ISSUE 47, FALL 2003 XCELL JOURNAL XILINX, INC.

Issue 47
Fall 2003

Xcell journal

THE AUTHORITATIVE JOURNAL FOR PROGRAMMABLE LOGIC USERS

Discover New Applications For Low-Cost Solutions

SOFTWARE

Xilinx ISE 6.1i Lowers Design Costs – Again

SUPPORT

MySupport.xilinx.com Named One of Top 10 Best Web Support Sites

Maximize Spartan-3 Designs with Xilinx Global Services Education and Support

APPLICATIONS

Performance + Time = Memory 3D Design Saves Money

Virtex-II FPGAs Enter Formula 1 Racing

How to Implement a DDR SDRAM Controller

NEW PRODUCT

New Platform Flash PROM Family Minimizes Costs and Board Space

COVER STORY

5,700 Spartan FPGAs Bring BioWall to "Life"

XILINX®

Add Bluetooth to your MicroBlaze-based designs.

(It doesn't get any easier than this.)

Now there's a faster, easier way to build and test a Bluetooth® system using embedded processors. It's the Embedded Bluetooth Development Kit from Memec Design, and it provides everything you need to add a Bluetooth interface to existing Xilinx® MicroBlaze™ designs. Developed with industry-leading partners, this unique kit can help you save weeks of development time in your challenging next-generation applications. To find out more, contact us at **www.insight.na.memec.com/xilinx_bluetooth**.

XILINX® The Programmable Logic Company℠ | Memec Design | stonestreet one | BROADCOM.

Insight MEMEC
Americas, Europe, Asia, & Japan

www.insight.na.memec.com/x-fest2003
Visit X-Fest 2003 to find out more.

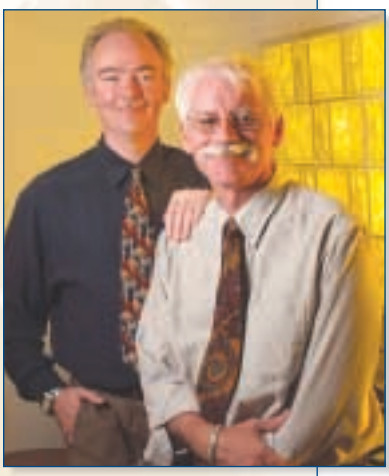

EDITOR IN CHIEF — Carlis Collins, editor@xilinx.com, 408-879-4519

MANAGING EDITOR — Tom Durkin, tom.durkin@xilinx.com, 530-271-0899

ASSOCIATE EDITOR — Jack Shandle

COPY EDITOR — Charmaine Cooper Hussain

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

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

ART DIRECTOR — Scott Blair

Xcell journal

Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124-3400
Phone: 408-559-7778
FAX: 408-879-4780
©2003 Xilinx, Inc.
All rights reserved.

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, RocketPHY, 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.

Xilinx Wins Again

Xilinx has built its reputation on the triumvirate of silicon, software, and support. And although silicon and software get the most news coverage, it's the support services team that holds it all together.

The *support.xilinx.com* website – and the personalized *mysupport.xilinx.com* portals – are accessed by more than 110,000 users each month, and are credited for saving customers substantial time to knowledge, which translates to faster time to market.

In fact, the *MySupport.xilinx.com* support service is so good that the Association of Support Professionals has named it to its list of "The Year's Ten Best Web Support Sites." Xilinx joins fellow industry leaders Dell, Hewlett Packard, Mentor Graphics, and Novell on the list of companies awarded for excellence in online service and support.

The Year's Ten Best Web Support Sites

To select this year's winners, judges evaluated companies on such criteria as usability, design, navigation, and knowledge base, as well as search implementation, interactive features, personalization, and major site development.

According to Jeffrey Tarter, executive director of the Association of Support Professionals, "Users are looking for highly interactive technical support services they can personalize to fit their needs and can access in real time. Xilinx is a leader in delivering the kind of online support customers are looking for."

