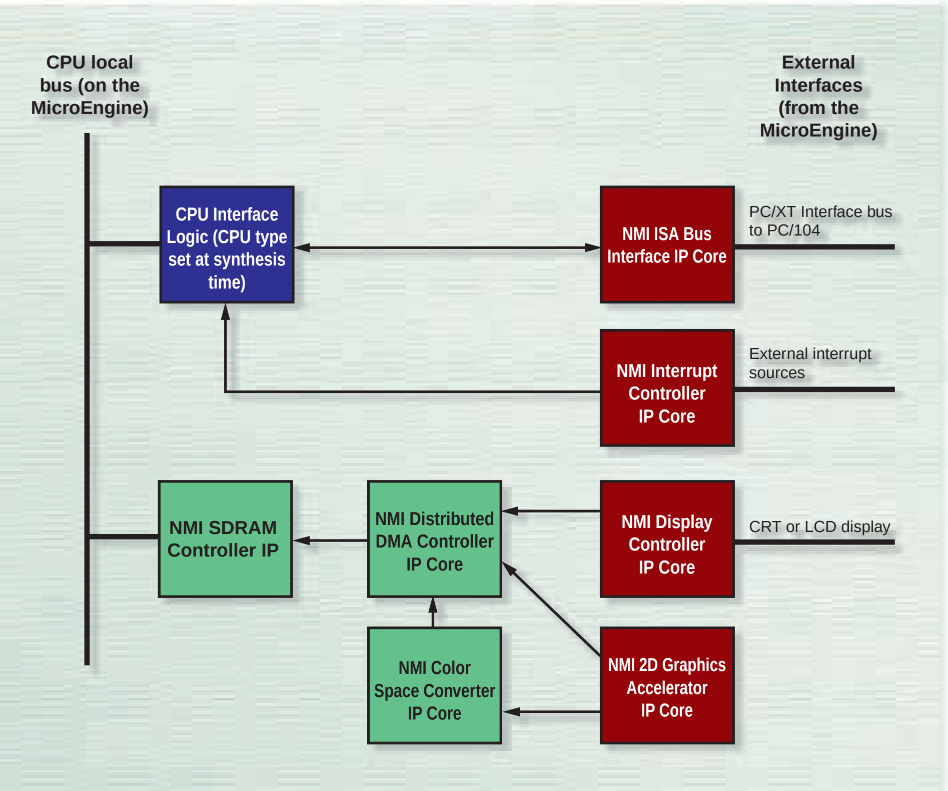
CPU local bus (on the MicroEngine)
External Interfaces (from the MicroEngine)
CPU Interface Logic (CPU type set at synthesis time)
NMI ISA Bus Interface IP Core
PC/XT Interface bus to PC/104
NMI Interrupt Controller IP Core
External interrupt sources
NMI SDRAM Controller IP
NMI Distributed DMA Controller IP Core
NMI Display Controller IP Core
CRT or LCD display
NMI Color Space Converter IP Core
NMI 2D Graphics Accelerator IP Core

Figure 2 - Example μEngine FPGA design

SDRAM), but which uses a Xilinx Spartan-II or Virtex-II FPGA to provide the system interface and peripheral functions. This arrangement provides a totally flexible core module, and it enables you to include precisely the peripherals and interfaces you need for your system. It also means that you can use the same basic board in a wide range of equipment.

In addition, the µEngine addresses the conceptually simple, yet practically more difficult, problem of designing microprocessor systems with high-speed external clocks and buses. The µEngine is, in its own right, a self-contained, pretested, high-performance microprocessor subsystem. All it needs to "run" is power.

In other words, all you need to implement your application is a baseboard containing a power supply and the specific interface logic to suit your application. In many cases, the baseboard can be relatively straightforward, and use lower technology design and less stringent manufacturing rules than the high-speed µEngine design.

The connection between the µEngine and the baseboard is made via an industry-standard, 144-pin, SODIMM connector that carries both power and logic signals. Eighty-eight pins of the interface are connected to the FPGA and are completely user-programmable. Twelve CPU-specific pins carry such dedicated functions as serial ports, ADCs, DACs, or USB, depending on the CPU deployed on the µEngine (Figure 2).

The image for the FPGA is held in the µEngine's flash memory and is completely reprogrammable. You can even place more than one FPGA image on a µEngine, enabling it to support multiple baseboards. It does this by means of a mechanism that identifies the type of baseboard it is plugged into at power-up and automatically loads the correct FPGA for the application.

In addition, the ability to isolate the CPU from the baseboard allows you to plug different CPU-based µEngines into the same baseboard. This is one of the µEngine architecture's greatest advantages. It enables upgrades of CPUs as well as the use of CPUs of varying performance levels, such as an application that requires modest graphics performance in an entry-level product and high performance in another, high-end product. For instance, the entry-level product might be based on a Hitachi SH3 (without FPU) 100 MHz µEngine (Figure 3) – and the high-end product might be based on a Hitachi SH4 (with FPU) 200 MHz µEngine (Figure 4). Both units would use the same baseboard.

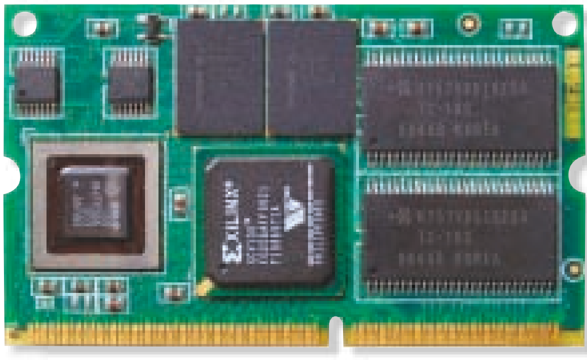

Figure 3 - Hitachi SH3 µEngine with Virtex XCV100

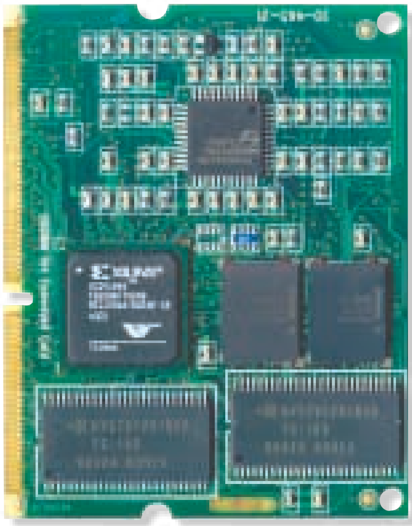

Figure 4 - Virtex XC2V1000 implemented on Hitachi SH4 µEngine

FPGA Intellectual Property

FPGA IP cores for the µEngine are available from many sources:

- NMI provides a wide range of proven IP (for example, PCI host bridge, display controllers, UART, frame grabber, 2D graphics accelerator).
- Xilinx LogiCORE™ IP
- Third-party IP
- Your own IP
- Custom-developed IP.

These elements can be freely mixed in the µEngine to produce the unique functionality required for any application. The NMI deployment module has many features that simplify integration into the final system. For instance, because most systems using high-performance CPUs are running an embedded operating system, such as Windows® CE.NET, we have provided Windows CE software drivers for all of our FPGA IP on the full range of µEngines.

What's more, you can populate a µEngine with various FPGA densities: 50K to 200K gates on Spartan-II FPGAs, 50K to 300K gates on Virtex-E FPGAs, and 250K to 1M gates on Virtex-II FPGAs. This variable gate population makes the µEngine the most cost-effective solution for almost any application, interface, peripheral, or accelerator mix.

Lastly, to ease portability from one FPGA device to another, our FPGA designs use only high-level description languages.

Conclusion

NMI developed the µEngine deployment module through imaginative use of FPGAs in CPU-based systems, creating a high-performance module that provides extraordinary hardware flexibility and upgradeability. The availability of FPGA IP and reference designs facilitates rapid and low-risk development of new products and applications, allowing companies to focus on adding value, rather than having to reinvent the core technology.

NMI offers several development platforms, enabling you to easily evaluate the µEngine and its associated IP.

Editor's note: Since this article was written, NMI Electronics was purchased by Intrinsyc Software Inc. For more information on the innovative use of Xilinx FPGAs in µEngines, visit www.intrinsyc.com/products/microengine/.

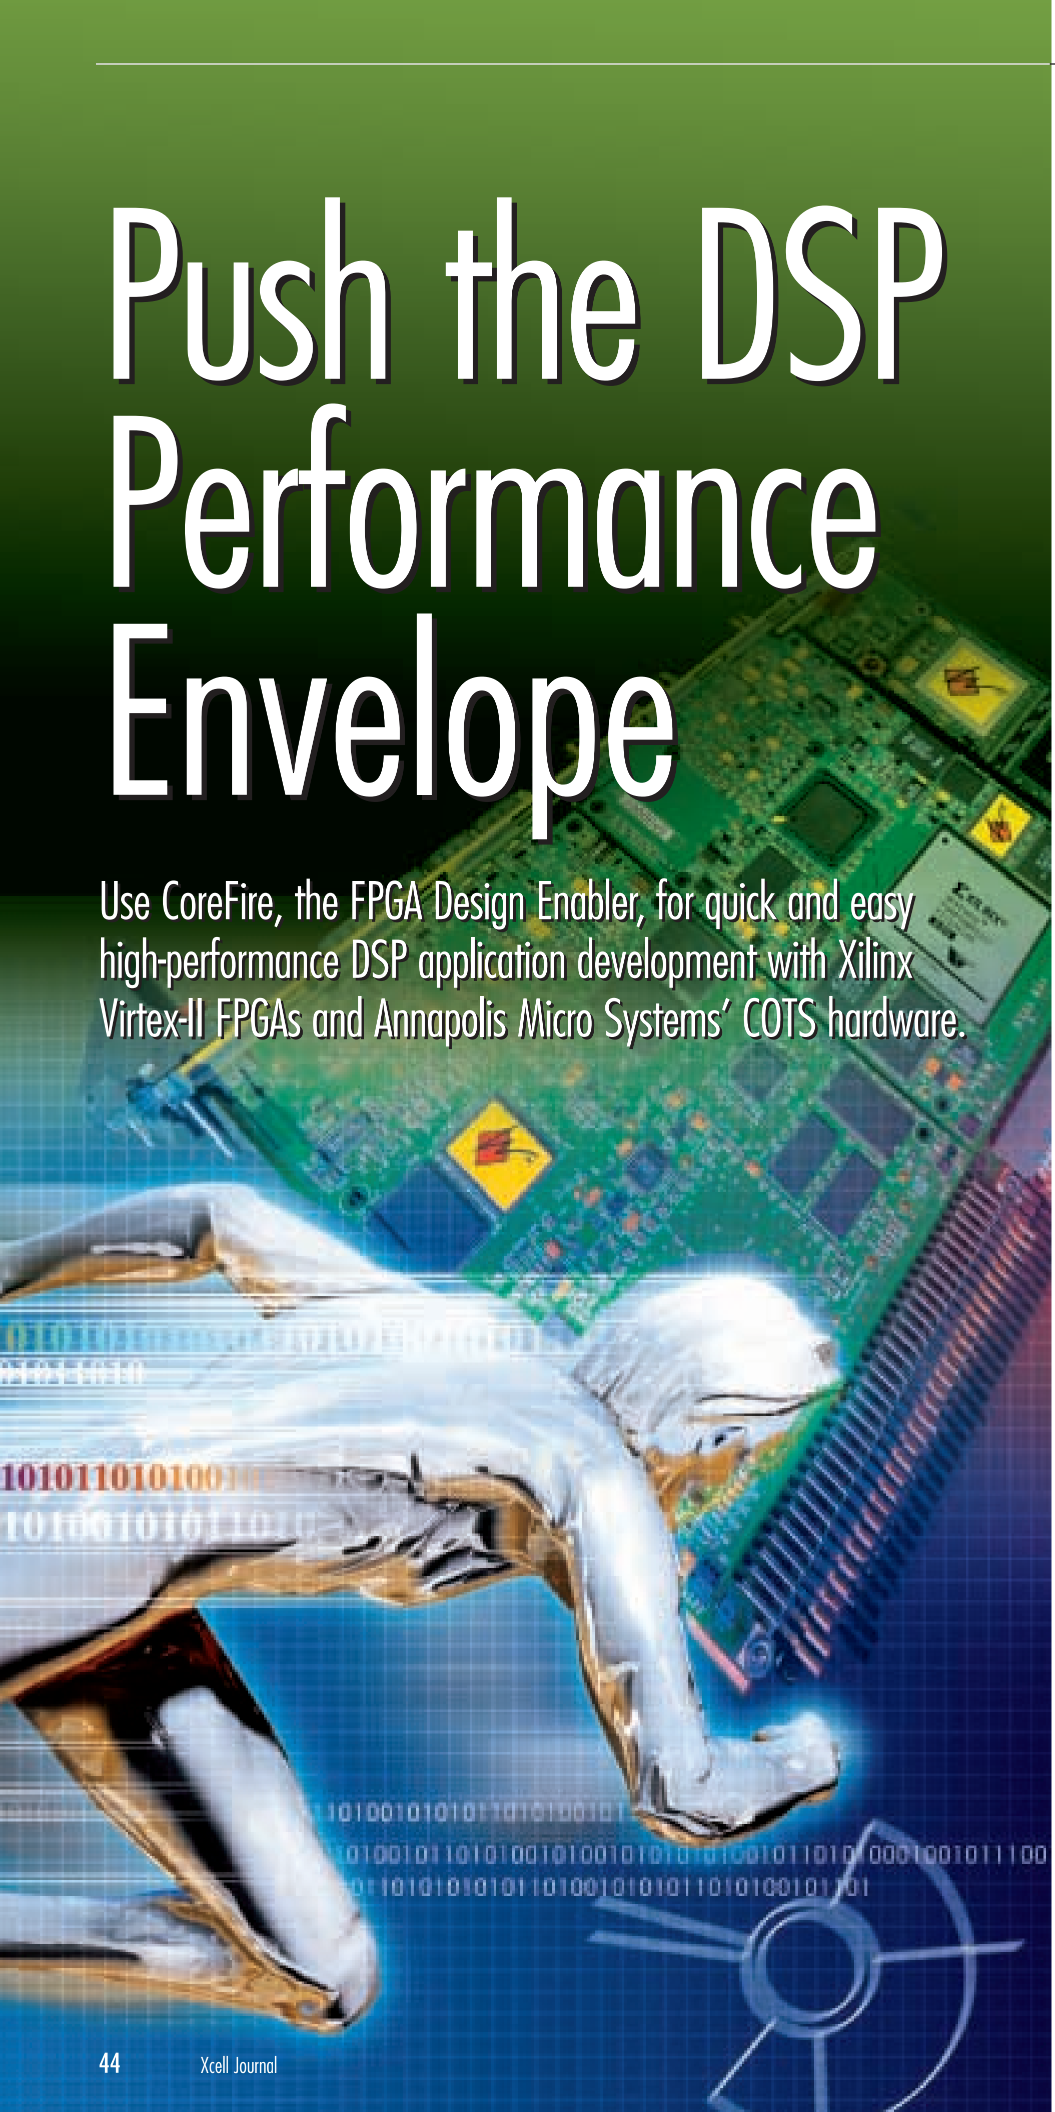

Push the DSP Performance Envelope

Use CoreFire, the FPGA Design Enabler, for quick and easy high-performance DSP application development with Xilinx Virtex-II FPGAs and Annapolis Micro Systems' COTS hardware.

by Jane S. Donaldson
President
Annapolis Micro Systems, Inc.
jdonald@annapmicro.com

DSP developers and their customers know FPGA-based processing outperforms conventional processors on a board-for-board comparison, resulting in significant improvements in processing speed, size, weight, power, and costs. Your FPGA design can be a customized parallel processing chip, specifically crafted for a particular application, accelerating the application to run in hardware and at hardware speeds far faster than could be achieved with software on a generic processor.

- Process data in real time, on site, saving all the time and money involved in data collection and off-site processing.
- Modify the processing by simply reconfiguring the chip (by downloading a different FPGA file) to fix bugs, to adapt to a new set of interface requirements, or to modify the processing in response to application input data or processed results.
- Deliver new applications in place, with no human on-site intervention, by any means of file transfer, including Internet, internal network, hard drive storage, smart card, or wireless modem.

To deploy your application quickly to meet customer demands, you need commercially available hardware with the latest Xilinx FPGAs – plenty of gates, ample memory, and fast standard I/O options like Fibre Channel 2 and 1.5 GHz A-to-D input.

You need a quick and easy way to develop, modify, and test your applications. With VHDL, Verilog, or schematics, even the most experienced ASIC designers need many months to develop applications using upwards of 40 million gates for a single VME or PCI slot.

You can jumpstart your DSP design process – saving time and money – by using the eighth-generation, commercial off-the-shelf (COTS), general purpose Xilinx Virtex™-II FPGA-based hardware from Annapolis Micro Systems, Inc. You can develop at the application level with the easy-to-learn CoreFire™ FPGA Design Enabler. It's loaded with high-performance IP modules, created for your use by FPGA application design experts.

Meet the Demand for Real-Time DSP Applications

Managing digital signal processing data in real time for applications like radar and image processing is very demanding. You need:

- High-speed real-time processing
- Very fast data rates
- A combination of complex and real data types
- Integer and floating point data representations and computation
- Variable and changing data path sizes

When you implement your digital signal processing application in a Xilinx Virtex-II FPGA, you build a customized, parallel-processing design that outperforms both general-purpose processors and digital signal processing chips. Some of the Virtex-II features that enable this very high performance are:

- Chip performance in excess of 300 MHz
- Multiple on-chip memory banks for vector-based processing
- High ratio of memory to logic
- Fast embedded multipliers
- 16 pre-engineered clock domains to support the multiple frequencies and multiple-phase requirements of complex system design.

To illustrate the requirements of DSP applications, we chose a radar signal processor that uses a 16-channel channelizer with a polyphase FIR filter and FFTs to divide the incoming data into multiple frequency channels for real-time processing. Refer to Figure 1 to see the data flow and processing required by this application.

The data comes into the system at 1200 MegaSamples/second with 8 bits per sample. This stream is broken into 16 data streams and processed with polyphase FIR filters and FFTs. The resulting data is again split into 16 channels, the nine most interesting of which are chosen for further processing. Channels 1-3 are sent into the first radar signal processing module, 4-6 are sent into the second radar signal processing module, and 7-9 are sent into the third radar signal processing module, all at 300 MB/s per channel.

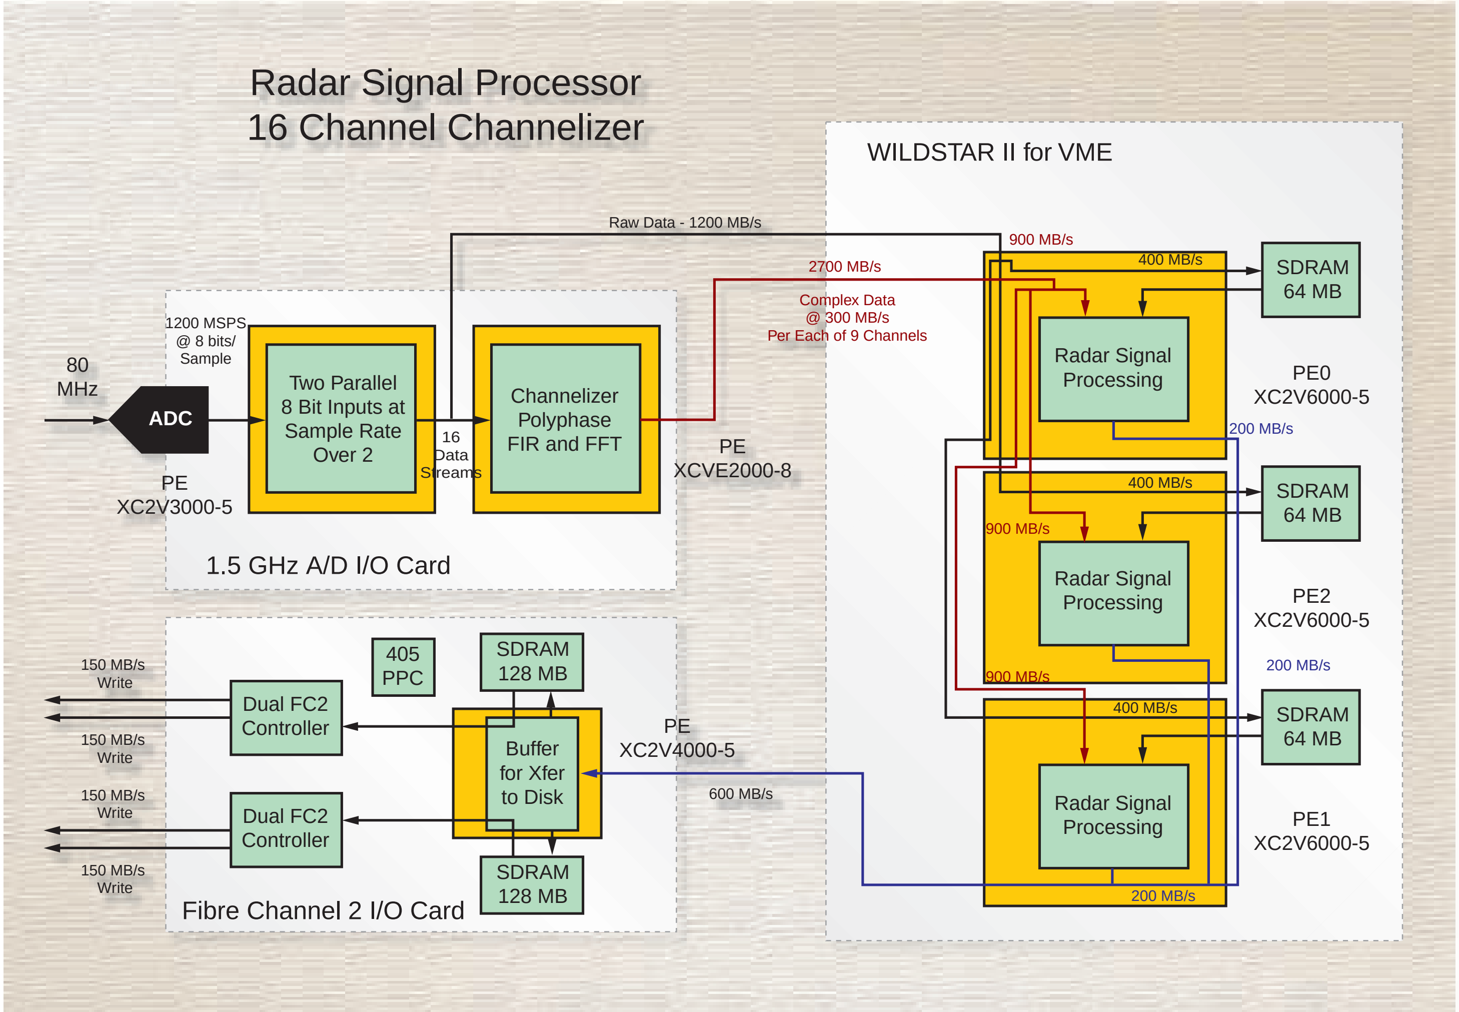

Figure 1 - Channelizer block diagram

The pre-channelizer raw data, at 1200 MB/s, is divided into three streams of 400 MB/s each. Each stream is stored in its own SDRAM block. The appropriate raw data is folded back into the radar signal processing with the channelizer-processed data. The radar signal processors perform filters, FFTs, and other DSP functions on the data.The final result is sent to buffer memory, and then out to disk at 600 MB/s.

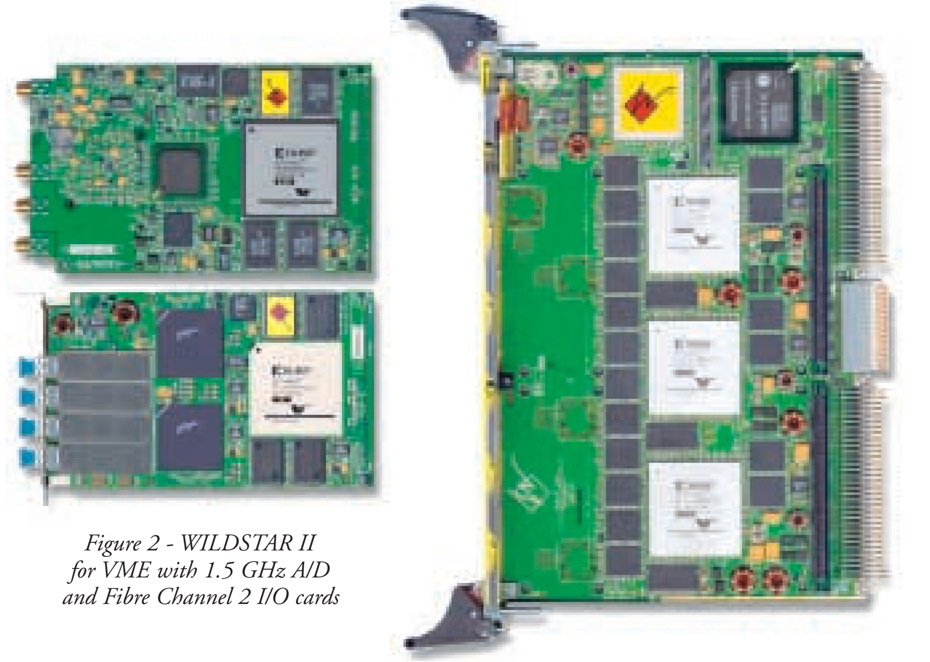

Figure 2 - WILDSTAR II for VME with 1.5 GHz A/D and Fibre Channel 2 I/O cards

Use COTS Hardware from Annapolis for Fast Deployment

On the right side of Figure 2 is the Annapolis WILDSTAR™ II VME board. This board is available with one, two, or three Virtex-II 6000 or 8000 FPGAs, with up to 72 MB of DDR2 SRAM in 18 banks, up to 384 MB of DDR SDRAM in three banks, and programmable flash memory for storing FPGA files for fast reconfiguration.

On the top left side of Figure 2 is the Annapolis 1.5 GHz A/D I/O card, which, for this application, plugs into the top slot of the WILDSTAR II card. This board comes with a MAX 104 or MAX 108 8-bit A/D converter, one Virtex-II 1000 or 3000, one Virtex-E 1000 or 2000, with up to 2 MB of DDR2 SRAM accessible by the Virtex-II bridge PE and up to 16 MB of ZBT SRAM in four banks accessible by the Virtex-E PE.

On the bottom left side of Figure 2 is the Annapolis Fibre Channel 2 I/O card, which, for this application, plugs into the bottom slot of the WILDSTAR II card. This board has four full duplex Fibre Channel 2 I/O channels, with peak rates of 200 MB each way per channel. The board comes with two QLogic ISP2312s, a Virtex-II 4000, 264 MB of DDR SDRAM in four banks, and an IBM PowerPC™ 405 running Linux.

Figure 3 shows the boards connected together and ready to fit into one slot in the VME chassis. The ADC on the 1.5 GHz A/D I/O card performs the analog input and A/D conversion. The Virtex-II and Virtex-E FPGAs on the 1.5 GHz A/D I/O card create parallel data streams and perform the channelizer function, using polyphase filters, FFTs, and other DSP functions, as well as data reduction.

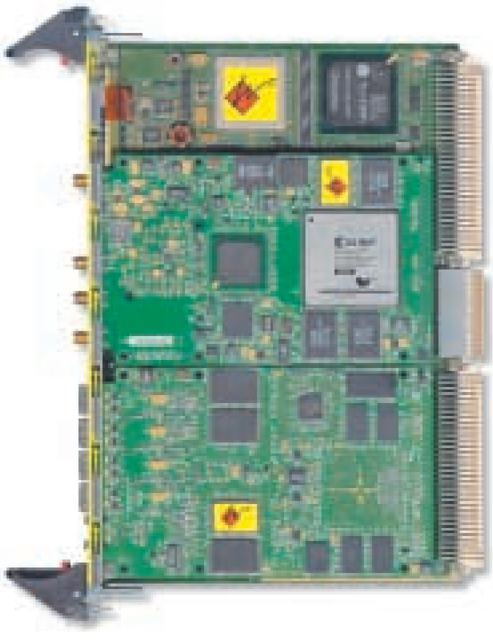

Figure 3 - WILDSTAR II for VME with I/O cards ready to insert in chassis

The channelized data and raw data are both split into three paths in the Virtex-II PE0 on the WILDSTAR II card. Each Virtex-II PE on the WILDSTAR performs radar signal processing functions. The Virtex-II PE1 on the WILDSTAR II card gathers and processes the results for output to the Fibre Channel 2 I/O card.

The Virtex-II FPGA on the Fibre Channel 2 I/O card accepts the data from the WILDSTAR II card, buffers it, and sends it out to disk via the four Fibre Channel 2 channels with the help of the QLogic and PowerPC chips.

Table 1 is a comparison of the system data transfer speeds provided by this Annapolis system to the data transfer speeds required by this channelizer application. You can see that the system easily meets the throughput requirements for the channelizer application.

These WILDSTAR II and I/O boards are the eighth-generation of Xilinx FPGA-based, high-performance processing boards produced by Annapolis Micro Systems. Annapolis continues to push the high-performance envelope, using latest-standard Xilinx FPGAs.

You Can Build Your Application Quickly and Easily with CoreFire

You can see that the channelizer application fits on the chosen WILDSTAR II system, so acquiring readily available hardware for your application will be easy.

The next step is to figure out how you will develop the host software and FPGA implementations for your application. Remember, this project stretches across three different printed circuit cards and six different FPGAs.

Data Path	Channelizer Data Transfer Speed Requirements	System Data Transfer Speed
Input to ADC Speed	80 MHz	75-150 MHz
ADC to PE Speed	1200 MegaSamples/s	80-1500 MegaSamples/s
A/D I/O to PE0 on WS II	3900 MB/s	4000 MB/s
WS II PE to its SDRAM	900 MB/s	950 MB/s
PE0 to PE1	1300 MB/s	4000 MB/s
PE0 to PE2	1500 MB/s	4000 MB/s
WS II PE2 to FC II I/O	600 MB/s	4000 MB/s

Table 1 - System data transfer speeds versus channelizer requirements

The classic VHDL methodology for implementing applications on FPGAs is difficult, and requires expert knowledge and countless months of painstaking work. You cannot wait months to deploy your product. You need a tool that will allow you to deploy your project within weeks, not months or years. You must be able to develop new application files rapidly and easily, as well as accommodate specification changes, functional additions, and algorithm development.

Using the CoreFire FPGA Design Suite from Annapolis, you can implement each of your algorithms in as little as a few hours. Use the standard WILDSTAR II C or Java API to write your host program. The CoreFire board support packages handle all the I/O, memory, and FPGA interfaces seamlessly, providing excellent performance. Refer to the CoreFire screen display in Figure 4.

CoreFire is a graphical user interface FPGA application development tool that allows you to build your application very quickly by dragging and dropping library elements onto the design window. Choose from more than 400 expertly crafted modules. Modify your input and output types, numbers of bits, and other variables by changing module parameters with pull-down menus. Move modules around on the screen and reconnect with a flick of the mouse.

The modules automatically provide correct timing and clock control. Insert debug modules to report actual hardware values for in-the-loop debugging. Hit the Build Button to check for errors and sizes and to build an encrypted EDIF file. Use the Xilinx ISE tool to place-and-route each FPGA design.

Modify and use the jar file created by the CoreFire build to load your new file into your WILDSTAR II and I/O card hardware. Use the CoreFire debugger to view and modify register and memory contents in the FPGA, and to step through the data flow of your design running in the real physical hardware.

Armed with your debug results, you will find it very easy to use the CoreFire design window to modify and rebuild your FPGA design until you are satisfied with the results. Use the CoreFire program to build and debug each of your FPGA designs, and then use the jar file and the WILDSTAR II API to develop your overall host program.

Conclusion

It is easy to push the DSP performance limits with Virtex-II FPGAs and Annapolis Micro Systems boards. Some of our customers have gone from initial inquiry to first deployment in as short a time as two months.

When you buy high-performance, world-class, Virtex-II-based off-the-shelf hardware from Annapolis, it is easy to build and modify applications. You have more time to fine-tune your algorithms. You can get prototypes up and running sooner, so you have more time to test market your product.

Final development is just as easy. You will be in the market far ahead of your competition, saving time and money. To learn more, contact Annapolis Micro Systems, Inc., at 410-841-2514, or visit our website at *www. annapmicro.com.*

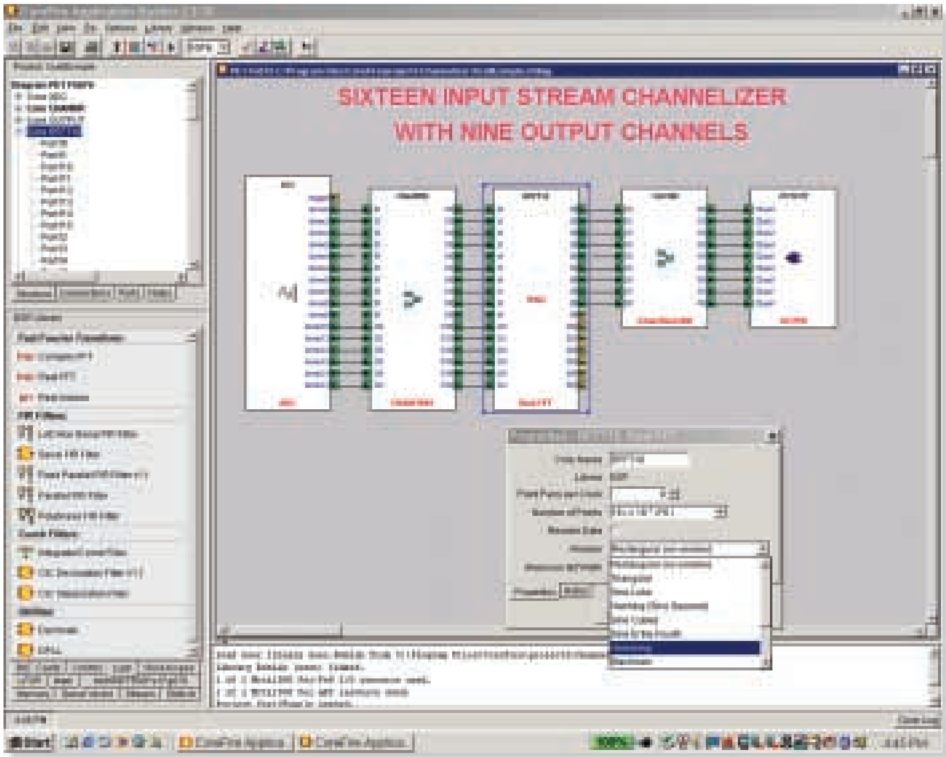

Figure 4 - CoreFire screen display

ISE 5.2i Further Reduces Your Design Costs

Xilinx industry-leading design tools provide a low-cost, low-risk, and high-performance logic solution.

by Mark Goosman / Lee Hansen
Product Marketing Managers
Xilinx, Inc.
mark.goosman@xilinx.com, lee.hansen@xilinx.com

Reducing project costs isn't new to most designers, but in tight economic times the pressure to bring project costs down becomes much more important. In his *Xcell Journal* Winter 2002 article titled, "When Total Cost Management Counts, Xilinx PLDs Pay Off" (*www.xilinx.com/publications/products/cool2/xc_tcm43.htm*), Eric Thacker described how programmable logic devices (PLDs) offer significant benefits in dynamic, rapidly changing markets.

With our ISE development systems and development options, Xilinx not only supports the benefits of PLDs but also offers additional cost savings. ISE 5.2i, the latest release of our design software, delivers a number of productivity technologies that shorten logic design flow, optimize design results, shorten implementation and verification cycles, and provide interactive design assistance. At the same time, ISE 5.2i enables you to realize even faster design performance. The end result to you is cost savings across your entire project.

The shorter design cycles and time-to-market advantages of FPGAs and CPLDs mean that you need less engineering resources. This allows you to make the best use of your staff when difficult economic conditions restrict your ability to hire more engineers. Our fast, efficient, and highly productive ISE software tools help you get the job done in less time, and they make each engineer more productive.

Free ISE WebPACK

The ISE WebPACK™ design suite is the ideal Web-downloadable desktop solution. It offers a complete development environment with modules from ABEL and HDL synthesis to device fitting and JTAG programming. ISE WebPACK tools are a subset of our award-winning ISE Foundation™ design tools, providing instant access to the ISE tools at no cost. By providing a design solution that is always up-to-date, with error-free downloading and single file installation, Xilinx has created a solution that allows instant productivity. Because ISE WebPACK development tools are available for download from the Xilinx website at *www.xilinx.com/ise/webpack5*, you can get started immediately on designs for leading Xilinx CPLDs and mid-density FPGAs.

This Web-downloadable design solution reduces your design costs by including all the tools you need to complete your design.

ISE WebPACK includes:

- **ModelSim Xilinx Edition (MXE-II) Starter Version**
  ModelSim® XE is a complete HDL simulation environment that has been optimized for programmable logic design, enabling you to quickly verify source code and functional and timing models of your design.
- **HDL Bencher**
  Within the ISE WebPACK toolset, the HDL Bencher™ test bench generator automatically imports the current HDL design file and creates an editable stimulus waveform by default.
- **StateCAD**
  The StateCAD® FSM wizard automates the state machine design process. You can specify complex state machines to quickly meet tough product requirements. The state machines can then be automatically translated to an HDL format you can include in your design flow.
- **ChipViewer**
  ChipViewer is a pre- and post-fit graphical utility to assign or view pin placement and implemented logic for all Xilinx CPLD devices. This removes the risks associated with changes late in the design process.
- **XPower**
  XPower is a graphical power-analysis tool. Total device power, power per-net, fitted, routed, partially routed, or unrouted designs can be easily analyzed.

Optimized Design Performance and Device Utilization

Xilinx ISE design tools have raised the industry standard for both design performance and device utilization. Through patented implementation algorithms, ISE allows you to achieve the fastest possible design performance. Compared with competitive solutions, designs can achieve better than 15% higher performance.

This performance edge means you can potentially target a lower cost device – leveraging faster performance from the software. Thus, you can hit your timing goals earlier, spending less time in the design flow.

For example, based on benchmark data, you can achieve 20% to 30% better performance in Virtex-II Pro™ designs using ISE than you can get from an offering from the leading competitor. In many cases, you can target your design to a slower speed grade device and still achieve targeted design performance.

ISE also reduces project costs by packing more logic into Virtex™-II devices, letting you fit your design in the smallest possible device. Advanced FPGAs are not solely made of look-up tables and flip-flops anymore. Today's logic fabrics are best described as "feature rich." This trend requires sophisticated algorithms in both synthesis and implementation tools, providing optimal performance and logic utilization by leveraging new hardware features.

Xilinx ISE development tools separate unrelated functions and assign them to different clusters (called a slice) on the fabric. This avoids conflicting placement constraint and guarantees optimal performance. As the device gets full, powerful algorithms pack unrelated logic into common clusters. This gradual process ensures that the device is utilized at its best, with minimal impact to design performance.

With competing FPGA solutions development tools, the packing of logic requires a special option. With this option turned on, packing is limited, because an unrelated LUT using its 4-inputs and a flip-flop cannot be merged together in any logic element – and the limited packing comes at a cost to design performance. Virtex-II logic utilization with ISE comes out 15% better than the nearest competitive offering.

In the Virtex-II fabric, the LUT and flip-flop can be used independently, without restrictions. In Stratix devices, a LUT cannot be used with its flip-flop in all circumstances, because one input pin of the LUT is shared with the path that has direct access to the flip-flop. By default, when the flip-flop is not fed by any logic, the LUT in that LE is unavailable to the rest of the design. As a remedy, Quartus II tools provide a register packing option (off by default) to enable the packing of LUTs along with the flip-flop. This still does not allow LUTs using 4-inputs to be packed, because the connectivity restriction is still present. Figure 1 shows LUT to flip-flop connectivity in both Virtex-II and Stratix devices.

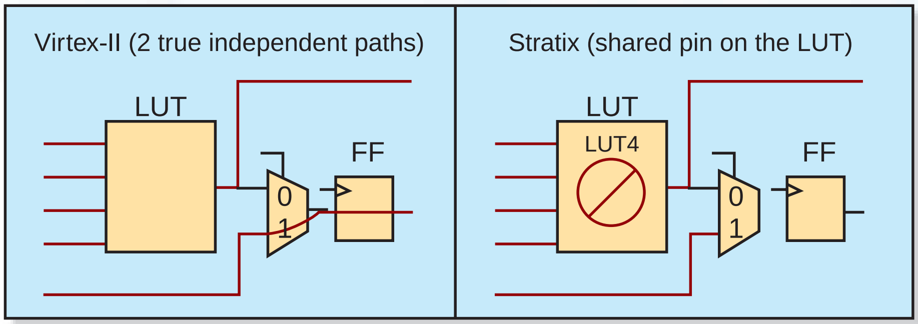
Virtex-II (2 true independent paths)
LUT
FF
0
1
Stratix (shared pin on the LUT)
LUT
LUT4
FF
0
1

Figure 1 - Connectivity LUT to flip-flop in Virtex-II and Stratix devices

Advanced Technology Streamlines the Design Flow

ISE is also packed with advanced software technology designed to accelerate the more time-consuming parts of the design and debug logic flow.

Incremental Design is a technology included in ISE that shortens design recompile times. By locking performance for areas of the design that don't need to change, Incremental Design lets you perform re-synthesis and re-place-and-route on only those pieces of the design that have to change. This reduction in time adds up fast in the crucial verification cycle, where debug changes are common.

The Xilinx ChipScope™ Pro integrated logic analyzer also delivers added productivity to the verification cycle. Through small, easy-to-place software debug cores, the ChipScope Pro tool allows you to monitor – in real time – any signal in the FPGA. This includes the IBM® PowerPC™ 405 peripheral bus in the advanced Virtex-II Pro FPGA. Design signals are captured and brought to the outside world through the FPGA JTAG programming port. This minimizes the amount of dedicated FPGA space and I/O pins required – as opposed to using more traditional ASIC and competing FPGA debug methodologies.

Additionally, signal monitor points can be changed through the ISE FPGA editor without having to re-compile the design, saving even more debug time. The ChipScope Pro analyzer cuts verification times dramatically, even when the device is on the board – or in the field.

Conclusion

For logic design, the true cost of the project includes much more than just device cost. Factors like development cost, project timelines, access to development tools, designer efficiency, ability to achieve device performance goals, and verification costs can have a big impact on the overall project cost.

Xilinx allows you to meet – or beat – your project budget through free ISE WebPACK development tools, other ISE configurations, a complete design environment, ISE's powerful implementation tools, robust verification technology, and more. As you evaluate various logic design solutions, look at the total costs associated with design tools and designer resources in addition to the cost of the device.

Prototype *All* Xilinx Devices In-System or Standalone – with MultiPRO Desktop Tool

Now you can program/configure Xilinx devices from your desktop with a single programming hardware solution.

by Michelle Badal
Marketing Manager, Configuration Solutions
Xilinx, Inc.
michelle.badal@xilinx.com

In today's competitive landscape, engineers need to have flexibility – it's the key to maximizing your design efforts. The MultiPRO Desktop Tool, our latest programming hardware, offers a number of features to provide you with the most flexible, low-risk, and cost-effective way to prototype all Xilinx devices.

MultiPRO Desktop Tool Increases Flexibility and Reduces Cost

MultiPRO Desktop Tool (Table 1), released in December, is a complete programming solution that enables you to realize the full potential of Xilinx programmable logic devices. Designed specifically to interface with a PC via parallel port (IEEE 1284), the MultiPRO Desktop Tool provides:

- In-system programming and configuration of all Xilinx devices
- Standalone programming of CoolRunner™-II CPLDs and XC18V00 ISP PROMs
- Comprehensive configuration mode support
- A low-cost solution by integrating desktop programmer and download cable functionality.

In-System Programming

With MultiPRO Desktop Tool, you can program and configure all Xilinx devices in-system. MultiPRO Desktop Tool's flexibility enables you to program your devices right at your desk by following these steps:

- Power MultiPRO with +5VDC, via an external AC power brick
- Use ISE iMPACT v5.1i SP3 or higher on your PC
- Interface the MultiPRO pod with the PC via a parallel port (IEEE 1284) (cable included)
- Connect the MultiPRO pod to the PCB with the ribbon cables (included)
- Run software to download bitstream and program/configure target FPGA, CPLD, or PROM (Figure 1).

Standalone Programming

MultiPRO Desktop Tool reduces the risks of prototyping by enabling you to perform standalone programming on CoolRunner-II CPLDs and XC18V00 ISP. As with in-system programming, you can program your devices right at your desk by following these steps:

- Power MultiPRO with +5VDC, via an external AC power brick
- Use ISE iMPACT v5.1i SP3 or higher on your PC
- Interface the MultiPRO pod with the PC via a parallel port (IEEE 1284) (cable included)
- Connect the MultiPRO pod to a CoolRunner-II CPLD or XC18V00 ISP PROM adapter (adapters available for all package types)
- Run software to program target CoolRunner-II or XC18V00 ISP PROM (Figure 2).

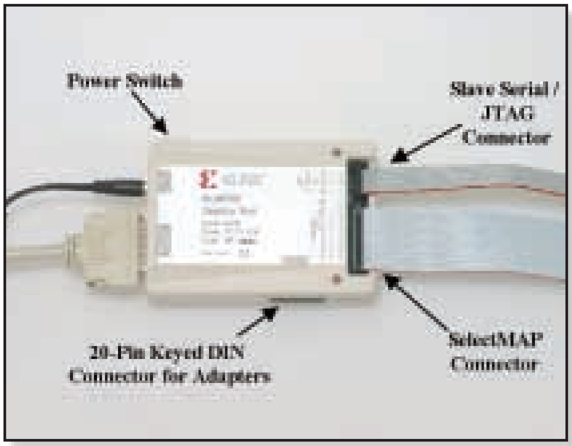
Power Switch
Slave Serial / JTAG Connector
20-Pin Keyed DIN Connector for Adapters
SelectMAP Connector

Figure 1 - MultiPRO with download cable ribbons

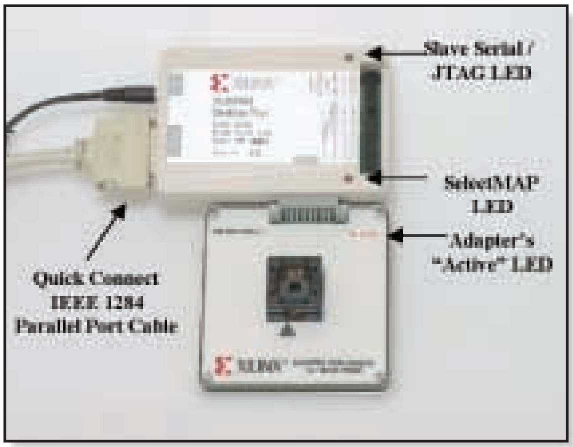

Figure 2 - MultiPRO with programming adapter

Comprehensive Configuration Mode Support

The comprehensive configuration mode support gives you the flexibility to choose the most suitable mode for your design, while using only one programmer hardware solution. MultiPRO Desktop Tool is supported by ISE iMPACT download software v5.1i SP3 or higher, and supports JTAG (IEEE 1149.1), Xilinx slave serial, and SelectMAP modes (Table 1).

MultiPRO Desktop Tool Features

- Provides standardization with J Drive™ IEEE 1532 programming engine
- Provides prototype debug environment with ChipScope™ ILA and ChipScope Pro compatible
- Automatically adapts to target I/O voltage.

Conclusion

MultiPRO Desktop Tool reduces the risks and costs of prototyping in a number of ways:

- It provides a single solution that meets all your desktop programming and download cable needs.
- It frees up more expensive resources designed for mass programming and configuring of devices.

MultiPRO Desktop Tool provides the most comprehensive programming solution to date from Xilinx. For more information on how to increase flexibility, reduce risk, and reduce system prototyping cost, visit: *www.xilinx.com/xlnx/xil_prodcat_product.jsp?title=csd_cables/.*

	Devices Supported	Software	Host/Interface	Configuration Mode
MultiPRO Desktop Tool	• Standalone programming of CoolRunner-II CPLDs and XC18V00 ISP PROMs • In-system programming of all Xilinx ISP PROMs and CPLDs • In-system configuration of all Xilinx FPGAs	• ISE iMPACT download software v5.1i SP3 or higher	• PC (Parallel Port - IEEE 1284)	• JTAG (IEEE 1149.1) mode • Slave Serial mode • SelectMAP mode • Desktop Programming mode

Table 1 - MultiPRO Desktop Tool

DLK Enables Cost-Effective Design

The new design environment from SENG digitale Systeme GmbH – based on FPGAs and standard parts – makes it easier to estimate development costs.