Co-winner Mentor Graphics also praised *MySupport.xilinx.com*: "Many of our customers have told us that the quality and coverage of *support.xilinx.com's* technical content allows them to readily understand complex FPGA architectures."

As we said, Xilinx support doesn't get as much publicity as Xilinx silicon and software, but it is just as important.

Don Mullikin, vice president of Xilinx Global Services, puts it this way: "Design teams tackling today's most complex system-design challenges need support that matches the caliber of the technical capabilities of our products."

Our award-winning support matches our award-winning hardware and software, so you get a complete design solution that is the best in the industry.

Tom Durkin

Tom Durkin
Managing Editor

Correction

In the summer 2003 edition of *Xcell Journal*, the URL for more information on RocketPHY™ transceivers was incorrectly reported in the article "New RocketPHY Transceiver Family Debuts at 10 Gbps." The correct URL is *www.xilinx.com/rocketphy/*. We regret this error and any inconvenience it may have caused.

FALL 2003, ISSUE 47

Xcell journal

COVER STORY

Biology Goes Digital

An array of 5,700 Spartan FPGAs brings the BioWall to "Life."

ISE 6.1i Continues to Lower Design Costs

The latest release of Xilinx programmable logic design software saves you time – and money

Performance + Time = Memory

By approaching FPGA designs as three-dimensional endeavors, you can radically reduce device size – and cost.

Formula 1 Racing: The Xilinx Advantage

The BMW WilliamsF1 team lifts the hood on its fifth-generation vehicle control system

46

To subscribe to the printed *Xcell Journal*, or to view the Web-based *Xcell Online* Journal, visit: *www.xilinx.com/publications/xcellonline/*

Innovation, Education, Exploration

From Programmable World in 2003 to landing on Mars in 2004, Xilinx leads the industry in new products and new applications.

by Wim Roelandts
CEO, Xilinx, Inc.

Xilinx products and services continue to expand the limits of possibility. We strive to make the best and most innovative products in the programmable logic industry – and deliver them as a complete solution that gives you everything you need to bring your product to market first. Here are a few of our innovations that may interest you.

Programmable World

On May 6, Xilinx launched Programmable World 2003 – a nine-month international initiative designed to help you understand and use the latest programmable logic solutions. It includes technical forums, exhibits, and workshops in dozens of cities around the world. By the time it ends in November, Programmable World 2003 will have reached more than 10,000 executives, system architects, engineers, and project managers worldwide. Co-sponsored by more than 30 industry leaders – including Agilent, Cadence, IBM, Intel, Texas Instruments, Mentor Graphics, and Wind River Systems – Programmable World 2003 started in North America, then went to Europe, on to China and Taiwan, and will finish in Japan and Korea in November.

Following the worldwide forums, which ended in August, we have been holding in-depth technical workshops, including multiple full-day courses on high-speed DSP, connectivity, processing, and system design.

At the Programmable World 2003 forum in Japan this summer, Dr. Tsugio Makimoto, chief technical officer of Sony, asserted that field programmability is "hitting its peak" from 1997 through 2007. Dr. Makimoto stated a second wave of digital consumer products will become the mainstay for future market expansion, just as personal computing products drove the first wave of the digital revolution. "Field programmability is a must in this new era of digital consumer goods," Dr. Makimoto said. "These products have short time-to-market windows, and product life cycles that come quickly and end dramatically. Flexible, agile solutions are critical."

Xilinx Goes to Mars

NASA recently launched two Mars Exploration Rover (MER) missions, which depend on the most advanced, radiation-tolerant Virtex™ FPGAs in critical applications for both the landers and rover vehicles. The Spirit MER was launched June 10 and the Opportunity MER was launched July 7.

Chosen because of their reprogrammability and high densities, the Virtex FPGAs serve as the main brain of the motor control boards. These FPGAs will be used to control the pyrotechnic devices that will cushion the touchdown of the landers in January 2004. The Xilinx FPGAs will also direct the motor control functions on the rovers, including controllers for the wheels, steering, and antenna gimbals.

The radiation-tolerant Virtex devices on the two MER missions have up to one million system gates and total ionizing dose guaranteed to 100 kRads(si). We will soon announce a next-generation family of radiation-tolerant Virtex-II products with up to six million system gates and total ionizing dose guaranteed to 200 kRads(si).

RocketPHY

In May, we introduced our new RocketPHY™ stand-alone, SONET-compliant, 10 Gbps transceivers. These devices are among the industry's first 10 Gbps products using CMOS process technology.

The RocketPHY transceivers are the first products from Xilinx that are not directly related to programmable logic devices, but they are designed to work with our FPGAs. These transceivers use the same high-speed I/O technology as our Virtex-II Pro™ X FPGAs in application-specific designs that solve many I/O challenges – at far less cost than any competing product in networking applications.

Conclusion

Xilinx continues to innovate in technology and penetrate new markets. Although our process technology, devices, and software are continually setting the highest standards for performance, quality, and reliability, we are also setting new standards in other critical areas that affect your life. You'll find that our products, our support services, and our educational programs work together to bring you a cohesive, faster time to market, and lowest cost path to your next product introduction. There simply is no easier, faster, or lower cost way to develop flexible new products than with Xilinx silicon, software, and support.

ISE 6.1i Continues to Lower Design Costs

The latest release of Xilinx programmable design software saves you time – and money.

by Lee Hansen
Product Solutions Marketing
Xilinx, Inc.
lee.hansen@xilinx.com

The newest version of Xilinx Integrated Software Environment programmable design tools – ISE 6.1i – is now available to registered in-maintenance Xilinx customers.

This new release of our industry-leading programmable design suite continues the Xilinx commitment to lower your design costs, to deliver ease of use that focuses on solving your specific engineering problems, and to expand support for low-cost Spartan™-3 FPGAs – and native support for the Red Hat™ Linux™ operating system.

Some of the new capabilities in ISE 6.1i include:

- Support for the largest Spartan-3 XC3S4000 and XC3S5000 devices
- Increased Spartan-3 device support in ISE WebPACK™ and ISE BaseX
- New ISE native Linux support
- Automatic Web software update
- PACE (Pinout and Area Constraints Editor) enhancements for CPLD, and PC board design assistance
- Project Navigator mixed-language and embedded design enhancements
- As much as 15% better performance than ISE 5.1i.

ISE Shrinks Costs

Spartan-3 Platform FPGAs are the world's lowest cost FPGA device family. ISE 6.1i software will help you lower the cost of using Spartan-3 FPGAs even further with these new features.

More Spartan-3 Devices in ISE

ISE WebPACK 6.1i, which is free to designers, has added support for the Spartan-3 XC3S200 and XC3S400 devices, at 200,000 and 400,000 system gates, respectively. ISE BaseX now supports the XC3S400 device as well.

High Density Spartan-3 Support

ISE Foundation™ and ISE Alliance Series™ 6.1i versions support the highest-density Spartan-3 devices, the XC3S4000 and XC3S5000 FPGAs at 4 million and 5 million gates, respectively. Now you have even greater flexibility in choosing Spartan-3 devices and ISE configurations.

ISE Supports Native Linux

ISE 6.1i is also the first ISE release that runs on native Red Hat Linux versions 7.3 and 8.0. The installation CDs for 32-bit Linux that come with your ISE 6.1i shipment will help you make the best use of your corporate programmable design platforms.

Automatic Web Update

Another new capability is automatic Web software update. Upon execution, ISE will notify you whenever a new service pack upgrade is available. And if selected, ISE will only download those parts of the service pack that apply to your unique installation. This feature saves you the time of identifying and updating your current software configuration – and minimizing the required disk space.

Ease of Use Engineers Demand

Many software design companies settle for a good look-and-feel graphical user interface as their standard for "ease of use."

ISE 6.1i goes beyond being just another pretty GUI. It focuses on solving engineering bottlenecks and design headaches that hinder your design process and progress.