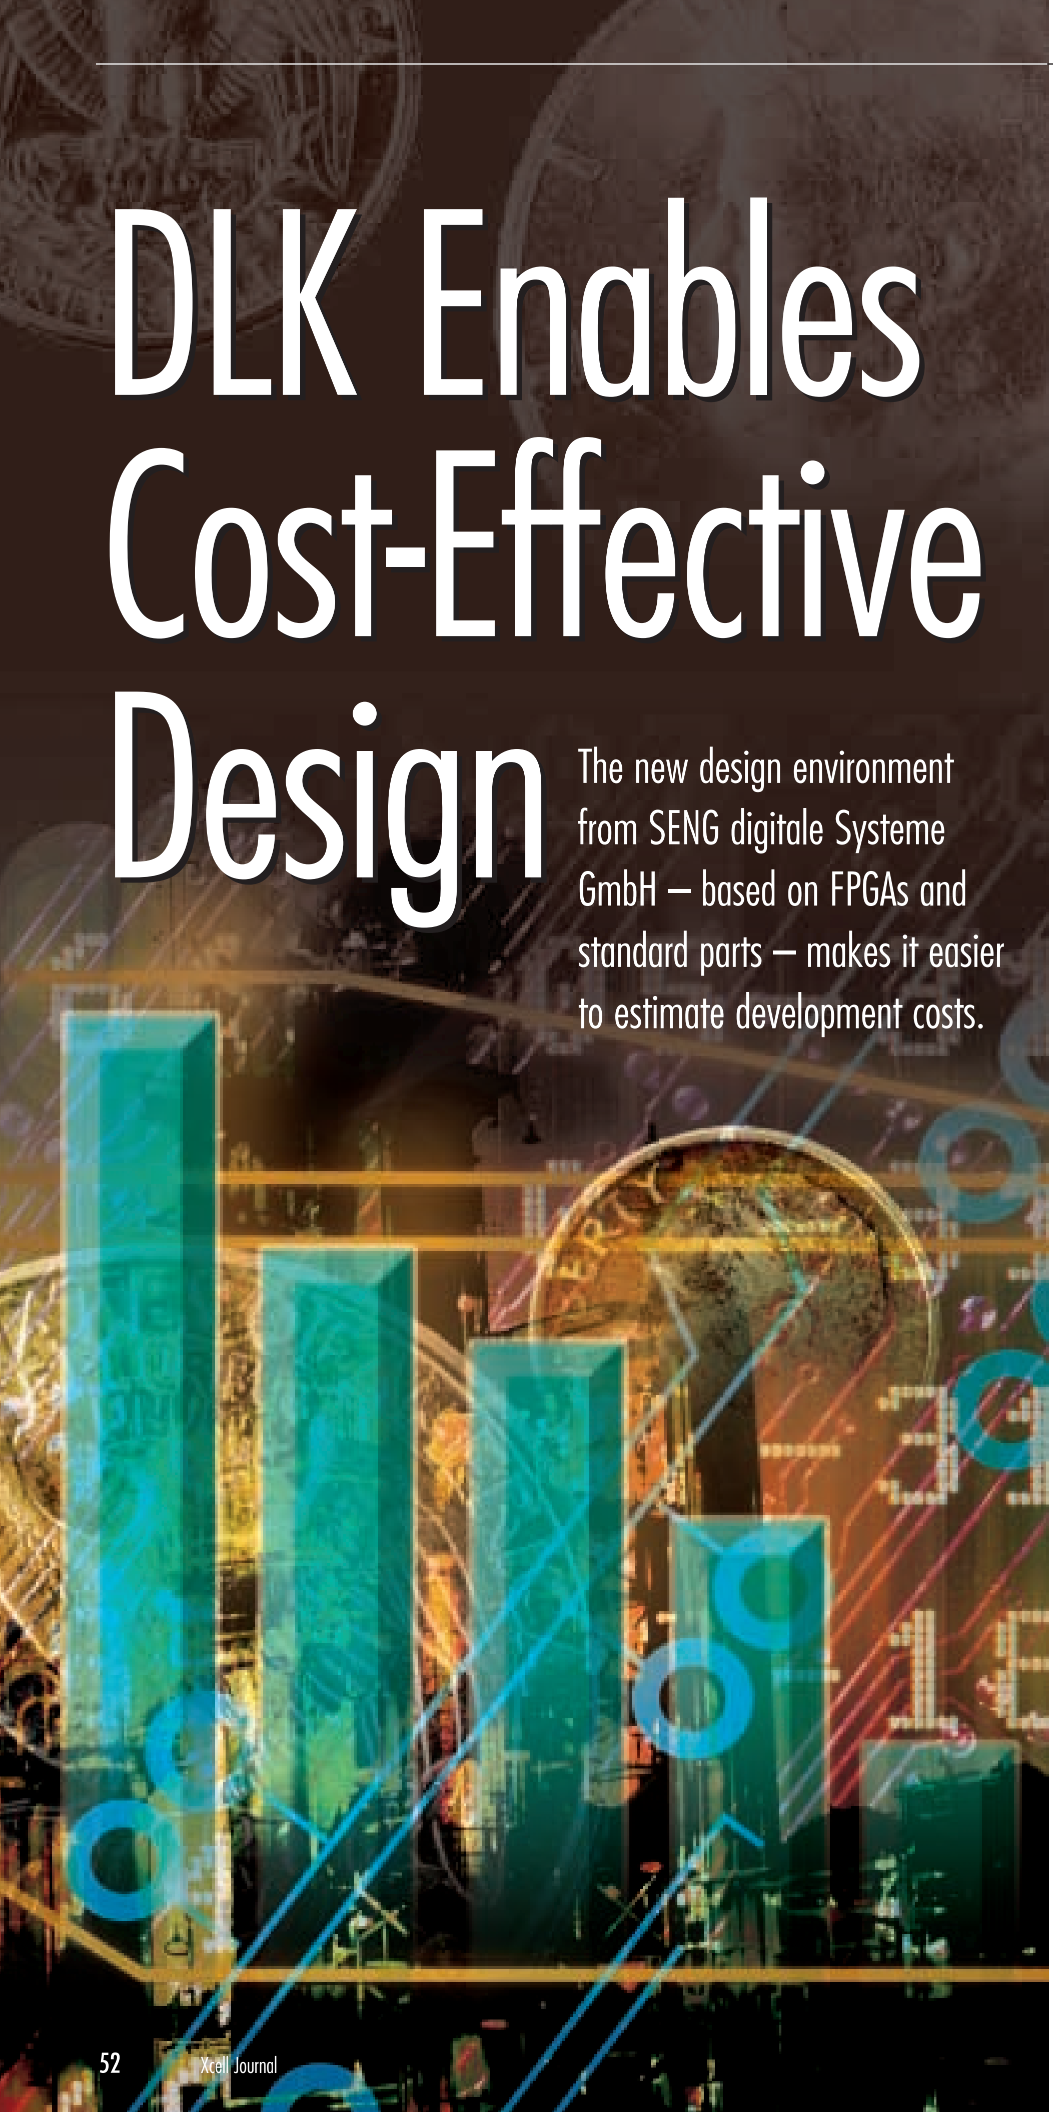

by Peter Seng
Managing Director
SENG digitale Systeme GmbH
peter@seng.de

Existing application-specific development and production solutions may be fine for high-volume markets, but for low- and mid-range product volumes, these solutions are most often either too expensive or otherwise not suitable. As a result, these types of products are usually built with off-the-shelf PC- or microcontroller board-based components. Although standard components are relatively low cost, they are rarely a perfect fit for your design specification and almost always require you to develop and add specialized hardware. Programmable FPGA-based hardware offers an attractive alternative, but so far, entrance barriers have been high, and development time and costs difficult to predict.

SENG digitale Systeme GmbH recently introduced a scalable development environment that makes it easy to prototype and produce programmable FPGA-based hardware using standard parts and at predictable costs. The core of the SENG environment is the Digital Logic Kernel (DLK), which consists of an FPGA + CPU + memory + PC interface. The DLK uses standard Xilinx parts, including FPGA and CPLD source codes, software, and design rules, thus eliminating the need to program equipment or preprogram parts before use. The DLK is, by default, a self-bootable device. To exchange data or for administration purposes, the DLK is accessed via an integrated PC parallel printer port interface. All you need to build, program, and service devices in the field is a PC with a parallel printer port (Figure 1).

The DLK Concept

The DLK is well suited for designing, building, and programming products containing any kind of hard or soft CPU with internal and external memory, and programmable logic. The system interconnect concept makes the DLK appropriate for several CPU, FPGA, and CPLD device families.

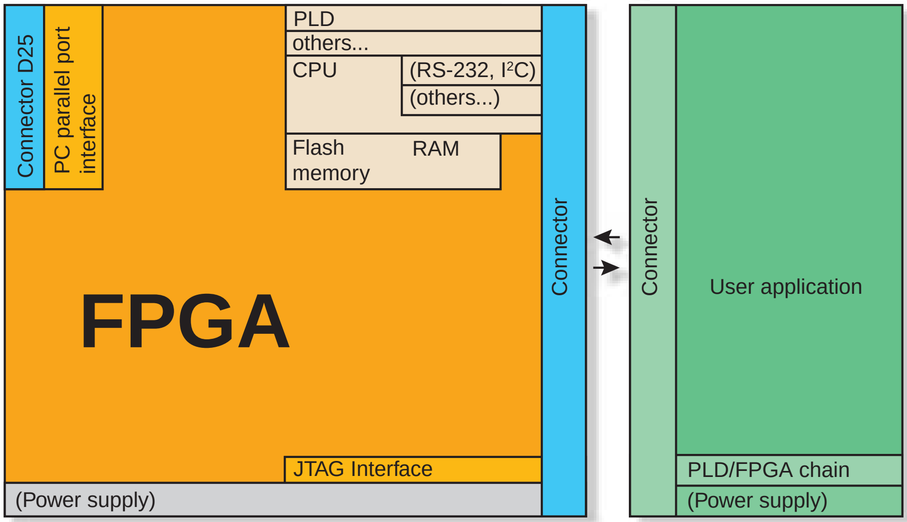

Figure 1 - DLK concept

Components

The DLK development kit consists of software and hardware source code, software, schematics, and a demonstration board. Components include:

- FPGA internal 8-bit bidirectional interface to the PC with parallel I/O bus structure (Figure 2)
- Driver and software for several operating systems, including Win 9x- and WinNT 4.0-based OS
- Flash CPLD-based state machines and memory
- FPGA internal emulation of JTAG programming interface
- Administration software running on PC for Windows™ OS, including dynamic link library, application program, and source code (Figure 3)
- Demonstration board (Figure 4)
- FPGA source codes (varies with grade of license)
- 8032 C sample source code containing LCD, UART, interrupt, and I²C routines.

DLK Design Flow

1. Design an application-specific board according to the DLK principles, or use one of the available DLK boards.

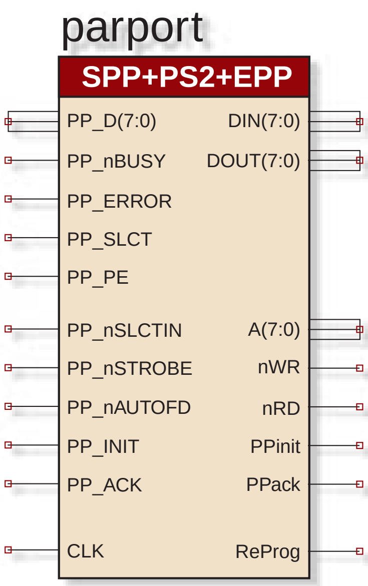

Figure 2 - PC parallel port interface, ISE schematic symbol

2. Rearrange FPGA, CPLD, and PC software according to your specific needs; compile.
3. Connect application-specific board to PC parallel port; configure it with JTAG emulation bitstream using DLK software.
4. Use iMPACT software to embed flash memory on board CPLD.
5. Develop application-specific logic and CPU programs; compile.
6. Download FPGA bitstream to the board and communicate with it using DLK routines.
7. Debug your application.
8. Store the final FPGA configuration bitstream and CPU program in flash memory; disconnect from PC. The application is ready.
9. For administration, upgrade, or communication purposes, just reconnect board to PC.

Conclusion

The DLK concept is an open system environment that uses standard tools and readily available elements, enabling development of digital products quickly and cost-effectively. This development system is especially suited for developing prototypes and for low- and mid-range product volumes. The system is compatible with Xilinx ISE design tools and Windows operating systems. The kit includes several real-world examples to minimize training. All source code is available.

DLK development kits, including demo boards, are available now, starting at $298. System and support is available from SENG digitale Systeme GmbH. Write to *info@seng.de* or visit *www.seng.de* for more information.

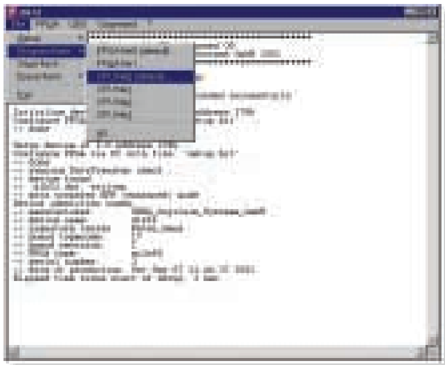

Figure 3 - DLK running on Windows

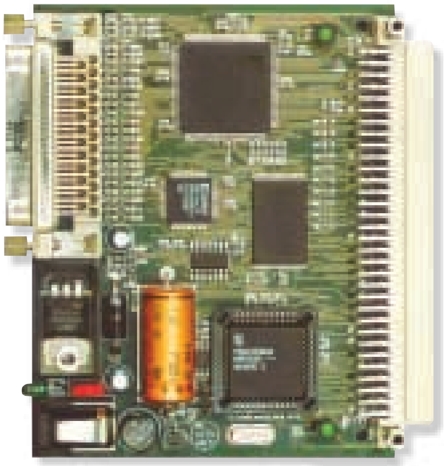

Figure 4 - DLK51 board

Use SpyGlass Predictive Analysis for Effective RTL Coding

Atrenta Inc.'s SpyGlass software uses a "look-ahead" engine based on fast-synthesis technology to help you identify potential problems – and fix them – early in the design process.

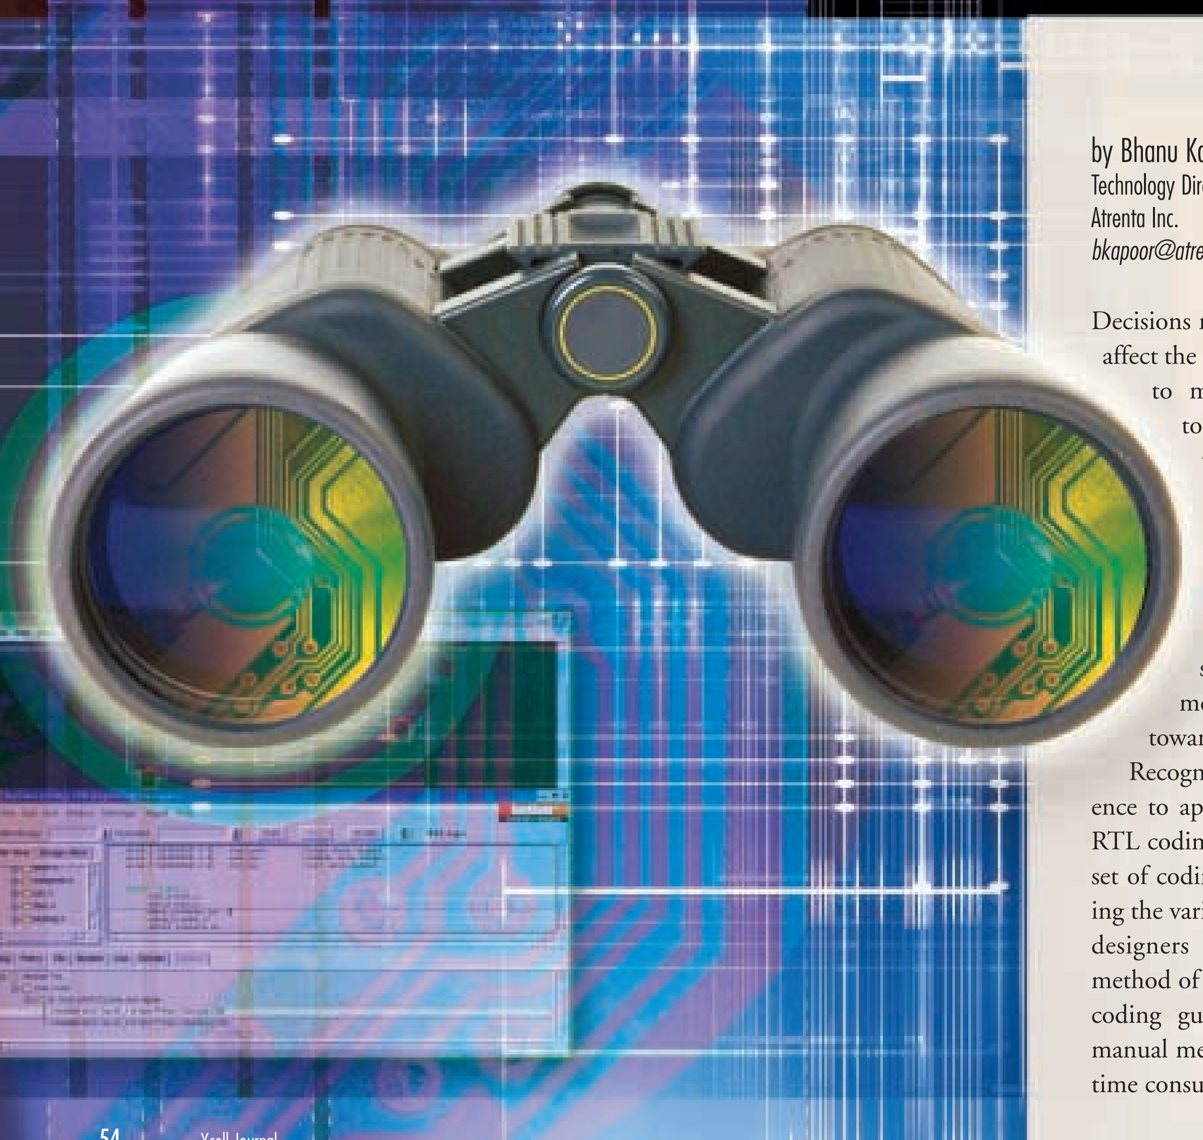

by Bhanu Kapoor
Technology Director
Atrenta Inc.
bkapoor@atrenta.com

Decisions made early in the design process affect the entire chip design process. Thus, to manage designs effectively, the tools you use for RTL (register transfer level/language) design must enable a stepwise refinement of the code by predicting likely downstream problems. For example, to achieve high performance, the guidelines for synchronous design must be met, as they play an important role toward meeting timing objectives.

Recognizing the importance of adherence to appropriate guidelines during the RTL coding process, Xilinx recommends a set of coding guidelines for designs targeting the various Xilinx device families. Most designers currently follow a manual method of design reviews to check to see if coding guidelines have been met. This manual method is both prone to error and time consuming. Automating adherence to

RTL coding guidelines, therefore, would greatly increase the timely success of projects targeting high-performance designs on Xilinx device families, such as the Virtex™-II series of platform FPGAs.

The SpyGlass™ Predictive Analyzer tool from Atrenta Inc. automates the process of meeting design guidelines through the use of an underlying predictive analysis technology. This approach performs detailed structural analysis on RTL aspects to check for coding styles, RTL-handoff, design reuse, clock/reset requirements, verification, timing guidelines, and much more. The SpyGlass tool employs a "look-ahead" engine based on fast-synthesis technology and a fast, built-in, cycle-based simulator to carry out such analyses. Such a look-ahead methodology only uses RTL code as input to the SpyGlass tool and does not require any vectors or assertions. It is therefore very easy to set up and run.

Policy-Based RTL Coding

The SpyGlass™ Predictive Analyzer is a comprehensive, policy-based system that defines, in a succinct and organized way, design policies that automatically point out time-consuming downstream issues. The SpyGlass tool then offers suggestions on ways to overcome these downstream issues during the RTL code development process. This helps you to meet your time-to-market target.

Policy

A policy is a collection of rules for a specific purpose, such as rules associated with a standard, a silicon vendor, or a specific design tool. Policies enhance user extensibility, allowing you to develop and manage customized groupings of rules more easily. The design methodology includes a set of policies that you can select during the process of RTL code development. Examples of such policies include lint, reuse, verification, timing, and testability.

Rule Groups

A collection of rules within a policy is termed a group. Typically, groups consist of rules addressing a particular area of interest in the RTL code. Groups are hierarchical, meaning that a group can contain other lower-level rule groups as well as individual rules. A group provides an additional level of modularity in applying policies to a given RTL design.

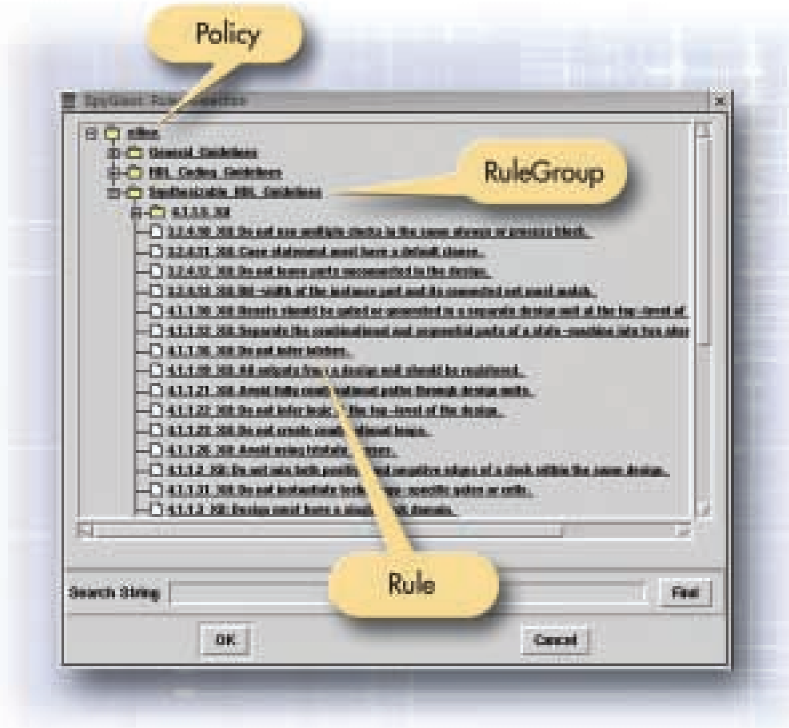

```
xilinx
  General_Guidelines
  RTL_Coding_Guidelines
  Synthesizable_RTL_Guidelines
    4.1.1.5_XIl
      3.2.4.10_XIl Do not use multiple clocks in the same always or process block.
      3.2.4.21_XIl Case statement must have a default clause.
      3.2.4.32_XIl Do not leave ports unconnected in the design.
      3.2.4.33_XIl Bit-width of the instance port and its connected net must match.
      4.1.1.30_XIl Resets should be gated or generated in a separate design unit at the top-level of
      4.1.1.32_XIl Separate the combinational and sequential parts of a state-machine into two alw
      4.1.1.15_XIl Do not infer latches.
      4.1.1.20_XIl All outputs from a design unit should be registered.
      4.1.1.21_XIl Avoid fully combinational paths through design units.
      4.1.1.23_XIl Do not infer logic at the top-level of the design.
      4.1.1.25_XIl Do not create combinational loops.
      4.1.1.26_XIl Avoid using tristate buffers.
      4.1.1.2_XIl Do not mix both positive and negative edges of a clock within the same design.
      4.1.1.31_XIl Do not instantiate technology specific gates or cells.
      4.1.1.3_XIl Design must have a single clock domain.
```

Figure 1 - Example of policy consisting of rule groups and rules

Rules

A rule is the most fundamental element in the policy-based management system. It describes a set of conditions that – when checked by the policy engine – result in an indication of a specific problem with the RTL code. Rules allow standard analysis of the RTL code. An example of a policy, rule groups, and rules in the SpyGlass system is shown in Figure 1. With this system, you can selectively turn on specific groups or specific rules within a group.

Policy Engine

The policy engine accelerates electronic product development by enabling development teams to capture, aggregate, distribute, and apply constraints and requirements early in the development cycle. The fast synthesis engine internally creates a structural view of the design and foresees downstream issues early in the development cycle, thereby eliminating errors at the earliest possible stage.

Although you can check many complex rules statically on the internal synthesized structural view, other rules require some understanding of the logic function of the design. This is particularly true for testability-related checks. In order to perform a testability check, you must use an evaluator. The evaluator in the policy engine is effectively a zero-delay, cycle-based simulator you can use to resolve functional design constraints, as well as carry out a simulation required to set up the design for testability analysis.

Policy implementation requires a traversal engine that works on the RTL netlist produced by fast synthesis. The SpyGlass engine generates basic primitives for rules that cover design traversal. The connectivity information, coupled with the traversal primitives, enable you to create rules that look for violations across the design hierarchy.

By applying a policy over a given RTL design code, you can obtain a wealth of information about the design – as well as any of violations of the defined rules. The SpyGlass interface highlights rule violations on the specific lines of RTL code. Not only does the SpyGlass predictive analysis tool offer an extensive help function to assist you in understanding the nature of the violation, but also, where

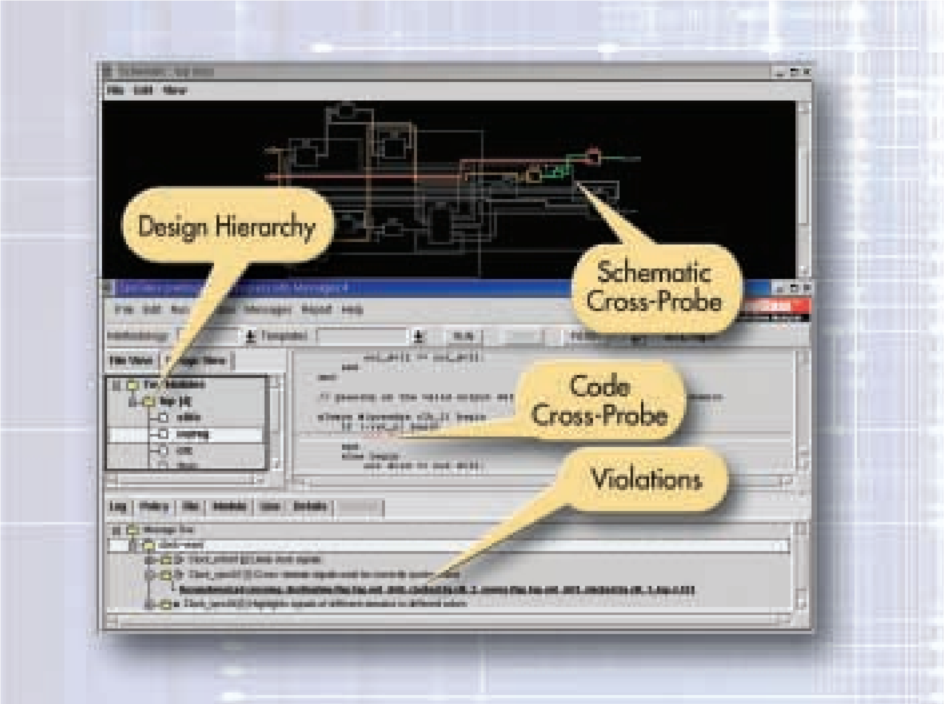

Figure 2 - Example of rule violation and associated debug windows

appropriate, it makes suggestions for resolving the problem.

Moreover, each violation is highlighted on a structural view of the design, which can give you additional debugging information for complex problems. And, as shown in Figure 2, cross-probing between the code and structural views reveals an example of a clock domain synchronization problem. Additionally, a comprehensive set of reporting mechanisms and violation management utilities – such as waiver to allow a specific rule violation in a given point in the design – further facilitate the RTL coding and assessment process.

Coding for Xilinx Devices

Xilinx recommends a set of coding guidelines for designs targeting the various Xilinx device families. These include general coding guidelines, as well as synthesizable HDL guidelines. Furthermore, synchronous design guidelines help you meet timing objectives in high-performance designs such as those targeting the Virtex-II FPGAs, which can accommodate as many as 8 million gates and operate at frequencies as high as 400 MHz.

Clock distribution also requires specific design code guidelines for successful implementation. For this reason, it is important to identify the clocks and resets in the design, as well as the clock-tree network, early in the code development process.

Many designs, especially in the networking domain, typically use multiple clock domains, a situation that presents challenging integration and verification issues. Signals that originate in one clock domain may be used in other clock domains. Correct chip operation can only be ensured if rules governing correct use of signals across asynchronous clock boundaries are followed.

The routability of design is another aspect that can benefit greatly by following appropriate coding guidelines. The number of I/Os, fanout of intermediate nodes, and widths of internal buses must meet specified limits for various parts of a given design.

Using the SpyGlass predictive analysis system, you can check for all of the above violations, problems, and issues during the RTL coding process itself. This saves many design iterations that might otherwise be needed to fix errors later in the design cycle.

Specifically, you can use the SpyGlass tool to check the following issues relevant to Xilinx FPGA designs:

- Single clock edge is used in the design to clock the data.
- Internally derived or generated clocks and set/resets are not used in the design.
- Appropriate synchronizers are used as signals cross clock-domain boundaries.
- The number of levels of logic between registers is less than a specified value.
- The fanout of any design node does not exceed a specified number.
- Unintended latch inferences are avoided in the code.
- If-then-else or case statements are not nested deeper than three levels.
- Naming conventions are followed while coding pipeline logic.
- Assess memory requirements and find out feasibility of conversion into regular flops, distributed memory, and block memory.
- For state machines, separate the next-state decoding and output decoding into two discrete processes or always blocks.
- Outputs from design units are registered.
- Identify dead code and floating nodes in the design.

Several of these checks are critical for meeting timing objectives in high performance designs.

Conclusion

We have described the elements of policy-based design methodology, enabled through SpyGlass predictive analysis, for effective RTL coding leading to efficient implementation of designs on Xilinx-based platforms. This approach performs detailed structural analysis on RTL code to check for coding styles, design reuse, clock/reset requirements, verification, timing guidelines, and more. SpyGlass software includes the most comprehensive coverage of Xilinx-recommended coding guidelines that are essential for reuse and efficient design implementation.

To learn more about SpyGlass predictive analysis, go to *www.atrenta.com.*

Enter the next dimension

As a professional engineer you know that electronics design involves many dimensions. Capturing design data, analyzing and verifying the circuit's performance, synthesizing VHDL for FPGA implementation - these are not fragmented processes, but multiple dimensions of a single design flow. That's why Altium has pioneered nVisage DXP - the first design capture system to go beyond current one-dimensional schematic capture systems.

With nVisage DXP you can capture the multiple dimensions of your design within a single application. nVisage includes multiple design entry methods that support both schematic and VHDL based entry. nVisage also gives you multiple analysis and verification tools, such as integrated XSpice/SPICE 3f5 circuit simulation, pre- and post-layout signal integrity analysis, and VHDL simulation. For circuit implementation, nVisage's extensive output options for board layout are complemented by powerful RTL-level VHDL synthesis for multiple FPGA target architectures, allowing you to design for both programmable devices and PCBs within the one application.

Why not envisage a new dimension in design and take a look at the nVisage multimedia demonstration at **www.nVisage.com**. You can also download a free 30-day trial version, from the website or order the trial CD by calling: **800-470-2686**.

nVisage DXP™

Multi-dimensional design capture for every engineer

Switch to nVisage from $1,495*

1039-ADXA-1

Altium™
Making Electronics Design Easier™

*Retail value $2,995. Receive up to 50% discount if you switch to nVisage DXP from another qualifying Schematic Capture package. nVisage, nVisage DXP, Design Explorer and Altium and their respective logos are trademarks or registered trademarks of Altium Limited or its subsidiaries.

DESIGN EXPLORER
DXP
TECHNOLOGY BY ALTIUM

Get RealFast RTOS with Xilinx FPGAs

Real-time operating systems implemented in Xilinx FPGAs enhance performance, improve predictability, simplify design, and lower system costs.

by Tommy Klevin
Product Manager
RealFast
tommy.klevin@realfast.se

Designers once used field programmable gate arrays (FPGAs) as the glue logic to interconnect discrete components on a printed circuit board (PCB). As we start a new era, though, we can now build complete systems in one FPGA chip.

To take advantage of the possibilities offered by Xilinx Platform FPGAs, we must consider the kind of operating system best suited for these one-chip systems, as well as other systems that contain more than one central processing unit (CPU) or one digital signal processor (DSP).

The RealFast company in Sweden has many years of experience in FPGA design, in operating systems, and in real-time systems development. We recently developed the Sierra 16 real-time operating system (RTOS), implemented in an FPGA, to enhance performance and predictability and to reduce system complexity. This article describes the Sierra 16 RTOS, explains what it can do, and explores the possibilities of putting operating system functions into hardware.

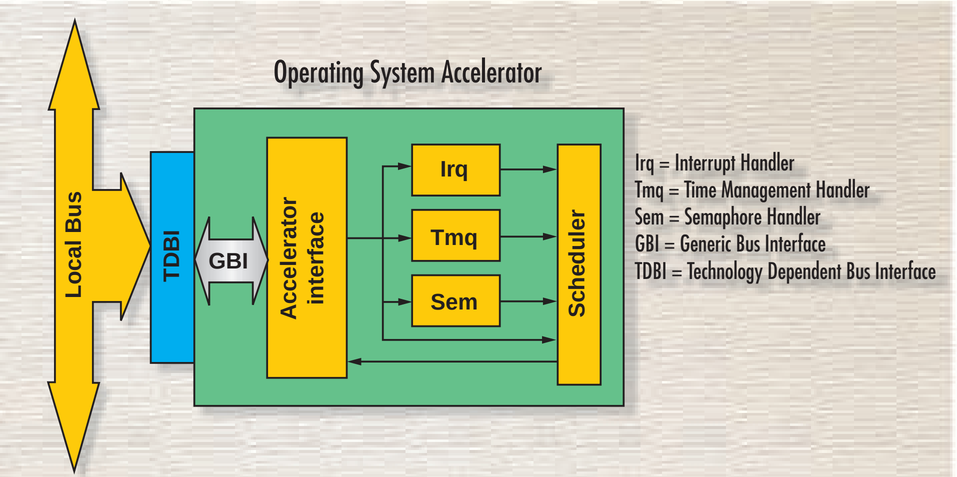
Operating System Accelerator
Local Bus
TDBI
GBI
Accelerator interface
Irq
Tmq
Sem
Scheduler
Irq = Interrupt Handler
Tmq = Time Management Handler
Sem = Semaphore Handler
GBI = Generic Bus Interface
TDBI = Technology Dependent Bus Interface

Figure 1 - Block schematics of the Sierra 16 HW-RTOS

Sierra 16 RTOS

Why implement an operating system in hardware? Is it not better to have the operating system in software as we are used to? Well, who would have imagined in the mid-1980s that mathematical operations would be performed by hardware in the CPU instead of software? A closer look at operating systems shows that many of the low-level primitives are similarly independent of the operating system. These primitives can be performed by either hardware or software.

Some example low-level primitives are:

- Task handling – create and schedule tasks
- Synchronization – semaphores, message passing
- Time handling – delay, timeout handling
- Interrupts.

What is the point of doing these operations in a hardware FPGA system instead of in software? Even in the new platform FPGAs, memory is limited if you do not add external memory. In small, cheap products, you probably want to squeeze your application into the available memory in the FPGA. If you use a traditional operating system, the operating system eats substantial memory, leaving less space for the application.

On the other hand, platform FPGAs contain logic that is just waiting to be used. By moving the operating system to logic, you save a lot of space in memory, and at the same time, you increase performance and achieve a fully predictable system.

The Sierra 16 RTOS has the same functionalities found in other traditional software real-time operating systems. As shown in Figure 1, the Sierra 16 RTOS has the following configuration:

- Handles as many as 16 tasks at eight priority levels
- Supports 16 semaphores
- Handles timing – delay, timers
- Supports eight external interrupts.

In the coming months, RealFast will complement the Sierra series with other configurations to suit both small and large systems.

Interrupt Handling

Interrupt handling is critical in most systems and is particularly critical in real-time systems. Interrupts introduce unpredictable behavior because of the difficulty in predicting when they will arrive. The Sierra 16 RTOS has an intelligent interrupt handler that makes it possible to achieve fully predictable system behavior, even though it is not possible to know exactly when the interrupts will arrive.

We treat the interrupt service routines (ISR) like any ordinary task with a certain priority in the system. As shown in Figure 2, when not taking care of any service, ISRs are in a wait state. When the interrupt arrives, the ISR is transferred to the ready state in the Sierra 16 RTOS scheduler. ISR tasks execute by priority. To switch between different tasks, the Sierra 16 RTOS interrupts the CPU with a task-switch interrupt. This is the only interrupt that interrupts the CPU. All mechanisms are handled by the Sierra 16 RTOS – from the physical pin, to scheduling, to interrupting the CPU so it can switch from present task code to the ISR code and start executing it.

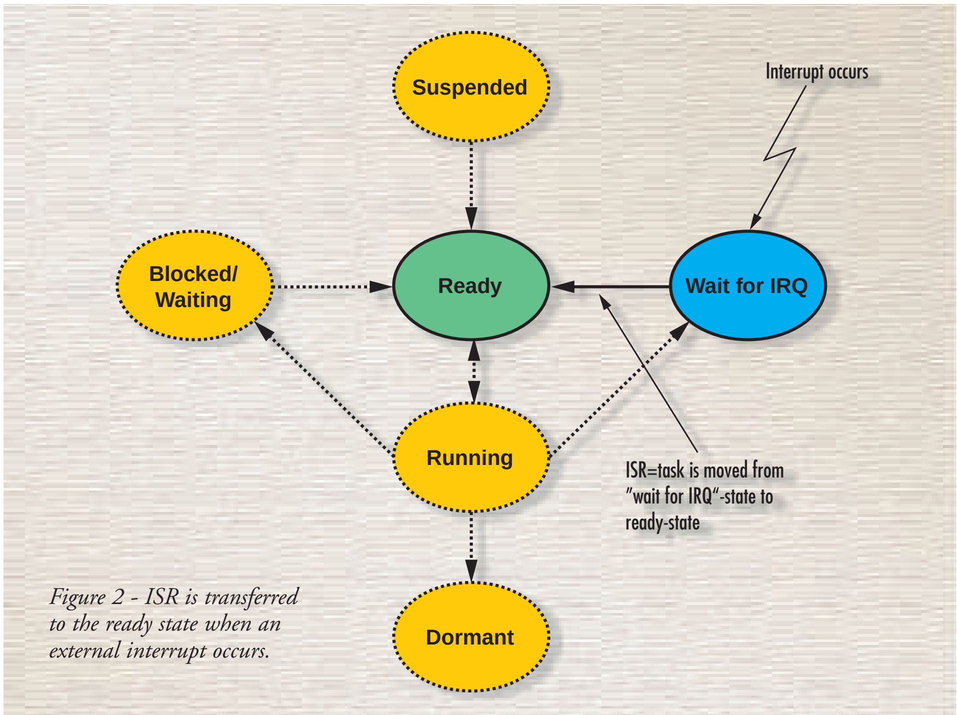

Figure 2 - ISR is transferred to the ready state when an external interrupt occurs.

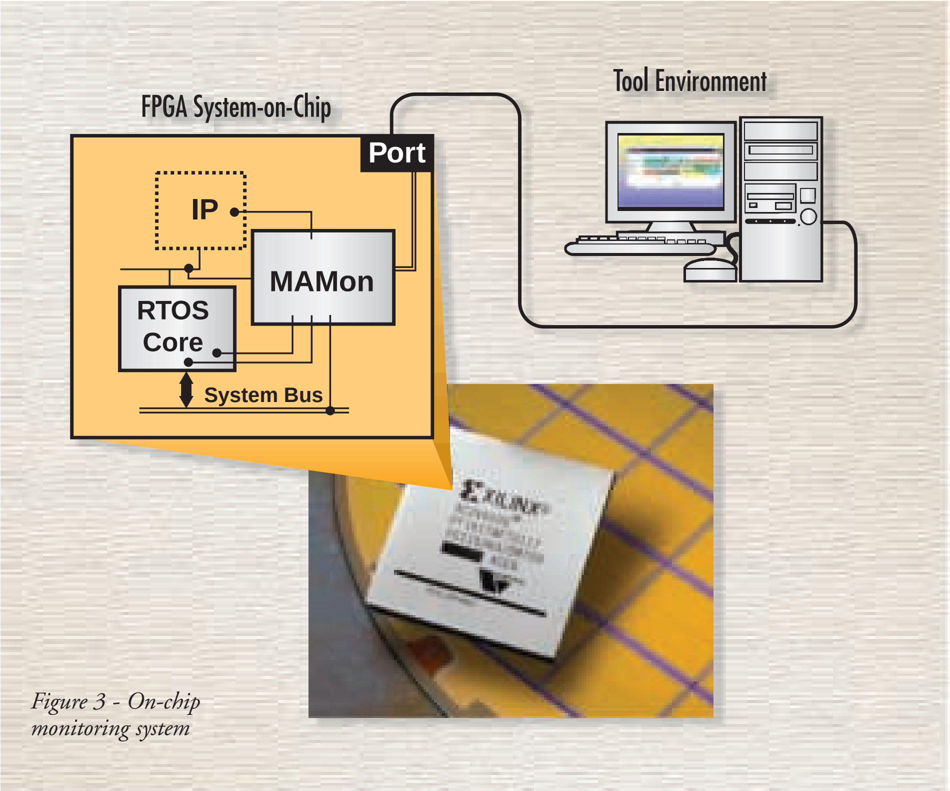

Figure 3 - On-chip monitoring system

Nonintrusive Monitoring and Debugging

Run-time observability in embedded system architectures is a requirement for testing, debugging, and validating design assumptions made about the behavior of the system and its environment. The classic approach to run-time observability is to apply monitoring – the process of detecting, collecting, and interpreting run-time information regarding the system's execution behavior.

When monitoring real-time systems, an important aspect is to minimize – or better yet, completely avoid – the intrusiveness of the monitor on the system's timing and execution properties. Failure to handle monitor intrusiveness may lead to probe effects that cause nondeterministic behavior in programs with race conditions and poor synchronization.

When we use a hardware RTOS (HW-RTOS) like the Sierra 16 RTOS, we avoid the intrusiveness problem. Because the Sierra 16 RTOS comprises a number of hardware components that know at each moment exactly what is going on in the system, this information can be extracted without any software probes normally used for extracting information. One intellectual property (IP) component available for the Sierra 16 RTOS is the Multiprocess Application Monitor (MAMon). As shown in Figure 3, this monitor is implemented in hardware and listens nonintrusively to the different parts inside the Sierra 16 RTOS.

Multiprocessor Solutions Made Easy

Loosely coupled systems and tightly coupled systems contain more than one CPU. In a loosely coupled system, all CPUs have their own memory for such things as program codes and data, and the CPUs communicate through a shared memory. In a tightly coupled system, often called Symmetrical Multiprocessor System (SMP), all CPUs share the same memory and execute code and data from this memory.

The problem with these kinds of systems, in particular SMP systems, is the difficulty in making efficient operating systems that use the efficiency of multiple CPUs. It takes complex algorithms to handle and synchronize multiple CPUs. Pure software solutions for operating systems with multiprocessor support become inefficient and do not give the desired performance boost compared to single CPU systems.

An RTOS implemented in hardware can, on the other hand, perform many parallel operations. This capability provides a completely new approach to solving the problem with multiprocessor systems. All algorithms that are complex and time-consuming in software are moved to hardware. All synchronization of tasks, such as semaphores, is handled by the HW-RTOS.

The Sierra 16 RTOS supports single CPU systems, but RealFast will soon release another HW-RTOS series that provides support for systems with two or more CPUs or DSPs.

MicroBlaze CPU + Sierra 16 HW-RTOS

A powerful but very inexpensive RTOS solution is to combine a Xilinx MicroBlaze™ software CPU with a RealFast Sierra 16 HW-RTOS. With a small 150K-gate Spartan™-II FPGA and (maybe) some external memory, you can create a complete and advanced real-time system at a very low cost.

In bigger Spartan-II and Virtex™-II FPGAs, it is even possible to exclude external memory, as the internal block RAMs will be big enough for both driver and application in many cases.

The software driver for the Sierra 16 has a minimal footprint, only a couple of kilobytes. This driver, together with the HW kernel, makes a very fast and predictable RTOS kernel. With a system running at 50 MHz, most of the system-calls are finished within 2 microseconds.

Conclusion

The new Xilinx Platform FPGAs offer new design approaches and new possibilities. Because logic has grown, we need to start solving problems in a parallel manner, rather than in the sequential manner that we are used to.

An operating system implemented in hardware is a step towards parallel processing, and we believe that these kinds of cost-effective solutions will meet the increasing demands of applications in the future. To learn more about the RealFast Sierra 16 HW-RTOS, visit *www.realfast.se/*.

Virtex-II Pro Platform FPGAs Deliver Proven Interoperability

With verified interoperability between specialty ASSPs and Xilinx Virtex-II series FPGAs, you can focus on designing and debugging your system rather than how to make all the parts work together.

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

Virtex™-II series Platform FPGAs include a rich set of system features such as embedded memories, DSP, and I/O connectivity with which you will design most blocks of your system. But your design might need to incorporate specialty ASSPs. These could be large memories, optical transceivers, analog filters, A/D and D/A converters, specialized DSP/network processors, connectors, and many others. Interoperability of all these devices on the printed circuit board (PCB) is a tough challenge.

Xilinx Virtex-II Pro™ and Virtex-II FPGAs solve these challenges by enabling you to connect to almost any ASSP without a hitch. This article will show you how.

Interoperability Can Make or Break Your Design

Making multiple disparate devices compatible – often from different vendors – can be an arduous task involving numerous design and debug iterations. Even if you do succeed in making all the devices work together, the inefficiency of such a design is likely to impose a performance penalty that prevents the overall design from achieving its maximum performance objectives. Unknown interoperability of devices increases the risk of cost overruns and slipped schedules, and can even jeopardize design completion. Starting with a Platform FPGA that delivers interoperability of the devices you will use in your design is critical to predicting project time, cost, and feasibility.

Xilinx Leads in Interoperability

Xilinx offers complete interoperability solutions with leading ASSPs across a broad range of interface standards. This ensures that the integration of third-party ASSPs will go smoothly, that the cost of design implementation will be predictable, and that your projects will succeed the first time through.

By choosing Virtex-II Platform FPGAs over alternatives that lack powerful capabilities for interoperability, you will lower the risk of slipped deadlines, significantly lower the cost of the project, accelerate design and development cycles, and simplify the expensive system debug process.

With its extensive network of partnerships with ASSP vendors through the Reference Design Alliance Program – the only such program in the programmable logic industry – Xilinx gives you the edge. Xilinx and ASSP vendors work closely to deliver fully interoperable solutions. This allows you to simplify system development and eliminates the tough challenges of making all the different devices on the board work together seamlessly. You can focus your creative energy on product differentiation and maximizing system performance.

Virtex-II Pro Delivers Proven Interoperability

With proven RocketIO™ serial and SelectIO™-Ultra parallel I/Os in Virtex-II Pro Platform FPGAs and IP cores for protocols, the Xilinx SystemIO solution provides ultimate connectivity with ASSPs. The Virtex-II Pro FPGA family was built on the Virtex-II family in a completely scalable framework – specifically to address interoperability and prevent problems common with ground-up competing architectures. Virtex-II Pro FPGAs inherit the entire Virtex-II SelectIO-Ultra technology, which directly addresses connectivity and interoperability with ASSPs using established parallel system interfaces. Virtex-II FPGAs have been widely used as an interface to bridge disparate ASSPs. Given that ASSP vendors have been working with SelectIO-Ultra technology for more than two years, Virtex-II Pro FPGAs give you unparalleled edge in interoperability over any competing FPGA in the world.

Given that ASSP vendors have been working with SelectIO-Ultra technology for more than two years, Virtex-II Pro FPGAs give you unparalleled edge in interoperability over any competing FPGA in the world.

Virtex-II Pro devices support dozens of emerging, established, and even proprietary connectivity standards. The embedded RocketIO serial transceivers enable the widest range of programmable serial bandwidth: 622 Mbps to 75 Gbps covering emerging serial standards such as Gigabit Ethernet, 10 Gigabit Ethernet XAUI, PCI Express™, SxI-5, TFI-5, Serial RapidIO™, InfiniBand™, and Fibre Channel.

SelectIO-Ultra parallel I/O system interfaces support such standards as SPI3 (POS PHY 3), SPI-4.1 (Flexbus 4), SPI4.2 (POS PHY 4), XGMII, RapidIO, PCI, PCI-X, CSIX, HyperTransport™, XSBI, and SFI-4 using 22 single-ended and six differential electric standards including LVTTL, LVCMOS, PCI, PCI-X, HSTL, SSTL, LVDS, LVPECL, and HyperTransport.

Checking silicon features and electrical I/O standard support is only the first step towards interoperability with ASSPs.

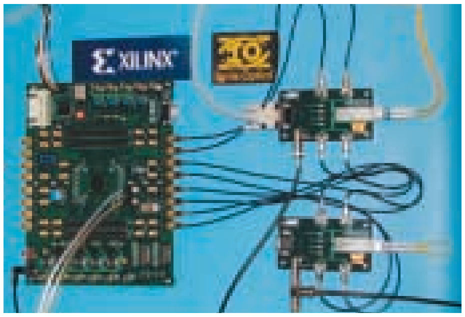

Figure 1 - Four RocketIO channels at 2.488 Gbps with Ignis Optics' four IGP-2000 OC-48 SFP optical transceivers

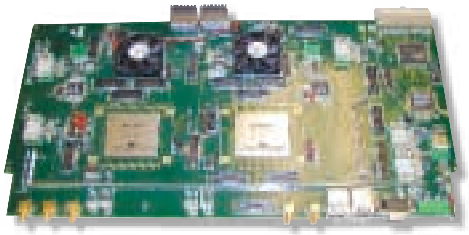

Figure 2 - SPI4.2 static and dynamic alignment mode HW verification completed with PMC-Sierra S/UNI-9953

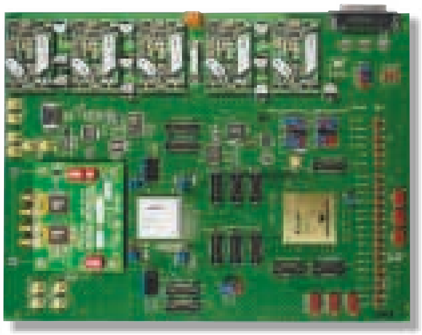

Figure 3 – SPI4.1 HW interoperability with AMCC Ganges-II device

Xilinx goes far beyond this by providing pre-verified interface and controller IP cores – jointly verified in hardware with ASSP vendors during several months of engineering testing to guarantee full interoperability. Reference designs and boards are also available to demonstrate true hardware interoperability, as shown in Figures 1, 2, and 3.

The list of ASSP devices interoperable with Virtex-II series FPGAs is shown in

Table 1. The latest list of interoperable devices, along with reference designs, is always available at *www.xilinx.com/company/reference_design/interop_solutions.htm*. These ASSPs include devices from Intel, AMCC, PMC-Sierra, Mindspeed Technologies, IDT, and other vendors.

In addition, Xilinx actively participates in interoperability events and testing activities. Recently, Xilinx submitted Virtex-II Pro FPGAs for interoperability testing to the 10 Gigabit Ethernet Consortium at the University of New Hampshire's InterOperability Lab, which was attended by 13 participating vendors. (*ftp://public.iol.unh.edu/pub/10gec/Oct02_GTP_release.pdf*).

Xilinx Lets You Choose the Best ASSPs

The Virtex-II Pro solution permits you to design-in the best ASSPs that fit your needs. ASSPs come with varying and often incompatible interface standards. Lack of extensive interoperability capabilities can limit design alternatives. For example, choosing a network processor ASSP supporting the HyperTransport standard could prevent you from choosing a security co-processor ASSP supporting only the RapidIO standard. Even if these two ASSPs matched your requirements perfectly, you may still be forced to choose less suitable ASSPs just because they work together.

With the Virtex-II Pro FPGAs' support for multiple system connectivity standards and proven ASSP interoperability, you can effortlessly bridge standards and ASSPs. Select the right ASSPs that best fit your product needs based on capability and price, not on how well they work together. The Xilinx leadership in interoperability gives you the clear advantage in differentiating your products from the competition.

Conclusion

Interoperability of disparate devices on the board is a critical parameter for design success. Your design productivity, performance, and even the probability of completion can be significantly improved by thoroughly verified interoperability among semiconductor devices and vendors. Xilinx and the leading ASSP vendors want you to be confident that you are in safe hands when working with their devices.

Start your next design on a robust platform that has a proven record of several years of diverse interoperability to its credit. Focus your creative energy on product function, performance, and differentiation, rather than worrying about whether the components will work together.

ASSP Vendor	ASSP Device	ASSP Device Type	Interface Standard	Reference Design
AMCC	Several devices	Several	ViX-v3 to SPI4	Yes
AMCC	Ganges-II ATM, SONET/SDH	OC-192, 4xOC-48 POS, Framer/Mapper	SPI4.1	Coming
Bay Microsystems	Montego	Network Processor	Host/Accountant	Yes
Broadcom Octal SerDes	BCM8040	Multi-rate	XAUI	Coming
Ignis Optics	2.5 Gbps SFP	Optical transceiver	2.488 Gbps per Channel	Yes
IDT	TeraSync FIFO	FIFOs	HSTL	Coming
Intel	IXF1810x Family 10Gb Ethernet	OC-192c POS/GFP & MAC/Framer	SPI4.2	Coming
Mindspeed Technologies	OptiPHY M29730	OC-192/STM-64, Quad OC-48/STM-16 POS Framer	SPI4.2	Yes
Netlogic Microsystem	NSE4256, IXP1200, SA110	CAM	NSE bus/nPu bus	Yes
PMC-Sierra (PM3386)	S/UNI-2xGE Ethernet Controller	Dual Gigabit	SPI3	Coming
PMC-Sierra	Xenon Family	10 Gb Ethernet	SPI4.2	Coming
PMC-Sierra	S/UNI 9953 POS, ATM, Ethernet	10 Gb PHY for alignment	SPI4.2 dynamic	Coming
TeraCross	TXQ1450, TXS1400	Scheduler, traffic generation & monitoring	Line Card Interface Control Link	Yes
Velio	VC1003	SONET backplane, gbE SerDes	LVDS	Yes
Velio	VC1021/1022	OC-48, OC-192 SerDes	LVDS	Yes
Velio	VC1061/1062	Storage Quad SerDes	10 GigE MAC	No
ZettaCom	ZTM202	Traffic Management	CAM, SRAM, CSIX	Yes

Table 1 - Virtex-II series interoperable solutions

Introducing The Platform for Programmable Systems.

VIRTEX-II PRO

Available now, the new Virtex®-II Pro Platform FPGAs herald an astonishing breakthrough in system-level solutions. With up to four IBM PowerPC™ 405 processors immersed into the industry's leading FPGA fabric, Xilinx/Conexant's flawless high-speed serial I/O technology, and Wind River System's cutting-edge embedded design tools, Xilinx delivers a complete development platform of infinite possibilities. The era of the programmable system is here.

The Power of Xtreme Processing

Each IBM PowerPC runs at 300+ MHz and 420 Dhrystone MIPS, and is supported by IBM CoreConnect™ bus technology. In addition, the FPGA fabric enables ultra-fast hardware processing, such as TeraMACs/s DSP applications. With Xilinx's unique IP-Immersion™ architecture, system architects can now harness the power of high-performance processors, along with easy integration of soft IP into the industry's highest performance programmable logic. Designers have absolute freedom to implement any design they can imagine. Never before has such performance been achieved enabling hardware accelerated processing and multiple processing in an off-the-shelf device.

IBM technology PowerPC processor™

Enabling a new development paradigm

For the first time ever, system designers can partition and re-partition their system between hardware and software at any time during the development cycle—even after the product has shipped. That means you can optimize the overall system, guaranteeing your performance target in the most cost-efficient manner. You can also debug hardware and software simultaneously at speed.

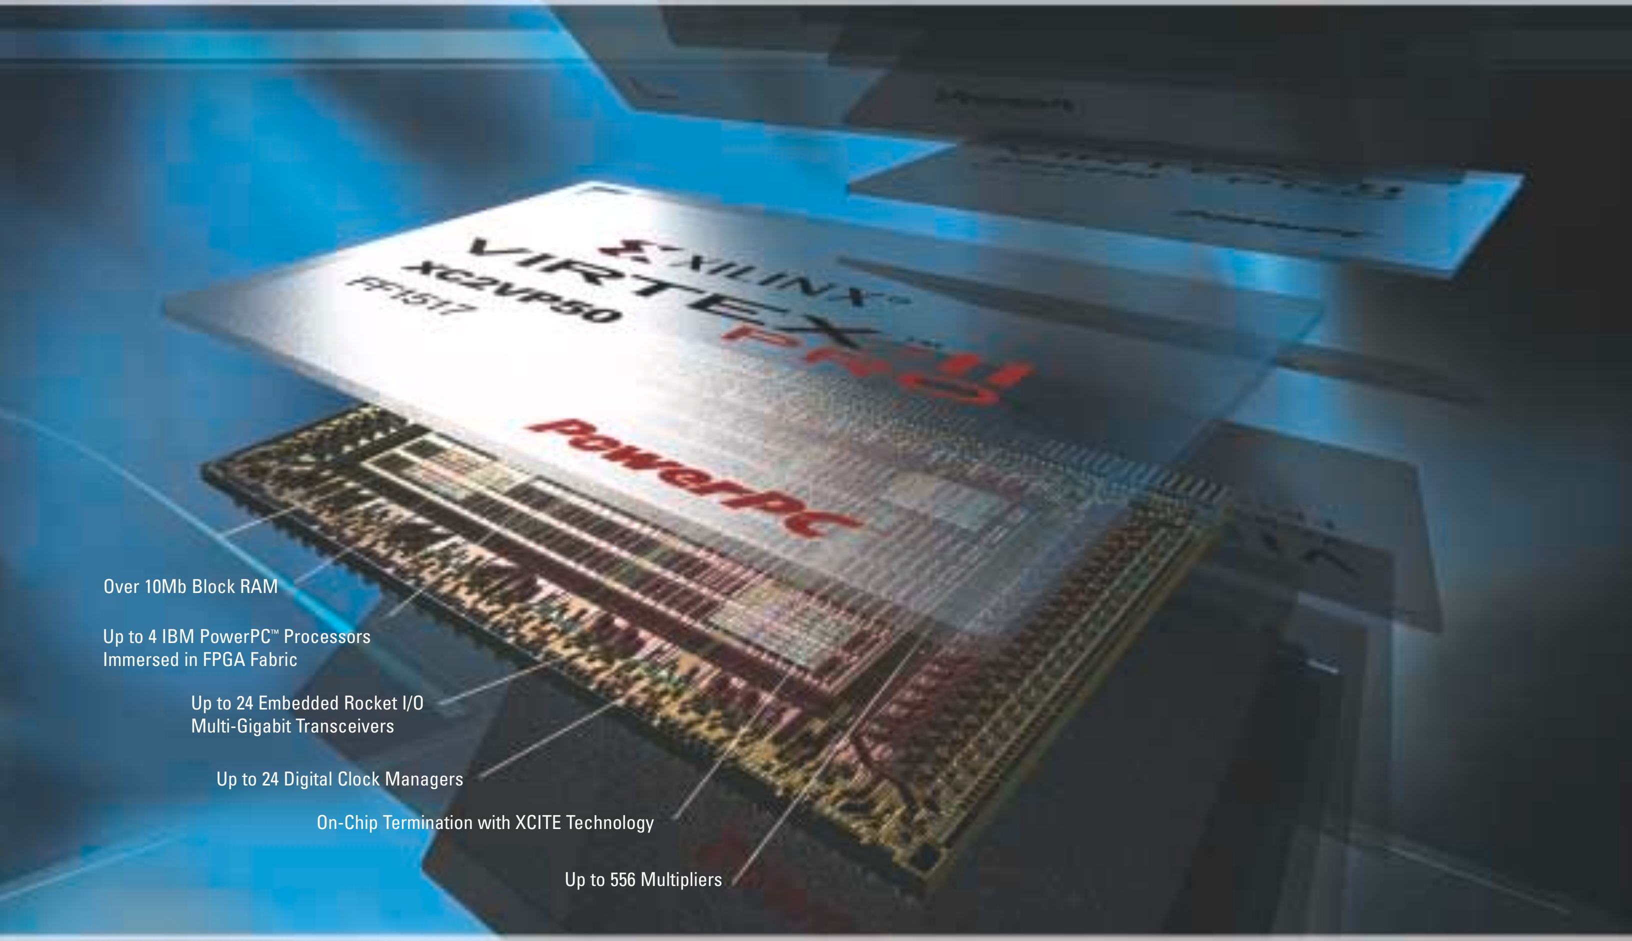
XILINX
VIRTEX-II PRO
XC2VP50
FF1517
PowerPC
Over 10Mb Block RAM
Up to 4 IBM PowerPC™ Processors Immersed in FPGA Fabric
Up to 24 Embedded Rocket I/O Multi-Gigabit Transceivers
Up to 24 Digital Clock Managers
On-Chip Termination with XCITE Technology
Up to 556 Multipliers

The ultimate Connectivity platform

The first programmable device to combine embedded processors along with 3.125 Gbps transceivers, the Virtex-II Pro series addresses all existing connectivity requirements as well as emerging high-speed interface standards. Xilinx Rocket I/O™ transceivers offer a complete serial interface solution, supporting 10 Gigabit Ethernet with XAUI, 3GIO, SerialATA, InfiniBand, you name it. And our SelectI/O™-Ultra supports 840 Mbps LVDS and other parallel methodologies. Think of it: up to 16 Rocket I/O 3.125 Gbps transceivers at your disposal, delivering the ultra-high bandwidth you need (40+ Gigabit per second total serial bandwidth) for real market challenges like optical networking, high-end broadcast, storage and DSP systems and so much more.

The Power of Integration

In a single off-the-shelf programmable device, system architects can take advantage of microprocessors, multi-gigabit transceivers, digital clock managers, highest density on-chip memory, on-chip termination and more. The result is a dramatic simplification of board layout, a reduced bill of materials, and unbeatable time to market.

Industry-Leading Tools from Wind River and Xilinx

Optimized for the PowerPC, Wind River's industry-proven embedded tools are the premier support for real-time microprocessor and logic designs. And driving the Virtex-II Pro FPGA is Xilinx's lightning-fast ISE 4.2i software, the most comprehensive, easy-to-use development system available.

See the new Virtex-II Pro Platform FPGA in detail. Visit ***www.xilinx.com/virtex2pro*** today and step into the era of the programmable system.

XILINX®

The Programmable Logic Company℠

www.xilinx.com/virtex2pro

FORTUNE® 2002
100 BEST COMPANIES TO WORK FOR

Spartan-IIE Family Grows

Building on a tradition of cost-effective speed and reliability, two new Xilinx Spartan-IIE devices offer enhanced flexibility and higher densities – at the lowest cost per I/O.

by Rufino Olay
Solutions Marketing Manager
Xilinx, Inc.
rufino.olay@xilinx.com

With today's challenging economic times, making cost-sensitive products such as plasma displays, set-top boxes, and broadcast video equipment requires a low-cost solution. Additionally, the integration of more features in digital consumer products often demands more pins than previously available. Thus, we were challenged to create a low-cost solution that addressed the need for a high pin count device.

Two new additions to the successful Spartan™-IIE family of devices meet these exacting criteria. The two devices, XC2S400E and XC2S600E, are low-cost, high-density, high-I/O devices that will allow you to target a wider spectrum of designs than you previously could with programmable logic.

Fourth Generation Spartan FPGAs

Since introducing the Spartan series more than four years ago, Xilinx has delivered four generations and shipped more than 40 million Spartan series FPGAs, with prices starting as low as $2.55 per device. With the Spartan series, you get the advantages of high I/O-count ASICs and gate arrays – and you get the added flexibility of a general purpose, programmable architecture. You also get the benefit of a proven architecture with the industry's fastest and most productive software tools, plus the most comprehensive offering of IP cores from Xilinx and third-party AllianceCORE™ vendors.

On November 18, 2002, Xilinx announced an extension to the Spartan product line to address customer demand for even higher density and higher I/O-count devices in the price ranges required for consumer applications. As shown in Table 1, the XC2S400E device (400,000 system gates and up to 410 I/Os) and the XC2S600E device (600,000 system gates and up to 514 I/Os) give you the ability to integrate more functionality into a smaller form factor and still meet your stringent budget requirements. With the Spartan-IIE family, you get:

- The lowest cost per I/O – The new Spartan-IIE devices give you more I/Os at much lower prices than any competing FPGA.
- Up to 514 I/Os – With the highest number of I/Os available in the low-cost segment of the FPGA industry, Spartan-IIE devices allow you to put higher density ASIC designs into FPGAs and still keep the benefits of reprogrammability.
- Four DLLs – The DLLs allow easy clock duplication, quick frequency adjustment, faster state machines using different clock phases, de-skewing of the incoming clock, and generation of fast setup and hold times or fast clock to outs.
- More than one billion MACs/sec per dollar – You can implement high performance DSP functionality at the lowest cost possible.

Device	XC2S50E	XC2S100E	XC2S150E	XC2S200E	XC2S300E	XC2S400E	XC2S600E
System Gates	50K	100K	150K	200K	300K	400K	600K
Logic Cells	1,728	2,700	3,888	5,292	6,912	10,800	15,552
Block RAM Bits	32K	40K	48K	56K	64K	160K	288K
Distributed RAM Bits	24K	37K	54K	73K	96K	150K	216K
DLLs	4	4	4	4	4	4	4
I/O Standards	19	19	19	19	19	19	19
Max Differential I/O Pairs	83	86	114	120	120	172	205
Max Single Ended I/O	182	202	265	289	329	410	514
Packages	TQ144	TQ144					
	PQ208	PQ208	PQ208	PQ208			
	FT256	FT256	FT256	FT256	FT256		
		FG456	FG456	FG456	FG456	FG456	FG456
						FG676	FG676

Table 1 - Spartan-IIE product matrix

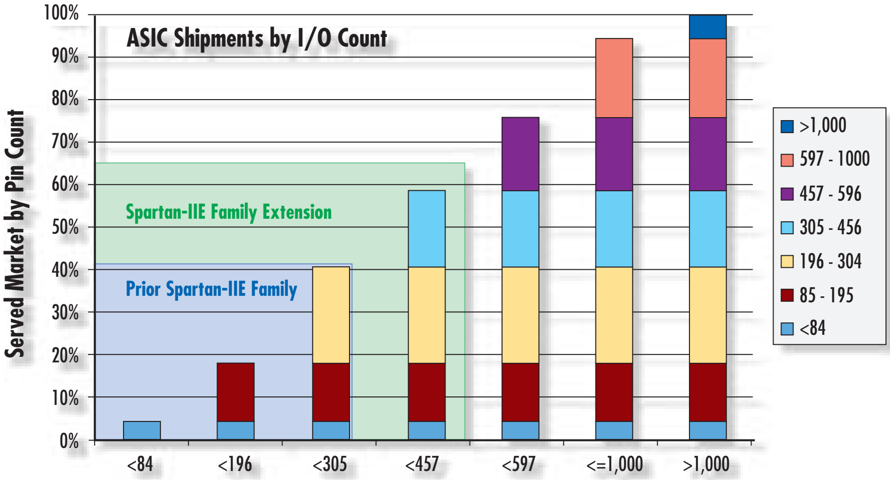

Source: Xilinx

Figure 1 - Spartan-IIE extensions meet customer demand for more I/O.

More I/Os for Digital Consumer Applications

Moving to advanced technologies has always enabled Xilinx to dramatically reduce costs and simultaneously bring larger density devices within the reach of many more cost-conscious customers. Now, with the two new Spartan-IIE FPGAs, this same advantage is being brought into the I/O arena.

Traditionally, consumer applications with more than 305 I/Os have required ASICs, as shown in Figure 1. But with the introduction of the XC2S400E and XC2S600E devices, you now have up to 410 and 514 I/Os, respectively, and with the added advantage of reprogrammability.

These two new FPGAs deliver a greater than 67% increase in I/O capacity over previous Spartan-IIE offerings, and up to 100% more I/Os than competing FPGAs in the same density ranges. Additionally, the I/Os can be configured as differential I/O pairs (up to 205), giving you LVDS performance up to 400 Mbps.

Supporting More I/O Standards

Today's designs are more complicated than ever, with the majority typically containing numerous I/O standards on a single PCB. With the Spartan-IIE FPGAs you can connect as many as 19 different I/O standards on a single chip. This flexibility gives you

	XC2S50E	XC2S100E	XC2S150E	XC2S200E	XC2S300E	XC2S400E	XC2S600E
TQ144	102	102					
PQ208	146	146	146	146	146		
FT256	182	182	182	182	182	182	
FG456		202	265	289	329	329	329
FG676						410	514

TQ144	2 devices for migration
PQ208	5 devices for migration
FT256	6 devices for migration
FG456	6 devices for migration
FG676	2 devices for migration

Table 2 - Density migration possibilities

the ability to bridge different I/O standards and protocols – and completely eliminates the need for costly bus transceivers.

Scaleable Footprints

Following the tradition of Xilinx FPGA families, the Spartan-IIE family continues to support density migration across common packages without changing the PC board footprint. The relative positions of VCC and GND remain constant across like packages, unlike competing low-cost FPGAs.

For example, when using a FG456 package, six density members of the Spartan-IIE family can be interchanged, providing outstanding flexibility for design revision, upgrade, or cost optimization, as shown in Table 2.

More RAM

The new Spartan-IIE devices differentiate themselves in the amount of available memory both in block RAM and distributed RAM.

The XC2S400 has four columns and the XC2S600E has six columns of block RAM, equating to 160K and 288K of block RAM, respectively. This is a greater than 4X increase in capacity over the previous largest density Spartan-IIE device. With this increase comes the possibility of storing more data, coefficients, FIFO functions, and larger general memory functions in your memory-hungry applications.

Xilinx continues to be the only FPGA supplier to offer distributed RAM, which is an ideal solution for designs that require multiple small, fast, and flexible memories situated close to the logic. As with all other Xilinx FPGA families, the 4-input LUT in Spartan-IIE devices can also be used as memory, where it can be configured as ROM, or single-port or dual-port synchronous RAM.

We've doubled the amount of distributed RAM in these new devices to 150K for the XC2S400E, and 216K for the XC2S600E.

Software and IP

The entire Spartan-IIE family is supported by the Xilinx ISE (Integrated Software Environment) tool set, which includes the industry's most advanced timing-driven implementation tools available for programmable logic design, along with design entry, synthesis, and verification capabilities (*www.xilinx.com/ise5/*).

There are also more than 200 IP cores, including PCI, DSP, and other predesigned and tested solutions (*www.xilinx.com/ipcenter/*) to get your designs up and running fast.

Processing Solutions on a Budget

By utilizing the Xilinx MicroBlaze™ 32-bit field programmable controller option with the Spartan series devices, you can create an easy-to-use, low-cost, customized processing solution. The MicroBlaze processor is the fastest, most powerful soft processor and peripheral solution on the market today for traditional 16-bit and 32-bit microprocessor and microcontroller applications.

Coupling the ISE and MicroBlaze solutions gives you a winning combination with the benefits of:

- Flexibility – Easily create a customized processor design that can be modified at any time during the design cycle.
- Guaranteed product availability – You can purchase the MicroBlaze source code and never have to worry about processor obsolescence.
- Reduce system cost – By integrating your entire processing solution within one device you not only save time and effort but you also reduce your bill of materials, inventory, and debug time.

Conclusion

To be successful in this tough, competitive marketplace, you need an inexpensive and flexible design solution. You also need fast, reliable performance and the lowest cost per I/O. With the expanded Spartan-IIE family there is no faster, safer, or lower cost way to develop next-generation consumer products.

To obtain a free Spartan-IIE Resource CD containing a wealth of information on the XC2S400E and XC2S600E Spartan-IIE devices, visit *www.xilinx.com/spartan2e*.

CoolRunner-II Solutions Save Money

With CoolRunner-II CPLD devices, great things come in small packages.

by Steve Prokosch
CPLD Product Marketing
Xilinx, Inc.
steve.prokosch@xilinx.com

When you're looking for a simple, low cost, and easy-to-use programmable logic device that incorporates multiple functions, think Xilinx CoolRunner™-II CPLDs. These versatile, nonvolatile devices can save you time and money on your next design by reducing board costs and redesigns.

Cost can be thought of in several different ways, depending on your point of view. For a buyer, it's the bottom line of a bill of materials. For a design engineer, it's time invested and looming deadlines.

Engineers also face tradeoffs, such as how fast a product can be designed with a minimal number of board layouts. Your success may rest in the decisions you make while trying to accomplish this goal. When making component choices, it pays to have built-in flexibility; with reprogrammable logic, you get the best dollar value as well as the ability to deliver products ahead of schedule.

Additionally, engineers must consider such factors as single-chip solutions, package size, density, versatility, flexible I/O structures, and the ability to modify pin functionality after placement on the board. By considering these items before parts selection, you can save costs and still maintain flexibility.

Single Chip Integration

If you have unlimited board space, a large stocking warehouse, and inexpensive test and assembly costs, some of these cost factors may not enter into the price equation. But if you're in a competitive marketplace, usually one or more of these items will be scrutinized:

- Power-efficient board size
- Minimum number of parts and suppliers
- Low assembly costs.

Board Size

Typically, the packaging of your product is defined by board size, which is driven by the number of components you need to get the job done. If you can squeeze out the required functionality and still stay within the power budget, you have met your goal.

Don't forget that cost can mean board space to some engineers and inexpensive parts to others. As shown in Figure 1, with CoolRunner-II you can select small BGA packages such as 56- or 132-ball chip scale packages (CSP) for high integration or flat pack (FP) packages for low-cost solutions. If you are concerned with board size, you may need unique CSP options. Xilinx also offers 0.5 mm to 0.8 mm ball spacing packages that can save you more than 50% when compared to similar I/O count FP package options.

Although the space savings from a 14 mm-by-14 mm, 100-pin FP package to an 8 mm-by-8 mm, 132-ball CSP package may seem trivial, consider the routing involved. With flat packs, all pins typically route outward from the package. With BGA packages, routing can be achieved by running traces between the adjacent solder balls. These packages also offer more options when using denser, multilayer PCBs. This may yield twice the routing efficiency of a comparable FP package, further reducing board space. Thus, the capability of these small packages goes well beyond the "wow factor" of their physical size.

Parts and Suppliers

Lower power consumption can also be achieved through reduced component counts. A single low-power CPLD device improves reliability by reducing the total number chance of cold or weak solder joints that may cause intermittent failures. Heat dissipation may be reduced. And more solder joints also increase the chance of manufacturing problems. The more solder joints, the higher the chances of developing manufacturing problems. Heat dissipation may be reduced through fully utilizing a single part instead of powering multiple parts that may not be fully utilized. These two factors can have a direct impact on customer service and customer reliability ratings

Figure 2 illustrates the many functions you can squeeze into a single CPLD and still get the low power operation you desire.

Maintaining multiple components for specific functions can lead to a nightmare for procurement. By expanding your supplier base, you increase demands on many different departments within your company and thus lengthen your time-to-market. These areas may include accounting, shipping and receiving, or component engineering. If you have a quality department, they may want reports on each individual device.

Furthermore, the more devices you specify, the higher the chance of encountering a production problem. It would be devastating to not be able to ship a multi-million-dollar product due to a $2 part on back order. By using more parts, you also run a higher risk of device obsolescence. This may not cause a delay in shipment, but it typically costs a board re-layout.

Assembly Costs

The more components shipped to your contract manufacturer, the more money you spend in shipping costs. Each electronic component placed on your board reflects a direct assembly charge. If your contract manufacturer charges you to stock devices, this will also add cost to your end product.

By keeping the component count down, you can dramatically reduce both direct expenses and the indirect cost of doing business.

Integration and Flexibility

If you need multiple I/O standards for unique memory devices or CMOS level translation, conversion devices may be necessary. Depending on your application, specialty memory devices may also be required.

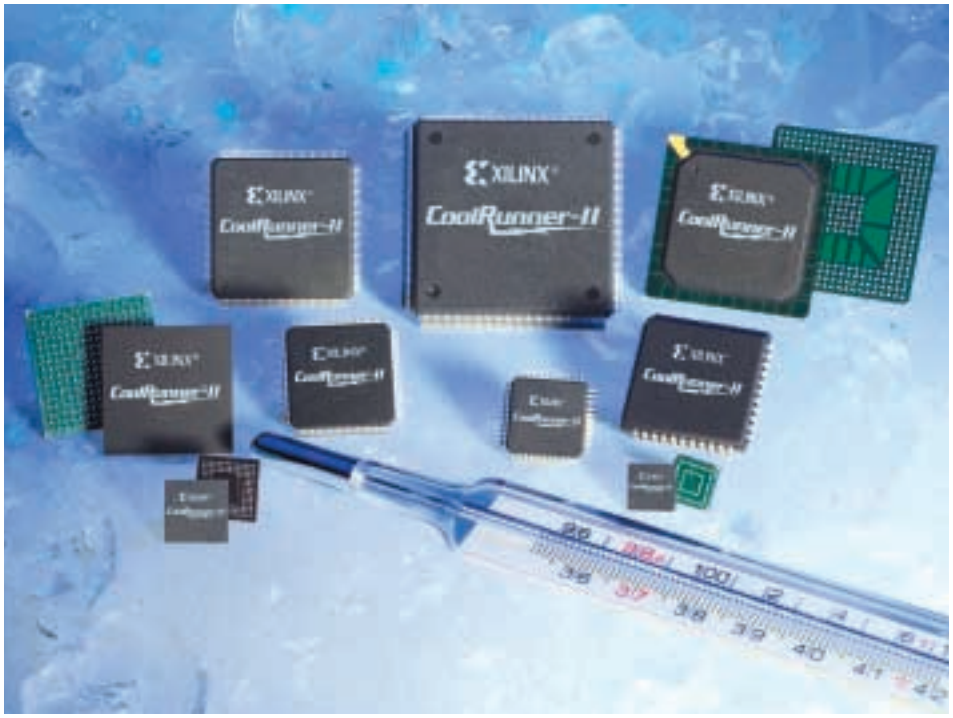

Figure 1 - CoolRunner-II CPLD package offerings

If your processor does not support HSTL or SSTL memory types, you may need to select voltage referenced to CMOS translators. In high-volume applications, these single-function translators can cost from $4 to $6 in 48-pin packages. The problem is, they only serve one purpose. If you don't use all of the pins, it's wasted board space and power. With a single CPLD, you get translation coupled with extra logic capabilities and the freedom to use pins for other purposes besides translation.

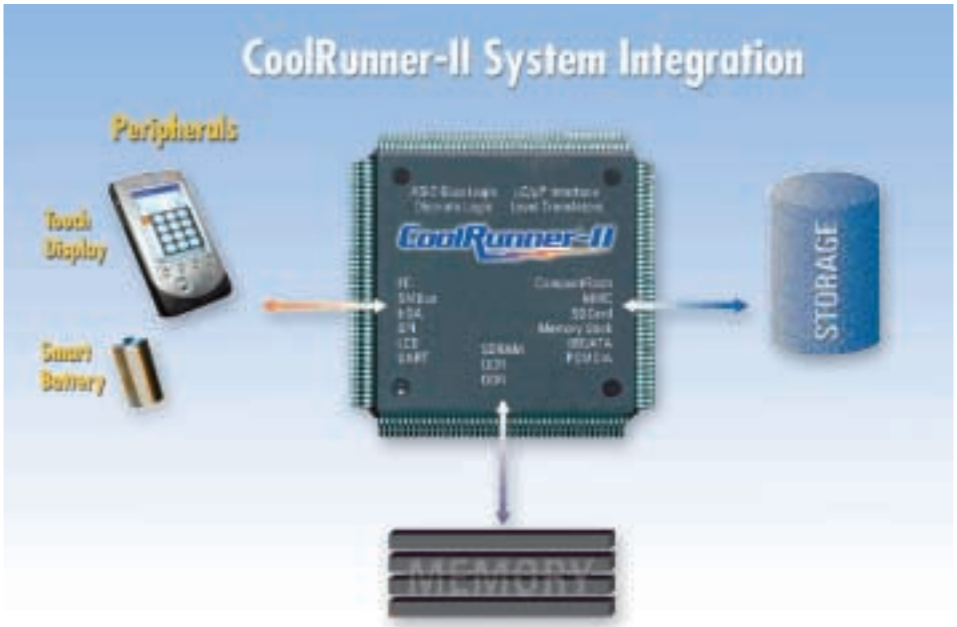

Figure 2 - Multiple functions in a single CPLD

Even if you don't use any specialty memory, what about legacy parts that use different voltage levels than your processor? Again, you have the choice of purchasing a single function device that can cost around $2 for a 48-pin package. A comparable pin count CPLD can cost half as much as this single function device – and again, give you more functionality. So if you need voltage level translation in the range of 1.5V to 3.3V, CoolRunner-II CPLDs can also provide this integrated function.

One specialty function that sometimes is not considered but may prevent board respins is input hysteresis. Schmitt trigger inverters can cost from $4 to $8 in 20-pin packages. These devices usually operate from 1.6V to 3.6V, which gives them a wide operating window. CoolRunner-II CPLDs have input hysteresis on every input pin. And because you can configure CoolRunner-II CPLD input buffers to any voltage from 1.5V to 3.3V, they also have a wide range of operation. In a head-to-head comparison, CoolRunner-II CPLDs can cost 75% less than a discrete Schmitt trigger device.

Also, by using a CPLD solution, you can enable the input hysteresis, if required; if not, just leave it disabled. Because you don't always know if you need hysteresis, this flexibility may save a new board layout.

Moreover, with features such as clock dividers and doublers (DualEDGE flip-flops), you can set up independent clock domains in CoolRunner-II CPLDs, thus eliminating the need for independent oscillators or crystals. The devices can handle fast-running sequential functions such as pulse width modulator, conversion functions (BCD to decimal), and serial communications functions.

Conclusion

Due to their multifunctional nature, Xilinx CPLDs can integrate many applications to save costs in your design. The high-performance, low-power CoolRunner-II CPLDs can reduce the number of board redesigns, minimize the total number of devices, and increase overall flexibility. This will have a direct impact on bringing your product to market faster.

To get you started with CPLDs, Xilinx offers multiple aids, including beginner tutorials with demo boards and reference designs that include detailed application notes with HDL code. Some design examples include SMBus, I^{2}C, SPI, and processor interfaces. You can also look at full-up reference designs, such as designing an MP3 player. Whatever your level of experience, Xilinx makes it easy to use reprogrammable logic.

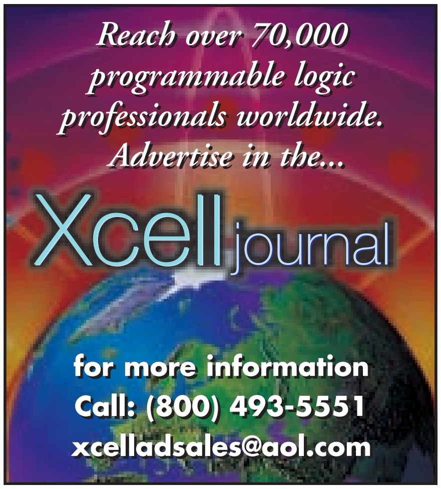

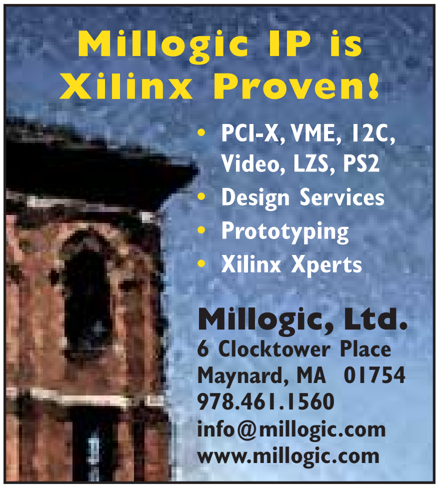

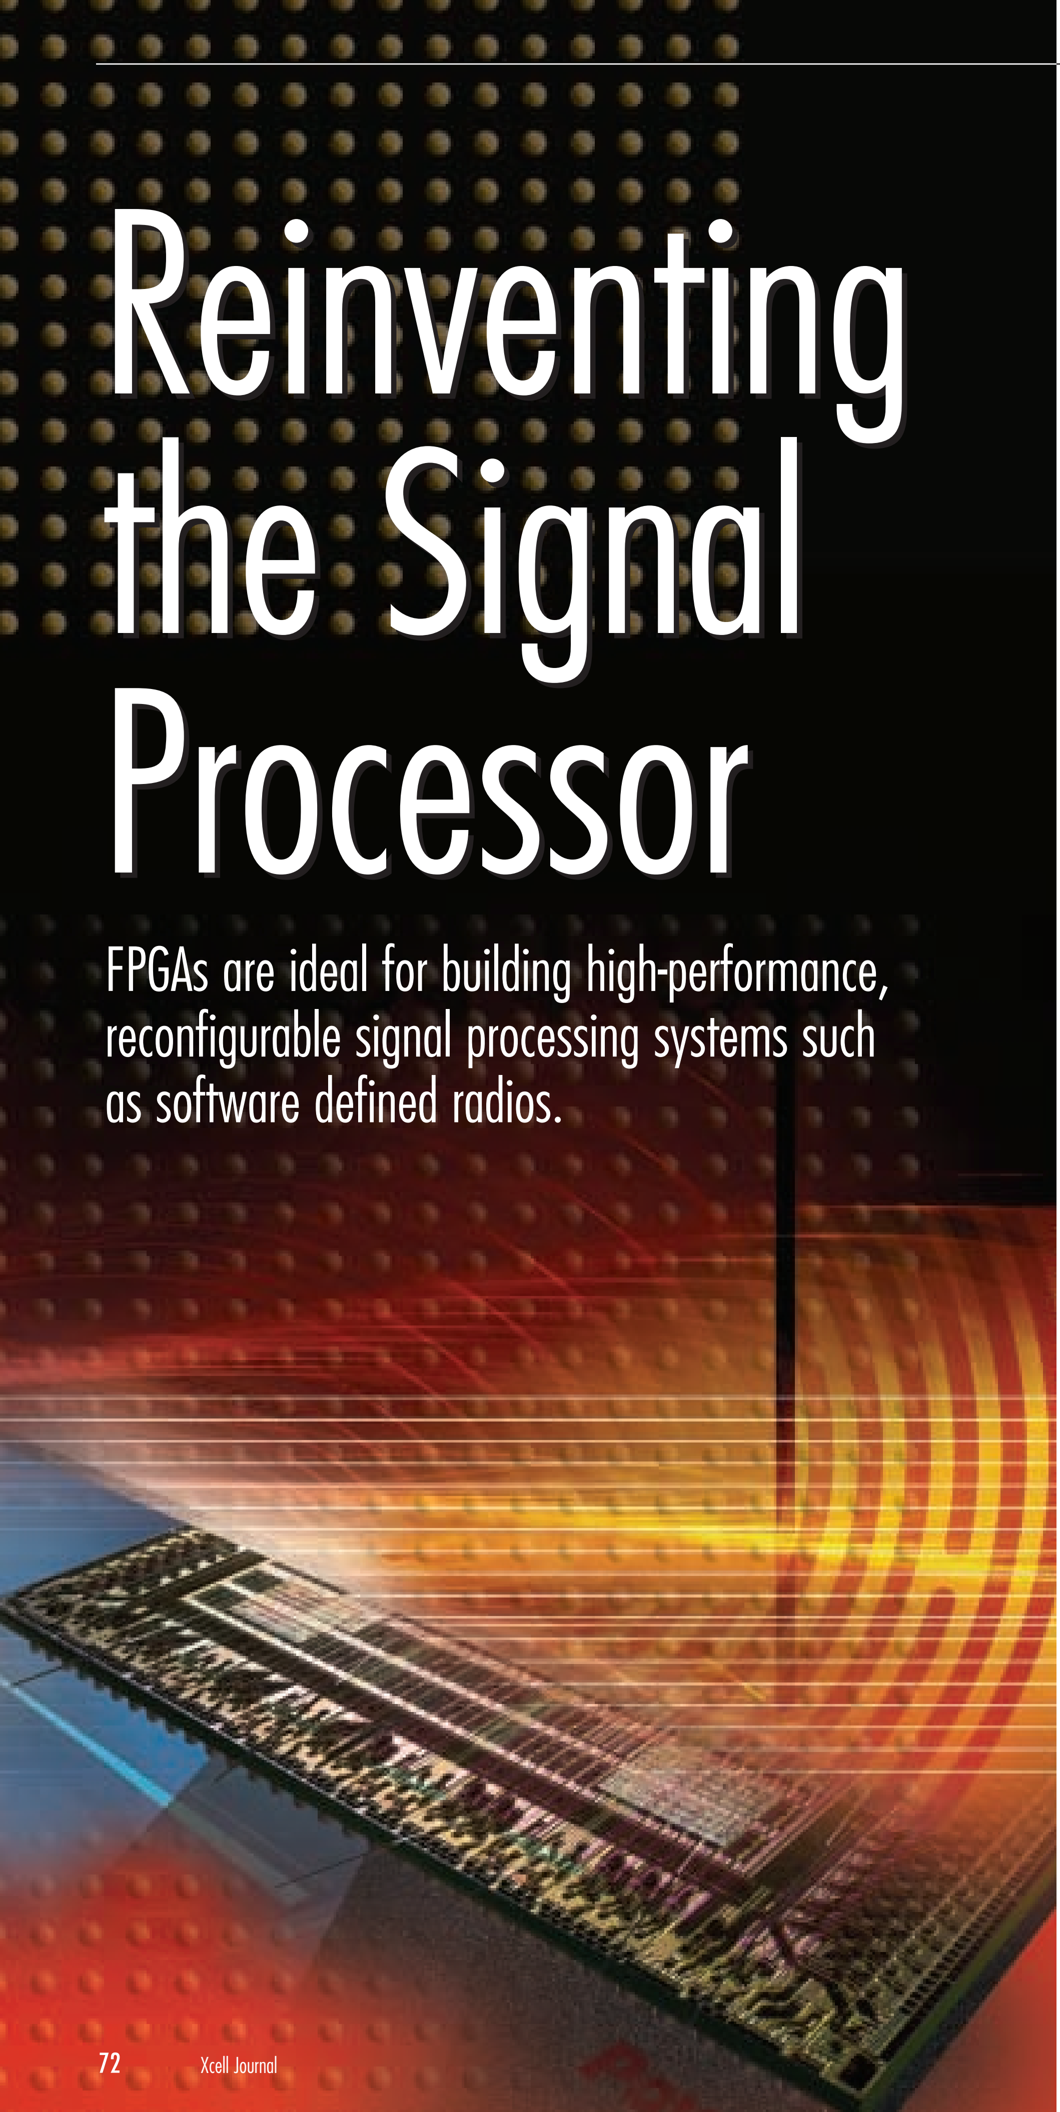

Reinventing the Signal Processor

FPGAs are ideal for building high-performance, reconfigurable signal processing systems such as software defined radios.

by Chris Dick, Ph.D.
Chief DSP Architect and Director
Signal Processing Engineering
Xilinx, Inc.
chris.dick@xilinx.com

The ultimate goal in software radio has been the realization of an agile radio that can transmit and receive at any carrier frequency using any protocol, all of which can be reprogrammed virtually instantaneously.

The Software Defined Radio Forum (SDRF) (*www.sdrforum.org*), an organization dedicated to supporting the development, deployment, and use of open architectures for advanced wireless systems, defines a software defined radio (originally coined by Joe Mitola in 1991 [1]) as radios that provide software control of a variety of modulation techniques. These include wide-band or narrow-band operation, communications security functions (such as hopping), and waveform requirements over a broad frequency range.

Figure 1 shows the architecture of a generic software radio. Smart antenna array technology is used for both the receive and transmit paths in the system. On the receive side, multiple high-bandwidth digitized antenna data is channelized, converted to baseband, and filtered – typically the sample rate is adjusted at this node. Other sections of the radio's physical layer (PHY) perform demodulation, synchronization, multiuser detection, adaptive interference cancellation, source decoding, forward error correction, beam forming, and adaptive equalization.

All of these computations present significant challenges for the radio PHY signal processing engine. Furthermore, much of the processing occurs at very high data rates.

Demands for Configurability and Agility

One of the driving objectives underlying SDR concepts is the desire to have a single hardware platform capable of servicing a number of radio environments. This type of reconfigurability could be used in several ways. For example, manufacturers developing infrastructure equipment or network operators building out a network could deploy a soft radio system in Europe configured to support Universal Mobile

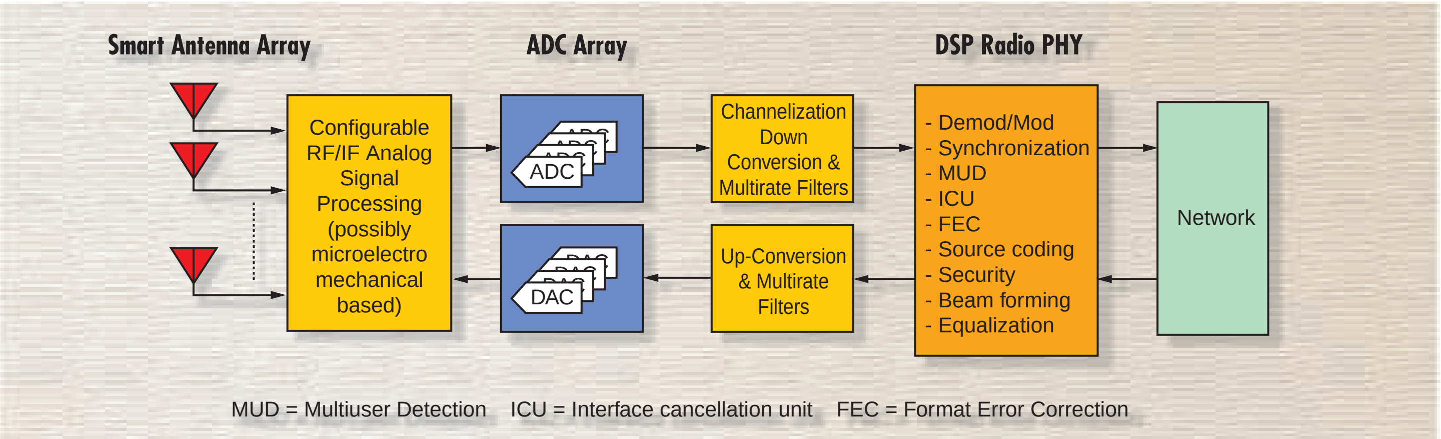
Smart Antenna Array
ADC Array
DSP Radio PHY
Configurable RF/IF Analog Signal Processing (possibly microelectro mechanical based)
ADC
DAC
Channelization Down Conversion & Multirate Filters
Up-Conversion & Multirate Filters
- Demod/Mod
- Synchronization
- MUD
- ICU
- FEC
- Source coding
- Security
- Beam forming
- Equalization
Network
MUD = Multiuser Detection
ICU = Interface cancellation unit
FEC = Format Error Correction

Figure 1 - Generic software radio architecture showing adaptive antenna array, analog signal processing, digitization, and the radio signal processing PHY.

Telecommunications System (UMTS) or Global System for Mobile Communication (GSM) standards, or operate the system in the U.S. with a Code Division Multiple Access (CDMA)2000 radio personality profile. This one system could also be operated as a multimode radio in an environment that employs both wideband and narrowband CDMA communications.

Radio agility is important in situations where standards are fluid. For example, consider the evolution of the 3GPP standard and the length of time required for that standard to stabilize. Agility is also important during transition periods. As we move from 2G to 3G mobile cellular systems, multiple standards such as Personal Digital Communication System (PCS), GSM, IS-95, Personal Handyphone System (PHS), DECT, EDGE, GPRS, IMT-2000, and CDMA2000 must all coexist.

Multistandard support will be a fact of life for the foreseeable future. When the 4G wireless network build-out is completed, multimode operation will be required to support third-generation wireless direct sequence spread spectrum (DSSS) and orthogonal frequency division multiplexing (OFDM), the modulation scheme most likely to be deployed in 4G systems. From a network operator perspective, base tranceiver station (BTS) configurability could be used to dynamically allocate radio resources. This might occur on the time-scale of hours in order to provide the highest quality of service to the subscriber base at any given time.

Both manufacturers and network operators could also use a configurable BTS to permit field upgrades or bug fixes to equipment already deployed, by supplying a new BTS profile using the Internet or a microwave link from a radio network controller to the BTS. Soft radios can also be viewed as a means to protect infrastructure investments by keeping radio hardware from becoming obsolete as new standards and techniques become available.

Economics and The Effect on Software Radio Development

Although commercial technology and economics have always been inextricably linked, significant changes are occurring in both of these domains that will alter the way electronic equipment, including software radios, is developed and deployed. From a purely technological perspective, *The International Technology Roadmap for Semiconductors (ITRS)* shows that Moore's Law will remain in effect for at least another 15 years, and that in the year 2016 devices will be produced on a 22 nanometer node.

Yet Patrick Gelsinger, Intel's chief technology officer, announced at last year's International Solid State Circuit's Conference (ISSCC 2001) plans for a 20 or 30 nanometer process in 2010, delivering a device consisting of 2 billion transistors operating at a clock frequency of 30 GHz. The estimated power consumption of the device would be 3 to 5 kW, or a power density of 1 kW/square centimeter, about the same as a rocket nozzle. This has obvious thermal implications that must be dealt with using techniques radically different from today's methods.

Instruction set architecture (ISA) signal processors share many similarities to general purpose processors. Architectural differentiates such as very long instruction word (VLIW), super-scalar extensions, and various types of predictive enhancements are really micro-architecture evolutions of the basic architecture credited to von Neumann and his colleagues in the 1940s and 1950s. As such, signal microprocessors have leveraged most of their performance via a raw increase in clock frequencies. For example, in the early 1980s the first fixed-point signal processors supported clock frequencies in the 5- to 10-MHz region. Current-generation high-end ISA Digital Signal Processors (DSPs) use 600-MHz clocks and are on a trajectory to the Giga-Hertz region. Obviously, this curve has the same thermal pitfalls described above.

Although FPGAs take advantage of Moore's Law (and other advanced process technology such as all-copper interconnect and low-K dielectric substrates) to provide increased clock frequencies over time, their primary mechanism for supplying performance is completely different than the ISA approach. FPGAs exploit the large amount of parallelism inherent in most signal processing algorithms. With as many as 556 embedded multipliers and 125,136 logic cells in the Xilinx Virtex-II Pro™ Platform FPGA, we can readily see how these devices can be viewed as a naturally parallel processing engine that can take advantage of the rich parallelism in a software radio PHY.

The software radio PHY is a complex signal processing system in which algorithmic and functional-level parallelism can be leveraged to realize a high-performance system that does not rely on raw speed for its performance. The multiplier array could be used to implement space-time processing in a receiver, while at the functional level multiple turbo convolutional decoders could be operating concurrently to support multiple users, each with a 2 Mbps data rate in a 3G environment.

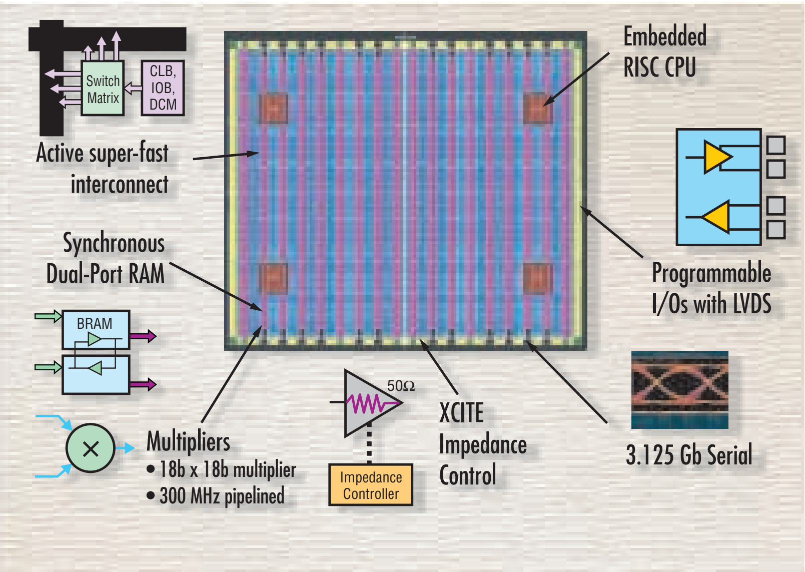

Figure 2 - Virtex-II Pro Platform FPGA showing the multiplier array for supporting parallel signal processing, multi-gigabit transceivers for inter-chip and inter-system connectivity, and embedded RISC processor technology for performing decision-oriented tasks and running a real-time operating system.

Manufacturers have made 60% more transistors available to circuit designers per area of silicon compared with what was available a year earlier. In contrast, the ratio at which designers are able to utilize transistors in circuits of any given tier of complexity has only been increasing at a rate of 20% per year [3].

This "design gap" is associated with performance supply and demand, but another aspect is methodology related. It is becoming an increasingly complex, time-consuming, and error-prone procedure to develop and verify a sophisticated ASIC. Furthermore, at a cost of $1-$2 million for mask set costs, it is becoming prohibitively expensive. ASIC development timelines are now spanning years, and may even extend beyond the window of opportunity for the intended product.

Harvard Business School Professor Clayton Christensen highlights that while price and performance are still important, there are signs that a seismic shift is taking place, leading to a new era where other factors – such as customization – matter more [4]. This is precisely where the FPGA fits in: it is the ultimate in customization.

FPGAs address the technical as well as business perspectives outlined above. Because they are off-the-shelf commodity items, companies can access state-of-the-art device technology with minimal NRE, and quickly build and deploy customized systems, achieving very short time-to-market while simultaneously maximizing first-to-market revenue streams. The Virtex-II Pro platform FPGA shown in Figure 2 is the cornerstone technology for building high-performance reconfigurable signal processing systems – which includes the PHY in an SDR. In conjunction with the logic fabric and active interconnect, this device has an array of embedded multipliers for supporting the most demanding of arithmetic tasks in a radio PHY. This particular FPGA family also offers integrated Power PC 405 technology, multi-gigabit transceivers and dynamic impedance matching capability on the device I/O ports that can be used to simplify printed circuit board design and manufacturing. Using a platform-based approach to system implementation, system designers can create product differentiates by implementing signal processing functions in the logic fabric as well as through embedded software running on a Power PC.

The DSP Dilemma

One approach to BTS implementation has been to employ a combination of ASIC and ISA DSPs. The ASIC technology is typically used to address the significant arithmetic requirements of the radio front-end, such as digital down conversion and channelization filters to support multicarrier W-CDMA or CDMA2000 standards. These functions are beyond the capabilities of even state-of-the-art ISA DSPs. Even though the ASICs used in this part of the system may offer some programmability, it is generally limited in nature and is certainly a departure from the intended philosophy of the fully configurable soft radio. DSP processors might be used for certain baseband functions such as source (de-) encoding, as with CELP codec. Reduced Instruction Set Code (RISC) processing resources in the system could also support the requirements of the higher levels in the protocol stack.

From a soft radio perspective, the ASIC/processor combination is poor partitioning from both a flexibility and efficiency standpoint. In recent years FPGAs have experienced hyper-growth in both arithmetic complexity and compute density (number of operations/unit area of FPGA) that can be achieved by current generation devices. What types of signal processing functions can be usefully realized by an FPGA?

Radio designers working with FPGA technology implement IF sampled receivers, channelizers of different varieties including classical digital down (and up) conversion (DDC and DUC) architectures, FFT-based polyphase transforms, multistage multirate polyphase decimators and interpolators, adaptive interference cancellers for DSSS channels, multiuser detection (MUD), and rake receivers (including acquisition and tracking). More recently, FPGAs have been used to construct space-time processors for advanced smart antenna systems. FPGAs are extremely adept and flexible at implementing FFTs, and this functionality has been used to construct OFDM modulators and demodulators.

FPGAs have also found extensive use in narrowband bandwidth-efficient Quadrature Amplitude Modulation (QAM) systems. In this environment they have been used to implement adaptive channel equalizers, digital timing recovery circuits, carrier recovery loops, frequency locked loops, and fractional rate change filters.

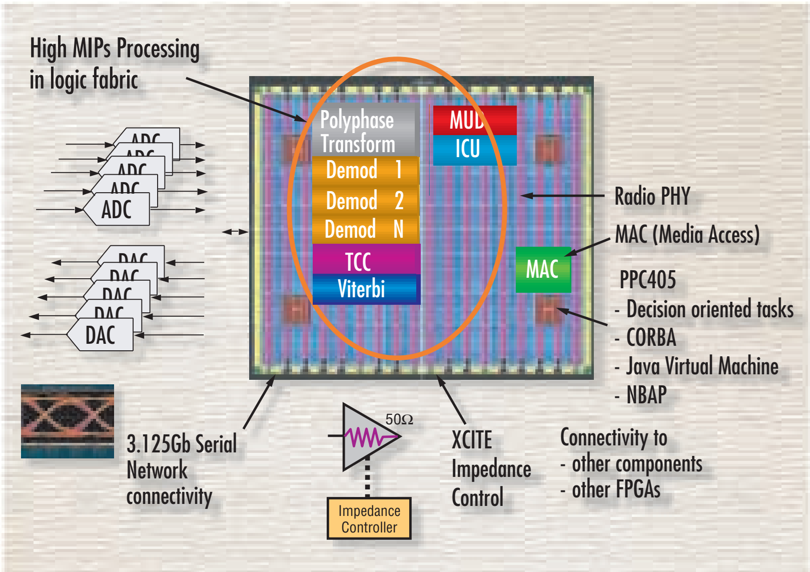

Figure 3 - Platform FPGA approach to software-defined radio realization. The high MIPs processing is implemented in the logic fabric, while decision-oriented and non-real-time tasks are provided as embedded software running on the Power PC. The multi-gigabit transceivers could be used for providing connectivity to the broader network.

FPGAs are also extensively used for forward error correction in communication systems. For example, OC-3 155 Mbps Viterbi decoders, Reed-Solomon decoders, Reed-Solomon decoding at OC-192 10 Gbps data rates, and (de-)interleavers operating at clock frequencies greater than 200 MHz are all achievable with current generation Virtex™-II FPGAs.

FPGAs are the ultimate device technology in terms of user customization. They allow system architects to perform area-performance tradeoffs and to therefore "right-size" the functional components in the system. FFTs with execution times in the microsecond to tens-of-microseconds are possible. In the context of an OFDM communication system, a small number of FPGA resources could be used to realize a (de-)modulator that supports a moderate data rate, or by using more resources an extremely high-performance high data-rate link could be realized.

With FPGA technology, control of the silicon is put back into the hands of the system developer rather than the chip architect – as is the case with an ISA signal processor. In fact, one way to view an FPGA is as a miniature silicon foundry with turnaround times of hours rather than months.

Leveraging these types of tradeoffs does not always mean that the engineering team has to construct the functional units from first principles. To facilitate rapid product development, many signal processing functions are available from the FPGA manufacturers themselves and from third-party intellectual property (IP) suppliers. FFTs, multirate filters, Viterbi decoders, and Reed-Solomon encoders and decoders are all available as pre-verified IP from Xilinx (*www.xilinx.com/xlnx/xil_prodcat_landingpage.jsp?title=Xilinx+DSP*).

One of the roadblocks to the widespread deployment of FPGA-based signal processing has been design methodology related. In the past, FPGA-DSP design has required signal processing and communication engineers to use tool flows and languages with which they are typically unfamiliar. The introduction of tools like System Generator for DSP (*www.xilinx.com/xlnx/xil_prodcat_product.jsp?title=system_generator*) has gone a long way to let engineers work in the language of the problem. In this case the system is developed using a visual dataflow paradigm in The Mathworks Simulink environment. The approach not only allows the design to be specified, simulated, and parameterized, but it also enables design reuse through the use of IP cores.

The Reconstruction of the Software Radio

The platform FPGA provides an opportunity for the radio architect to reinvent the system. Instead of having a radio card that is responsible for the DSP heavy lifting at the front-end of a soft radio system, and then passing this partially processed data over a VME or PCI-X bus to a baseband processor, multiple functions could be integrated into one or a small number of platform FPGAs.

As shown in Figure 3, compute-intensive tasks in the radio PHY could be implemented in the FPGA logic fabric, while more decision- and control-oriented tasks are run as embedded software on the Power PC™. This embedded processor could even be used to run a Node B application protocol (NBAP) for a BTS, as a java virtual machine, or even provide CORBA support.

Advances in analog-to-digital converter technology are still required to support the high dynamic range requirements of wideband radio front-ends that offer true multimode global operability. Advances are also required in the area of configurable high-bandwidth analog signal processing for realizing the RF and IF stages of a radio. Micro-electromechanical systems (MEMs) appear to be a promising technology for addressing in-system configurable analog signal processing. As this technology matures and is combined on a single platform with digital functionality, the ideal of a completely configurable radio will move closer to reality.

The significant computation demands of the SDR PHY have been largely satisfied by highly parallel signal processing platforms realized using recent generation FPGA technology from companies like Xilinx. To complement the device technology, an increased emphasis on Signal Processing IP libraries and design methodologies such as System Generator for DSP are taking on renewed roles to provide a solution to the challenges presented by the software radio.

(Note: A useful primer for engineers and executives interested in developing products in the SDR application space can be found on the SDRF's web site, *www.sdrforum.org/sdr_primer.html.*)

References

[1] Mitola, J., III. 2000. *Software Radio Architecture: Object Oriented Approaches to Wireless Systems Engineering.* John Wiley and Sons.

[2] Allan, A., D. Edenfeld, W. H. Joyner Jr., A.B. Kahng, M. Rodgers, and Y. Zorian. 2002. 2001 Technology Roadmap for Semiconductors. IEEE *Computer Magazine* 35 (1):42-53.

[3] Verlinden, M.C., S.M. King, and C.M. Christensen. 2002. Seeing Beyond Moore's Law. *Semiconductor International* 25 (8):50-60.

[4] Bass, M.J., and C.M. Christensen. The Future of the Microprocessor. 2002. *IEEE Spectrum Magazine*: 34-39.

How to Make Smart Antenna Arrays

The Nallatech BenADIC card combines a 20-channel data acquisition system with Xilinx XtremeDSP technology and Virtex FPGAs for high-performance digital signal processing.

by Malachy Devlin, Ph.D.
Chief Technology Officer
Nallatech
m.devlin@nallatech.com

Wireless communication has created a continuing demand for increased bandwidth and better quality of service. With the ever-increasing number of mobile network subscribers, available capacity is becoming more of a premium.

"Smart" antenna arrays are one way to accommodate this increasing demand for bandwidth and quality. These antenna arrays provide numerous benefits to service providers. However, the processing requirements for smart antenna arrays are many orders of magnitude greater than those for single antenna implementations.

In this article, we will describe how smart antenna arrays work and present a new product from Nallatech™ that combines a 20-channel data acquisition system with an FPGA computing fabric for handling the high-performance digital signal processing (DSP) operations. We also show you how this combined product is integrated into a scalable system using Xilinx Internet Reconfigurable Logic (IRL™) technology for remote configuration and control of the system using Nallatech's field upgrade systems environment (FUSE™) software.

Focus Power with Smart Antennas

Figure 1 shows a conventional antenna as omnidirectional. It radiates and receives information equally in all directions. This equal distribution leads to power being transmitted to, but not received by, the user. This wasted power becomes potential interference to other users or to other base stations in other cells. Interference and noise reduce the signal-to-noise ratio used by the detection and demodulation operations, resulting in poor signal quality.

To overcome the problems associated with omnidirectional arrays, smart antennas focus all transmitted power to the user and only "look" in the direction of the user for the received signal. This ensures that the user receives the optimum quality of service and maximum coverage for a base station. An intermediate step to this ideal is using directional antennas that divide the 360-degree coverage into sectors. As shown in Figure 2, four directional antennas can each cover approximately 90 degrees.

Instead of using individual antennas, we can create a smart antenna array and add further processing intelligence to the data received or transmitted with this array. Smart antenna arrays enable us to direct beams in specific directions through electronic or software control.

Two types of smart antenna arrays are switched-beam arrays and adaptive arrays. As shown in Figure 3, switched-beam arrays comprise a number of predefined beams. The control system switches among the beams and selects the beam that provides the maximum signal response.

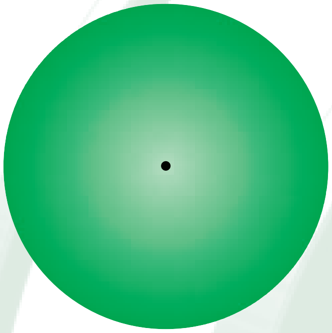

Figure 1- Omnidirectional antenna

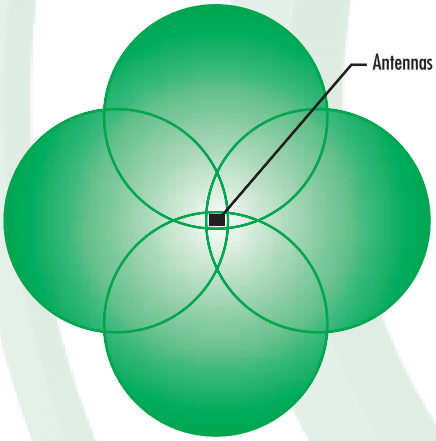

Figure 2 - Sectorized antenna with four sectors

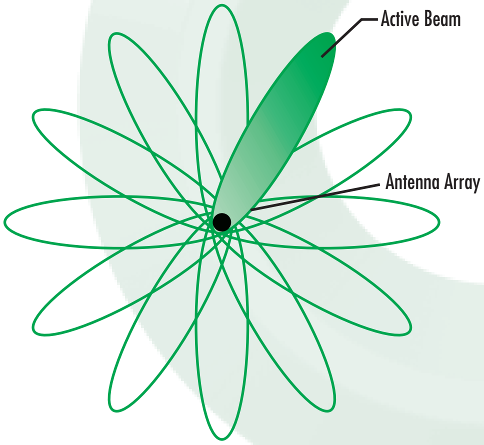

Figure 3 - Switched antenna array with active beam highlighted

Adaptive antenna arrays, on the other hand, incorporate more intelligence into their control system than do switched-beam arrays. Adaptive antennas monitor their environment and, in particular, the response of the data path between the user and the base station. This information is then used to adjust the gains of the data received or transmitted from the array to maximize the response for the user. With adaptive antenna arrays, the control system has full flexibility and determines how the gains of the arrays are adjusted. By adjusting the gains in this way, the control system can – in addition to maximizing the gain from a particular user – also attenuate the signal from an interfering source, such as from another user or from multipath signals. Therefore, as shown in Figure 4, adaptive arrays maximize the signal-to-interference-plus-noise ratio (SINR) and not just the signal-to-noise ratio (SNR).

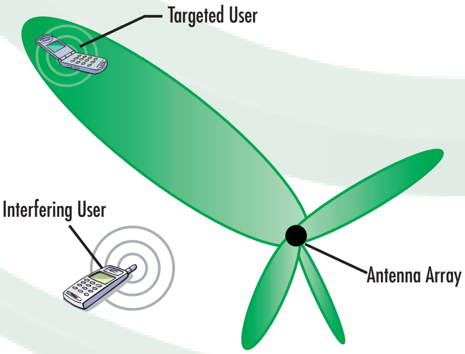

Figure 4 - Adaptive antenna array enhancing the SINR

This dynamic adaptation of the antenna array response provides focused beams to specific users and a new mechanism for multiuser access to the base station. Conventionally, multiple users are separated when accessing the base station by using different frequencies, as in frequency division multiple access (FDMA). FDMA is used in advanced mobile phone services (AMPS) and total access communications systems (TACS). FDMA is also used in time, as in time division multiple access (TDMA) for global systems for mobile communications, interim standard 136 (IS-136), or code division multiple access (CDMA), which is used in third generation (3G) systems.

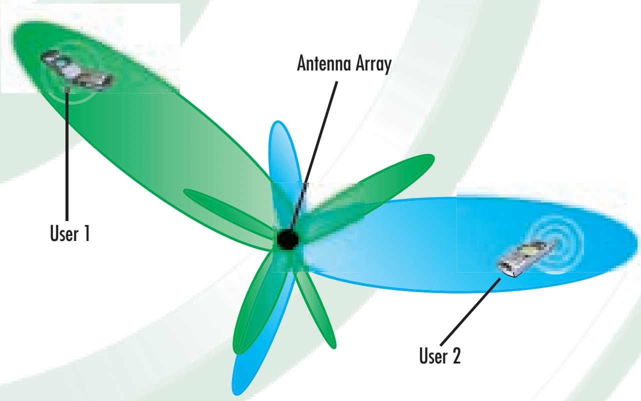

Figure 5 - SDMA allows two users to access the same base station on the same frequency.

As shown in Figure 5, by using smart antenna arrays, we can now use space division multiple access (SDMA). In this case, users may use the same frequency, time, or code allocations over the air interface and only be separated spatially. This enables SDMA to be a complementary scheme to FDMA, TDMA, and CDMA, and SDMA thus provides increased capacity within congested areas.

Smart Antenna Processing

A fully adaptive antenna array implementation requires a considerable increase in processing requirements. Previously, we had a single stream of data coming from a single antenna; now, we have multiple data streams to process. As shown in Figure 6, the data flow diagram for a beam-forming application is not a single input data stream. We now have N data streams that must be processed from the N antenna elements.

The fundamental operation carried out in adaptive arrays is to pass the data stream from each antenna through an adaptive finite impulse response (FIR) filter. Note that in narrowband applications, the adaptive FIR filters simplify to a single weight vector. The processing requirements increase, however, with each beam processed.

If we consider a simple example where we have four antennas and a narrowband system, such that the adaptive filters result in a single multiplication, we can see that the processing requirements approach one-half billion multiple accumulates (MACs) per second, for a sample rate of 105 mega samples per second. This sample rate is for a single beam and does not include the processing requirements for the adaptive update algorithm. This amount of processing does not seem unreasonable for performance in a DSP processor. However, if we want to support multiple beams and achieve finer beams by increasing the number of antennas, we could quickly exhaust the processing capability of a standard processor architecture as we reach processing requirements of several billion MACs per second.

By using FPGAs, we have powerful DSP devices for handling these high-performance requirements at sampled data rates. Furthermore, we can take advantage of the FPGA flexibility for directly handling acquisition control and other DSP functions, such as digital down-conversion, demodulation, and matched filtering.

20-Channel Data Acquisition

Figure 7 shows the Nallatech BenADIC™ data acquisition card, which can simultaneously capture data from 20 sources at a sustained rate of 105 mega samples per second. The analog inputs have a 250 MHz bandwidth and the data is digitized at 14 bits resolution. The card produces 3.675 gigabytes of digitized data every second, or the equivalent of 5.4 audio CDs, for processing.

The large number of tightly coupled input channels makes the BenADIC card ideal for processing smart antenna arrays. As shown in Figure 8, the 20 input channels are partitioned into five groups of four channels. The analog-to-digital converters (ADC) in each of these groups are connected to their own Xilinx Virtex™-E FPGAs. This enables local processing of the four channels. Thus, the architecture can be

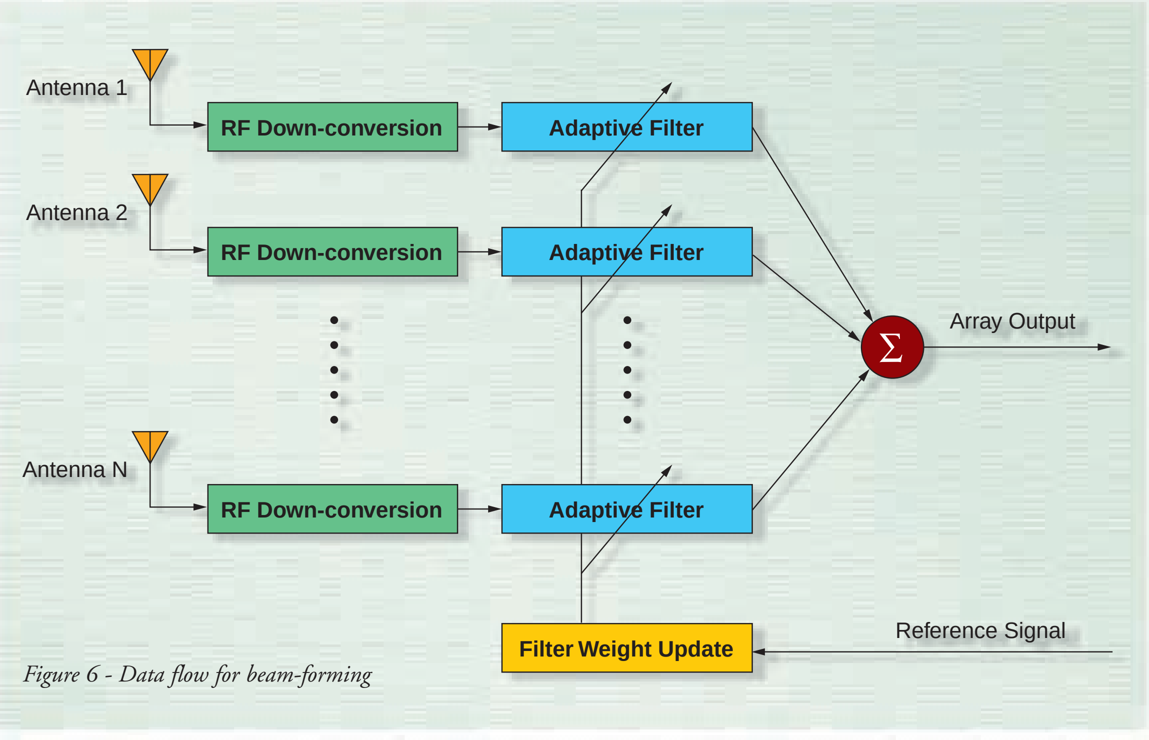

Figure 6 - Data flow for beam-forming

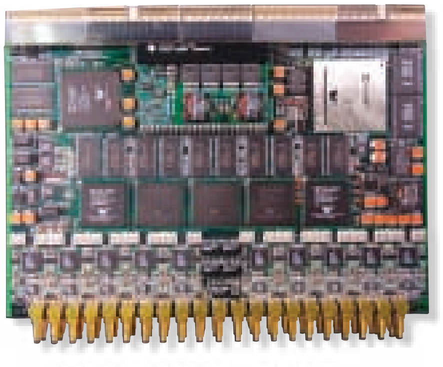

Figure 7 - Xilinx Virtex-E FPGAs on a BenADIC 20-channel, 14-bit data acquisition card

arranged to handle five antenna arrays, each with four antennas within the array. Alternatively, the high-speed internal buses enable these groups to be interconnected to handle an array of 20 antennas.

In addition to the channel group FPGAs, a large FPGA can handle further processing and communicate with the compact PCI (cPCI) backplane and the PCI bus (via the interface FPGA). Communications over the cPCI backplane allow data transfer to other cards in the system, such as to the Nallatech DIME-II™-based BenADIC card, which can accept DIME-II modules and provide more than 50 million system gates with the Virtex-II family for Xilinx XtremeDSP™ operations.

The BenADIC card is compatible with the tools and cores produced by the Xilinx DSP group and includes the powerful System Generator tool. By using System Generator, you can develop and verify your algorithms within MATLAB™ and Simulink™ environments. You can then synthesize and implement your design for the FPGA. This implementation exercise has been carried out successfully for the BenADIC using System Generator.

Xilinx FPGAs Allow for Software-Defined Radio

The great thing about Xilinx FPGAs is their ability to be reprogrammed on the fly and to give hardware different personalities based on the application. Nallatech has been implementing dynamically reconfigurable FPGA systems for a number of years. The BenADIC card is Nallatech's newest product.

The BenADIC card is compatible with the Nallatech FUSE software environment, which provides the capability to selectively and dynamically change the operation of FPGAs in the BenADIC card or other FPGA-based systems, including modular DIME systems. This ability to dynamically update a system leads to the definition of a software-defined radio where the receiver characteristics are controlled via software. By using the FUSE software, this control can be handled locally over a PCI or remotely via TCP/IP, for example.

Conclusion

The BenADIC card from Nallatech couples the power of Xilinx FPGAs with a highly integrated 20-channel data acquisition system on a single card. It is ideally suited for handling large, smart antenna arrays or several smaller arrays.

The combination of the FPGA performance and flexibility enables the realization of advanced DSP algorithms, which in turn opens the possibility for deploying advanced wireless interfaces. Deploying advanced wireless interfaces provides users with a better quality of service and gives service providers greater capacity.

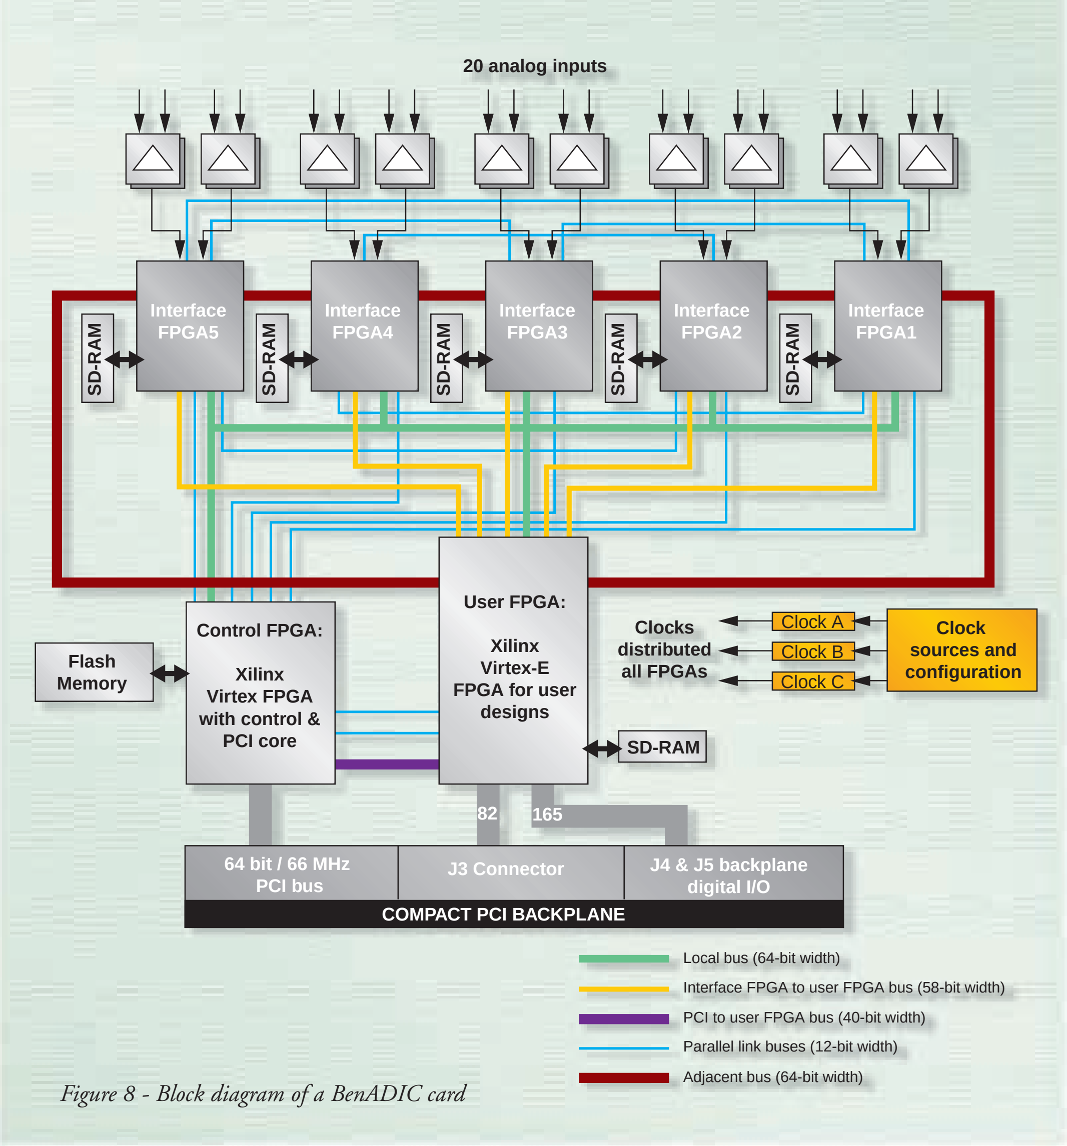

20 analog inputs
Interface FPGA5
Interface FPGA4
Interface FPGA3
Interface FPGA2
Interface FPGA1
SD-RAM
SD-RAM
SD-RAM
SD-RAM
SD-RAM
Flash Memory
Control FPGA: Xilinx Virtex FPGA with control & PCI core
User FPGA: Xilinx Virtex-E FPGA for user designs
Clocks distributed all FPGAs
Clock A
Clock B
Clock C
Clock sources and configuration
SD-RAM
82
165
64 bit / 66 MHz PCI bus
J3 Connector
J4 & J5 backplane digital I/O
COMPACT PCI BACKPLANE
Local bus (64-bit width)
Interface FPGA to user FPGA bus (58-bit width)
PCI to user FPGA bus (40-bit width)
Parallel link buses (12-bit width)
Adjacent bus (64-bit width)

Figure 8 - Block diagram of a BenADIC card

FPGAs Have the Multiprocessing I/O Infrastructure to Meet 3G Base Station Design Goals

Two-dimensional fabric efficiently links arrays of processors inside Virtex-II Pro devices to enable parallel processing of data.

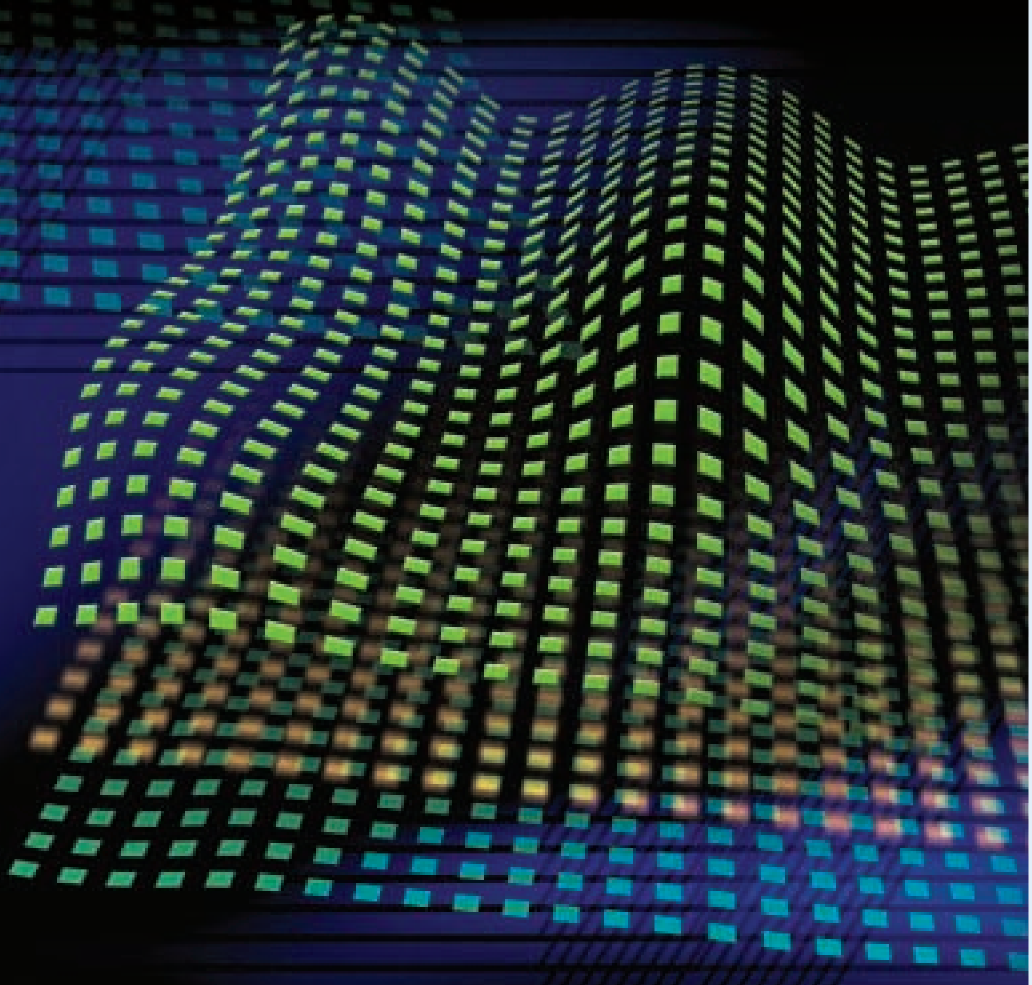

by Peter Galicki, CEO
CrossBow Technologies Inc.
peter.galicki@crossbowip.com

With increased data traffic and new multiuser detection and adaptive beamforming algorithms, data processing requirements of 3G base stations will increase by as much as 100 times relative to current equipment. This increase in processing capacity must be matched by low power consumption, as the new picocell base stations mounted on building sides will not be using forced air cooling.

Arrays of small and specialized processors (Figures 1 and 2) will provide a power-efficient method of increasing performance, more so than can be obtained by increasing the features of larger, general-purpose super processors.

The evolution of current standards and introduction of new standards currently force base station operators to perform frequent upgrades to their wireless infrastructure, often requiring board replacements. To reduce field maintenance, 3G equipment must be upgradable without board swapping. The high cost of 3G equipment often leaves wireless infrastructure manufacturers with thin profit margins; costs will have to come down to enable large-scale deployment.

Finally, OEMs cannot abandon their current design methods to start designing 3G equipment from scratch. They must be able to reuse semiconductor IP, code, and development tools to hit market windows and to obtain a return on investment.

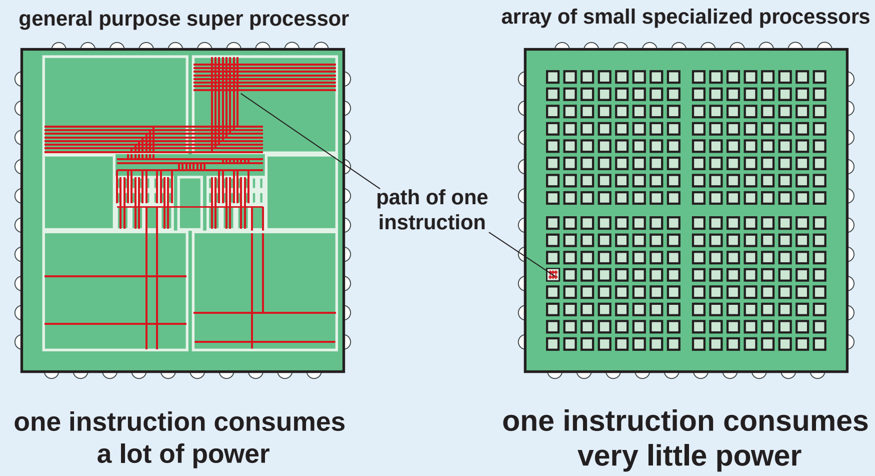
general purpose super processor
array of small specialized processors
path of one instruction
one instruction consumes a lot of power
one instruction consumes very little power

Figure 1 – Arrays of small specialized processors reduce power consumption.

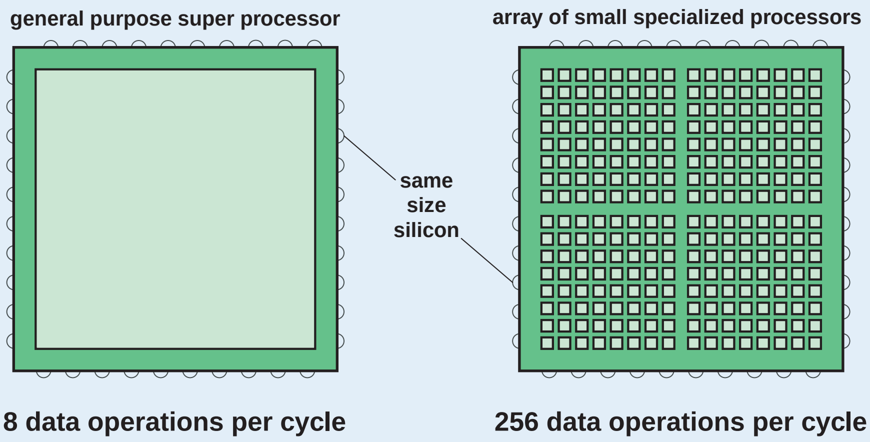
general purpose super processor
array of small specialized processors
same size silicon
8 data operations per cycle
256 data operations per cycle

Figure 2 – Arrays of small specialized processors increase performance.

Meeting these seemingly exclusive goals requires a combination of top-notch process technology, combined with comprehensive component library and efficient data communication methods, inside-chip and chip-to-chip. Reaching 3G design goals hinges on achieving the right balance between the size and number of data processing components to keep most of the chip busy at all times, while reducing the overall distance that the data has to travel inside ICs. System efficiency is heavily influenced by design partitioning, optimization of individual data processing components, and streamlining data flow between components. To keep data processing components busy, inter-processor data transfers must have low latency and be precisely deterministic – otherwise components will waste valuable processing cycles while waiting for data.

Low latency requires the removal of data communication bottlenecks by spreading the data flow over the entire area of a chip. Transfer determinism is achieved by a combination of low latency and a uniform data communications structure inside the chip.

Data Processing Elements

Each 3G chip is likely to contain hundreds of processing elements, many representing autonomous processors with their own data processing flows, control flows, memory, and communications ports. Some data processing flows may be augmented with dedicated DSP blocks. Virtex™-II FPGAs support DSP functions and a MicroBlaze™ soft processor. Virtex-II Pro™ devices also feature embedded IBM PowerPC™ processors. Depending on the task at hand, smaller processors may be better suited for simpler functions, and larger processors may be a better fit for more complex algorithms.

In order to work in parallel, processors must be able to easily communicate with each other. An efficient way for processors to communicate is through a fabric dispersed across the entire design that looks to individual processors like conventional memory (Figure 3). This approach enables each processing element to be developed and verified individually, yet easily exchange data with other processors.

Two-Dimensional Data Communications

An effective data interconnect fabric must support low latency and deterministic data transfers occurring simultaneously among multiple processing elements. It must also be flexible and scalable to allow for the addition of new elements or the removal of unwanted elements without affecting the rest of the design. Finally, it should be compatible with existing processors and be as easy to use as accessing memory.

Memory-Like Interface

Using conventional bus cycles to transfer data between processors dispenses with exotic and hard-to-implement communications peripherals and protocols in favor of a simple memory-like interface. As shown in Figure 4, 2D-fabric from CrossBow appears to processors as a memory-mapped peripheral on an IBM CoreConnect™ bus. PowerPC and MicroBlaze processors can issue conventional read and/or write bus cycles to their local 2D-fabric peripherals to communicate with other processors on the chip (Figures 5 and 6). The payload for each transfer is derived from the data bus. The destination location and the initial direction of travel are derived from the address bus. The transfers are totally transparent to the sending and receiving processors, launching transfers with write cycles and terminating transfers with read cycles.

Routing of data from source to destination, as well as arbitration with other data traffic, is performed autonomously by the interconnected 2D-fabric peripherals.

A 2D Array of Data Transport Links

Efficient 3G designs will feature global communication fabrics using single sets of lines to transfer all kinds of data, including payloads, control words, and configuration data. Duplication of data transfer lines reduces overall system efficiency.

As shown in Figure 7, 2D-fabric peripherals of adjacent processors are interconnected with a single mesh of horizontal and vertical data transport links. Individual bus cycles are autonomously converted to small packets that travel between source processors and destination processors through chains of 2D-fabric peripherals of the intermediate processors along the way. Short point-to-point links reduce power consumption. Small packets with single word payloads reduce data transfer latencies, enabling data and control packets to share common transfer lines. The same lines can also be used for system initialization and configuration.

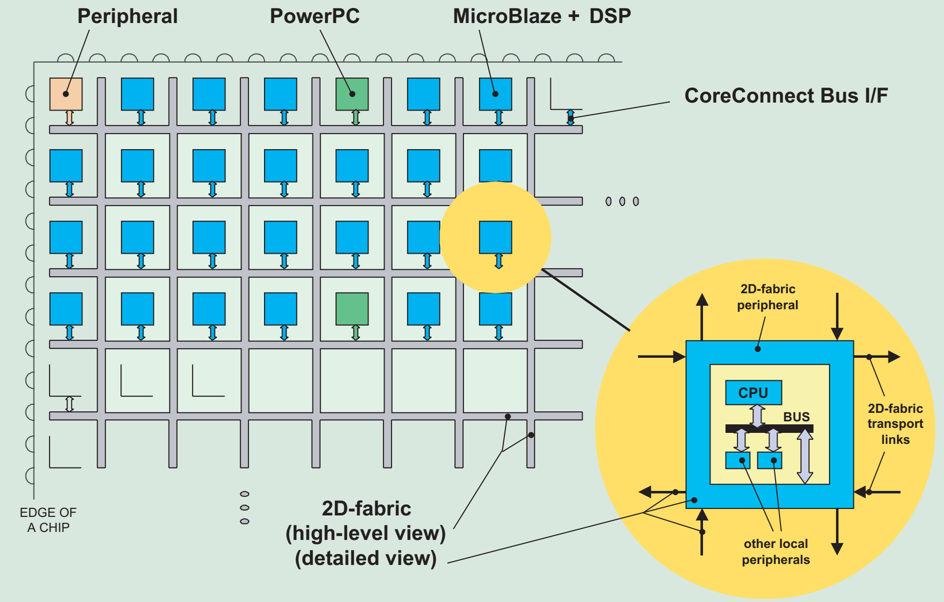

Figure 3 – Two-dimensional I/O fabric enables arrays of processors to work in parallel.

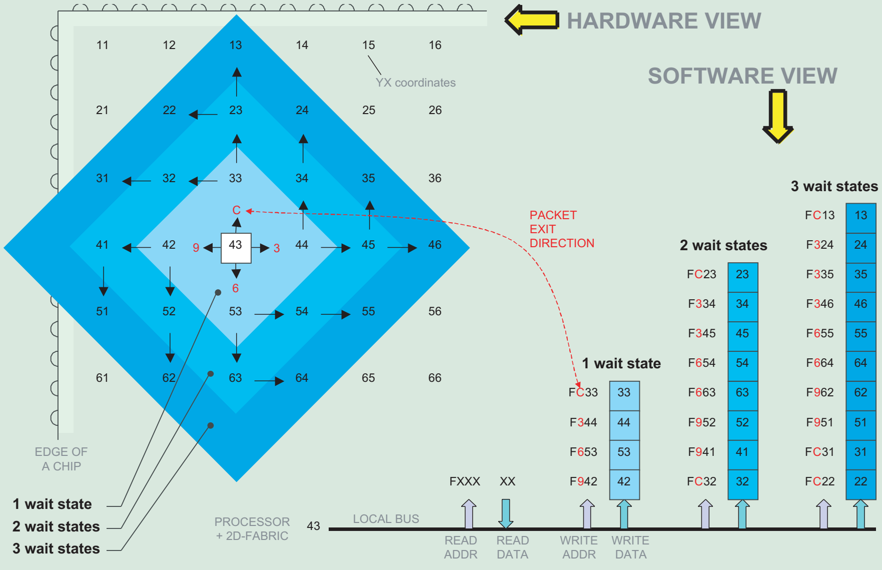

Figure 4 – Deterministic worst-case transfer latencies look like memory wait-states.

Scalability

Scalability is an important requirement for the design effort and product field upgrades. Constantly changing standards may require adding or removing processors late in the design cycle or even after field deployment.

In the past, adding or removing processors has always been difficult when using centralized DMAs for movement of data. In any centralized I/O structure, removing or adding new components is likely to affect other system components. Two-dimensional I/O structures are much less sensitive to design changes. Adding another processor to a chip is as simple as wrapping it with a 2D-fabric peripheral and connecting the respective data transport links to the existing fabric. This can be easily done without affecting any hardware or software already in place.

Low Latency and Deterministic Data Transfers

In computing environments where hundreds of processors are simultaneously exchanging data, how can you guarantee that any one of those transfers is going to arrive at its destination no later than a fixed amount of time? Buses, crossbars, and other centralized I/O structures force all data traffic through one central location, creating huge traffic jams. Two-dimensional I/O structures, however, can easily guarantee data delivery by spreading out data traffic across the design. As shown in Figure 7, a two-dimensional data transport grid dispersed across the entire design area removes communication bottlenecks to allow individual transfers to complete on time, without interfering with other transfers.

Individual processors must use worst-case transfer latency when planning data transfers. Although it is acceptable for data to wait to be transferred, processors waiting for data are wasting precious processing cycles. Total transfer latency depends on the worst-case latency across one processing node and the number of intermediate processing nodes between the source and destination nodes.

Worst-Case Latency Across One Processing Node

A 50 ns packet latency across one node represents the time elapsed from when the packet started entering the node to the time when it started exiting that node. A packet delay time is the time from when it starts entering the node to the time when it completely emerges. Thus, a 100 ns packet delay time is 50 ns latency plus another 50ns for the packet to fully emerge from the node.

If packets exiting from a given output port can arrive from three different sources, the worst-case latency for any one packet is 250 ns. This is equal to the best-case latency of 50 ns plus two packet-delay slots of 100 ns each.

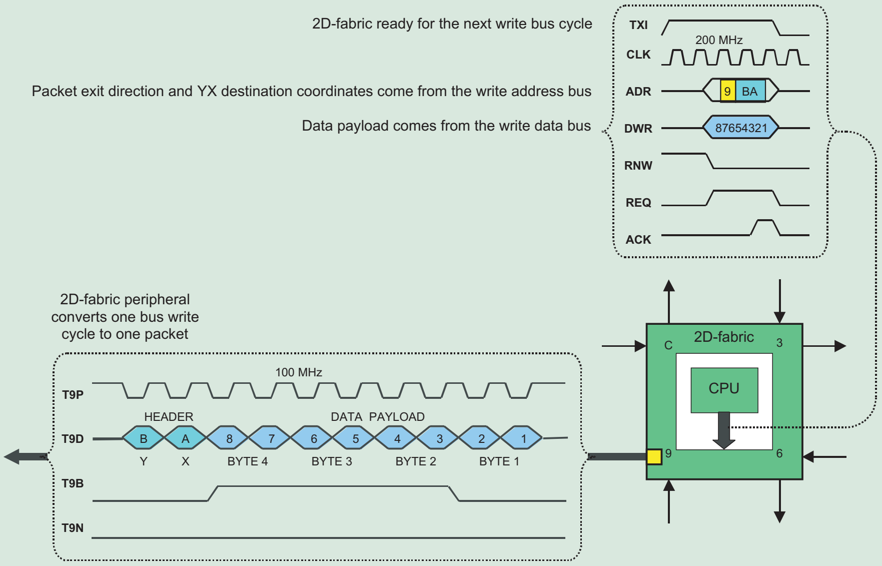
2D-fabric ready for the next write bus cycle
Packet exit direction and YX destination coordinates come from the write address bus
Data payload comes from the write data bus
TXI
200 MHz
CLK
ADR
9
BA
DWR
87654321
RNW
REQ
ACK
2D-fabric peripheral converts one bus write cycle to one packet
100 MHz
T9P
HEADER
DATA PAYLOAD
T9D
B
A
8
7
6
5
4
3
2
1
Y
X
BYTE 4
BYTE 3
BYTE 2
BYTE 1
T9B
T9N
2D-fabric
C
3
CPU
9
6

Figure 5 – Conventional write bus cycles launch data packets.

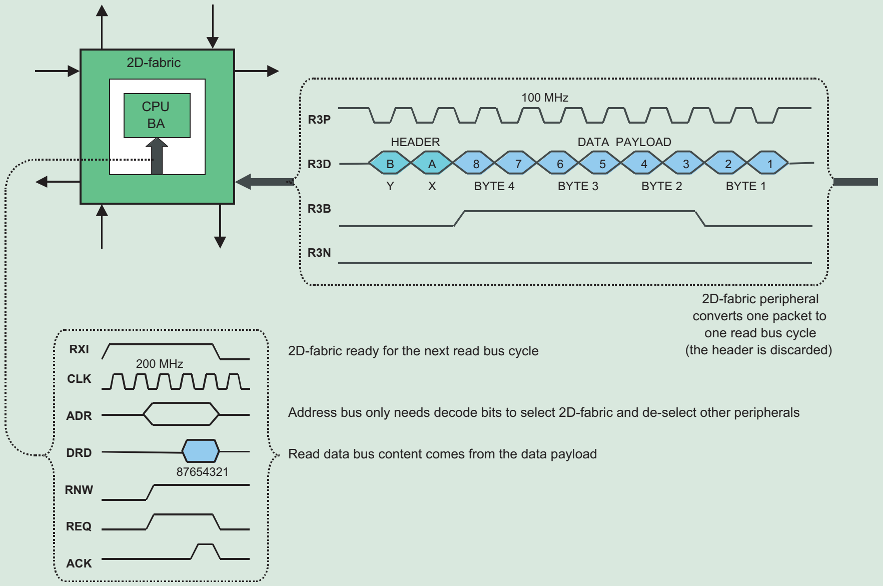
2D-fabric
CPU
BA
100 MHz
R3P
HEADER
DATA PAYLOAD
R3D
B
A
8
7
6
5
4
3
2
1
Y
X
BYTE 4
BYTE 3
BYTE 2
BYTE 1
R3B
R3N
2D-fabric peripheral converts one packet to one read bus cycle (the header is discarded)
RXI
200 MHz
CLK
ADR
DRD
87654321
RNW
REQ
ACK
2D-fabric ready for the next read bus cycle
Address bus only needs decode bits to select 2D-fabric and de-select other peripherals
Read data bus content comes from the data payload

Figure 6 – Conventional read bus cycles receive data packets.

Worst-Case Latency Across Several Processing Nodes

If the worst-case latency for crossing of one processing node is 250 ns, the worst-case latency for the entire transfer chain of two nodes, for example, would amount to 500 ns. Thus, if a packet is launched from a source processor two nodes away from its destination, it will take it a maximum of 500 ns to arrive at its destination processor, regardless of any other data traffic in the system (Figure 8).

Total Latency

Because 2D-fabric appears to processors as if it were memory, and because transfer latency increases with the geographical distance from the source, processors can treat transfer latency as memory wait states for the purpose of scheduling the transfers. In a fully deterministic way, the further you go, the more wait states will be required to complete a transfer (Figure 4).

Although actual latency for the above example is most likely to be closer to the best-case latency of 100 ns, the worst-case latency should always be used when planning data transfers between processors. In some I/O fabrics, worst-case latency can be further reduced by launching packets in specific routing directions to avoid interference with other packets, thus reducing the number of packet delay slots from two to one, or even down to zero.

As shown in Figure 4, 2D-fabric allows processors to easily determine the worst-case transfer latency for any destination inside the chip by simply counting the number of intermediate nodes. 2D-fabric also enables packets to be launched in any one of four possible directions by encoding exit directions in the address field of each data write cycle.

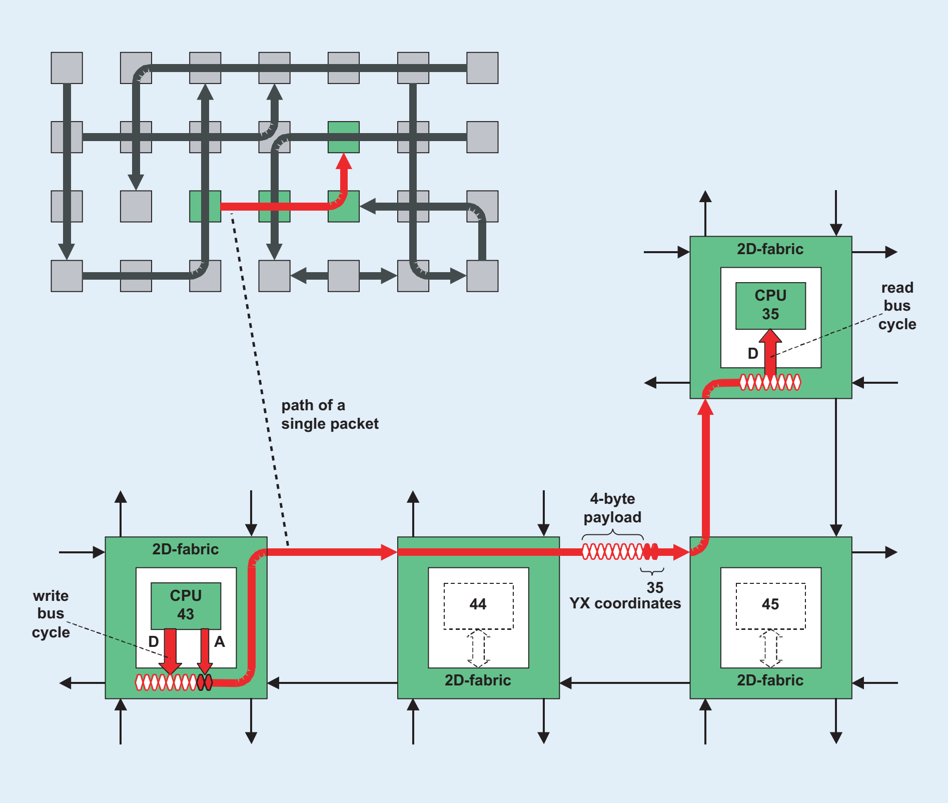

Figure 7 – Non-blocking multiple simultaneous data transfers

Conclusion

Two-dimensional inter-processor interfaces enable fast, easy, and efficient data communications among hundreds of data processing elements of 3G functions implemented inside Virtex-II FPGAs. In addition to 3G, two-dimensional I/O also benefits voice-over-packet, routers, medical imaging, radar, and sonar applications. Linking processing elements with 2D-fabric increases system performance by enabling multiple processors to process data in parallel. At the same time, 2D-fabric reduces power consumption by minimizing the total distance that data has to travel inside chips.

And because it looks to the processors like conventional memory, 2D-fabric does not force system programmers to change their programming methods to benefit from higher performance. Serial programming code investment is preserved, because each processing element has only one processor.

Finally, system designers can now drastically increase processing throughput and I/O bandwidth while retaining current processor architectures and design tools.

For more information on the 2D-fabric parallel-processing interface, go to *www.xilinx.com/products/logicore/alliance/crossbow/crossbow.htm.*

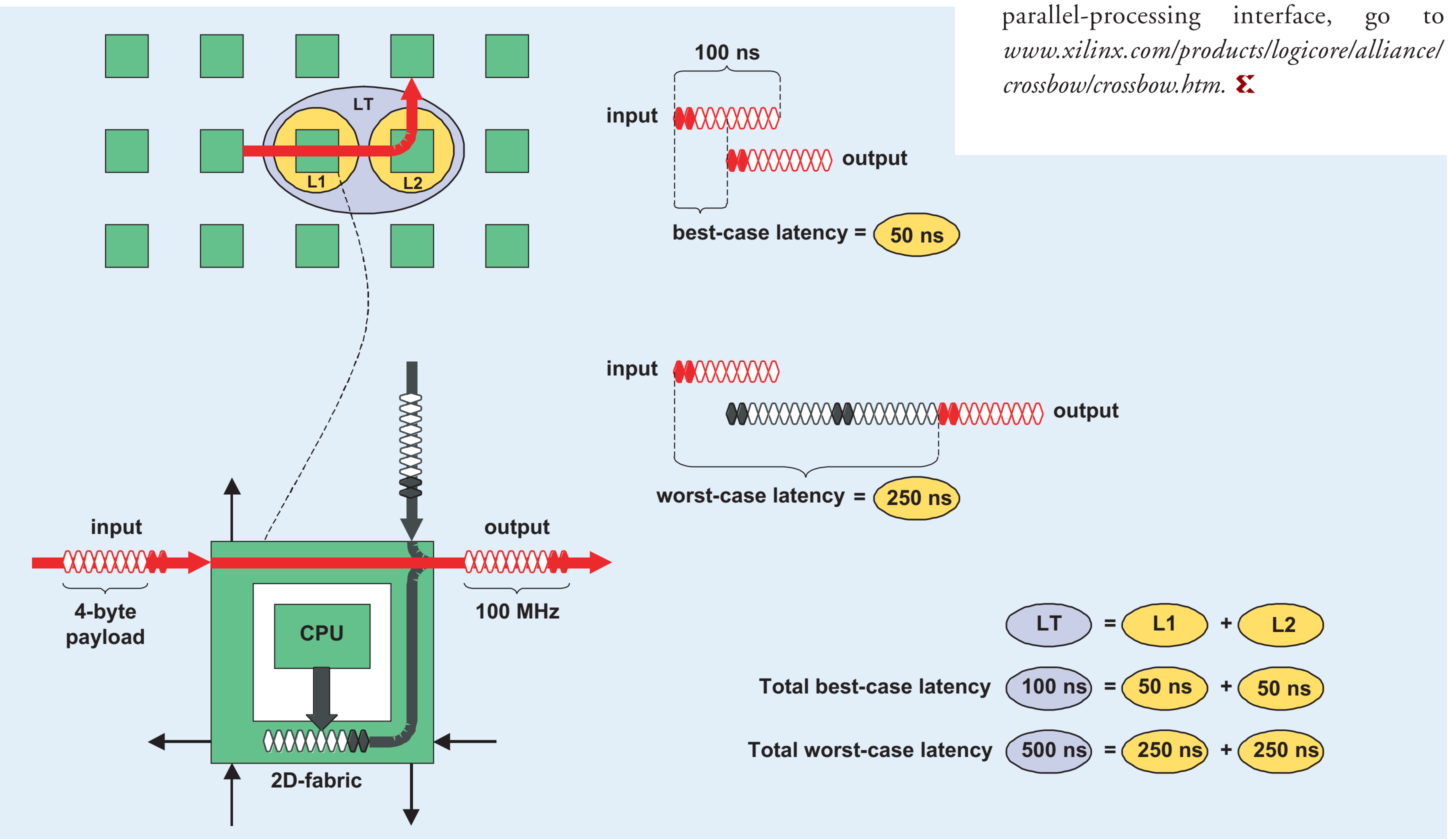
LT
L1
L2
100 ns
input
output
best-case latency = 50 ns
input
output
worst-case latency = 250 ns
input
4-byte
payload
output
100 MHz
CPU
2D-fabric
LT = L1 + L2
Total best-case latency 100 ns = 50 ns + 50 ns
Total worst-case latency 500 ns = 250 ns + 250 ns

Figure 8 – Latency can be determined by counting intermediate nodes.

PCI

AWESOME ASIC Prototyping Solutions

PCI-X

Cool Emulation Platforms

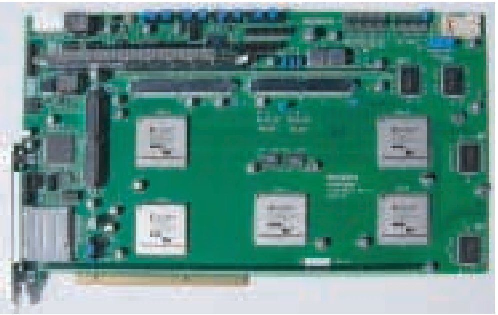

DN3000K10S

Single VirtexII™ PWB

- 2V4000/6000/8000
- 2 PLL's for high speed clock distribution
- 3 512k x 36 SSRAM's
- 1 GB SDRAM DIMM (PC133)
- FPGA configuration using SmartMedia Card
- Sun, Win2000/NT/98, LINUX Drivers and Utilities

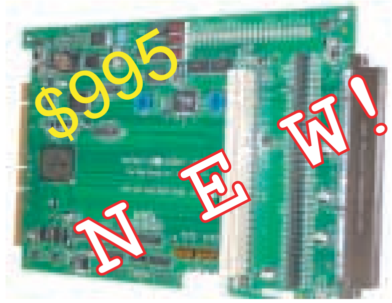

DN3000K10

Logic Emulation PWB with the *Hot* new Xilinx VirtexII™ FPGA's.

- Up to 5 2V8000's per PWB -- 466,000 FF's!
- 2 PLL's for high speed clock distribution
- 4 512k x 36 SSRAM's
- 1 GB SDRAM DIMM (PC133)
- Monstrous amounts of interconnect
- FPGA configuration using SmartMedia Card

PCI-X Extender

64-bit/133MHz Active Extender

- Tundra TSI320 Dual-mode PCIX-to-PCIX bridge
- Active PCI Extender with TWO 32/64-bit slots
- Easy access to all PCI-X signals
- PCI/PCI-X Universal 5V/3.3V I/O
- Secondary PCI Frequencies as low as 10MHz
- Numerous LEDs show system status
- +3.3V not needed on host

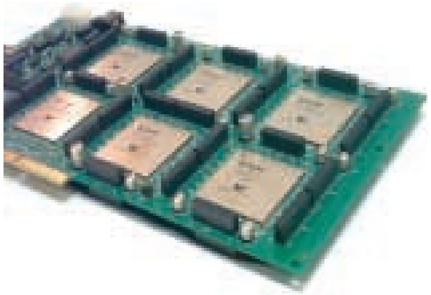

DN2000K10

Logic Emulation PWB based on the Xilinx VirtexE series FPGA's

- Up to 6 Xilinx Virtex or Virtex-E FPGAs
- Stackable to 4 or more PWBs!
- 32/64 bit Universal PCI

Design Consulting Services Available for any ASIC or FPGA projects!

PCI Extender

64-bit/66MHz Active Extender

- Intel® 21154 PCI-to-PCI bridge
- Active PCI Extender with THREE 32/64-bit slots
- Easy access to all PCI signals
- 5V → 3.3V or 3.3V → 5V signal translation
- Frequency division of PCI Clock
- Numerous LEDs show system status
- +3.3V not needed on host

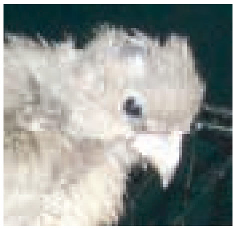

Mr. Freaky Says:

"Cheepest Way To a Working Design!"

- White-Backed Mousebird

A Sample of Our Services:

- FPGA/ASIC/Board Design
- Turnkey Verilog/VHDL Design and Verification
- ASIC → FPGA Design Translation
- HW Emulation and Testing
- Algorithmic Acceleration

Mr. Lazy Says:

"I hate searching for bugs; these guys made my life easy!"

- Blue Naped Mousebird

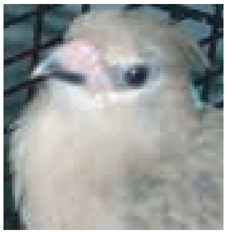

The DINI Group

1010 Pearl Street, Suite 6
La Jolla, CA 92037-5165
http://www.dinigroup.com
Voice: (858) 454-3419
FAX: (858) 454-1728
Email: sales@dinigroup.com

The DiNi Group

400M+ ASIC Gates SOLD!

Bluetooth Wireless Technology Gets BOOST Lite Processor for Virtex FPGAs

For easy and fast access to Bluetooth wireless technology, NewLogic offers the BOOST Lite processor and development board for Xilinx Virtex FPGA products.

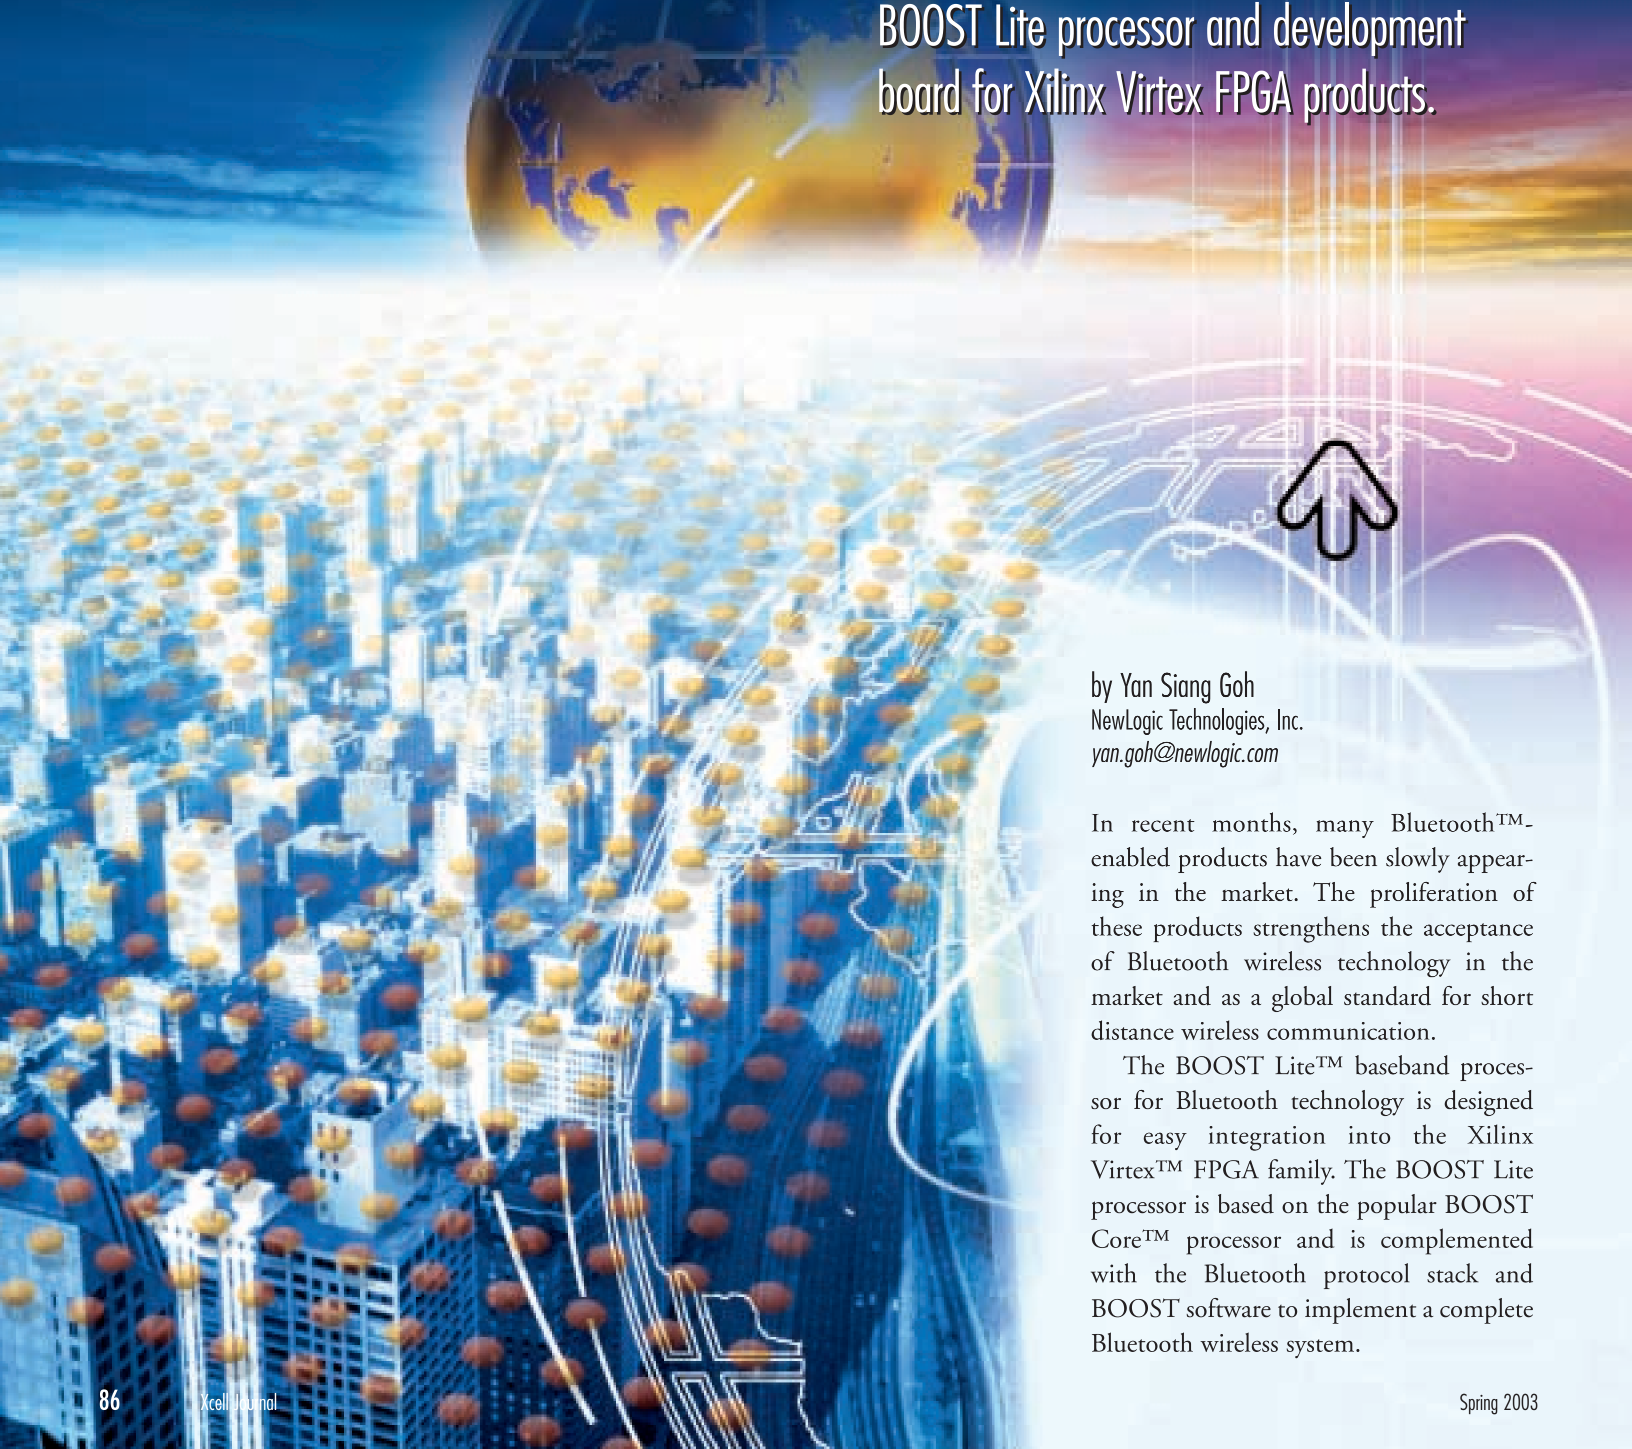

by Yan Siang Goh
NewLogic Technologies, Inc.
yan.goh@newlogic.com

In recent months, many Bluetooth™-enabled products have been slowly appearing in the market. The proliferation of these products strengthens the acceptance of Bluetooth wireless technology in the market and as a global standard for short distance wireless communication.

The BOOST Lite™ baseband processor for Bluetooth technology is designed for easy integration into the Xilinx Virtex™ FPGA family. The BOOST Lite processor is based on the popular BOOST Core™ processor and is complemented with the Bluetooth protocol stack and BOOST software to implement a complete Bluetooth wireless system.

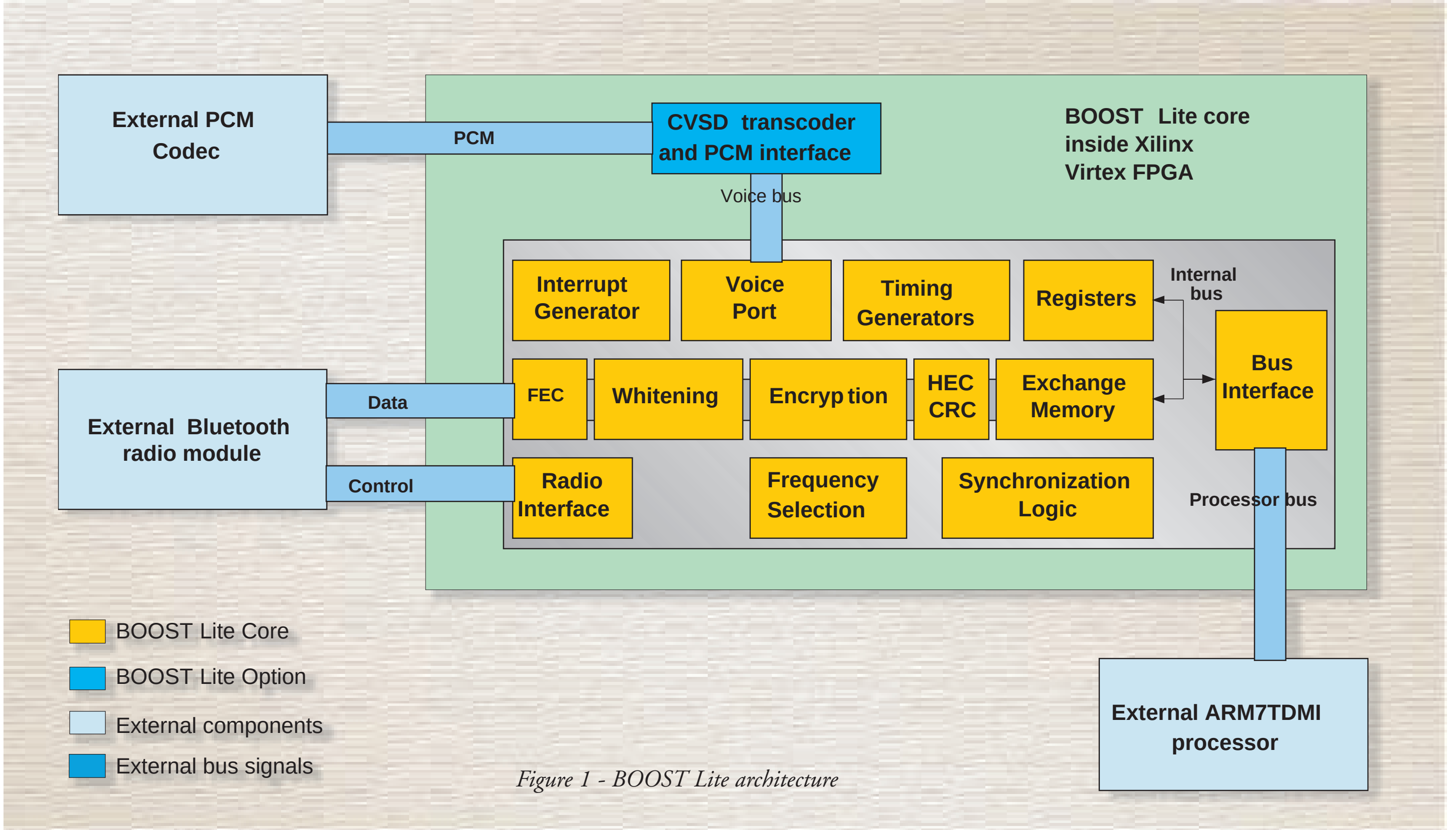

Figure 1 - BOOST Lite architecture

BOOST Lite Core Features

- Based on BQB version 1.1 qualified BOOST Core processor
- Interface to most major Bluetooth radios in the market
- Supports co-existence with 802.11b, piconet, and scatternet operation
- Supports Bluetooth low power modes
- Supports all data and voice packet types
- ARM™ processor AMBA-type interface
- Optimized for the Xilinx Virtex family.

BOOST Lite Architecture

The BOOST Lite core has a fixed bus interface to an external ARM processor and external Bluetooth radio (see Figure 1). Some external RAM and ROM (as well as EPROM, EEPROM, and so on) are necessary to host the BOOST software, which is the Bluetooth protocol stack. The CVSD encoder (available as an option) and a voice coder are necessary to support voice operation. For data applications, it is possible to input/output a data stream from a UART.

The BOOST Lite core interfaces to a fast processor bus or the AMBA-type bus. This bus ensures that data is moved quickly between the processor and the exchange memory, which is accommodated internally on the Xilinx Virtex FPGA.

From the architecture, you can see that the BOOST Lite processor is designed as an “out of the box” Bluetooth wireless solution. The BOOST Lite processor can be integrated with any other third-party system, such as a printer or test equipment core, that requires Bluetooth functionality.

Furthermore, you can use the BOOST Lite processor for fast prototyping before committing to a ASIC/ASSP design cycle. This enables an easy progression from an FPGA prototype to a silicon solution.

NewLogic offers an option for BOOST Lite users to upgrade to the BOOST Core processor as a full, source code version of a Bluetooth baseband processor.

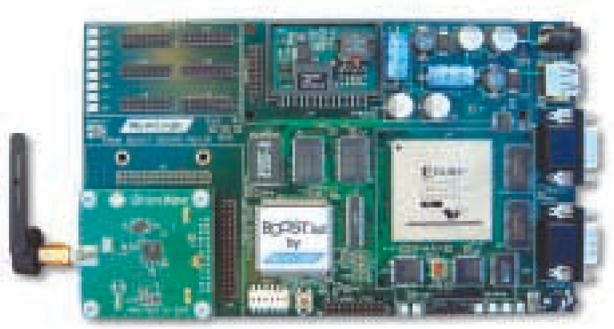

Figure 2 - BOOST Lite development board

A BOOST Lite development board is also available (see Figure 2). The board is supplied with an ARM7TDMI™ processor, external Bluetooth radio, and a Xilinx Virtex FPGA as the BOOST Lite core.

Conclusion

With the availability of BOOST Lite processors for the Virtex FPGA family, the process of designing a Bluetooth-enabled system using Xilinx Virtex FPGAs is greatly simplified. For more information on the BOOST™ family and the WiLD™ 802.11 WLAN technology family, visit NewLogic’s website at *www.newlogic.com*.

Decode MPEG-2 Video with Virtex FPGAs

Amphion's CS6651 video decoder enables the decompression of video streams in real time.

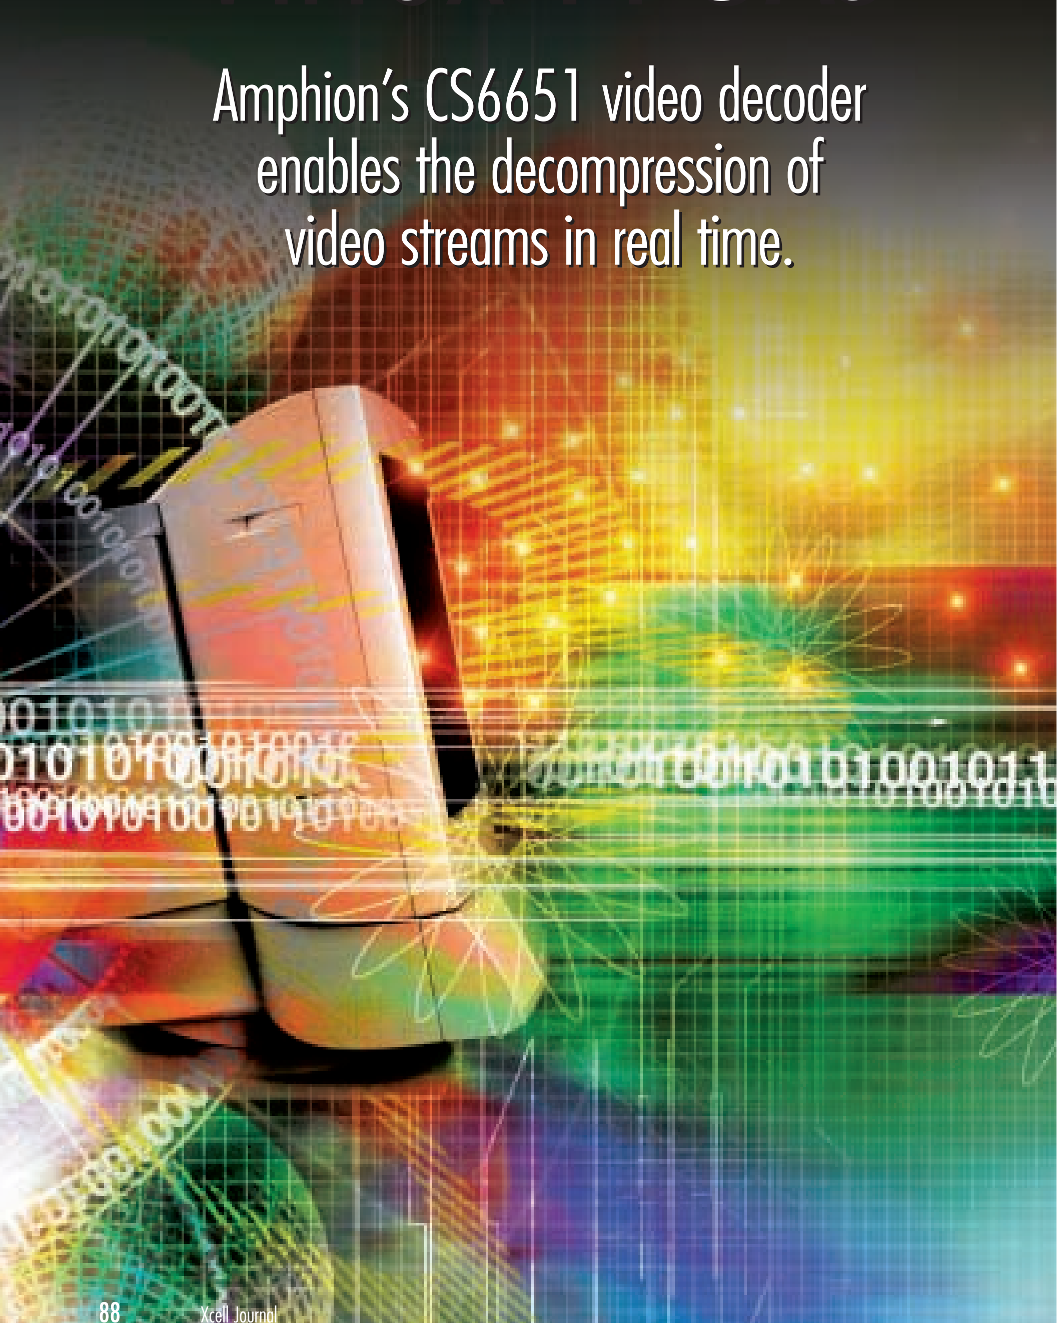

by Rick Richmond
Senior Design Engineer
Amphion
rick@amphion.com

MPEG-2 is the digital video paradigm of today. It is at the heart of the Digital Video Broadcast (DVB) and Advanced Television Systems Committee (ATSC) standard as well as high-definition digital television systems and DVD-video, which has seen incredible market growth in recent years.

The widespread adoption of these applications and systems, coupled with considerable investment by broadcasters and distributors, indicate that MPEG-2 is going to be around for a good while to come – despite the emergence of new, even better video compression algorithms. Future consumer digital applications, in which audio, video, and data networking technologies converge, are certain to need built-in MPEG-2 video capability.

Yet MPEG-2 video is fairly complex and computationally intensive to decode. The main features of the algorithm are discrete cosine transform (DCT)-based compression and motion estimation techniques. Until now, decoder implementations for digital STBs (Set Top Boxes) and DVD-video players had been the domain of ASIC implementations or software running on very powerful processors.

As the sophistication of products like STBs grows, designs will require ever-increasing flexibility and ever-decreasing development time scales. Now, using the Xilinx Virtex™ series of FPGAs and a new intellectual property (IP) core from Amphion, building MPEG-2 video into your designs is simple. Amphion's CS6651 video decoder, which incorporates an integrated external SDRAM memory controller and display direct memory access (DMA), is a great example of the Platform FPGA capability of the Xilinx Virtex series. This solution allows decoding of MP@ML MPEG-2 video with NTSC or PAL frame rates and resolutions.

The MPEG-2 Video Algorithm

MPEG-2 MP@ML video provides a generic video compression solution for applications such as satellite, terrestrial, and cable television (in DVB and ATSC formats), as well as optical storage (DVD-video) at NTSC, PAL, and SECAM resolution and frame rates.

MPEG-2 video sequences are composed of three different types of pictures:

- Intra coded pictures (I-pictures), which are compressed using DCT-based techniques
- Predictive coded pictures (P-pictures), which use motion compensation to predict the current picture from a past reference picture
- Bidirectional predictive coded pictures (B-pictures), which use predictions from past and future reference pictures. B-type pictures are not themselves used as reference pictures.

DCT-based Compression

In MPEG-2 MP@ML video, each picture is broken down into 16x16-sized blocks of luminance samples, called macro blocks. These blocks are further divided into 8x8 blocks. Each macro block has six blocks in total: four luminance blocks and two sub sampled chrominance blocks (a 4:2:0 chrominance format). In the encoding process, a two-dimensional, 8-point DCT is applied to each 8x8 block; the resulting coefficients are quantized using a 64-element quantization matrix.

This process reduces amplitude and increases the number of zero-value coefficients. The quantized DCT coefficients are reordered in a zigzag fashion into a one-dimensional stream, effectively grouping together runs of zero-valued coefficients interspersed with non-zeros. This stream of run-level pairs is then encoded using Huffman-style variable length codes (VLCs) based on a statistical model.

Motion Compensation

In predictive and bidirectionally predictive coded pictures, each macro block may have a number of pairs of motion vectors. These specify the horizontal and vertical displacement from the current position at which the stored reference picture best resembles the current macro block. The difference (if any) between the motion-compensated prediction and the actual image is coded using the DCT-based techniques described above. The sum of the predicted samples and the prediction error give the final reconstructed macro block.

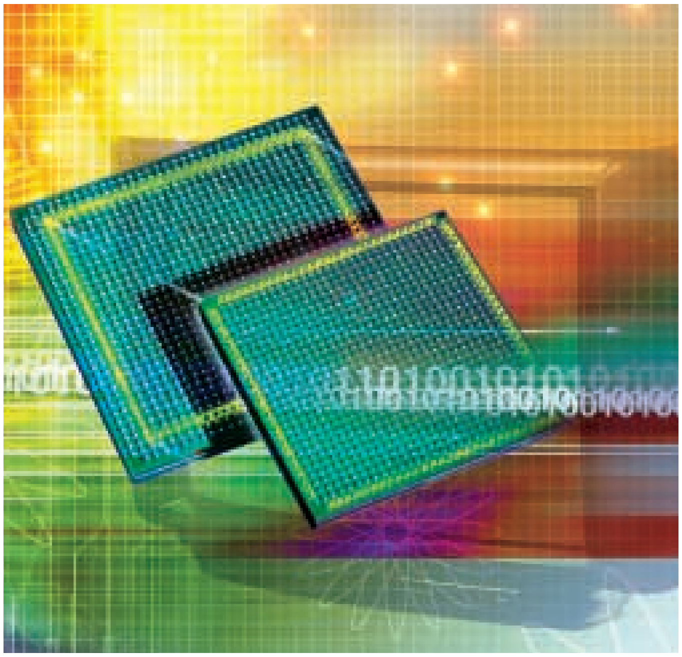

Decoding MPEG-2 Video

As devices and applications grow in complexity, the Amphion CS6651 MPEG-2 video decoder IP core can be implemented by taking advantage of the capabilities of Xilinx Platform FPGAs.

As part of a demonstration system, the core is implemented in a Xilinx Virtex XCV800 device. Including extra interfacing glue logic and PAL/NTSC video encoder driver logic, the core consumes fewer than 8,000 slices and 26 block RAMs. The implementation benefited greatly from the following features of the target device:

- Ample high-performance block RAM
- Fast I/Os
- Extensive logic resources.

The functional blocks and simplified interfaces of the Amphion CS6651 MPEG-2 video decoder IP core are shown in Figure 1. The core requires a minimum clock speed of 27 MHz to maintain MP@ML decoding rates. Video elementary streams are accepted into the core via the byte-wide ES_Data input port on the elementary stream interface.

Parser

The front end of the core is the video elementary stream parser. It searches the syntax of the incoming stream for start codes at which decoding may commence. The parser extracts the various encoding parameters from the headers, which are used to direct subsequent decoding. The remaining variable-length encoded picture data is passed onto the VLC decoder.

VLC Decoder

Here, the Huffman-style VLC picture data is decoded. The outputs of this block include DCT block run-level codes and motion vectors for motion compensation of each macro block.

Run-Level Decoding and Inverse Quantization

The run-level decoder converts run-level codes from the VLC decoder into complete blocks of 64 quantized DCT coefficients. These coefficients are then converted from zigzag scan order to natural row order before being dequantized. Virtex block RAMs are used to support the scan conversion operation and the storage of custom quantization matrices, which may be sent in the elementary stream headers.

Inverse DCT

This unit performs the computationally intensive inverse DCT (IDCT) on the dequantized 8x8 blocks of DCT coefficients. Making use of the high-speed on-chip block RAM, the IDCT unit is capable of streaming data continuously, transforming in 64 clock cycles an 8x8 block of DCT coefficients into an 8x8 block of luminance or chrominance samples or prediction errors.

Motion Compensation and Picture Reconstruction

For each macro block in a P-picture or B-picture, the motion compensation unit takes the decoded motion vectors from the VLC decoder and translates them into row and column coordinates for the prediction samples in the reference picture. The frame store memory then requests these samples and retrieves them via the SDRAM interface. The samples are combined with other prediction samples, if necessary, to complete the prediction for the macro block.

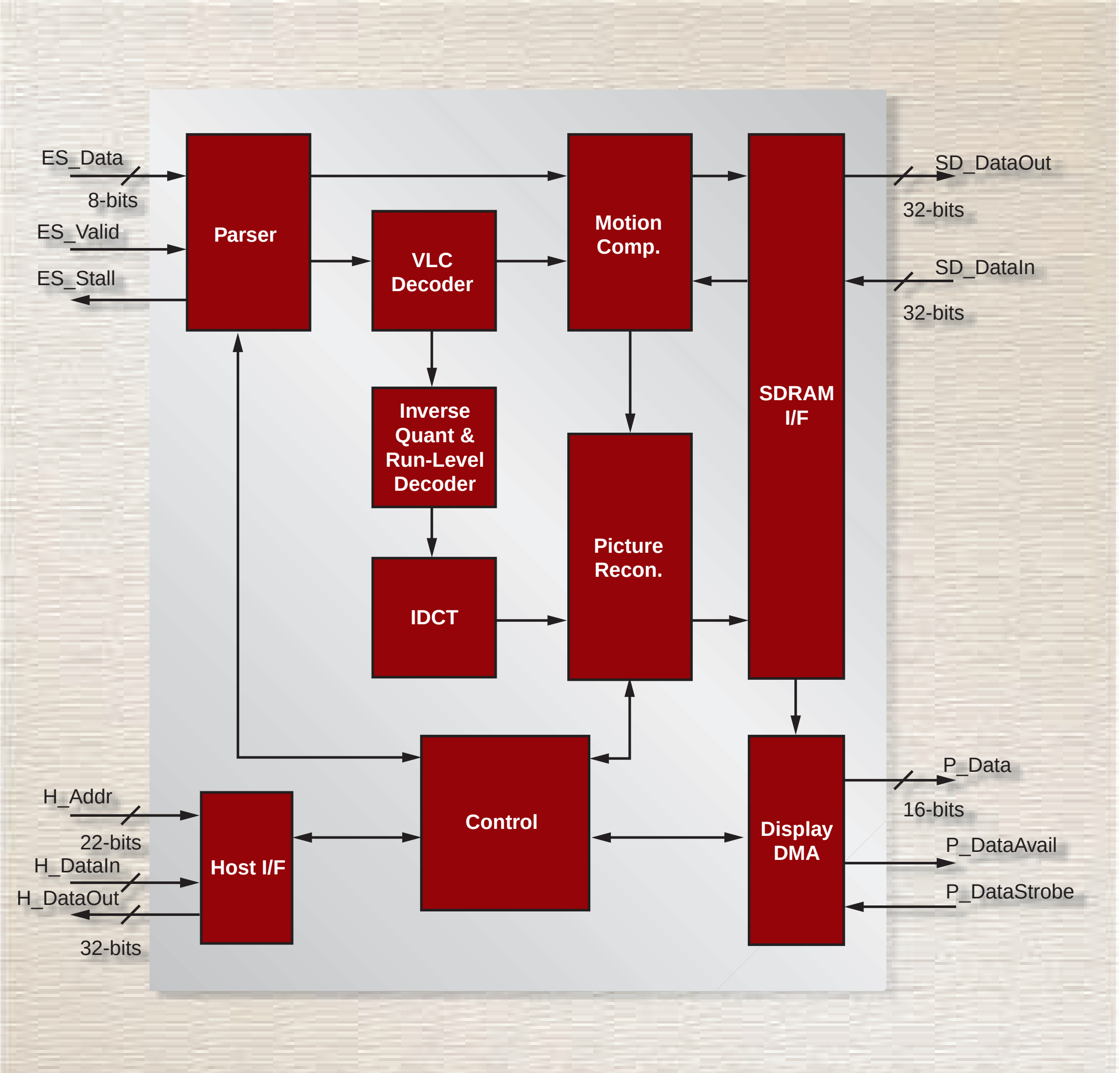

Figure 1 - Simplified block diagram of Amphion CS6651 MPEG-2 video decoder IP core

The final stage of decoding is to add the prediction samples to the prediction error corrections from the inverse DCT unit and write the reconstructed samples into the frame store memory. If a block has no predicted samples, then the samples from the inverse DCT are the final samples and are passed straight through. Both motion compensation and picture reconstruction employ Virtex block RAMs for buffering predictions, samples, or prediction errors and final reconstructed samples.

SDRAM Interface

The frame store memory can be implemented using a commodity PC100 or better 64-Mb SDRAM part. The SDRAM interface handles the mapping of row and column motion compensation prediction requests and reconstructed sample writes into linear memory addresses. Motion compensation places particular memory bandwidth demands on the decoder implementation. To achieve adequate decoding performance, this unit must then arbitrate between the other functions, such as the display DMA, which also access the frame store.

Display DMA

The display DMA unit retrieves decoded samples from the frame-store memory line-by-line for display. This unit has a configurable double-byte output interface for luminance and chrominance samples and can also perform chrominance upsampling to a 4:2:2 format in the vertical direction. This interface provides a number of handshake signals and flags (not shown in Figure 1) to easily allow for the addition of extra logic, to create sync pulses suitable for connecting the decoder to a NTSC or PAL video encoder chip.

Host Interface and Control

Access to internal control, status, and video stream parameter registers within the core is provided via the host interface. Simple 32-bit read/write access to the frame store is also available. In the demonstration system, a host processor controls the core to perform special effects modes such as pause and fast-forward in response to user commands.

Conclusion

MPEG-2 video decoding promises to be an important feature of many future products. In addition to the Amphion CS6651 MPEG-2 video decoder IP core implemented on a single Xilinx Virtex Platform FPGA, Amphion is developing even more sophisticated FPGA-based decoders.

For the latest information about Amphion's advanced decoders, visit *www.amphion.com/video.html.*

Bughunters @ Siemens

Siemens has developed a powerful diagnostic tool for its high-volume telephone switches. Advanced FPGA technology and high-productivity EDA tools from Xilinx have enabled sophisticated test strategies.

by Alfred Fuchs, MSEE
Project Manager
Siemens CES Design Services
www.ces-designservices.com
alfred.fuchs@siemens.com

Gerhard R. Cadek, Pd.D.
EDA Trainer
Oregano Systems
cadek@oregano.at

Most people are unaware that hundreds of microprocessors are involved in placing a cellular telephone call. Although this may be a surprising fact for the average cell phone user, mobile communications experts are even more astonished that such a complex system is stable and seemingly robust.

For despite a well-defined development methodology and quality assurance measures, bugs are sure to exist in any microprocessor code or software. Particularly annoying are sporadic bugs, which show up only under special, rare circumstances at intervals as long as months apart. Typically, such bugs are associated with the dynamics of the multiprocessor system.

Of course, there are plenty of debug tools to analyze software, but applying them may change the dynamics of the system and can therefore mask the problem.

This is where the Siemens Hardware (HW) Tracer comes into its own in identifying software bugs before they can cause problems in the field. Based on Xilinx Virtex™-E FPGAs, HW-Tracer is a standalone, rack-based data capture, analysis, and debug tool (Figure 1).

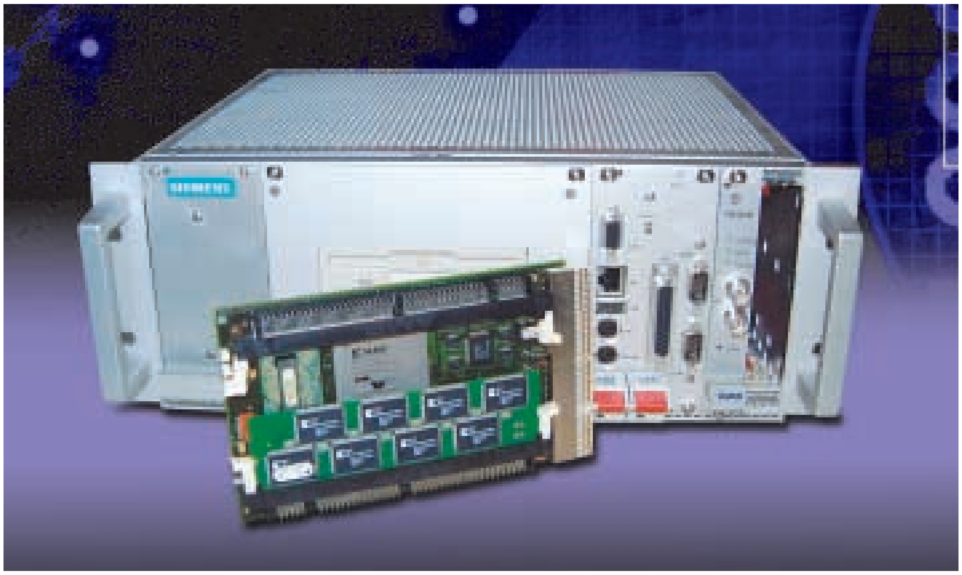

Figure 1 - Siemens "Bughunter" HW-Tracer with Virtex-E XCV2000E FPGA

Hardware Tracer

The HW-Tracer's task is to acquire, filter, and qualify all bus cycles and additional monitor signals. These signals can be accessed via interfaces on the front panels of every processor and memory board and are recorded in trace memory together with a timestamp that is 2 million cycles deep.

The tracer can log the complete activity of the system in real time without interfering with its operation. Moreover, it can also trigger specific events due to a tight coupling with the Coordination Processor's (CP) operating system.

HW-Tracer is essentially a logic analyzer with many proprietary extensions. Because of the number, high speed, and complexity of these extensions, we developed the HW-Tracer as an in-house tool for Siemens. For example, it tracked processes (whose addresses were determined at runtime) and compiled a view containing only program jumps.

To meet the demanding requirements of the new CP generation, HW-Tracer has to cope with a peak data rate of 2 gigabytes per second. Actual cycle rates can even be higher, as data bursts are transferred in a compressed format.

To use the HW-Tracer's two channels independently for processor or memory, we dynamically reconfigured the FPGA that acts as the heart of the tool. Reconfiguring the hardware to complete two different trace tasks halved the cost of the hardware.

The Technical Challenge

A number of key design considerations decided the final hardware implementation. Because of their reprogrammability and dynamic reconfigurability, we selected Virtex-E FPGAs.

The tasks the HW-Tracer had to perform included:

- Analyze nearly 200 input signals at a system speed of 75 MHz
- Capture a high number of 32-bit buses with a data path as wide as 144 bits, which would stress routing resources
- Cope with signal integrity issues on the board level.

Power aspects also had to be taken into account in the overall design.

We conducted a feasibility study before we developed the new HW-Tracer. Specifically, we investigated whether FPGA technology was capable of dealing with the necessary system performance. Our study showed that the Xilinx Virtex-E family of FPGAs was the only technology available that could meet the requirements. We chose Virtex-E FPGAs because:

- Virtex-E devices have up to 804 user I/Os and a 130 MHz internal performance level.
- There were more than enough routing resources to capture the 32-bit busses with widths as wide as 144 bits. With densities ranging from 58K to 4M system gates, a cascade chain for wide-input functions, and dedicated carry logic for high-speed arithmetic functions, the Virtex-E FPGAs were ideal for this task.
- The programmable SelectIO™ standards on Virtex-E devices were able to handle signal integrity issues at the board level, supporting 20 high-performance interface standards with as many as 804 single-ended I/Os or 344 differential I/O pairs.
- The fast, high-density, 1.8V Virtex-E family is designed for low-power operation.

The circuit design took a massive pipelining approach, mathematically speaking. In addition to a design complexity of about 400K gates, which was mainly determined by the logic in the data path, there was also complex control logic to be implemented. A unique feature in

Xilinx FPGAs is a mode called SRL16 (Shift Register LUT), which can be used to increase the effective number of flip-flops per configurable logic block (CLB) by a factor of 16. (Adding flip-flops enables fast pipelining.)

The state machines required to control the operation modes of the HW-Tracer could not be pipelined and threatened to limit the achievable system speed. We have found, however, that arithmetic performance is not the bottleneck in today's leading-edge FPGAs.

Choosing the Design Flow

Despite the extreme performance needed, we decided to code the design in VHDL, which is technology-independent at a high level and controls the implementation via just the constraints in the synthesis and place-and-route tools. We have successfully used this automated design flow already in the past for several FPGA designs. It prevents design issues and saves time when modifying the design.

Only the memory modules had to be instantiated; using the Xilinx CORE Generator™ tool, this proved to be quick and easy to implement.

It was clear that we would have to utilize simulation as if we were developing an ASIC. Simulation (including a gate-level simulation with the VHDL output of Xilinx Alliance Series™ software) added to the overall design time but paid off at the end, when the FPGA was delivered with only two very minor bugs detected.

Striving for Timing Closure

We selected Mentor Graphics' Leonardo Spectrum™ software to synthesize the design and Alliance Series software for the back-end task. Both tools enabled us to implement a completely automated design flow. This was important because in past designs, timing closure could not be guaranteed; code changes tended to result in new worst-case paths that violated the tight timing constraints. We have experienced this timing constraints problem when designing with ASICs.

The most recent releases of the Leonardo Spectrum synthesis tool address this issue successfully by interacting with Xilinx Alliance software, and thus, reducing the number of design iterations. We have seen improvements on the critical paths as high as 15%.

We carried out many implementation runs to incrementally move toward our performance goal. In our project, we started the development based on preliminary timing data; later updates of the timing data entailed some refinements of the constraints.

We benefited a great deal from the stability and performance of Xilinx design software. We achieved turnaround times of three to four hours for the whole implementation task, which is excellent when taking into account the complexity and tough timing constraints.

Management Views

The CP's bring-up-phase depended heavily on the availability of HW-Tracer, which was developed in parallel. To address this risk, we added several additional verification steps to the process.

The test equipment included in the VHDL system simulation comprised:

- Two main memory units, each with a million-gate ASIC containing two embedded CPU-cores
- Several different processing units, each with a million-gate ASIC and two CPUs
- Up to 16 peripheral models and an ATM-controller.

In the context of the system simulation, the CPU's firmware was already verified.

"The verification of the HW-Tracer at the virtual, simulated level ensured that it could be immediately used for testing the physical prototype of the CP unit. This is a great advantage over ASIC designs. Using FPGA simulation tools shortened our design time greatly," said Johann Notbauer, Siemens CES Design Services' technical director for ASIC and FPGA design.

Friedrich Wilhelm, technical director and specialist for proprietary test tools at Siemens, added, "The high logic densities, flexible high-performance interface standards, and pipelining capabilities of Virtex-E combined with the powerful design tools made the choice of FPGA easy."

Emulation

The application software of the HW-Tracer had to be thoroughly verified before its first use. In this environment, the Tracer design was stimulated by a pattern generator, which was included into the FPGA design using internal Virtex-E block memories. The patterns were again derived from the system simulation. In the final product the pattern generator is used for a "learn mode," which helps users familiarize themselves with HW-Tracer before being connected to an actual CP unit for test and debug.

Internet Reconfigurable

In order to facilitate troubleshooting on any of the switches installed throughout the world, the HW-Tracer was designed as a portable device and squeezed into a 3U-CompactPCI chassis.

The FPGAs were loaded by the embedded PC, which took the configuration data from the hard disk. Embedded software as well as FPGA design files are accessible on a dedicated homepage – where future updates and upgrades can be downloaded.

Conclusion

The powerful combination of Mentor Graphics' Leonardo Spectrum synthesis with Alliance software and Virtex-E hardware from Xilinx enabled Siemens to bring its new range of telephone switches to market. Virtex-E FPGAs were used at the heart of the HW-Tracer, which was used to debug and test the Coordination Processor project. The software combination allowed FPGA design iterations to be completed quickly without affecting the tight timing constraints. We chose Virtex-E devices for their high logic densities, high system speeds, flexible system I/O, and ability to perform fast pipelining utilizing the SRL16 mode.

As one customer put it, "The HW-Tracer is the most important debug tool in the CP project."

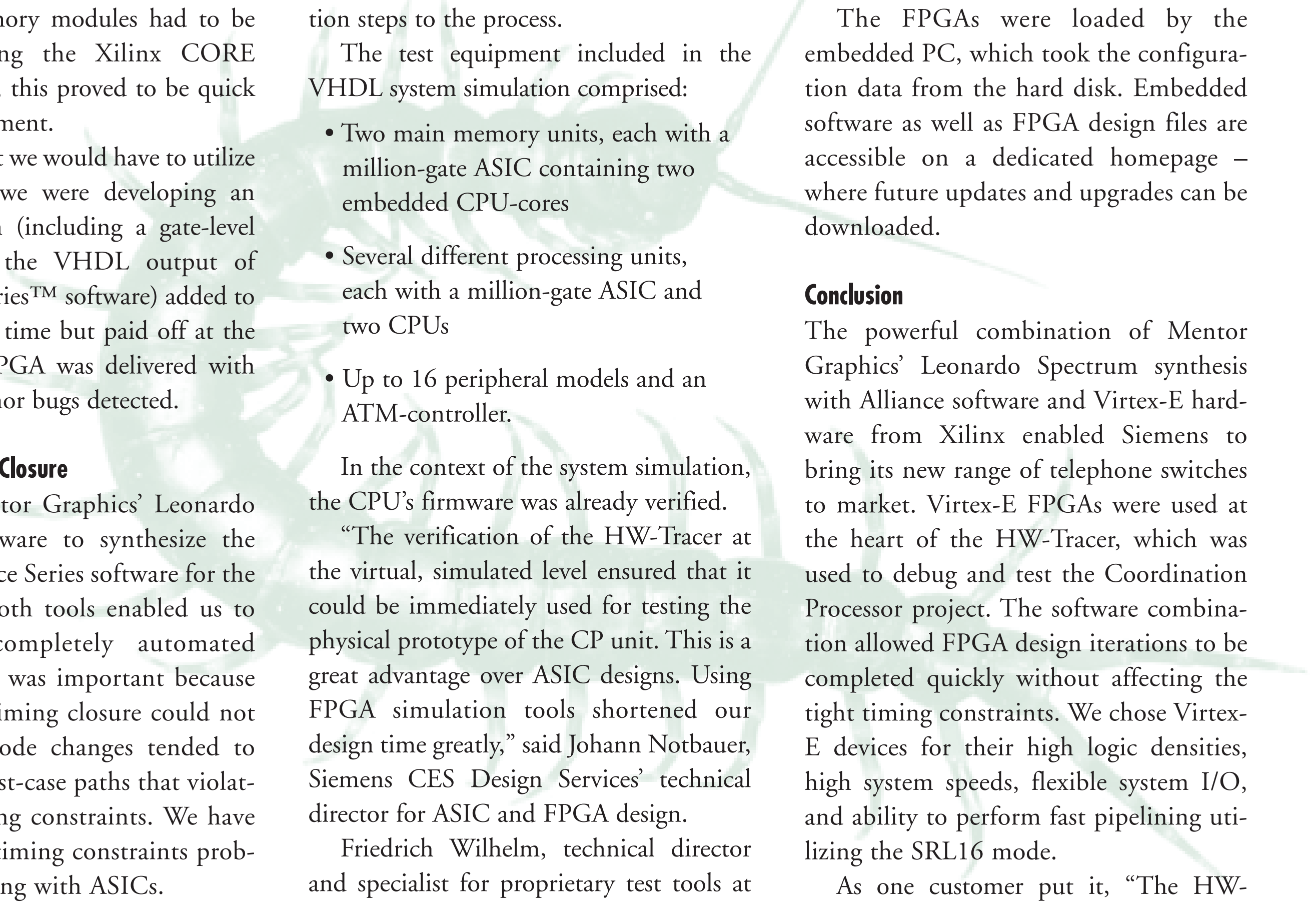

support.xilinx.com: Answers You Need When You Need Them

Need help? Simply go to *support.xilinx.com* on the World Wide Web.

by Doug Horne
Product Marketing & Web Development
Global Services Division
Xilinx, Inc.
Doug.Horne@xilinx.com

A robust, comprehensive, online resource, *support.xilinx.com* stands ready to answer all your design questions 24 hours a day, 7 days a week, 365 days a year (366 in a leap year). With customized access to such features as Answers Search, Software Manuals, Tech Tips, Forums, Problem Solvers, techXclusives, WebCase, Software Updates, and Agents, you can get the answer you need when you need it (Figure 1):

- Answers Databases: Use our advanced search or answer browser tools to easily access the latest answers to technical issues from our huge database of more than 4,000 answers, indexed by logical category.

- Software Manuals: View manuals in PDF, compressed PDF, or in convenient search-enabled HTML format.
- Tech Tips: Get the latest technical information about development tools, device families, interface tools, and Virtex-II Pro™ FPGAs.
- Forums: Collaborate with other designers in discussion groups or chat rooms; join the popular newsgroup *comp.arch.fpga*.
- Problem Solvers: Get instant help with installation and configuration, PCI applications, and JTAG implementation. This interactive tool uses a series of questions to diagnose and troubleshoot your configuration or installation problem automatically, saving you hours of work.
- techXclusives: Expert Xilinx application engineers will keep you informed of the hottest issues and latest techniques.
- WebCase: Use our WebCase to manage hotline cases. You can check status, add notes, and even close a case -on the Web, anytime.

Let's Get Personal

MySupport.xilinx.com puts the powerful tools of *support.xilinx.com* to work for you with an easy-to-use interface that lets you create your own personalized page (Figure 2). Whenever you log in, you will automatically be notified of new downloads, such as service packs and IP update releases, relevant to your design. Stock quotes, bookmarks (for Xilinx and non-Xilinx websites), and Answer Records trails are handy features that add a personal touch to *MySupport.xilinx.com*. Registration is easy and free!

MySupport makes it easy to stay on top (Figure 3):

- Support News: Get the latest in support-related news and developments on *support.xilinx.com*; techXclusives are also featured in this section.

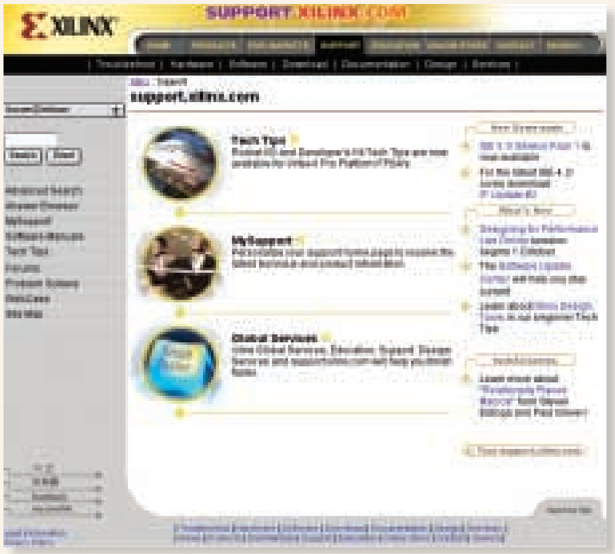

Figure 1 - support.xilinx.com: the best in online support

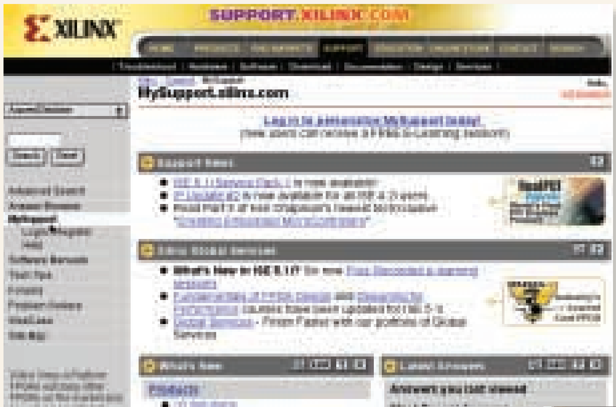

Figure 2 - MySupport.xilinx.com, the information you need when you need it

- Xilinx Global Services: Here you'll find all the latest news on the Xilinx Global Services Division's offerings, designed to help you "finish faster." You can customize your news alerts according to your interests.
- MyAlerts: This personalized service gives you the latest product news and Answer Records pertaining to your specific needs.
- What's New: View the latest products, devices, design tools, and system solutions from the Xilinx product catalog.
- Latest Answers: View the answer records you visited, make note of the most recent answers published by Xilinx engineers, and see what visitors are saying.
- Tech Tips: This is the best resource for hot issues and tips that will get you up and running quickly.
- My Bookmarks: Create links to your favorite websites and Answer Records, or search Queries on *support.xilinx.com*.
- Stock Quotes: Watch your favorite stock ticker symbols, latest price, change, and last trade information.

Do It Today

Take advantage of the powerful tools at your disposal by visiting *support.xilinx.com* today. To sign up for *MySupport.xilinx.com*, visit *mysupport.xilinx.com*, or click on the MySupport logo on the *support.xilinx.com* home page. Questions can be directed to Doug Horne at 408-626-6317 or *doug.horne@xilinx.com*.

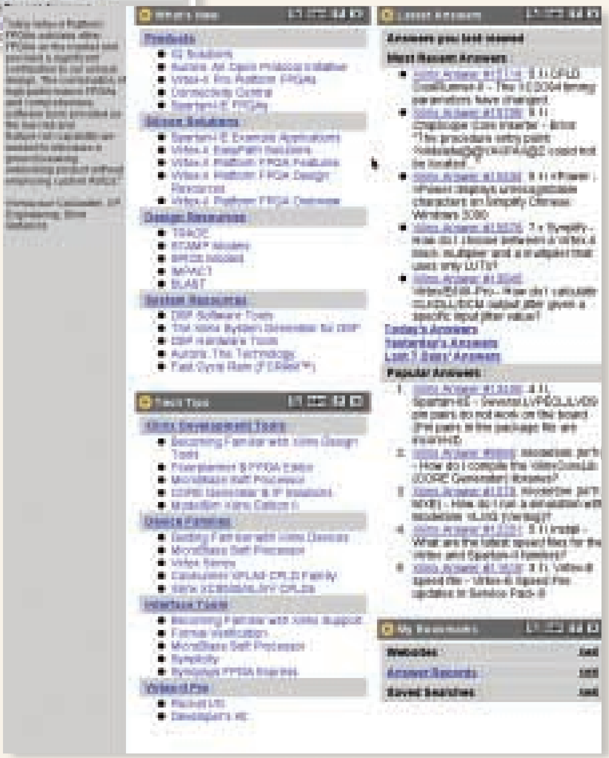

Figure 3 - Keeping track of issues is easy.

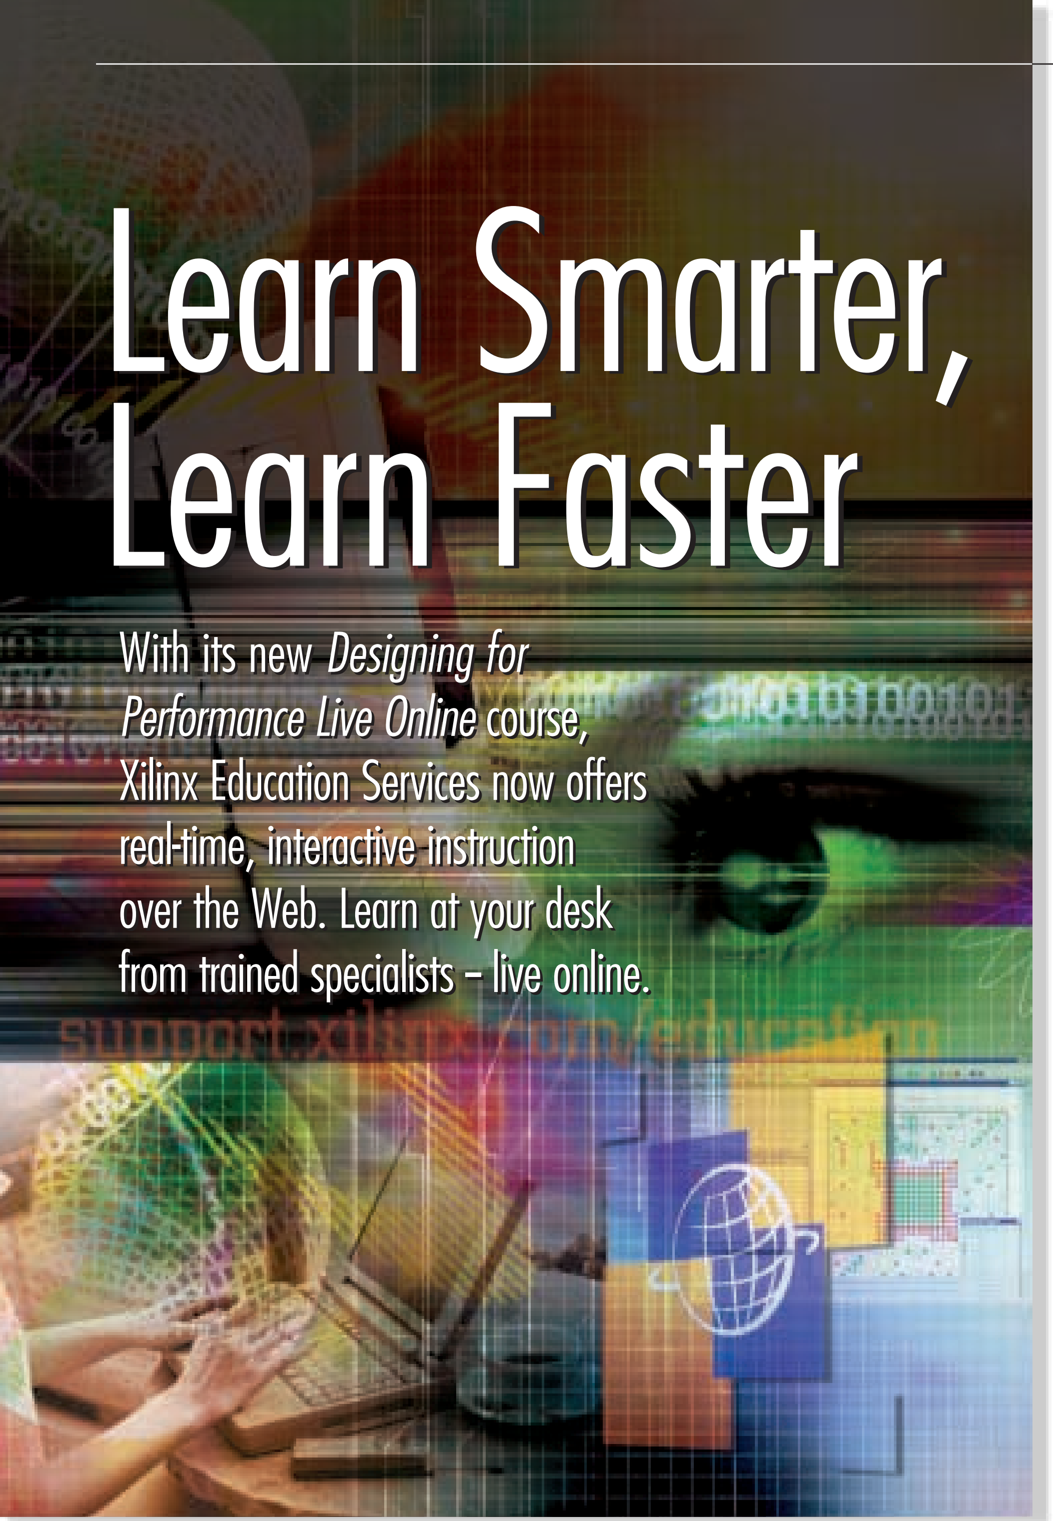

Learn Smarter, Learn Faster

With its new *Designing for Performance Live Online* course, Xilinx Education Services now offers real-time, interactive instruction over the Web. Learn at your desk from trained specialists – live online.

by Cindy Andruss
Technical Writer/Production Editor
Xilinx, Inc.
cindy.andruss@xilinx.com

Reduced training budgets, travel restrictions, and explosive advances in programmable logic technology make it harder than ever to stay on the high side of the learning curve. Combined with demands to do more in less time, your job as a design engineer – or as a manager of a design team – requires an innovative training solution that will save you time and money.

In addition to developing state-of-the-art logic devices and software, Xilinx also continues to pioneer educational opportunities that reduce engineers' time to knowledge and increase their proficiency in using Xilinx FPGA design tools. The latest offering from Xilinx Education Services, *Designing for Performance Live Online*, is an education package that combines the best of live instruction with none of the inconveniences and lost opportunity costs of travel to off-site training centers.

Learning Without Luggage

Design engineer Lukose Ninan was just getting into Xilinx FPGA design when he found out about our first online FPGA training, *e-Series I*. "I wanted a refresher course," said Ninan of ComSonics Inc. "*e-Series I* is exactly what I wanted to get up to speed on Xilinx design, and I got the information I needed without having to travel."

When Mike Schell of Convergent Design Inc. learned about the online program, he enrolled immediately. "The number one reason I prefer e-learning classes is convenience," said Schell, a design consultant. "I have deadlines to meet, and it would be lost time if I had to travel to a class."

Ninan and Schell are among a growing number of engineers and managers who use online e-learning programs to acquire or enhance existing skills (see "Benefits for Both Managers and Engineers"). According to a report from Merrill Lynch, employees in more than half of U.S. corporations used e-learning training programs in 2000. Here at Xilinx, online course enrollments have steadily increased since our initial e-learning offerings began in 2000.

Virtual Education

The *Designing for Performance Live Online* package consists of five one-hour lecture modules and four two-hour lecture-and-lab modules (see "*Designing for Performance Live Online* Modules"). This series of modules was selected from the popular *Designing for Performance* course. Each module is delivered live on the World Wide Web.

The entire package is available for $900 USD or nine training credits, and you can register and pay for the package online. If you prefer, you may purchase modules individually.

The modules are scheduled sequentially, two per week, over a five-week time period. Some students say these smaller units of content provide a more lasting learning effect compared with the average content retention rate for intensive, all-day seminars.

"It's not necessary to have all the information in an eight-hour class," Schell said. "That's what I really like about the program. You have a day or two to absorb the

Benefits for Both Managers and Engineers

As a manager, you will benefit from *Designing for Performance Live Online* if:

- You plan to hire engineers throughout the year and would like to offer this packaged training to new hires as they come on board.
- You do not have a sufficient number of engineers at any one time to warrant an on-site learning program.
- You do not want to send the engineers away for training because of time or budget restrictions.
- You want accessible, consistent, and affordable training for your engineers who have experience with Xilinx ISE software tools but who need to enhance their knowledge of FPGAs.

As a design engineer, you will learn how to:

- Use synchronous design techniques to improve performance.
- Design synchronization circuits to improve design reliability.
- Write HDL code to efficiently target Virtex-II architecture resources.
- Generate customized cores using the CORE Generator™ system.
- Estimate power consumption using the XPower utility.
- Pinpoint design bottlenecks by interpreting timing reports.
- Apply advanced timing constraints to meet your performance objectives.
- Improve design performance by using advanced implementation options.

information before the next session. It's easier to learn that way."

The predetermined schedule lets you lock in your dates ahead of time. If by chance you miss a session, the sequence is repeated every five weeks. You may even begin the series with almost any of the first five modules, but units 6 through 9 are somewhat dependent upon each other and should be taken sequentially. Once you have completed *Designing for Performance Live Online*, you'll receive a Certificate of Completion to add to your record of continuing education credits.

A Classroom at Your Fingertips

As the name promises, *Designing for Performance Live Online* is delivered by a live instructor – no recordings. At the beginning of each interactive session, the instructor will review labs and address any questions you may have before moving on to the lecture.

Real-time, synchronous training sessions create an environment where you can ask questions and hold collaborative discussions. Also, you can use the chat feature to ask questions of other engineers in the class, or pose your questions to the instructor.

All you need to begin *Designing for Performance Live Online* are a Web-enabled computer and a telephone line. You will log on to an online server (provided by Toolwire) to run the lab exercises. Before enrolling in this class, you must pretest your system to ensure that it will perform the labs in the Toolwire virtual environment. Go to *support.xilinx.com/support/training/usingtoolwire.htm* and follow the steps for testing your network, connection speed, firewall compatibility, and installation of the Citrix ICA client software. Once you have pretested your system, you will be able to connect to the Toolwire remote Windows 2000 desktop. If any issues arise, simply contact the registrar at 1-877-959-2527 for support.

After you register for *Designing for Performance Live Online*, Xilinx will send you a series of e-mails containing the URL address and phone number for each session, along with lab requirements and instructions. Ten minutes before each session begins, you can log in and download the presentation materials and lab documentation so you will be prepared for class.

Conclusion

The *Designing for Performance Live Online* series delivers a convenient, advanced training solution for design engineers who have taken the Xilinx Fundamentals course or who have equivalent knowledge of Virtex™-II architecture, software tool flow, and global timing constraints. This *Live Online* course focuses on enhancing your knowledge of the latest Xilinx FPGA design tools and techniques. To learn more about *Designing for Performance Live Online* or other Xilinx e-learning courses, visit the Xilinx Education Services website at *support.xilinx.com/support/education-home.htm* or call the registrar at 877-XLX-CLAS (877-959-2527).

Designing for Performance Live Online Modules*

1. FPGA Design Techniques
2. HDL Coding Style (and lab)
3. Synthesis Techniques (and lab)
4. CORE Generator™ System (and lab)
5. Xpower
6. Achieving Timing Closure
7. Timing Groups and Offset Constraints
8. Path-Specific Timing Constraints (and lab)
9. Advanced Implementation Options

*Modules are subject to change to stay current with emerging technology. For the most up-to-date information, visit the Xilinx Education Services website at *support.xilinx.com/support/education-home.htm.*

Xilinx Technology Enabled Instant Deployment of Real-PCI Express

Xilinx delivered the world's first PCI Express product on the same day the specification was finalized.

By Xilinx Staff

Last July – on the same day the PCI Special Interest Group (SIG) announced the final PCI Express specification – Xilinx delivered the world's first PCI Express intellectual property (IP) core: Real-PCI™ Express. The solution expedited the implementation of PCI Express for Xilinx customers by 12 to 18 months, demonstrating the power of programmable logic over ASIC technology.

PCI Express is the successor to the legacy peripheral component interconnect local bus standard established by Intel. The new PCI Express standard is targeted at the desktop, mobile, server, storage, and embedded communications markets.

"PCI Express takes PCI to another level, with a high-speed, scalable, serial architecture that provides exciting new I/O options for system partitioning designs and form factors," said Tony Pierce, PCI-SIG chairman.

The PCI Express Core

The Xilinx Real-PCI Express core uses the proven RocketIO™ 3.125 Gbps serial transceivers on Xilinx Virtex-II Pro™ FPGAs – the only devices on the market capable of implementing the new specification.

The Real-PCI Express interface can be used to maximize performance and feature quality in high-performance workstations as well as consumer gaming devices. Designers can use the core to design high performance PCI Express systems using Xilinx Smart-IP™ technology to meet critical 2.5 GHz timing requirements. The core reaches a line speed of 2.5 Gbps, utilizing the features of the RocketIO multi-gigabit transceivers, such as clock data recovery, 8B/10B encoding, 3.125 Gbps SerDes, transmit/receive FIFOs, and CRC.

The Logic Advantage

The programmability and serial transceiver capability of Virtex-II Pro FPGAs allowed Xilinx to develop the core simultaneously with the definition of the PCI Express specification as it evolved.

The shipment of the Real-PCI Express core at the same time a final specification was released allowed designers to begin prototyping PCI Express solutions immediately, well ahead of any ASIC-based implementations, according to Cary Snyder, a noted industry expert and analyst.

"By using the programmability and serial transceiver capability of the Virtex-II Pro device, Xilinx was able to develop its core in parallel with its participation in the definition of the specification – a true testament to the capability and benefits of programmable systems," Snyder added.

Always Looking Forward

Real-PCI Express is currently compatible with both the protocol and electrical requirements of the v1.0 base PCI Express specification, but Xilinx is continuing to participate as an active developer in the PCI Express Advanced Switching working group to develop the communications extension for PCI Express. Xilinx plans to incorporate PCI Express with advanced switching into its Virtex™-II series FPGAs.

License Price and Availability

Real-PCI Express is available now as a Xilinx LogiCORE™ product under the terms of SignOnce™ IP license and is priced at $25,000. Once purchased, it may be configured and downloaded from the Xilinx website at *www.xilinx.com/pciexpress/*

For more information and to purchase Virtex-II Pro FPGAs, visit *www.xilinx.com/platform/*

Information about licensing and other Xilinx LogiCORE products is available on the Xilinx IP Center at *www.xilinx.com/ipcenter/*

For complete information about Real-PCI Express, visit *www.xilinx.com/systemio/*

Xilinx Events and Tradeshows

Xilinx participates in numerous trade shows and events throughout the year to help you stay informed. This is a perfect opportunity to meet our silicon and software experts, ask questions, see demonstrations of new products, and discuss the latest trends. You'll meet people just like you and you'll see how they are using programmable logic to solve their technical challenges.

Worldwide Event Schedule

January 28-29	**Platform Conference**	San Jose, CA
January 30-31	**EDSF 2003**	Pacifico, Yokohama
February 17-21	**3GSM**	Cannes, France
February 19-20	**Platform Conference**	Taipei, Taiwan
February 23-25	**FPGA**	Monterey, CA
February 25-27	**Wireless Systems Conference**	San Jose, CA
March 3-11	**IIC China Spring**	Shanghai, Beijing, Shenzhen
March 17	**Synopsys Users Group**	San Jose, CA
April 1-3	**Global Signal Processing Expo**	Dallas, TX
April 8-11	**FCCM**	Napa, CA
April 23-25	**Embedded Systems Conference**	San Francisco, CA
June 16	**Embedded Processor Forum**	San Jose, CA
June 2-4	**40th Design Automation Conference**	Anaheim, CA

For more information about Xilinx Worldwide Events,
please contact one of the following Xilinx team members or see our website at:
www.xilinx.com/events/

For North American shows, contact Jennifer Waibel at: *jennifer.waibel@xilinx.com*
For European shows, contact Andrew Stock at: *andrew.stock@xilinx.com*
For Japanese shows, contact Yumi Homura at: *yumi.homura@xilinx.com*
For Asia Pacific shows, contact Mary Leung at: *mary.leung@xilinx.com*

FUZZY LOGIC

Want to win a new Handspring PDA? Heard any good jokes lately?

In each issue of the *Xcell Journal* we give you many golden nuggets of information about programmable logic. So, we thought it would be fun to bury a few "diamonds" for you to find as well – something to challenge you and give you a diversion from your usual routine. Some of these diamonds are buried deep and are difficult to dig up, some lie there on the surface, easy to find.

Plus, we thought you might enjoy your colleagues' humor.

With a funny joke and a keen eye, you have a good chance of winning one of five Handspring™ Visor Pro™ PDAs.

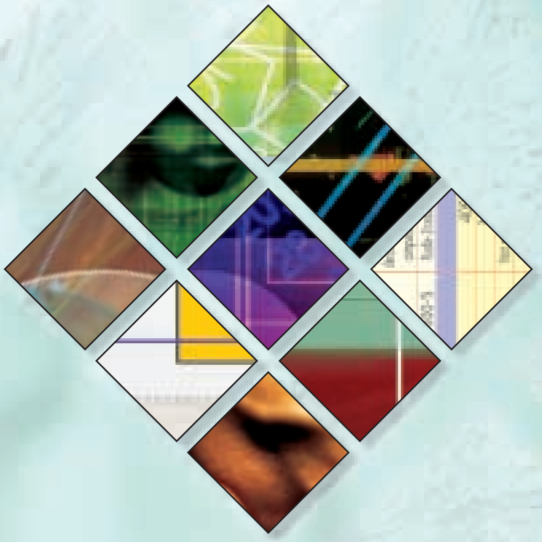

Write down the page number where you find each diamond, add up the nine page numbers, and send us the total along with your joke. You could win a new Handspring PDA!

Contest Rules

There are nine diamonds shown below. Each comes from a page somewhere in this issue of Xcell. Your job is to find the page where each diamond originated. Add up the nine page numbers where you find each diamond, to give you a total - that's your answer. Send us your answer along with your funniest clean joke, (one that is acceptable for printing).

The five winning entries will be the ones with the correct answers and the funniest jokes (as judged by the Xcell editors). The five winners will receive a new Handspring Visor Pro PDA. Only one entry per person; Xilinx employees and contractors, and their families, are not eligible. The entry deadline is May 1, 2003. The winners and their jokes will be announced in the next issue of Xcell.

Send your entry to editor@xilinx.com, with the following words in the subject line:

Xcell Diamonds - 123
(Replace 123 with your answer).

In the body of your e-mail send us your funniest clean joke.

Here's a joke from us to get you started:

There are only 10 kinds of engineers in the world:

Those who understand binary numbers, and those who don't.

Xilinx CPLD Product Selection Matrix

PRODUCT SELECTION MATRIX

						I/O Features		Speed			Clocking	
	System Gates	Macrocells	Product terms per Macrocell	Input Voltage Compatible	Output Voltage Compatible	Max. I/O	I/O Banking	Min. Pin-to-Pin Logic Delay (ns)	Commercial Speed Grades (fastest to slowest)	Industrial Speed Grades (fastest to slowest)	Global Clocks	Product Term Clocks per Function Block
CoolRunner-II Family — 1.8 Volt												
XC2C32	750	32	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	33	1	3.5	-3 -4 -6	-6	3	17
XC2C64	1500	64	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	64	1	4	-4 -5 -7	-7	3	17
XC2C128	3000	128	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	100	2	4.5	-4 -6 -7	-7	3	17
XC2C256	6000	256	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	184	2	5	-5 -6 -7	-7	3	17
XC2C384	9000	384	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	240	4	6	-6 -7 -10	-10	3	17
XC2C512	12000	512	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	270	4	6	-6 -7 -10	-10	3	17
CoolRunner XPLA3 Family — 3.3 Volt												
XCR3032XL	750	32	48	3.3/5	3.3	36		5	-5 -7 -10	-7 -10	4	16
XCR3064XL	1500	64	48	3.3/5	3.3	68		6	-6 -7 -10	-7 -10	4	16
XCR3128XL	3000	128	48	3.3/5	3.3	108		6	-6 -7 -10	-7 -10	4	16
XCR3256XL	6000	256	48	3.3/5	3.3	164		7.5	-7 -10 -12	-10 -12	4	16
XCR3384XL	9000	384	48	3.3/5	3.3	220		7.5	-7 -10 -12	-10 -12	4	16
XCR3512XL	12000	512	48	3.3/5	3.3	260		7.5	-7 -10 -12	-10 -12	4	16
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	3	18
XC9572XV	1600	72	90	2.5/3.3	1.8/2.5/3.3	72	1	5	-5 -7	-7	3	18
XC95144XV	3200	144	90	2.5/3.3	1.8/2.5/3.3	117	2	5	-5 -7	-7	3	18
XC95288XV	6400	288	90	2.5/3.3	1.8/2.5/3.3	192	4	6	-6 -7 -10	-7 -10	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	3	18
XC9572XL	1600	72	90	2.5/3.3/5	2.5/3.3	72		5	-5 -7 -10	-7 -10	3	18
XC95144XL	3200	144	90	2.5/3.3/5	2.5/3.3	117		5	-5 -7 -10	-7 -10	3	18
XC95288XL	6400	288	90	2.5/3.3/5	2.5/3.3	192		6	-6 -7 -10	-7 -10	3	18

PACKAGE OPTIONS AND USER I/O

		CoolRunner-II						CoolRunner XPLA3						XC9500XV				XC9500XL			
	Body Size	XC2C32	XC2C64	XC2C128	XC2C256	XC2C384	XC2C512	XCR3032XL	XCR3064XL	XCR3128XL	XCR3256XL	XCR3384XL	XCR3512XL	XC9536XV	XC9572XV	XC95144XV	XC95288XV	XC9536XL	XC9572XL	XC95144XL	XC95288XL
PLCC Packages (PC)																					
44	16.5 x 16.5 mm	33	33					36	36					34	34			34	34		
PQFP Packages (PQ)																					
208	28 x 28 mm				173	173	173				164	172	180				168				168
VQFP Packages (VQ)																					
44	12 x 12 mm	33	33					36	36					34	34			34	34		
64	12 x 12 mm																	36	52		
100	16 x 16 mm		64	80	80				68	84											
TQFP Packages (TQ)																					
100	14 x 14 mm														72	81			72	81	
144	20 x 20 mm			100	118	118				108	120	118*				117	117			117	117
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 7mm							36	40					36	38			36	38		
144	12 x 12 mm									108						117				117	
280	16 x 16 mm										164						192				192
BGA Packages (BG) — wire-bond standard BGA (1.27 mm ball spacing)																					
256	27 x 27 mm																				192
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 Fineline BGA (1.0 mm ball spacing)																					
256	17 x 17 mm																192				192
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.

Important: Verify all Data with Device Data Sheet and Product Availability with your local Xilinx Rep

🟩 Automotive products are highlighted: -40C to +125C ambient temperature for CPLDs

Xilinx IQ Solutions for Automotive Intelligence

Xilinx Virtex FPGA Product Selection Matrix

			Virtex-II Pro (1.5V)										Virtex-II (1.5V)										
			XC2VP2	XC2VP4	XC2VP7	XC2VP20	XC2VP30	XC2VP40	XC2VP50	XC2VP70	XC2VP100	XC2VP125	XC2V40	XC2V80	XC2V250	XC2V500	XC2V1000	XC2V1500	XC2V2000	XC2V3000	XC2V4000	XC2V6000	XC2V8000
Pins	Body Size	I/O's	204	348	396	564	692	804	852	996	1164	1200	88	120	200	264	432	528	624	720	912	1104	1296
Chip Scale Packages — wire-bond chip-scale BGA (0.8 mm ball spacing)																							
144	12 x 12 mm												88	92	92								
BGA Packages (BG) — wire-bond standard BGA (1.27 mm ball spacing)																							
575	31 x 31 mm																328	392	408				
728	35 x 35 mm																			516			
FGA Packages (FG) — wire-bond fine-pitch BGA (1.0 mm ball spacing)																							
256	17 x 17 mm		140	140									88	120	172	172	172						
456	23 x 23 mm		156	248	248										200	264	324						
676	27 x 27 mm					404	416	416										392	456	484			
FFA Packages (FF) — flip-chip fine-pitch BGA (1.0 mm ball spacing)																							
672	27 x 27 mm		204	348	396																		
896	31 x 31 mm				396	556	556										432	528	624				
1152	35 x 35 mm					564	644	692	692											720	824	824	824
1148*	35 x 35 mm							804	812														
1517	40 x 40 mm								852	964											912	1104	1108
1704	42.5 x 42.5 mm									996	1040	1040											
1696*	42.5 x 42.5 mm										1164	1200											
BFA Packages (BF) — flip-chip fine-pitch BGA (1.27 mm ball spacing)																							
957	40 x 40 mm																		624	684	684	684	

Note: Within the same family, all devices in a particular package are pin-out (footprint) compatible.
Virtex-II packages FG456 and FG676 are also footprint compatible.
Virtex-II packages FF896 and FF1152 are also footprint compatible.
* The FF1148 and FF1696 packages support higher number of user I/O and zero RocketIO™ multi-gigabit transceivers.
Important: Verify all Data with Device Data Sheet (http://www.xilinx.com/partinfo/databook.htm)

Numbers indicated in the matrix are the maximum number of user I/O's for that package and device combination, I/Os for RocketIO MGTs are not included in this table.

	Virtex-II Pro Package Configurations with Available RocketIO Transceiver Blocks									
Package	XC2VP2	XC2VP4	XC2VP7	XC2VP20	XC2VP30	XC2VP40	XC2VP50	XC2VP70	XC2VP100	XC2VP125
FG256	4	4								
FG456	4	4	8							
FG676				8	8	8				
FF672	4	4	8							
FF896			8	8	8					
FF1152				8	8	12	16			
FF1148						0*	0*			
FF1517							16	16		
FF1704								20	20	24
FF1696									0*	0*

Note: * FF1148 and FF1696 packages support higher number of user I/O and zero RocketIO multi-gigabit transceivers

			CLB Resources				Memory Resources			DSP	Clock Resources		I/O Features				Speed							
		System Gates (see note 1)	CLB Array (Row X Col)	Number of Slices	Logic Cells (see note 2)	CLB Flip-Flops	Max. Distributed RAM Bits(kbits)	# 18 kbits Block RAM	Total Block RAM (kbits)	# 18x18 Dedicated Multipliers	DCM Frequency (min/max)	# DCM Blocks (see note 3)	Digitally Controlled Impedance	Maximum Differential I/O Pairs	Max. I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Serial PROM Family	System ACE	Config. Memory (Bits)	RocketIO™ Transceiver Blocks	PowerPC Processor Blocks	Virtex-II Series EasyPath Solution (see note 4)
Platform FPGAs	Virtex-II Pro Family — 1.5 Volt																.13um Nine Layer Copper Process							
	XC2VP2	*	16 x 22	1,408	3,168	2,816	44	12	216	12	24/420	4	YES	100	204	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	-5 -6 -7	-5 -6	ISP/OTP	ISP	1.31M	4	0	
	XC2VP4	*	40 x 22	3,008	6,768	6,016	94	28	504	28	24/420	4	YES	172	348		-5 -6 -7	-5 -6			3.01M	4	1	
	XC2VP7	*	40 x 34	4,928	11,088	9,856	154	44	792	44	24/420	4	YES	196	396		-5 -6 -7	-5 -6			4.49M	8	1	
	XC2VP20	*	56 x 46	9,280	20,880	18,560	290	88	1,584	88	24/420	8	YES	276	564		-5 -6 -7	-5 -6			8.21M	8	2	
	XC2VP30	*	80 x 46	13,696	30,816	27,392	428	136	2,448	136	24/420	8	YES	372	644		-5 -6 -7	-5 -6			11.36M	8	2	✔
	XC2VP40	*	88 x 58	19,392	43,632	38,784	606	192	3,456	192	24/420	8	YES	396	804		-5 -6 -7	-5 -6			15.56M	0** or 12	2	✔
	XC2VP50	*	88 x 70	23,616	53,136	47,232	738	232	4,176	232	24/420	8	YES	420	852		-5 -6 -7	-5 -6			19.02M	0** or 16	2	✔
	XC2VP70	*	104 x 82	33,088	74,448	66,176	1,034	328	5,904	328	24/420	8	YES	492	996		-5 -6 -7	-5 -6			25.60M	16 or 20	2	✔
	XC2VP100	*	120 x 94	44,096	99,216	88,192	1,378	444	7,992	444	24/420	12	YES	572	1,164		-5 -6 -7	-5 -6			33.65M	0** or 20	2	✔
	XC2VP125	*	136 x 106	55,616	125,136	111,232	1,738	556	10,008	556	24/420	12	YES	644	1,200		-5 -6 -7	-5 -6			42.78M	0**,20, or 24	4	✔
	Virtex-II Family — 1.5 Volt																.15um Eight Layer Metal Process							
	XC2V40	40K	8 x 8	256	576	512	8	4	72	4	24/420	4	YES	44	88	LDT-25, LVPECL-33, LVDS-33, LVDS-25, LVDSEXT-33, LVDSEXT-25, BLVDS-25, ULVDS-25, LVTTL, LVCMOS33, LVCMOS25, LVCMOS18, LVCMOS15, PCI33, PCI66, PCI-X, GTL, GTL+, HSTL I, HSTL II, HSTL III, HSTL IV, SSTL2I, SSTL2II, SSTL3 I, SSTL3 II, AGP, AGP-2X	-4 -5 -6	-4 -5	ISP/OTP	ISP	0.4M			
	XC2V80	80K	16 x8	512	1,152	1,024	16	8	144	8	24/420	4	YES	60	120		-4 -5 -6	-4 -5			0.6M			
	XC2V250	250K	24 x16	1,536	3,456	3,072	48	24	432	24	24/420	8	YES	100	200		-4 -5 -6	-4 -5			1.7M			
	XC2V500	500K	32 x 24	3,072	6,912	6,144	96	32	576	32	24/420	8	YES	132	264		-4 -5 -6	-4 -5			2.8M			
	XC2V1000	1M	40 x 32	5,120	11,520	10,240	160	40	720	40	24/420	8	YES	216	432		-4 -5 -6	-4 -5			4.1M			
	XC2V1500	1.5M	48 x 40	7,680	17,280	15,360	240	48	864	48	24/420	8	YES	264	528		-4 -5 -6	-4 -5			5.7M			
	XC2V2000	2M	56 x 48	10,752	24,192	21,504	336	56	1,008	56	24/420	8	YES	312	624		-4 -5 -6	-4 -5			7.5M			
	XC2V3000	3M	64 x 56	14,336	32,256	28,672	448	96	1,728	96	24/420	12	YES	360	720		-4 -5 -6	-4 -5			10.5M			✔
	XC2V4000	4M	80 x 72	23,040	51,840	46,080	720	120	2,160	120	24/420	12	YES	456	912		-4 -5 -6	-4 -5			15.7M			✔
	XC2V6000	6M	96 x 88	33,792	76,032	67,584	1,056	144	2,592	144	24/420	12	YES	552	1,104		-4 -5 -6	-4 -5			21.9M			✔
	XC2V8000	8M	112 x 104	46,592	104,832	93,184	1,456	168	3,024	168	24/420	12	YES	554	1,108		-4 -5				29.1M			✔

Note: 1. System Gates include 20-30% of CLBs used as RAM
2. Logic cell = (1) 4 Input (LUT) Look Up Table + Flip Flop + Carry Logic.
3. DCM – Digital Clock Management
4. Virtex-II Series EasyPath solution available to provide a no risk, no effort cost reduction path for volume production.
* System gate count not meaningful for Virtex-II Pro devices with immersed special blocks such as PowerPC processors and multi-gigabit transceivers.
** The FF1148 and FF1696 packages support higher number of user I/O and zero RocketIO multi-gigabit transceivers.
Important: Verify all Data with Device Data Sheet (http://www.xilinx.com/partinfo/databook.htm)

VIRTEX-II PRO EASYPATH SOLUTIONS

VIRTEX-II EASYPATH SOLUTIONS

Xilinx Spartan FPGAs

PRODUCT SELECTION MATRIX

		CLB Resources					BLK RAM		CLK Resources					I/O Features				Speed				
	System Gates (see note 1)	CLB Array (Row X Col)	Number of Slices	Logic Cells (see note 2)	CLB Flip-Flops	Max. Distributed RAM Bits	# Block RAM	Block RAM (kbits)	# Dedicated Multipliers	DLL Frequency (min/max)	# DLL's	Frequency Synthesis	Phase Shift	Digitally Controlled Impedance	Number of Differential I/O Pairs	Max. I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Serial PROM Family		Config. Memory (Bits)
Spartan-IIE Family — 1.8 Volt																		**.18/.15um Six Layer Metal Process**				
XC2S50E	50K	16 x 24	768	1,728	1,536	24K	8	32K	NA	25/320	4	YES	YES	NA	83	182	LVTTL,LVCMOS2, LVCMOS18, PCI33, PCI66, GTL, GTL+, HSTL I, HSTL III, HSTL IV, SSTL3 I, SSTL3 II, SSTL2 I, SSTL2 II, AGP-2X, CTT, LVDS, BLVDS, LVPECL	-6 -7	-6	ISP	OTP	0.6M
XC2S100E	100K	20 x 30	1,200	2,700	2,400	37K	10	40K	NA	25/320	4	YES	YES	NA	86	202		-6 -7	-6			0.9M
XC2S150E	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	NA	25/320	4	YES	YES	NA	114	265		-6 -7	-6			1.1M
XC2S200E	200K	28 x 42	2,352	5,292	4,704	73K	14	56K	NA	25/320	4	YES	YES	NA	120	289		-6 -7	-6			1.4M
XC2S300E	300K	32 x 48	3,072	6,912	6,144	96K	16	64K	NA	25/320	4	YES	YES	NA	120	329		-6 -7	-6			1.9M
XC2S400E	400K	40 x 60	4,800	10,800	9,600	150K	40	160K	NA	25/320	4	YES	YES	NA	172	410		-6 -7	-6			2.7M
XC2S600E	600K	48 x 72	6,912	15,552	13,824	216K	72	288K	NA	25/320	4	YES	YES	NA	205	514		-6 -7	-6			4.0M
Spartan-II Family — 2.5 Volt																		**.22/.18um Six Layer Metal Process**				
XC2S15	15K	8 x 12	192	432	384	6K	4	16K	NA	25/200	4	YES	YES	NA	NA	86	LVTTL, LVCMOS2, PCI33 (3.3V & 5V), PCI66 (3.3V), GTL, GTL+, HSTL I, HSTL III, HSTL IV, SSTL3 I, SSTL3 II, SSTL2 I, SSTL2 II, AGP-2X, CTT	-5 -6	-5	ISP	OTP	0.2M
XC2S30	30K	12 x 18	432	972	864	13.5K	6	24K	NA	25/200	4	YES	YES	NA	NA	132		-5 -6	-5			0.4M
XC2S50	50K	16 x 24	768	1,728	1,536	24K	8	32K	NA	25/200	4	YES	YES	NA	NA	176		-5 -6	-5			0.6M
XC2S100	100K	20 x 30	1,200	2,700	2,400	37.5K	10	40K	NA	25/200	4	YES	YES	NA	NA	196		-5 -6	-5			0.8M
XC2S150	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	NA	25/200	4	YES	YES	NA	NA	260		-5 -6	-5			1.1M
XC2S200	200K	28 x 42	2,352	5,292	4,704	73.5K	14	56K	NA	25/200	4	YES	YES	NA	NA	284		-5 -6	-5			1.4M
Spartan-XL Family — 3.3 Volt																						
XCS05XL	5K	10 x 10	100	238	200	3.1K	NA	NA	NA	NA	NA	NA	NA	NA	NA	77	TTL, LVTTL, CMOS, LVMOS, PCI	-4 -5	-4	ISP	OTP	0.05M
XCS10XL	10K	14 x 14	196	466	392	6.1K	NA	NA	NA	NA	NA	NA	NA	NA	NA	112		-4 -5	-4			0.09M
XCS20XL	20K	20 x 20	400	950	800	12.5K	NA	NA	NA	NA	NA	NA	NA	NA	NA	160		-4 -5	-4			0.18M
XCS30XL	30K	24 x 24	576	1,368	1,152	18.0K	NA	NA	NA	NA	NA	NA	NA	NA	NA	192		-4 -5	-4			0.25M
XCS40XL	40K	28 x 28	784	1,862	1,568	24.5K	NA	NA	NA	NA	NA	NA	NA	NA	NA	224		-4 -5	-4			0.33M

PACKAGE OPTIONS AND USER I/O

			Spartan-IIE (1.8V)							Spartan-II (2.5V)						Spartan-XL (3.3V)				
			XC2S50E	XC2S100E	XC2S150E	XC2S200E	XC2S300E	XC2S400E	XC2S600E	XC2S15	XC2S30	XC2S50	XC2S100	XC2S150	XC2S200	XCS05XL	XCS10XL	XCS20XL	XCS30XL	XCS40XL
Pins	**Body Size**	**I/O's**	182	202	265	289	329	410	514	86	132	176	196	260	284	77	112	160	192	224
PLCC Packages																				
84	30 x 30 mm															61	61			
PQFP Packages (PQ)																				
208	28 x 28 mm		146	146	146	146	146				132	140	140	140	140			160	169	169
240	32 x 32 mm																		192	192
VQFP Packages (VQ)																				
100	14 x 14 mm									60	60					77	77	77	77	
TQFP Packages (TQ)																				
144	20 x 20 mm		102	102						86	92	92	92				112	113	113	
Chip Scale Packages — wire-bond chip-scale BGA (0.8 mm ball spacing)																				
144	12 x 12 mm									86	92						112	113		
280	16 x 16 mm																		192	224
FGA Packages (FT) — wire-bond fine-pitch thin BGA (1.0 mm ball spacing)																				
256	17 x 17 mm		182	182	182	182	182	182												
FGA Packages (FG) — wire-bond fine-pitch BGA (1.0 mm ball spacing)																				
256	17 x 17 mm											176	176	176	176					
456	23 x 23 mm			202	265	289	329	329	329				196	260	284					
676	27 x 27 mm							410	514											
BGA Packages																				
256	27 x 27 mm																		192	205

Note: 1. System Gates include 20-30% of CLBs used as RAM
2. Logic Cell is defined as a 4 input LUT and a register

Important: Verify all Data with Device Data Sheet (http://www.xilinx.com/spartan)

Numbers indicated in the matrix are the maximum number of user I/O's for that package and device combination.

🟩

Automotive products are highlighted:
-40C to +125C junction temperature for FPGAs

Xilinx IQ Solutions for Automotive Intelligence

Xilinx IQ Solutions

Part Number	Speed	Package	Voltage	Description
XC9500XL CPLD				
XC9536XL	10 ns/100 MHz	VQ44, VQ64	3.3V	36 Macrocells (800 Gates), ISP, JTAG, Bus Hold & I/P Hysteresis
XC9572XL	10 ns/100 MHz	VQ64, TQ100	3.3V	72 Macrocells (1,600 Gates), ISP, JTAG, Bus Hold & I/P Hysteresis
CoolRunner XPLA3				
XCR3032XL	10 ns/100 MHz	VQ44	3.3V	32 Macrocells (800 Gates), Low Power, Slew Rate Control, ISP & JTAG
XCR3064XL	10 ns/100 MHz	VQ44, VQ100	3.3V	64 Macrocells (1,600 Gates), Low Power, Slew Rate Control, ISP & JTAG
XCR3128XL	10 ns/100 MHz	VQ100, TQ144	3.3V	128 Macrocells (3,200 Gates), Low Power, Slew Rate Control, ISP & JTAG
XCR3256XL	10 ns/100 MHz	TQ144, PQ208	3.3V	256 Macrocells (6,400 Gates), Low Power, Slew Rate Control, ISP & JTAG
XCR3384XL	10 ns/100 MHz	PQ208	3.3V	384 Macrocells (9,600 Gates), Low Power, Slew Rate Control, ISP & JTAG
XCR3512XL	10 ns/100 MHz	PQ208	3.3V	512 Macrocells (12,800 Gates), Low Power, Slew Rate Control, ISP & JTAG
CoolRunner-II				
XC2C32	6 ns/145 MHz	VQ44	1.8V	32 Macrocells (800 Gates), 6 I/O Standards, Slew Rate Control, Clock Doubler, Bus Hold, I/P Hysteresis. Ultra low power.
XC2C64	7.5 ns/127 MHz	VQ44, VQ100	1.8V	64 Macrocells (1,600 Gates), 6 I/O Standards, Slew Rate Control, Clock Doubler, Bus Hold, I/P Hysteresis. Ultra low power.
XC2C128	7.5 ns/127 MHz	VQ44, VQ100	1.8V	128 Macrocells (3,200 Gates), 9 I/O Standards, Slew Rate Control, Clock Doubler, Clcok Divider, CoolClock, DataGate, Bus Hold, I/P Hysteresis. Ultra low power.
XC2C256	7.5 ns/127 MHz	VQ100, TQ144	1.8V	256 Macrocells (6,400 Gates), 9 I/O Standards, Slew Rate Control, Clock Doubler, Clcok Divider, CoolClock, DataGate, Bus Hold, I/P Hysteresis. Ultra low power.
XC2C384	10 ns/100 MHz	TQ144, PQ208	1.8V	384 Macrocells (9,600 Gates), 9 I/O Standards, Slew Rate Control, Clock Doubler, Clcok Divider, CoolClock, DataGate, Bus Hold, I/P Hysteresis. Ultra low power.
XC2C512	10 ns/100 MHz	PQ208	1.8V	512 Macrocells (12,800 Gates), 9 I/O Standards, Slew Rate Control, Clock Doubler, Clcok Divider, CoolClock, DataGate, Bus Hold, I/P Hysteresis. Ultra low power.

Note: See page 96 for CPLD IQ devices Package Options and User I/O.

Part Number	Speed Grade	Package	Voltage	Description
Spartan-XL				
XCS05XL	-4	VQ100	3.3V	Low cost FPGA with power down pin, 5V tol I/O, 5,000 Gate, 238 logic cells, 100 CLBs.
XCS10XL	-4	VQ100	3.3V	Low cost FPGA with power down pin, 5V tol I/O, 10,000 Gate, 466 logic cells, 196 CLBs.
XCS20XL	-4	TQ144, PQ208	3.3V	Low cost FPGA with power down pin, 5V tol I/O, 20,000 Gate, 950 logic cells, 400 CLBs.
XCS30XL	-4	TQ144, PQ208	3.3V	Low cost FPGA with power down pin, 5V tol I/O, 30,000 Gate, 1,368 logic cells, 576 CLBs.
XCS40XL	-4	PQ208, BG256	3.3V	Low cost FPGA with power down pin, 5V tol I/O, 40,000 Gate, 1,862 logic cells, 784 CLBs.
Spartan-II				
XC2S15	-5	TQ144	2.5V	High volume FPGA, on-chip RAM, 16 I/O standards, 15,000 Gate, 432 logic cells, 96 CLBs, 4 block RAM blocks, 4 DLLS.
XC2S30	-5	TQ144, PQ208	2.5V	High volume FPGA, on-chip RAM, 16 I/O standards, 30,000 Gate, 972 logic cells, 216 CLBs, 6 block RAM blocks, 4 DLLS.
XC2S50	-5	TQ144, PQ208, FG256	2.5V	High volume FPGA, on-chip RAM, 16 I/O standards, 50,000 Gate, 1,728 logic cells, 384 CLBs, 8 block RAM blocks, 4 DLLs.
XC2S100	-5	TQ144, PQ208, FG256	2.5V	High volume FPGA, on-chip RAM, 16 I/O standards, 100,000 Gate, 2,700 logic cells, 600 CLBs, 10 block RAM blocks, 4 DLLs.
XC2S150	-5	PQ208, FG256	2.5V	High volume FPGA, on-chip RAM, 16 I/O standards, 150,000 Gate, 3,888 logic cells, 864 CLBs, 12 block RAM blocks, 4 DLLs.
XC2S200	-5	PQ208, FG456	2.5V	High volume FPGA, on-chip RAM, 16 I/O standards, 200,000 Gate, 5,292 logic cells, 1,176 CLBs, 14 block RAM blocks, 4 DLLs.
Spartan-IIE				
XC2S50E	-6	TQ144, PQ208, FT256	1.8V	High volume FPGA, on-chip RAM, 19 I/O standards, 50,000 Gate, 1,728 logic cells, 384 CLBs, 8 block RAM blocks, 4 DLLs.
XC2S100E	-6	TQ144, PQ208, FT256	1.8V	High volume FPGA, on-chip RAM, 19 I/O standards, 100,000 Gate, 2,700 logic cells, 600 CLBs, 10 block RAM blocks, 4 DLLs.
XC2S150E	-6	PQ208, FT256	1.8V	High volume FPGA, on-chip RAM, 19 I/O standards, 150,000 Gate, 3,888 logic cells, 864 CLBs, 12 block RAM blocks , 4 DLLs.
XC2S200E	-6	PQ208, FT256	1.8V	High volume FPGA, on-chip RAM, 19 I/O standards, 200,000 Gate, 5,292 logic cells, 1,176 CLBs, 14 block RAM blocks, 4 DLLs.
XC2S300E	-6	PQ208, FG456	1.8V	High volume FPGA, on-chip RAM, 19 I/O standards, 300,000 Gate, 6,912 logic cells, 1,536 CLBs, 16 block RAM blocks, 4 DLLs.
XC2S400E	-6	FT256, FG456, FG676	1.8V	High volume FPGA, on-chip RAM,19 I/O standards, 400,000 Gate,10,800 logic cells, 2,400 CLBs, 40 block RAM blocks, 4DLLs.
XC2S600E	-6	FG456, FG676	1.8V	High volume FPGA, on-chip RAM,19 I/O standards, 600,000 Gate,10,800 logic cells, 3,456 CLBs, 72block RAM blocks, 4DLLs.

Note: See page 93 for Spartan IQ devices Package Options and User I/O.

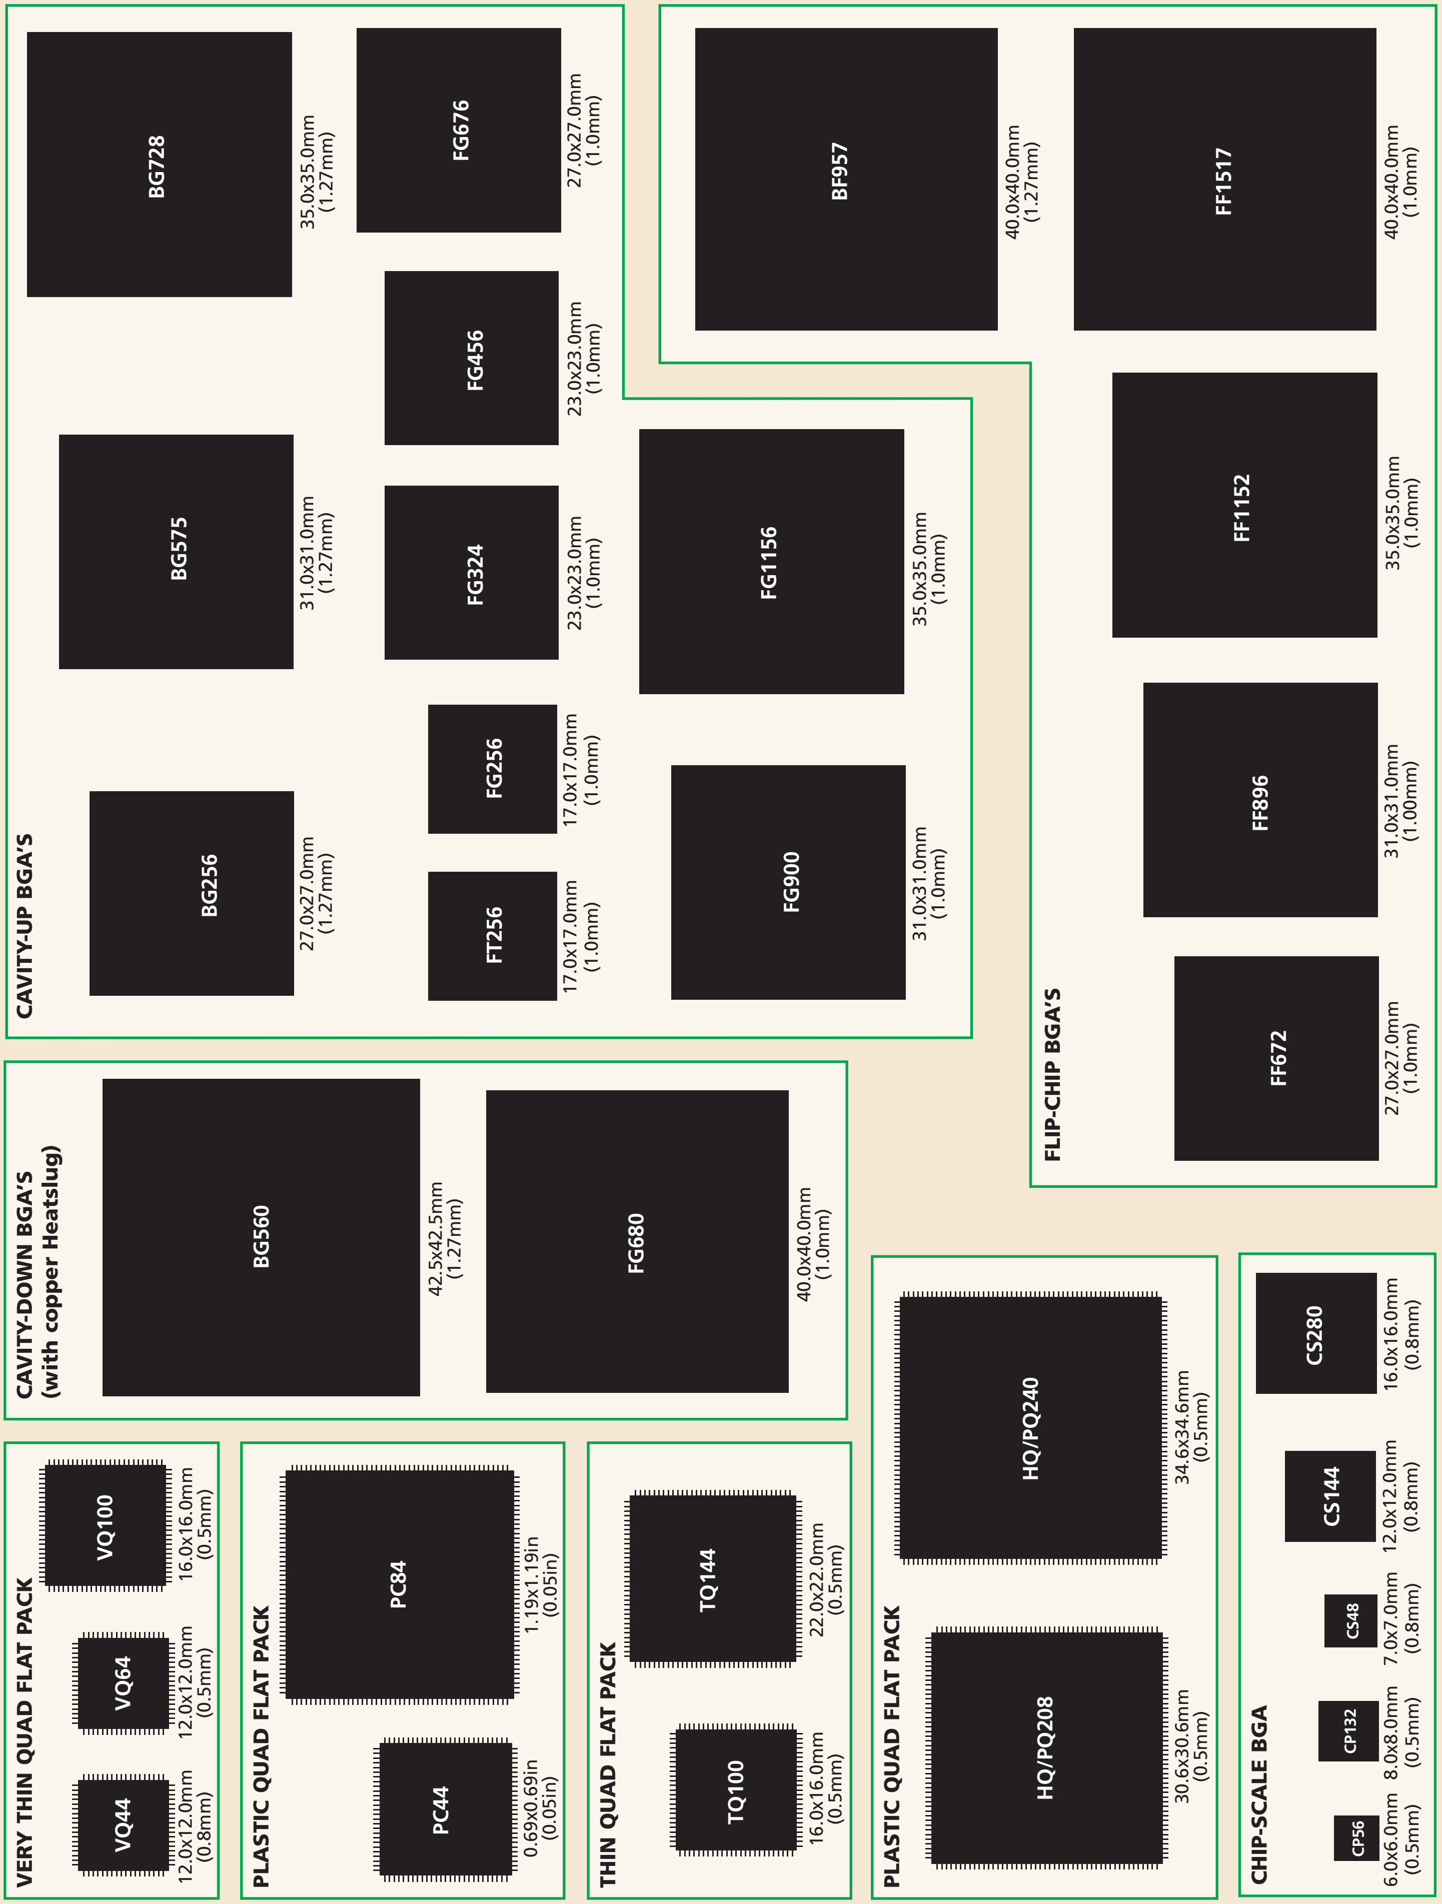
CAVITY-UP BGA'S
BG728
35.0x35.0mm
(1.27mm)
FG676
27.0x27.0mm
(1.0mm)
FG456
23.0x23.0mm
(1.0mm)
BG575
31.0x31.0mm
(1.27mm)
FG324
23.0x23.0mm
(1.0mm)
FG1156
35.0x35.0mm
(1.0mm)
FG256
17.0x17.0mm
(1.0mm)
BG256
27.0x27.0mm
(1.27mm)
FT256
17.0x17.0mm
(1.0mm)
FG900
31.0x31.0mm
(1.0mm)
FLIP-CHIP BGA'S
BF957
40.0x40.0mm
(1.27mm)
FF1517
40.0x40.0mm
(1.0mm)
FF1152
35.0x35.0mm
(1.0mm)
FF896
31.0x31.0mm
(1.00mm)
FF672
27.0x27.0mm
(1.0mm)
CAVITY-DOWN BGA'S
(with copper Heatslug)
BG560
42.5x42.5mm
(1.27mm)
FG680
40.0x40.0mm
(1.0mm)
VERY THIN QUAD FLAT PACK
VQ100
16.0x16.0mm
(0.5mm)
VQ64
12.0x12.0mm
(0.5mm)
VQ44
12.0x12.0mm
(0.8mm)
PLASTIC QUAD FLAT PACK
PC84
1.19x1.19in
(0.05in)
PC44
0.69x0.69in
(0.05in)
THIN QUAD FLAT PACK
TQ144
22.0x22.0mm
(0.5mm)
TQ100
16.0x16.0mm
(0.5mm)
PLASTIC QUAD FLAT PACK
HQ/PQ240
34.6x34.6mm
(0.5mm)
HQ/PQ208
30.6x30.6mm
(0.5mm)
CHIP-SCALE BGA
CS280
16.0x16.0mm
(0.8mm)
CS144
12.0x12.0mm
(0.8mm)
CS48
7.0x7.0mm
(0.8mm)
CP132
8.0x8.0mm
(0.5mm)
CP56
6.0x6.0mm
(0.5mm)

Xilinx Configuration Storage Solutions

System ACE	Memory Density	Number of Components	Min board space	Compression	FPGA Config. Mode	Multiple Designs	Software Storage	Removable	IRL Hooks	Max Config. Speed	Non-Volatile Media
SystemACE CF	up to 8 Gbit	2	25 cm^2	No	JTAG	Unlimited	Yes	Yes	Yes	30 Mbit/sec	CompactFlash
SystemACE MPM	16 Mbit 32 Mbit 64 Mbit	1	12.25 cm^2	Yes	SelectMAP (up to 4 FPGA) Slave-Serial (up to 8 FPGA chains)	Up to 8	No	No	Yes	152 Mbit/sec	AMD Flash Memory
SystemACE SC	16 Mbit 32 Mbit 64 Mbit	3	Custom	Yes	SelectMAP (up to 4 FPGA) Slave-Serial (up to 8 FPGA chains)	Up to 8	No	No	Yes	152 Mbit/sec	AMD Flash memory

PROM	Density	PD8	VO8	SO20	PC20	PC44	VQ44	Core Voltage	I/O Voltage	
									2.5V	3.3V
In-System Programming (ISP) Configuration PROMs										
XC18V256	256Kb			Y	Y		Y	3.3V	Y	Y
XC18V512	512Kb			Y	Y		Y	3.3V	Y	Y
XC18V01	1Mb			Y	Y		Y	3.3V	Y	Y
XC18V02	2Mb					Y	Y	3.3V	Y	Y
XC18V04	4Mb					Y	Y	3.3V	Y	Y
One-Time Programmable (OTP) Configuration PROMs										
XC17V01	1.6Mb		Y	Y	Y			3.3V	Y	Y
XC17V02	2Mb				Y	Y	Y	3.3V	Y	Y
XC17V04	4Mb				Y	Y	Y	3.3V	Y	Y
XC17V08	8Mb					Y	Y	3.3V	Y	Y
XC17V16	16Mb					Y	Y	3.3V	Y	Y

FPGA	PROM Solution	PD8	VO8	SO20	PC20	PC44	VQ44	Core Voltage	I/O Voltage	
									2.5V	3.3V
OTP Configuration PROMs for Spartan-II/IIE										
XC2S50E	XC17S50A	Y	Y	Y				3.3V	Y	Y
XC2S100E	XC17S100A	Y	Y	Y				3.3V	Y	Y
XC2S150E	XC17S200A	Y	Y				Y	3.3V	Y	Y
XC2S200E	XC17S200A	Y	Y				Y	3.3V	Y	Y
XC2S300E	XC17S300A						Y	3.3V	Y	Y
XC2S15	XC17S15A	Y	Y	Y				3.3V	Y	Y
XC2S30	XC17S30A	Y	Y	Y				3.3V	Y	Y
XC2S50	XC17S50A	Y	Y	Y				3.3V	Y	Y
XC2S100	XC17S100A	Y	Y	Y				3.3V	Y	Y
XC2S150	XC17S150A	Y	Y	Y				3.3V	Y	Y
XC2S200	XC17S200A	Y	Y				Y	3.3V	Y	Y

FPGA	PROM Solution	PD8	VO8	SO20	PC20	PC44	VQ44	Core Voltage	I/O Voltage	
									3.3V	5.0V
OTP Configuration PROMs for Spartan-XL										
XCS05XL	XC17S05XL	Y	Y					3.3V	Y	Y
XCS10XL	XC17S10XL	Y	Y					3.3V	Y	Y
XCS20XL	XC17S20XL	Y	Y					3.3V	Y	Y
XCS30XL	XC17S30XL	Y	Y					3.3V	Y	Y
XCS40XL	XC17S40XL	Y		Y				3.3V	Y	Y

Xilinx Home Page
http://www.xilinx.com

Xilinx Online Support
http://www.xilinx.com/support/support.htm

Xilinx IP Center
http://www.xilinx.com/ipcenter/index.htm

Xilinx Education Center
http://www.xilinx.com/support/education-home.htm

Xilinx Tutorial Center
http://www.xilinx.com/support/techsup/tutorials/index.htm

Xilinx WebPACK
http://www.xilinx.com/sxpresso/webpack.htm

Xilinx Software

	Feature	ISE *WebPACK*	ISE *BaseX*	ISE *Foundation*	ISE *Alliance*
Devices	Virtex™Series	**Virtex-E:** V50E – V300E **Virtex-II:** 2V40 – 2V250 **Virtex-II Pro:** 2VP2	**Virtex:** V50 – V300 **Virtex-E:** V50E – V300E **Virtex-II:** 2V40 – 2V250 **Virtex-II Pro:** 2VP2	ALL	ALL
	Spartan™ II/IIE Families	ALL (except XC2S400E and XC2S600E)	ALL	ALL	ALL
	CoolRunner™ XPLA3 / CoolRunner-II	ALL	ALL	ALL	ALL
	XC9500™ Series	ALL	ALL	ALL	ALL
Services	Educational Services	Yes	Yes	Yes	Yes
	Design Services	Sold as an Option	Sold as an Option	Sold as an Option	Sold as an Option
	Support Services	Web Only	Yes	Yes	Yes
Design Entry	Schematic Editor	Yes	Yes	PC Only	No
	HDL Editor	Yes	Yes	Yes	Yes
	State Diagram Editor	Yes	Yes	PC Only	No
	CORE Generator System	No	Yes	Yes	Yes
	PACE (Pinout and Area Constraint Editor)	Yes	Yes	Yes	Yes
	Architecture Wizards DCM – Digital Clock Management MGT – Multi-Gigabit Transcievers	No	Yes	Yes	Yes
	3rd Party RTL Checker Support	Yes	Yes	Yes	Yes
	Xilinx System Generator for DSP	No	Sold as an Option	Sold as an Option	Sold as an Option
Embedded System Design	GNU Embedded Tools GCC – GNU Compiler GDB – GNU Software Debugger	Yes	Yes	Yes	Yes
	WindRiver Xilinx Edition Development Tools Diab C/C++ Compiler SingleStep Debugger visionPROBE II target connection	No	Sold as an Option	Sold as an Option	Sold as an Option
Synthesis	Xilinx Synthesis Technology (XST)	Yes	Yes	Yes	No
	Synplicity Synplify/Pro	Integrated Interface	Integrated Interface	Integrated Interface (PC Only)	Integrated Interface (PC Only)
	Synplicity Amplify Physical Synthesis Support	Yes	Yes	Yes	Yes
	Leonardo Spectrum	Integrated Interface	Integrated Interface	Integrated Interface	Integrated Interface
	Synopsys FPGA Compiler II	EDIF Interface	EDIF Interface	EDIF Interface	EDIF Interface
	ABEL	CPLD	CPLD	CPLD (PC Only)	No
Implementation	iMPACT	Yes	Yes	Yes	Yes
	FloorPlanner	Yes	Yes	Yes	Yes
	Xilinx Constraints Editor	Yes	Yes	Yes	Yes
	Timing Driven Place & Route	Yes	Yes	Yes	Yes
	System ACE Configuration Manager	Yes	Yes	Yes	Yes
	Modular Design	Yes	Yes	Yes	Yes
	Timing Improvement Wizard	Yes	Yes	Yes	Yes
Board Level Integration	IBIS Models	Yes	Yes	Yes	Yes
	STAMP Models	Yes	Yes	Yes	Yes
	LMG SmartModels	Yes (Available from Synopsys)	Yes (Available from Synopsys)	Yes (Available from Synopsys)	Yes (Available from Synopsys)
	HSPICE Models*	Yes	Yes	Yes	Yes
Verification	HDL Bencher™	Yes	Yes	PC Only	No
	ModelSim® Xilinx Edition (MXE II)	ModelSim XE II Starter**	ModelSim XE II Starter**	ModelSim XE II Starter**	ModelSim XE II Starter**
	Static Timing Analyzer	Yes	Yes	Yes	Yes
	ChipScope PRO	No	Sold as an Option	Sold as an Option	Sold as an Option
	FPGA Editor with Probe	No	Yes	Yes	Yes
	ChipViewer	Yes	Yes	Yes	Yes
	XPower (Power Analysis)	Yes	Yes	Yes	Yes
	3rd Party Equivalency Checking Support	Yes	Yes	Yes	Yes
	SMARTModels for PPC and Rocket I/O	No	Yes	Yes	Yes
	3rd Party Simulator Support	Yes	Yes	Yes	Yes
Platforms		PC (MS Windows 2000/MS Windows XP)	PC Only (MS Windows 2000/MS Windows XP)	PC (MS Windows 2000/MS Windows XP)Sun Solaris, Linux	PC (MS Windows 2000/MS Windows XP)Sun Solaris, Linux
IP/CORE	For more information on the complete list of Xilinx IP products, visit the Xilinx IP Center at http://www.xilinx.com/ipcenter				

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

** MXE II supports the simulation of designs up to 1 million system gates and is sold as an option. For more information, visit the Xilinx Design Tools Center at **www.xilinx.com/ise**

Xilinx Global Services

Xilinx Global Services

support.xilinx.com

- Sign up for personalized email alerts @mysupport.xilinx.com
- Search our knowledge database
- Troubleshoot with Problem Solvers
- Consult with engineers in Forums

Xilinx Productivity Advantage Program (XPA)

The XPA provides customers with everything they need to improve their designs — Software, Education and Support Services, IP cores and demo boards — in one package. A single purchase order allows their designers to get what they need, when they need it, without the worry of ordering each piece of the solution separately.

Two types of XPAs are available; custom XPA and Prepacked "XPA Seat". "XPA Seat" is primarily for individuals or small design teams.

http://support.xilinx.com/support/gsd/xpa_program.htm

Xilinx Design Services (XDS)

- XDS provides extensive FPGA hardware and embedded software design experience backed by industry recognized experts and resources to solve even the most complex design challenge.
- **System Architecture Consulting** — Provide engineering services to define system architecture and partitioning for design specification.
- **Custom Design Solutions** — Project designed, verified, and delivered to mutually agreed upon design specifications.
- **IP Core Development, Optimization, Integration, Modification, and Verification** — Modify, integrate, and optimize customer intellectual property or third party cores to work with Xilinx technology. Develop customer-required special features to Xilinx IP cores or third party cores. Perform integration, optimization, and verification of IP cores in Xilinx technology.
- **Embedded Software** — Develop complex embedded software with real-time constraints, using hardware/software co-design techniques.
- **Conversions** — Convert ASIC designs and other FPGAs to Xilinx technology and devices.

Part Number	Product Description	Duration	Availability
Education Services		**hours**	
FPGA13000-5-ILT	Fundamentals of FPGA Design	8	Now
FPGA23000-5-ILT	Designing for Performance	16	Now
FPGA33000-5-ILT	Advanced FPGA Implementation	16	Now
LANG11000-5-ILT	Introduction to VHDL	24	Now
LANG21000-5-ILT	Advanced VHDL	16	Now
LANG12000-5-ILT	Introduction to Verilog	24	Now
PCI8000-4-ILT	PCI CORE Basics	8	Now
PCI28000-4-ILT	Designing a PCI System	16	Now
ASIC25000-5-ILT	Designing for Performance for the ASIC User	24	Now
DSP2000-3-ILT	DSP Implementation Techniques for Xilinx FPGAs	24	Now
DSP-10000-5-ILT	DSP Design Flow	24	Now
RIO22000-5-ILT	Designing with RocketIO MGT	16	Now
PROMO-5004-5-ILT	FPGA Essentials	15	Now
PROMO-5003-5-LEL	Designing for Performance, Live Online	9	Now
EMBD-21000-5-ILT	Embedded Systems Development	16	Now
Platinum Technical Service			
SC-PLAT-SVC-10	1 Seat Platinum Technical Service w/10 education credits	N/A	Now
SC-PLAT-SITE-50	Platinum Technical Service site license up to 50 customers	N/A	Now
SC-PLAT-SITE-100	Platinum Technical Service site license for 51-100 customers	N/A	Now
SC-PLAT-SITE-150	Platinum Technical Service site license for 101-150 customers	N/A	Now
Titanium Technical Service			
PS-TEC-SERV	Titanium Technical Service (minimum 40 hours)	N/A	Now
Design Services			
DC-DES-SERV	Design Services Contract	N/A	Now
Xilinx Productivity Advantage			
DS-XPA	Custom XPA Packaged Solution	N/A	Now
DS-ISE-ALI-XPA	XPA Seat, ISE Alliance	N/A	Now
DS-ISE-FND-XPA	XPA Seat, ISE Foundation	N/A	Now

Titanium Technical Service

- Dedicated application engineer
- Design Flow methodology coaching
- Contract based service
- Factory escalation process
- Timing Closure expertise
- Service at customer site or Xilinx site

Platinum Technical Service

- Access to Senior Applications Engineers
- Dedicated Toll Free Number*
- Priority Case Resolution
- Ten Education Credits
- Electronic Newsletter
- Formal Escalation Process
- Service Packs and Software Updates
- Application Engineer to Customer Ratio, 2x Gold Level

**Toll free number available in US only, dedicated local numbers available across Europe*

Education Services Contacts

North America: 877-XLX-CLAS (877-959-2527)
http://support.xilinx.com/support/training/training.htm

Europe: +44-870-7350-548
eurotraining@xilinx.com

Japan: +81-03-5321-7750
http://support.xilinx.co.jp/support/education-home.htm

Asia Pacific: +03-5321-7711
http://support.xilinx.com/support/education-home.htm

Design Services Contacts

North America & Asia:
Richard Fodor: 408-626-4256
Mike Barone: 512-238-1473

Europe:
Alex Hillier: +44-870-7350-516
Martina Finnerty: +353-1-4032469

designservices@xilinx.com

XPA Contacts

North America: 800-888-FPGA (3742)
fpga.xilinx.com

Europe:
Stuart Elston: +44-870-7350-532

Titanium Technical Service Contacts

North America:
Telesales: 1-800-888-3742

Europe:
Stuart Elston: +44-870-7350-632

Japan:
81-3-5321-7730
or japantitanium@xilinx.com

Xilinx IP Selection Guide

										Implementation Example				
Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Occ	MHz	Device	Key Features	Application Examples
Communication & Networking														
3G FEC Package	Xilinx	LogiCORE		V-II	V-E	V							Viterbi Decoder, Turbo Codec, Convolutional Encoder	3G Wireless Infrastructure
8b/10b Decoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		0.4%	160	XC2V40-5	Industry std 8b/10b en/decode for serial data transmission	Physical layer of Fiber Channel
8b/10b Encoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		0.4%	160	XC2V40-5	Industry std 8b/10b en/decode for serial data transmission	Physical layer of Fiber Channel
AWGN - Additive White Gaussian Noise	Xilinx	LogiCORE	V-IIP	V-II						93.0%	245	XC2V80-5	Probability density function (PDF) deviates less than 0.2 percent from the Gaussian PDF for |x| < 4.8_ and is obtained from a closed-form expression	BER measurements for Forward Error Correction cores such as Reed-Solomon, Viterbi Decoder,Turbo Convolutional code or Turbo Product code
ADPCM, 16 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II		V				89%	16	XCV150-6	Supports G.721 G.723 G.726 G.726a G.727 G.727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM, 256 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E					66%	30	XCV400E-8	Supports G.721 G.723 G.726 G.726a G.727 G.727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM, 512 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II		V				49%	6	XCV400E-8		
ADPCM, 768 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II		V				89%	50	XC2V500-5	Supports G.721 G.723 G.726 G.726a G.727 G.727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM, 1024 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E			S-II	S	89%	50	XC2V500-5	Supports G.721 G.723 G.726 G.726a G.727 G.727a, u-law, a-law	DECT, VOIP, cordless telephony
AES Decryption core family	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V		S-II		36%	103	XC2V250-5		eCommerce, Banking, Video phones, PDA, satellite communications
AES Encryption Core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		16%	131	XC2V250-5	Supports ECB, OFB, CFB, CBC modes; Supports 128, 192 and 256-bit keys	Transaction and secure communications; surveillance, storage and embedded systems
AES Encryption core family	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V		S-II		26%	134	XC2V250-5		eCommerce, Banking, Video phones, PDA, satellite communications
Bit Stream Analyzer and Data Extractor	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	71%	67	XCV50-6	Data syntax analysis of IP, MPEG, ATM	ATM, IP, MPEG
Bluetooth Baseband Processor	NewLogic GmbH	AllianceCORE		V-II		V						XC2V1000-4	Compliant to Bluetooth v1.1, BQB qualified software for L2CAP, LHP, HC1, voice support	Bluetooth applications
Bluetooth Hardware Baseband Controller	Wipro, Ltd.	AllianceCORE		V-II		V				69%	25	XC2V1500-4		
Cell Assembler	Paxonet Communications	AllianceCORE				V			S		60	XC4005XL-1	Octet wide operation, HEC computation, cell scrambling	ATM adapter cards, routers, switches
Cell Delineation	Paxonet Communications	AllianceCORE				V			S		40	XC4010XL-9	Octet wide operation, HEC verification, cell scrambling	ATM adapter cards, rtouters, switches
Convolutional Encoder	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	2%	144	XCV50-6	code rate, gen. vectors, CMSTR length customizable	Error correction
Convolutional Encoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		10%	26	XC2V40-6	k from 3 to 9, puncturing from 2/3 to 12/13	3G base stations, broadcast, wireless LAN, cable modem, xDSL, satellite com, uwave
CRC10 Generator and Verifier	Paxonet Communications	AllianceCORE				V			S	22%	20	XCS30-4	Separate generator and verifier blocks, compatible with ITU-T I.363 for AAL3/AAL4	ATM, SONET, and Ethernet
CRC32 Generator and Verifier	Paxonet Communications	AllianceCORE				V			S	44%	29	XCS30-4	Separate generator and verifier blocks, compat with ITU-T I.363 for AAL5	ATM, SONET, and Ethernet
DES3 Encryption	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		15%	167	XC2V1000-5		
DES and 3DES Cryptoprocessor	inSilicon Corporation	AllianceCORE				V		S-II	S		48	XC2S150-6	Compliant with ANSI X9.52, 128-bit key or two independent 64-bit keys	Secure communication, data storage
DES and 3DES encryption engine	Memec Core	AllianceCORE							S	79%	25	XCS20-4	NIST certified, supports EBC, CBC, CFB, and OFB	Secure communication, data storage
DES Cryptoprocessor	inSilicon Corporation	AllianceCORE				V		S-II	S		94	XC2S100-6	NIST certified, supports ECB, CBC, CFB, and OFB	Secure communication, data storage
DES Encryption	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		93%	204	XC2V40-5		
Distributed Sample Descrambler	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	14%	74	XCV50-6	Compliant with ITU-T I.432. Parameterizable data width, cell & header length	ATM PHY layer
Distributed Sample Scrambler	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	9%	104	XCV50-6	Compliant with ITU-T I.432 scrambler. Parameterizable data width, cell length, header length	ATM PHY layer
DVB Satellite Modulator	Memec Core	AllianceCORE				V					45-70	XCV50-4	Conforms to ETSI EN 300 421 v1.1.2, selectable convolutional code rate	Digital broadcast, microwave transmitter
Ethernet MAC, 1 Gigabit Full Duplex	Alcatel Technology Licensing Group	AllianceCORE		V-II						46%	31	XC2V1500-5	Single/multi-mode fiber optics; 802.3x full duples flow control; SERDES for GMII; Auto-negotiation for 1000BASE-X	NICs, routers, switches, hubs
EEthernet MAC, 1 Gigabit Half/Full duplex with GMII or Gigabit serial PHY	Xilinx	LogiCORE	V-IIP	V-II	V-E		S-IIE			23%	125 (GMII) or 1.25 Gbps	XC2V1000-4	IEEE 802.3-2000 compliant, support 8-bit GMII interface or integrated PCS/PMA interface, supporting 1000BASE-X application	GbE Network Interface Cards (NICs), Edge switches and terabit routers – packet based line cards, iSCSI line cards, SPI-3 (POS-PHY L3) to Gb Ethernet and other bridges
Ethernet MAC, 10 Gigabit Full Duplex with XGMII or XAUI	Xilinx	LogiCORE	V-IIP	V-II						25%	156.25 DDR for XGMII or 4 channels of 3.125 Gbps Rocket I/O transceivers	XC2V3000-5	Designed to IEEE 802.3ae, version D4.1 support both 32-bit XGMII parallel interface or XAUI interface, supports 10 GBASE-X, WAN/LAN functionality, Statistics gathering	Layer 2 swtiches/hubs, test equipment, bridge to POS PHY4, iSCSI line cards
Ethernet MAC, 10/100	Alcatel Technology Licensing Group	AllianceCORE		V-II				S-II		13%	31	XC2V1500-5	Single/multi-mode fiber optics; 10/100 MII PHY, 10Base-T, 100Base-T/TX/FX/4; RMON & Etherstats	NICs, routers, switches, hubs, printers
Ethernet MAC, 10/100	Paxonet Communications	AllianceCORE				V		S-II	S	45%	50	XCV150-4	IEEE 802.3 compliant RMON, MIBs stats, MII support	Ethernet switched, hub, NICS
SPI-4.1 (Flexbus 4) Interface Core, 1-Channel	Xilinx	LogiCORE		V-II						12%	200	XC2V1000 FG456-5	OIF SPI-4 Phase 1 and Flexbus4 compliant. Fully HW interoperable with AMCC OC-192 framers.	Line card: terabit routers & optical switches
SPI-4.1 (Flexbus 4) Interface Core, 4-Channel	Xilinx	LogiCORE		V-II						27%	200	XC2V1000 FG456-5	OIF SPI-4 Phase 1 and Flexbus4 compliant. Fully HW interoperable with AMCC OC-192 framers.	Line card: terabit routers & optical switches
HDLC Controller Core, 1-Channel	Xilinx	LogiCORE		V-II	V-E	V		S-II		15%	115	XC2V250	16/32-bit frame seq, 8/16-bit addr insert/delete, flag/zerop insert/detect	X.25, POS, cable modems, frame relay switches, video conferencing over ISDN
HDLC Controller, 1-Channel	Memec Core	AllianceCORE				V		S-II		95%	77	XC2S15-5	16/32-bit frame seq, 8/16-bit addr insert/delete, flag/zerop insert/detection	X.25, Frame Relay, B/D-Channel
HDLC Controller Core, 32-Channel	Xilinx	LogiCORE		V-II	V-E	V		S-II		34%	81	XC2V250	32 full duplex, CRC-16/32, 8/16-bit address insertion/deletion	X.25, POS, cable modems, frame relay switches, video conferencing over ISDN
HDLC-PPP for Packet over SONET	Paxonet Communications	AllianceCORE				V		S-II		75%	80	XC2S150-6	RFC1619 (IP&IPX) POS, 16/32 bit FCS generation and verification, stats	Bridges, switches, WAN links
Interleaver/De-interleaver	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		30%	208	XC2V40-6	Block & convolutional, width up to 256 bits, 256 branches	Broadcast, wireless LAN, cable modem, xDSL, satellite com,uwave nets, digital TV, CDMA2000
Interleaver/Deinterleaver	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	21%	73	XCV50-6	Block & convolutional support. param features. 3GPP, UMTS, GSM, DVB compliant	Channel coding in telecom/wireless, broadcast
Inverse Multiplexer for ATM	ModelWare, Inc.	AllianceCORE		V-II	V-E					100%	31	XC2V1000-4	Supports up to 32 links/32 groups; UTOPIA L2 PHY & ATM I/Fs, supports IDCR	Access systems, multi-service switches, DSLAMs, Basestation controllers
Network Processor	IP Semiconductors A/S	AllianceCORE		V-II	V-E						80	XC2V1500-5	Total Solution requires this core + SPEEDAnalyzer ASIC, 2.5 Gbps full duplex wire speed; network processor (NPV), low power an low device count than competing network processors.	Networking, edge and access, Switches and routers
Noisy Transmission Channel Model	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	23%	100	XCV50-6	Programmable noise generation profile	Noise emulation in transmission channel
SPI-3 (POS-PHY L3) Link Layer Interface, 1-256Ch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		9%	175	XC2V1000 FG456-4	OIF SPI-3 (POS-PHY L3) compliant. Fully HW interoperable with PMC-Sierra OC-48 framers.	Line cards, iSCSI cards, gigabit routers and switches
SPI-3 (POS-PHY L3) Link Layer Interface, 1-Ch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		10%	200	XC2V1000 FG456-4	OIF SPI-3 (POS-PHY L3) compliant. Fully HW interoperable with PMC-Sierra OC-48 framers.	Line cards, iSCSI cards, gigabit routers and switches
SPI-3 (POS-PHY L3) Link Layer Interface, 2-Ch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		55%	200	XCV50E-8	OIF SPI-3 (POS-PHY L3) compliant. Fully HW interoperable with PMC-Sierra OC-48 framers.	Line cards, iSCSI cards, gigabit routers and switches
SPI-3 (POS-PHY L3) Link Layer Interface, 4-Ch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		15%	200	XC2V1000 FG456-4	OIF SPI-3 (POS-PHY L3) compliant. Fully HW interoperable with PMC-Sierra OC-48 framers.	Line cards, iSCSI cards, gigabit routers and switches
SPI-3 (POS-PHY L3) Physical Layer Interface	Xilinx	LogiCORE			V-E					52%	104	XCV50E-8	OIF SPI-3 (POS-PHY L3) compliant.	Line cards, iSCSI cards, gigabit routers and switches
SPI-4.2 (POS-PHY L4) Multi-Channel Interface	Xilinx	LogiCORE	V-IIP	V-II						29%	400 DDR	XC2V3000 FG676-5	OIF SPI-4 Phase 2 compliant. Fully HW interoperable with PMC-Sierra and Mindspeed OC-192 framers. Supports dynamic alignment.	Line cards, switches, routers and optical switches
SPI-4.2 (POS-PHY L4) Multi-Channel Interface	Xilinx	LogiCORE		V-II						11%	100 DDR	XC2V3000 FG676-5	OIF SPI-4 Phase 2 compliant. Operates at OC-48 line rate	Line cards, switches, routers and optical switches

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Implementation Example			Key Features	Application Examples
										Occ	MHz	Device		
Communication & Networking (continued)														
SPI-4.2 (POS-PHY L4) to SPI-4.1 (Flexbus 4) Bridge	Xilinx	LogiCORE		V-II						35%	200 on FB4, 350 DDR on SPI-4.2, 175 internal	XC2V3000 FG676-5	OIF SPI-4 Phase 1& 2 compliant. Fully HW interoperable with AMCC, PMC-Sierra and Mindspeed OC-192 framers.	Line cards, switches, routers and optical switches
SPI-4.2 (POS-PHY L4) to XGMII (10GE MAC) Bridge	Xilinx	LogiCORE		V-II						53%	350 DDR on SPI-4.2, 156 DDR on XGMII	XC2V3000	OIF SPI-4 Phase 1& 2 compliant. Fully HW interoperable with AMCC, PMC-Sierra and Mindspeed OC-192 framers.	Line cards, switches, routers and optical switches
Reed-Solomon Decoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II	S	40%	98	XC2V250-6	Std or custom coding, 3-12 bit symbol width, up to 4095 symbols, error & erasure decoding	Broadcast, wireless LAN, cable modem, xDSL, satellite com,uwave nets, digital TV
Reed-Solomon Decoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II		V				97%	61	XC2V500-5	parameterizable, RTL available	Error correction, wireless, DSL
Reed Solomon Decoder	Amphion Semiconductor Ltd.	AllianceCORE				V			S	50%	50	XCV100-4		Error correction
Reed Solomon Decoder	Memec Core	AllianceCORE				V		S-II	S	83%	73	XCV50-6	Customizable, >580 Mbps	Error correction
Reed-Solomon Encoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II	S	42%	180	XC2V40-6	Std or cust coding, 3-12 bit width, up to 4095 symbols with 256 check symb.	Broadcast, wireless LAN, cable modem, xDSL, satellite com,uwave nets, digital TV
Reed Solomon Encoder	Amphion Semiconductor Ltd.	AllianceCORE				V			S	11%	82	XCV50-4		Error correction
Reed Solomon Encoder	Memec Core	AllianceCORE				V		S-II	S	12%	113	XCV50-6	Customizable, > 900 Mbps	Error Correction
SDLC Controller	CAST, Inc.	AllianceCORE		V-II		V				11%	158	XC2V1000-5	Like Intel 8XC152 Global Serial Channel, Serial Comm., HDLC apps, telecom	Embedded systems, professional audio, video
SHA-1 Encryption processor	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		24%	79	XC2V500-6	SHA-1 algorithm compliant	Secure comms, video suveillance, data storage, financial transactions
T1 Framer	Virtual IP Group	AllianceCORE											D4, ESF, SLC-96 formats. For XC4000.	DSI trunk, PBX I/F
Turbo Convolutional Decoder - 3GPP Compliant	Xilinx	LogiCORE		V-II	V-E	V				80%	40	XC2V500	3GPP specs, 2 Mbps, BER=10-6 for 1.5dB SNR	3G Wireless Infrastructure
Turbo Decoder, 3GPP	SysOnChip, Inc.	AllianceCORE		V-II		V					66	XC2V500-5	3GPP/UMTS compliant, IMT-2000, 2Mbps data rate	Error correction, wireless
Turbo Decoder, DVB-RCS	iCoding Technology, Inc.	AllianceCORE		V-II	V-E	V				44%	71	XC2V2000-5	DVB-RCS compliant, 9Mbps, data rate, switchable code rates and frame sizes	Error correction, wireless, DVB, Satellite data link
Turbo Decoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II		V				70%	65	XC2V2000-5	3GPP/UMTS compliant, >2Mbps data rate	Error correction, wireless
Turbo Convolutional Encoder - 3GPP Compliant	Xilinx	LogiCORE		V-II	V-E	V				65%	60	XC2V250	Compliant w/ 3GPP, puncturing	3G Wireless Infrastructure
Turbo Encoder, DVB-RCS	iCoding Technology, Inc.	AllianceCORE		V-II	V-E	V				2%	69	XC2V2000-5	DVB-RCS compliant, 9Mbps, data rate, switchable code rates and frame sizes	Error correction, wireless, DVB, Satellite data link
Turbo Encoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II		V		S-II		48%	120	XC2V80-5	3GPP/UMTS compliant, upto 4 interleaver laws	Error correction, wireless
Turbo Product Code Decoder	Xilinx	LogiCORE	V-IIP	V-II						31%	285	XC2V40-5	Fully compliant with IEEE 802.16 and 802.16a standards	Error Correction, LMDS, MMDS
Turbo Product Code Encoder	Xilinx	LogiCORE	V-IIP	V-II						64%	150	XC2V1000-5	Fully compliant with IEEE 802.16 and 802.16a standards	Error Correction, LMDS, MMDS
UTOPIA level 2 slave inteface	Paxonet Communications	AllianceCORE											Cell handshake in SPHY mode, 8/16 bit operation, internal FIFO, detects runt cells	ATM PHY layer
UTOPIA Level-2 PHY Side RX Interface	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	8%	53	XCV50-6	Protocol conversion from Pb (RACE BLNT) to UTOPIA L2, 8/16 bit operation	ATM PHY layer
UTOPIA Level-2 PHY Side TX Interface	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	10%	61	XCV50-6	Protocol conversion from UTOPIA L2 Pb (RACE BLNT), 8/16 bit operation	ATM PHY layer
UTOPIA Level-3 ATM Receiver	inSilicon Corporation	AllianceCORE				V		S-II	S					
UTOPIA Level-3 ATM Transmitter	inSilicon Corporation	AllianceCORE				V		S-II	S					
UTOPIA Level-3 PHY Receiver	inSilicon Corporation	AllianceCORE				V		S-II	S					
UTOPIA Level-3 PHY Transmitter	inSilicon Corporation	AllianceCORE				V		S-II	S					
UTOPIA Master	Paxonet Communications	AllianceCORE				V		S-II	S				SPHY, MPHY, HEC processing, round robin polling, ind. transmitter receiver	ATM PHY layer
UTOPIA Slave	Paxonet Communications	AllianceCORE				V		S-II	S	26%	79	XCV50-4	Cell handshake in SPHY mode, 8/16 bit operation, 32 bit FIFO interface, detects runt cells	ATM PHY layer
Viterbi Decoder, IEEE 802-compatible	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		35%	157	XC2V500-5	Parameterizable source code with constraint length(k)=7, G0=171, G1=133; or G0=133, G1=171, includes BestState Logic, ability to change code rate and traceback depth on the fly, supports Trellis Coded Modulation, IEEE802.11a/16a compatible, reaches OC3 (155Mbps) and higher	L/MMDS, broadcast equip, wireless LAN, cable modem, xDSL, sat com, uwave nets
Viterbi Decoder, General Purpose	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		38%	128	XC2V500-5	Puncturing, serial & parallel architecture, dynamic rate change, paramterized constraint length, soft/hard decision with programmable number of soft bits, dual rate decoder, erasure pins for external puncturing, compatible with standards such as DVB ETS, 3GPP2, IEEE802.16, Hiperlan, Intelsat IESS-308/309	3G base stations, broadcast, wireless LAN, cable modem, xDSL, satellite com, uwave, CDMA2000
Viterbi Decoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II	V-E			S-II		74%	91	XC2V500-5	Radix-2/radix4 architectures, BER, depuncturing. Code rate, contraint length parameterizable	Data transmission, wireless
Digital Signal Processing														
Bit Correlator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					4096 taps, serial/parallel input, 4096 bits width	
Cascaded Integrator Comb (CIC) Filter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					32 bits data width, rate change from 8 to 16384	
Comb Filter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Cascaded Integrator Comb Filter	
CORDIC	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					Polar to rectangular, rectangular to polar, sin & cos, sinh & cosh, atan & atanh, square root	Digital receivers
Digital Down Converter (DDC)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		47%	116	XC2V500-5	Configurable datapath comprising mixer, DDS, optional CIC filter, optional polyphase decimators, 25 to 108db DDS spurious free dynamic range	Wireless & wireline communication systems such as software defined radios, digital receivers, cable modems, BPSK, QPSK, QAM demodulators, spread spectrum communication systems, CDMA2000 & 3G basestations
Direct Digital Synthesizer (DDS)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		12%	245	XC2V80-6	8-65K samples, 32-bits output precision, phase dithering/offset	
FFT/IFFT, 1024-Point Complex	Xilinx	LogiCORE			V-E	V								
FFT/IFFT for Virtex-II, 1024-Point Complex	Xilinx	LogiCORE		V-II						62%	41us, 100MHz	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
FFT/IFFT, 16-Point Complex	Xilinx	LogiCORE		V-II	V-E	V								
FFT/IFFT for Virtex-II, 16-Point Complex	Xilinx	LogiCORE		V-II						37%	123ns 130MHz	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
FFT/IFFT, 256-Point Complex	Xilinx	LogiCORE		V-II	V-E	V								
FFT/IFFT for Virtex-II, 256-Point Complex	Xilinx	LogiCORE		V-II						54%	7.7us, 100MHz	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
FFT/IFFT, 32-Point Complex	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		29%	110	XC2V500-6	Complex FFT, Forward and Inverse transform. Supports bit precisions from 2-32 bits. Embedded memory	
FFT/IFFT, 64-, 256-, 1024-Point Complex	Xilinx	LogiCORE	V-IIP	V-II						32%	140	XC2V1500-6	64-,256-,1024-point programmable point size, forward and inverse transform, 16-bit complex data, 18-bit phase factors, 2's complement, built-in memories, programmable data scaling, 140 MHz, 1024-point transform - 7.31 us, 256-point transform - 1.83 us, 64-point transform - 0.46 us	
FFT/IFFT, 64-Point Complex	Xilinx	LogiCORE			V-E	V								
FFT/IFFT for Virtex-II, 64-Point Complex	Xilinx	LogiCORE		V-II						38%	1.9us, 100MHz	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
FIR Filter, Distributed Arithmetic (DA)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					32-bit input/coeff width, 1024 taps, 1-8 chan, polyphase, online coeff reload	
FIR Filter, Distributed Arithmetic Parallel	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Arithmetic FIR Filter	
FIR Filter, Distributed Arithmetic Serial	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Arithmetic FIR Filter	
FIR Filter, MAC	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		16%		XC2V250	Single rate, Polyphase Decimator, Polyphase Interpolator	3G base stations, wireless communications, image filtering
LFSR, Linear Feedback Shift Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					168 input widths, SRL16/register implementation	
Oscillator, Dual-Channel Numerically Controlled	Xilinx	LogiCORE			V-E	V	S-IIE	S-II	S				Not recommended for new designs. Suggested replacement: Direct Digital Synthesizer	
Oscillator, Numerically Controlled	Xilinx	LogiCORE			V-E	V	S-IIE	S-II	S				Not recommended for new designs. Suggested replacement: Direct Digital Synthesizer	
Time-Skew Buffer, Nonsymmetric 16-Deep	Xilinx	LogiCORE							S				Not recommended for new designs.	
Time-Skew Buffer, Nonsymmetric 32-Deep	Xilinx	LogiCORE							S				Not recommended for new designs.	

										Implementation Example			
Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Occ	MHz	Device	Key Features	Application Examples
Digital Signal Processing (continued)														
Time-Skew Buffer, Symmetric 16-Deep	Xilinx	LogiCORE							S					
TMS32025 DSP Processor core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		66%	63	XC2V500-5		
Math Functions														
1s and 2s Complement	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Bus Gate or Twos Complementer	
Accumulator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256s bit wide	
Adder Subtracter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256s bit wide	
Divider, Pipelined	Xilinx	LogiCORE		V-II	V-E	V	S-IIE	S-II	S				32-bit input data width, multiple clock per output	
Floating Point Adder	Digital Core Design	AllianceCORE		V-II		V		S-II			66	XC2V250-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
Floating Point Divider	Digital Core Design	AllianceCORE		V-II		V		S-II			53	XC2V250-5	Full IEEE-754 compliance, 15 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
Floating Point Multiplier	Digital Core Design	AllianceCORE		V-II	V-E	V		S-II			74	XC2V250-5	Full IEEE-754 compliance, 7 pipelines,32x32 mult, Single precision real format support	DSP, Math, Arithmetic apps.
Floating Point Square Comparator	Digital Core Design	AllianceCORE		V-II		V		S-II			91	XC2V80-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps.
Floating Point Square Root Operator	Digital Core Design	AllianceCORE		V-II	V-E			S-II			66	XC2V250-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
Floating Point to Integer Converter	Digital Core Design	AllianceCORE		V-II		V		S-II			66	XC2V250-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
Integer to Floating Point Converter	Digital Core Design	AllianceCORE		V-II		V		S-II			73	XC2V250-5	Full IEEE-754 compliance, double word input, 2 pipelines, Single precision real output	DSP, Math, Arithmetic apps
Integrator	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Cascaded Integrator Comb Filter	
Multiplier, Constant Coefficient	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Multiplier	
Multiplier, Constant Coefficient - Pipelined	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Multiplier	
Multiplier, Dynamic Constant Coefficient	Xilinx	LogiCORE			V-E	V							Not recommended for new designs. Suggested replacement: Multiplier	
Multiplier, Parallel - Area Optimized	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Multiplier	
Multiplier for Virtex, Variable Parallel	Xilinx	LogiCORE			V-E	V	S-IIE	S-II					Not recommended for new designs. Suggested replacement: Multiplier	
Multiply Accumulator (MAC)	Xilinx	LogiCORE	V-IIP	V-II		V	S-IIE	S-II					Input width up to 32 bits, 65-bit accumulator, truncation rounding	
Multiply Generator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					64-bit input data width, constant, reloadable or variable inputs, parallel/sequential implementation	
Registered Adder	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Adder Subtracter	
Registered Loadable Adder	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Adder Subtracter	
Registered Loadable Subtracter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Adder Subtracter	
Registered Scaled Adder	Xilinx	LogiCORE							S				Not recommended for new designs.	
Registered Serial Adder	Xilinx	LogiCORE							S				Not recommended for new designs.	
Registered Subtracter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Adder Subtracter	
Scaled-by-One-Half Accumulator	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Accumulator	
Sine Cosine Look Up Table	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II	S	12%	270	XC2V40-6	3-10 bit in, 4-32 bit out, distributed/block ROM	
Square Root	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: CORDIC	
Twos Complementer	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					Input width up to 256 bits	
Memories & Storage Elements														
ATM Utopia Level 2	Xilinx	LogiCORE	V-IIP											
Block Memory, Dual-Port	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits, 2-13K words	
Block Memory, Single-Port	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits, 2-128K words	
Content Addressable Memory (CAM)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-512 bits, 2-10K words, SRL16	
CAM, for Internet Protocol	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	19%	49	XCV50-6		IP routers
Distributed Memory	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-1024 bit, 16-65536 word, RAM/ROM/SRL16, opt output regs and pipelining	
FIFO, Asynchronous	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits, 15-65535 words, DRAM or BRAM, independent I/O clock domains	
FIFO, Synchronous	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Synchronous FIFO supporting Spartan-II/Spartan-IIE	
FIFO, Synchronous	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits, 16-256 words, distributed/block RAM	
Pipelined Delay Element	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Memory using SRL16 based memory type	
Registered ROM	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Memory	
Registered Single Port RAM	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Memory	
Microprocessors, Controllers & Peripherals														
10/100 Ethernet MAC w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Evaluation Core (Available w/EDK) or High Value Core to be bought separately ($)	Networking, communications, processor apps
10/100 Ethernet MAC Lite w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Evaluation Core (Available w/EDK)	Networking, communications, processor apps
16-bit proprietary RISC processor	Loarant Corporation	AllianceCORE		V-II	V-E	V		S-II			91	XC2V500-5	44 opcodes, 64-K word data, program, Harvard arch.	Control functions, State machines, Coprocessor
16-Word-Deep Registered LUT	Xilinx	LogiCORE							S					
16450 UART	Virtual IP Group	AllianceCORE							S		60	XC2S50-6	Independently controlled transmit, receive and data interrupts. 16X clock.	Serial data applications, modems
16450 UART w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Evaluation Core (Available w/EDK) or High Value Core to be bought separately ($)	Processor applications
16550 UART w/FIFOs	Virtual IP Group	AllianceCORE							S		16	XCS20-4	Prog. Data width, parity, stop bits. 16X internal clock, FIFO mode, false start bit detection	Serial data applications, modems
16550 UART w/FIFOs & synch interface	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		57%	87	XC2V80-5	Supports 16450 and 16550a synchronous s/w; programmable serial interface	Serial and modem applications
16550 UART w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Evaluation Core (Available w/EDK) or High Value Core to be bought separately ($)	Processor applications
2901 microprocessor slice	CAST, Inc.	AllianceCORE				V		S-II		19%	36	XC2S50-6	Eight function ALU, 4 status flags- Carry, Overflow, Zero and Negative	Simple microcontroller applications
2910A microprogram controller	CAST, Inc.	AllianceCORE				V		S-II	S	12%	63	XCV50-6	Based on AMD 2910a	High-speed bit slice design
68000 compatible microprocessor	CAST, Inc.	AllianceCORE		V-II		V				89%	32	XC2V500-5	MC68000 Compatible	Embedded systems, professional audio, video
80186 compatible processor	eInfochips Pvt. Ltd.	AllianceCORE		V-II		V				77%	49	XC2V1000-5	i80186 compatible plus enhanced mode	Industrial, wireless, communications, embedded
8051 compatible microcontroller	CAST, Inc.	AllianceCORE				V		S-II		52%	68	XCV200E-8	80C31 instruction set, 8 bit ALU, 8 bit control, 32 bit I/O ports, two 16 bit timer/counters, SFR I/F	Embedded systems, telecom
8051 compatible microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			73	XC2V250-5	80C31 instruction set, RISC architecture 6.7X faster than standard 8051	Embedded systems, telecom, video
8051 compatible microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			90	XC2V250-5	80C31 instruction set, high speed multiplier, RISC architecture 6.7X faster than standard 8051	Embedded systems, telecom, video
8051 compatible microcontroller	Dolphin Integration	AllianceCORE				V		S-II	S	33%	29.8	XCV300-6	12X faster (average) and code compatible wrt legacy 8051, verification bus monitor, SFR interface	Telecom, industrial, high speed control
8051 compatible microcontroller	Dolphin Integration	AllianceCORE		V-II	V-E			S-II		39%	38	XC2V1000-5	8X faster (average) and code compatible wrt legacy 8051, verification bus monitor, SFR interface, DSP focused	DSP, Telecom, industrial, high speed control
80515 high-speed 8-bit RISC microcontroller	CAST, Inc.	AllianceCORE				V				90%	42	XCV200E-8	RISC implementation, 8 bit ALU, 8 bit control, 32 bit I/O, 16 bit timer/counters, SFR I/F, ext. memory I/F	High speed embedded systems, audio, video
8052 compatible microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			71	XC2V250-5	80C31 instruction set, high speed multiplication and division ,RISC architecture 6.7X faster than standard 8051	Embedded systems, telecom, video
80530 8-bit microcontroller, compact	CAST, Inc.	AllianceCORE				V		S-II		88%	51	XC2S150-6	32 bit I/O, 3 counters, interrupt controller, SFR interface, dual data pointer	Low cost embedded systems, telecom
80530 8-bit microcontroller	CAST, Inc.	AllianceCORE				V				81%	66	XCV200E-8	32 bit I/O, 3 counters, 27-bit watchdog timer, 3-priority interrupt controller, SFR interface	Embedded systems, telecom

										Implementation Example				
Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Occ	MHz	Device	Key Features	Application Examples
Microprocessors, Controllers & Peripherals (continued)														
80C51 compatible RISC microcontroller	CAST, Inc.	AllianceCORE			V-E	V		S-II		75%	34	XC2S150-6	12X faster, SRF I/F	Embedded systems
8237 programmable DMA controller	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		40%	34	XC2S100-6	4 independent DMA channels (expandable), software DM	Microprocessor based systems
8250 UART	Memec Core	AllianceCORE							S	58%	10	XCS10-4	DC to 625K baud	Serial communications
8254 programmable interval timer/counter core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		74%	68	XC2V80-5	Status feedback, counter latch, Square wave mode, 6 counter modes, binary/BCD count, LSB/MSB R/W	Event counterr, baud rate generator
8254 programmable timer	Virtual IP Group	AllianceCORE				V		S-II	S					
8254 programmable timer/counter	eInfochips Pvt. Ltd.	AllianceCORE		V-II				S-II		11%	51	XC2V1000-4	Status red, Six prog, counter modes, Intel8254 like	Processor I/O interface
8255 programmable I/O controller	eInfochips Pvt. Ltd.	AllianceCORE		V-II				S-II		2%	175	XC2V1000-5	Three 8-bit parallel ports, 24 programmable IO lines, 8-bit bidi data bus	Processor I/O interface
8255 programmable peripheral interface	Memec Core	AllianceCORE							S	64%	8	XCS05-4	Bit set/reset support	Embedded systems
8255 programmable peripheral interface	Virtual IP Group	AllianceCORE				V		S-II	S		227	XCV50E-8	Three 8-bit peripheral ports, 24 programmable I/O lines, 8-bit bidi data bus	Processor I/O interface
8255A peripheral interface	CAST, Inc.	AllianceCORE				V		S-II		10%	227	XCV50E-8	Three 8-bit peripheral ports, 24 programmable IO lines, 8-bit bidi data bus	Processor I/O interface
8256 multifunction microprocessor support controller	Memec Core	AllianceCORE							S	89%	10	XCS20-4	Baud rate generator for 13 common baud rates, parallel I/O ports, prog. timer/counters	Communication, embedded systems
8259 programmable interrupt controller	Virtual IP Group	AllianceCORE							S			XC2S50-6	8 vectored priority interrupts, all 8259/A modes programmable- e.g., special mask, buffer	Real-time interrupt based uP designs
8259A programmable interrupt controller	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		90%	142	XC2V40-5	8 vectored priority interrupts, all 8259/A modes programmable- e.g., special mask, buffer	Real-time interrupt based uP designs
8279 programmable keyboard display interface	Memec Core	AllianceCORE							S	46%	8	XCS20-4	8 char keyboard FIFO, 2-key lockout, n-key rollover, 4-16 char display	Embedded systems interface
Arbiter and Bus Structure w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Infrastructure Core (includes device drivers), Available w/EDK	Processor applications
Arbiter and Bus Structure w/PLB interface	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Infrastructure Core (includes device drivers), Available w/EDK	PowerPC embedded system design
ARC 32-bit configurable RISC processor	ARC International plc	AllianceCORE			V-E			S-II		89%	37	XC2S150-6	4 stage pipeline, 16 single cycle instructions10, 3 interrupt exception levels, 24 bit stack pointer	32 bit processing, DSP
ATM Utopia Level 2 Master and Slave w/OPB Interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Evaluation Core (Available w/EDK) or High Value Core to be bought separately ($)	Networking, communications, processor apps
ATM Utopia Level 2 Master and Slave w/PLB Interface	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Evaluation Core (Available w/EDK) or High Value Core to be bought separately ($)	Networking, communications, processor apps
BRAM Controller w/LMB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Memory Controller Core, Available w/EDK	Processor applications
BRAM Controller w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Memory Controller Core, Available w/EDK	Processor applications
BRAM Controller w/PLB interface	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Memory Controller Core, Available w/EDK	PowerPC embedded system design
BSP Generator (SW only)	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Infrastructure Core (includes device drivers), Available w/EDK	PowerPC embedded system design
DCR Bus Structure	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (DCR), Infrastructure Core (includes device drivers), Available w/EDK	Processor applications
External Memory Controller (EMC) w/OPB interface (Includes support for Flash, SRAM, ZBT, System ACE)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Memory Controller Core, Available w/EDK	Processor applications
External Memory Controller (EMC) w/PLB interface (Includes support for Flash, SRAM, ZBT, System ACE)	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Memory Controller Core, Available w/EDK	PowerPC embedded system design
GPIO w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Available w/EDK	Processor applications
HDLC Controller (Single Channel) w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Evaluation Core (Available w/EDK) or High Value Core to be bought separately ($)	Processor applications
I2C w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Evaluation Core (Available w/EDK) or High Value Core to be bought separately ($)	Networking, communications, processor apps
Interrupt Controller (IntC) w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Peripheral Core (includes device drivers, RTOS adaptation layers), Available w/EDK	Processor applications
Interrupt Controller (IntC) w/DCR interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (DCR), Peripheral Core (includes device drivers, RTOS adaptation layers), Available w/EDK	Processor applications
IPIF Address Decode w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), IP Interface Module, Available w/EDK	Processor applications
IPIF DMA w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), IP Interface Module, Available w/EDK	Processor applications
IPIF Interrupt Controller w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), IP Interface Module, Available w/EDK	Processor applications
IPIF Master/Slave Attachment w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), IP Interface Module, Available w/EDK	Processor applications
IPIF Read/Write Packet FIFO w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), IP Interface Module, Available w/EDK	Processor applications
IPIF Scatter/Gather w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), IP Interface Module, Available w/EDK	Processor applications
IPIF Slave Attachment w/PLB interface	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), IP Interface Module, Available w/EDK	Processor applications
Java processor, 32-bit	Digital Communications Technologies, Ltd.	AllianceCORE		V-II		V		S-II		33%	40	XC2V1000-5	32bit data, 24 bit address, 3 Stage pipeline, Java/C dev. tools	Internet appliance, industrial control, HAVi multimedia, set top boxes
Java processor, configurable core	Derivation Systems, Inc.	AllianceCORE		V-II		V				38%	20	XC2V1000-5	32b data/address optional DES	Internet appliance, industrial control
JTAG UART w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Peripheral Core (includes device drivers, RTOS adaptation layers), Available w/EDK	Processor applications
Memory Test Utility (SW only)	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Infrastructure Core (includes device drivers), Available w/EDK	PowerPC embedded system design
MicroBlaze Soft RISC Processor	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			150	Virtex-II Pro (-6)	32-bit Soft Processor Core, Available w/EDK	Networking, communications
MicroBlaze Source Code	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			150	Virtex-II Pro (-6)	High Value Core to be bought separately ($)	Networking, communications
ML300 VxWorks BSP (SW only)	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Infrastructure Core (includes device drivers), Available w/EDK	PowerPC embedded system design
OPB2DCR Bridge	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (DCR), Infrastructure Core (includes device drivers), Available w/EDK	Processor applications
OPB2OPB Bridge (Lite)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Infrastructure Core (includes device drivers), Available w/EDK	Processor applications
OPB2PCI Full Bridge (32/33)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Infrastructure Core (includes device drivers), Available w/EDK	Processor applications
OPB2PLB Bridge	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Infrastructure Core (includes device drivers), Available w/EDK	PowerPC embedded system design
PIC125x fast RISC microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			126	XC2V80-5	PIC 12c4x like, 2X faster, 12-bit wide instruction set, 33 instructions	Embedded systems, telecom, audio and video
PIC1655x fast RISC microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			140	XC2V80-5	S/W compatible with PIC16C55X, 14-bit instruction set, 35 instructions	Embedded systems, telecom, audio and video
PIC165X compatible microcontroller	CAST, Inc.	AllianceCORE		V-II		V		S-II		61%	128	XC2V80-5	Microchip 16C5X PIC like	Embedded systems, telecom
PIC165x fast RISC microcontroller	Digital Core Design	AllianceCORE		V-II	V-E	V		S-II			126	XC2V80-5	PIC 12c4x like, 2X faster, 12-bit wide instruction set, 33 instructions	Embedded systems, telecom, audio and video
PLB2OPB Bridge	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Infrastructure Core (includes device drivers), Available w/EDK	PowerPC embedded system design
PowerPC Bus Master	Eureka Technology	AllianceCORE				V				3%	80	XCV400-6		
PowerPC Bus Slave	Eureka Technology	AllianceCORE			V-E					10%	80	XCV400E-8		
SDRAM Controller	Eureka Technology	AllianceCORE				V				37%	91	XCV50-6		
SDRAM Controller	NMI Electronics, Ltd.	AllianceCORE				V		S-II			137	XCV50-6	SDRAM refresh, customizable	Embedded systems using SDRAMs
SDRAM Controller, 200 MHz	Rapid Prototypes, Inc.	AllianceCORE				V		S-II						
SDRAM Controller, DDR	Memec Core	AllianceCORE		V-II	V-E			S-II		7%	133	XC2V1000-5	DDR SDRAM burst length support for 2,4,8 per access, supports data 16,32, 64, 72.	Digital video, embedded computing, networking

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Implementation Example			Key Features	Application Examples
										Occ	MHz	Device		
Microprocessors, Controllers & Peripherals (continued)														
SPI Master and Slave w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Peripheral Core (includes device drivers, RTOS adaptation layers), Available w/EDK	Networking, communications, processor apps
SDRAM Controller w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Memory Controller Core, Available w/EDK	Processor applications
SDRAM Controller w/PLB interface	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Memory Controller Core, Available w/EDK	PowerPC embedded system design
Systems Reset Module	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	CoreConnect Bus (PLB), Infrastructure Core (includes device drivers), Available w/EDK	PowerPC embedded system design
Timebase/Watch Dog Timer (WDT) w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Peripheral Core (includes device drivers, RTOS adaptation layers), Available w/EDK	Processor applications
Timer/Counter w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Peripheral Core (includes device drivers, RTOS adaptation layers), Available w/EDK	Processor applications
UART, generic compact	Memec Core	AllianceCORE		V-II	V-E	V	S-IIE	S-II	S	15%	73	XC2S50E-6	UART and baud rate generator	Serial data communication
UART Lite w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	CoreConnect Bus (OPB), Peripheral Core (includes device drivers, RTOS adaptation layers), Available w/EDK	Processor applications
V8-uRISC 8-bit RISC Microprocessor	ARC International plc	AllianceCORE									16	XC4000E	Proprietary 8-bit processor, 8 bit ALU, 16 bit stack pointer, 33 opcodes, 4 addr. Modes, 2 user opcodes	Embedded systems, 8-bit processing apps.
VxWorks Board Support Package (BSP)	Xilinx	LogiCORE	V-IIP								150	Virtex-II Pro (-6)	Interfaces HW and WindRiver RTOS	Embedded system design
Z80 compatible programmable timer/counter core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		9%	167	XC2V500-5		Programmable frequency divider, Pulse counter, pulse generator, interrupt controller
Z80 peripheral I/O controller core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		49%	76	XC2S50-6		Programmable , dual-port device, keyboards, printers, paper table readers
Z80CPU Microprocessor	CAST, Inc.	AllianceCORE		V-II		V		S-II		56%	72	XC2V500-5	Zilog Z80 compatible, 8-bit processor	Embedded systems, Communications
Standard Bus Interfaces														
Arbiter	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	24%	33	XCV50-6	Two priority classes - strong/weak, access counters	General purpose bus arbitration
I2C Bus Controller	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		21%	103	XC2S100-6	Philips I2C 1.1; supports master tx/rx and slave tx/rx modes	Embeded microcontroller and communications
I2C Bus Controller Master	Digital Core Design	AllianceCORE		V-II	V-E			S-II	S		143	XC2V40-5	I2C-like, multi master, fast/std. modes	Embedded systems
I2C Bus Controller Slave	Digital Core Design	AllianceCORE		V-II	V-E			S-II	S		157	XC2V40-5	I2C-like, Slave	Embedded
I2C Bus Controller Slave Base	Digital Core Design	AllianceCORE		V-II	V-E			S-II	S		187	XC2V40-5	I2C-like, Slave	Embedded Systems
I2C Two-Wire Serial Interface Master-Only	Memec Core	AllianceCORE		V-II	V-E	V	S-IIE	S-II		16%	83	XC2S50E-6	I2C-like, multi master fast/std. modes	Embedded microprocessor systems, I2C peripherals
I2C Two-Wire Serial Interface Master-Slave	Memec Core	AllianceCORE		V-II	V-E	V	S-IIE	S-II		28%	108	XC2V80-4	I2C-like, multi master fast/std. modes	Embedded microprocessor systems, I2C peripherals
PCI 64-bit/66-MHz master/target interface	Eureka Technology	AllianceCORE		V-II	V-E					22%	66	XC2V1000-5		
PCI host bridge	Eureka Technology	AllianceCORE		V-II	V-E					22%	33	XC2V1000-5		
PCI32 Interface Design Kit (DO-DI-PCI32-DKT)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		6%	66	XC2V1000 FG456-5	Includes PCI32 board, driver development kit, and customer education 3-day training class for US & Canada locations	PC boards, CPCI, Embedded, hiperf video, gb ethernet
PCI32 Interface, IP Only (DO-DI-PCI32-IP)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		6%	66	XC2V1000 FG456-5	v2.3 compiliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC add-in boards, CPCI, Embedded
PCI32 Single-Use License for Spartan (DO-DI-PCI32-SP)	Xilinx	LogiCORE					S-IIE	S-II		12%	66	XC2S200 PQ208-6	v2.3 compiliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC add-in boards, CPCI, Embedded
PCI64 & PCI32, IP Only (DO-DI-PCI-AL)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		6 - 7%	66	XC2V1000 FG456-5	v2.3 compliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC boards,CPCI,Embedded,hiperf video,gb ethernet
PCI64 Interface Design Kit (DO-DI-PCI64-DKT)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		7%	66	XC2V1000 FG456-5	v2.3 compliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly, customer education 3-day training class for US & Canada locations	PC boards,CPCI,Embedded,hiperf video,gb ethernet
PCI64 Interface, IP Only (DO-DI-PCI64-IP)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		7%	66	XC2V1000 FG456-5	v2.3 compliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC boards,CPCI,Embedded,hiperf video,gb ethernet
XAPP653: Virtex-II Pro 3.3V PCI Reference Design	Xilinx	Reference Design	V-IIP										Application note describes how to create compliant PCI3.3V designs with V-II Pro	
PCI-X 64/133 Interface for Virtex-II (DO-DI-PCIX64-VE) Includes PCI 64 bit interface at 33 MHz	Xilinx	LogiCORE		V-II						30%	133	XC2V1000 FG456-5	PCI-X 2.0 mode1 comp, 64/32-bit, 133 MHz PCI-X initiator and target IF, PCI 2.3 comp, 64/32-bit, 33 MHz PCI initiator and target IF, 3.3 V PCI-X at 33-66 MHz, 3.3 V PCI at 0-33 MHz	Server,Embedded,gb ethernet, U320 SCSI,Fibre Ch,RAID cntl,graphics
PCI-X 64/66 Interface for Virtex-E (DO-DI-PCIX64-VE) Includes PCI 64 bit interface at 33 MHz	Xilinx	LogiCORE			VE					30%	66	XCV300E-8	PCI-X 2.0 mode1 comp, 64/32-bit, 66 MHz PCI-X initiator and target IF, PCI 2.3 comp, 64/32-bit, 33 MHz PCI initiator and target IF, 3.3 V PCI-X at 33-66 MHz, 3.3 V PCI at 0-33 MHz	comm systems, SAN, clustered servers, Ultra 3 SCSI/Fibre Ch RAID, multi-port Gb
RapidIO 8-bit port LP-LVDS Phy Layer (DO-DI-RIO8-PHY)	Xilinx	LogiCORE	V-IIP	V-II						24%	250	XC2V1000 FG456-5	RapidIO Interconnect v1.2 compliant, verified with Motorola's RapidIO bus functional model v1.5	Routers, switches, backplane, control plane, data path, embedded sys, high speed interface to memory and encryption engines, high end video
RapidIO Logical (I/O) and Transport Layer (DO-DI-RIO8-LOG)	Xilinx	LogiCORE	V-IIP	V-II						20%	62.5	XC2V1000 FG456-5	RapidIO Interconnect v1.1 compliant, verified with Motorola's RapidIO bus functional model v1.5	Routers, switches, backplane, control plane, data path, embedded sys, high speed interface to memory and encryption engines, high end video
RapidIO Phy Layer to PLB Bridge reference design	Xilinx	Reference Design	V-IIP										World's first PowerPC with a RapidIO interface. The RapidIO physical layer core and the CoreConnect PLB are tightly integrated. Available to RapidIO PHY layer customers for free	Intelligent RapidIO peripherals, route exception processor, enterprise storage channel processor
Serial Protocol Interface (SPI) Slave	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE			14%	80	XC2S50E-6		Embedded microprocessor boards, and SOCs, audio/video, home and automotive radio
USB 1.1 Device Controller	Memec Core	AllianceCORE		V-II	V-E			S-II			12	XC2V1000-5	Compliant with USB1.1 spec., Supports VCI bus, Performs CRC, Supports 1.5 Mbps & 12 Mbps	Scanners, Printers, Handhelds, Mass Storage
USB 1.1 Function Controller	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		92%	48	XC2S200-6	USB 1.1; up to 31 endpoints; suspend and resume power mgmt; remote wake-up	Embedded systems, communications
Backplanes and Gigabit Serial I/O														
XAPP 649: SONET Rate Conversion in Virtex-II Pro Devices	Xilinx	Reference Design	V-IIP								2.488 Gbps per channel		Designed to I/F 16-bit SONET datapath to the RocketIO transceiver inside the FPGA @ 155.52 MHz	
XAPP 651: SONET and OTN Scramblers/Descramblers	Xilinx	Reference Design	V-IIP								2.488 Gbps per channel		SONET/OTN descrambler and scrambler are implemented in V-IIP fabric	
XAPP 652: Word Alignment and SONET/SDH Framing	Xilinx	Reference Design	V-IIP								2.488 Gbps per channel		Circuit to detect SONET A1A2 transitions	
WP160: Emulating External SERDES Devices with Embedded RocketIO Transceivers	Xilinx	White Paper	V-IIP										Discusses emulating Mindspeed CX27C201, Vitesse single channel VSC7123 and quad channel VSC7216-01, and Texas Instruments TLK2501 devices	
Aurora 401: single lane, 32-bit interface	Xilinx	Reference Design	V-IIP							~500 slices	3.125 Gbps per channel	Any	Reference design implements the Aurora link-layer protocol for a single channel. Uses LocalLink specification on the Aurora design back-end	Aurora is an ideal solution for backplanes and chip-to-chip interconnections
Video & Image Processing														
2D discrete/inverse cosine transform	Barco-Silex	AllianceCORE		V-II	V-E	V		S-II		77%	133	XC2V250-5		Picture and video, archiving, digital television compression and transmission, teleconference
2D Forward Discrete Cosine Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		42%	83	XC2V500-5		
2D Inverse Discrete Cosine Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		45%	95	XC2V500-5		
Block-based 2D Discrete Wavelet Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		62%	52	XC2V250-5		
Combined 2D Forward/Inverse Discrete Cosine Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		65%	95	XC2V500-5		
Combined 2D Forward/Inverse Discrete Wavelet Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		55%	95	XC2V500-5		
Compact Video Controller	Xylon d.o.o.	AllianceCORE		V-II		V		S-II			88	XC2V250-4	Single & double panel, LCD/CRT support, 4 gray, 256 colors	Video phone, Set-top box, PDA display
DCT/IDCT Forward/Inverse Discrete Cosine Transform, 1-D	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					8-32 pt FDCT, IDCT with 8-24 bits for coeff & input	JPEG, MPEG, H.261, H.263
DCT/IDCT Forward/Inverse Discrete Cosine Transform, 2-D	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		32%	140	XC2V1000-5	8x8 parameterized FDCT, IDCT & IEEE 1180-1990 compliant IDCT	JPEG, MPEG, H.261, H.263

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Implementation Example			Key Features	Application Examples
										Occ	MHz	Device		
Video & Image Processing (continued)														
Discrete Cosine Transform	eInfochips Pvt. Ltd.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		35%	83	XC2V1000-5	8-bit input/12-bit output precision; 76 clock cycle latency	Video coding and security, medical imaging, scanners, copiers, digital still cameras
Fast JPEG color image decoder	Barco-Silex	AllianceCORE		V-II		V				78%	56	XC2V1000-4	Conforms to ISO/IEC Baseline 10918-1, color, multi-scan, Gray-Scale	Video editing, digital camera, scanners
Fast JPEG gray scale image decoder	Barco-Silex	AllianceCORE		V-II	V-E	V				68%	73	XC2V1000-4	Conforms to ISO/IEC Baseline 10918-1, Gray-Scale	Video editing, digital camera, scanners
FIDCT Forward/Inverse Discrete Cosine Transform	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II		77%	78	XCV200-6	DCT for 8X8, 16X16, IDCT IEEE1180-1990 compliant	JPEG, MPEG, H.26X
Huffman Decoder	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		22%	25	XC2V1000-5		
JPEG encoder/decoder	inSilicon Corporation	AllianceCORE				V		S-II	S		20	XCV400E-8	Conforms to ISO/IEC Baseline 10918-1, 4 quantization tables, 4 Huffman tables. Stallable	Video editing, digital camera, scanners
Line-based programmable forward DWT	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		145%	51	XC2V250-5		
Longitudinal Time Code Generator	Deltatec S.A.	AllianceCORE				V				36%	80	XC2S15-5	SMTPE/EBU compliant, PAL/NTSC, lock-on external video reference	Audio/Video recording and editing equipment Motion JPEG Codec
core V1.0	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V					40	Virtex-II		
Motion JPEG Decoder core V1.0	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V					42	Virtex-II		
Motion JPEG Encoder core V2.0	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V					70	Virtex-II		
RGB2YCrCb Color Space Converter	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		29%	96	XC2S50E-7	8-bit I/O, 10-bit coeff, 13-bit internal precision; 5-cycle latency; fully synchronous.	Digital RGB to TV output conversion, image filtering, machine vision, still and video image processing.
RGB2YCrCb Color Space Converter	Perigee, LLC	AllianceCORE				V		S-II	S		202	XCV100E-8	One clock cycle throughput	HDTV, real time TV output modulation
YCrCb2RGB Color Space Converter	Perigee, LLC	AllianceCORE				V		S-II	S			XCV100E-8	One clock cycle throughput	HDTV, real time video
Basic Elements														
Binary Counter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					2-256 bits output width	
Binary Decoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					2-256 bits output width	
Bit Bus Gate	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Bit Gate	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Bit Multiplexer	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
BUFE-based Multiplexer Slice	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
BUFT-based Multiplexer Slice	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Bus Gate	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Bus Multiplexer	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					IO widths up to 256 bits	
Comparator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
FD-based Parallel Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
FD-based Shift Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-64 bits wide	
Four-Input MUX	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Bit Multiplexer or Bus Multiplexer	
LD-based Parallel Latch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Parallel-to-Serial Converter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: FD-based Shift Register	
RAM-based Shift Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide, 1024 words deep	
Register	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: FD-based Parallel Register	
Three-Input MUX	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Bit Multiplexer or Bus Multiplexer	
Two-Input MUX	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Bit Multiplexer or Bus Multiplexer	

Please visit the Xilinx IP Center for the most up-to-date information at **www.xilinx.com/ipcenter**

Eureka Technology offers a wide range of silicon proven IP cores for SoC designs. By using our free web-based **SoCDesigner**™ software, our customers are able to combine the benefits of using pre-verified silicon proven IP cores with the ability of customizing the SoC design according to their exact specifications. Our rich repertoire of IP cores enable our customers to build a wide varieties of system controller/CPU companion chips to support different CPU and bus standards such as **ARM**™, **PowerPC**™, **ARC**™, **MIPS**™, **SH**™**2/3/4**, **PCI**™, **Cardbus**™ and **PCMCIA**™.

To learn more about our silicon proven IP cores and SoCDesigner, please visit **www.eurekatech.com/socdesigner**

Eureka Technology

Call today for more information
tel: +1 650 960 3800
email: info@eurekatech.com
www.eurekatech.com
4962 El Camino Real, Los Altos,
CA 94022 USA

New
Up to 600K Gates and 514 I/Os

Got-a-Lotta-Pins. Not-a-Lotta-Price.

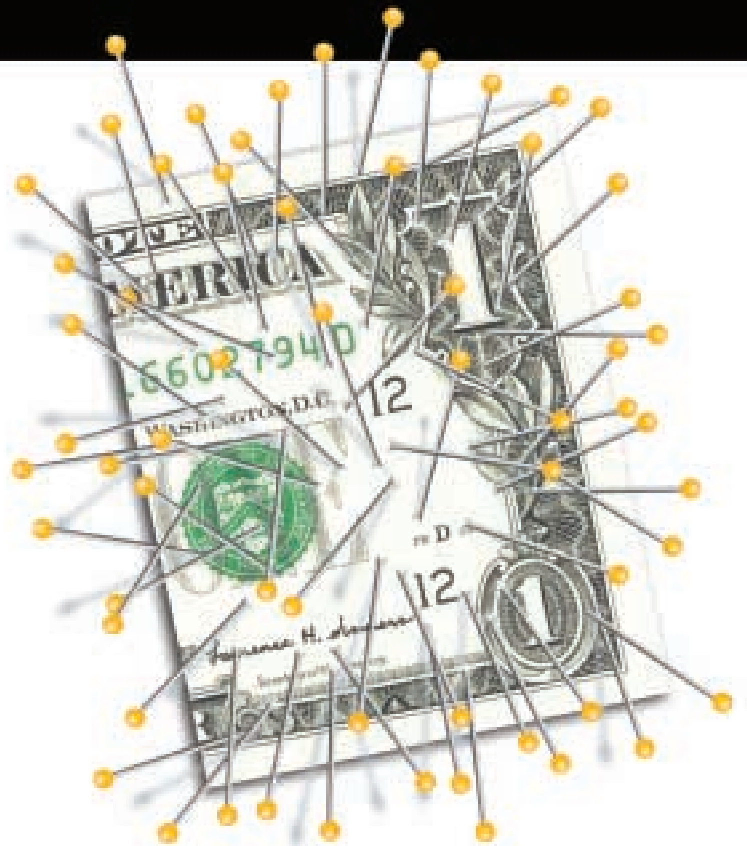

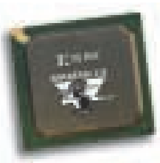

With up to 514 I/Os, the newly extended Spartan®-IIE family offers the lowest cost per pin in the industry. That's why Spartan-IIE FPGAs are the first choice of designers seeking a low-cost, higher density solution for high-volume applications.

All the I/O. All the density. No compromises.

Spartan-IIE FPGAs achieve minimum die size without sacrificing I/Os. With double the competitive I/O count, and densities ranging from 50,000 to 600,000 system gates, the Spartan-IIE series is the only true solution to ASIC headaches.

Fits your budget. Hits your market.

Spartan-IIE FPGAs lead the way in quick-turn, cost-sensitive markets like consumer, digital video, home networking, automotive and much more. Supporting 19 I/O standards and driven by Xilinx's proven, lightning-fast ISE 5.1i software, there's no better low-cost solution shipping today.

SPARTAN®-IIE
FPGA

Visit ***www.xilinx.com/spartan2e/*** today and find out how you can get all the pins you need . . . at the price you want.

XILINX®
The Programmable Logic Company℠
www.xilinx.com/spartan2e/

FORTUNE® 2003
100 BEST COMPANIES TO WORK FOR

PN 0010707