PACE Enhancements

Figure 1 shows an example of PACE (Pinout and Area Constraints Editor), introduced with ISE 5. PACE delivers pin definition and area management in an easy-to-use, graphically oriented environment. You can speed your design flow faster and easier with PACE.

PACE now offers CSV (comma separated value) file import and export. This capability offers you new flexibility in PC board design, including the ability to create pin tables in Microsoft™ Excel™ spreadsheets and import those into PACE. If the pin tables are modified, they can be exported back to the Excel workbook using the CSV format. This export/import capability eases the job of integrating the logic device into the board layout.

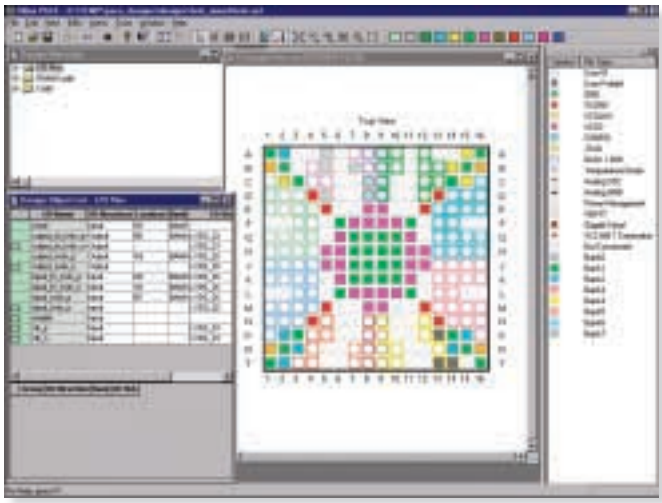

Figure 1 – Pinout and Area Constraints Editor (PACE)

PACE can also import and export VHDL and Verilog™ HDL files, which allow PACE to define I/O from the HDL port definitions – or write top-level HDL starting templates.

New design check capabilities help you predict output problems. PACE contains a new package flight-time display that graphically shows pin-delay time that is based on pin-to-pad estimates. PACE also checks for simultaneous switched outputs to prevent common high-drive strength outputs that could potentially create ground bounce signal problems.

The new version of PACE also supports an enhanced auto-floorplanning capability, which lets you identify area groups using PACE. Once a logic area group is identified, the ISE place-and-route tools create the floorplan, which saves you more design time.

Expanded Project Navigator

Project Navigator, the ISE design and project manager, has also been enhanced in version 6.1i. Project Navigator now supports mixed-language Verilog and VHDL design for customers using Synplicity® Corp.'s Synplify® tool suite or XST (Xilinx Synthesis Technology) software for their synthesis solutions. This new flexibility allows managers to mix the best possible design source code for any particular project. This, in turn, allows you to more easily and quickly mix and match your purchased IP with your own in-house design expertise regardless of design language.

Project Navigator also now links to the Xilinx EDK (Embedded Design Kit) XPS project manager supporting MicroBlaze™ and Virtex-II Pro™ embedded processor designs. This new integration shows an embedded project entity along with the design logic, and launches XPS when double-clicked, offering the first in a new series of upcoming enhancements that will bring Xilinx logic and embedded programmable design tools closer together.

Still the Fastest

ISE 6.1i continues to deliver the fastest programmable device performance available. Enhancements to our lightning-quick ProActive Timing Closure implementation technology now deliver up to 15% better performance over ISE 5.1i software.

The new INPUT_JITTER timing constraint lets an engineer describe system jitter and clock edge uncertainty. With more timing constraints, high-speed design rules, and local clocking options than any other programmable vendor, ISE 6.1i gives you the ability to design high-speed memory interface timing and double-data-rate local clock designs accurately.

Conclusion

ISE continues to define the standard of logic design. By concentrating on cost, productivity, and ease of use, ISE is delivering the tools necessary for programmable systems design that helps you squeeze the most out of your logic device.

To find out more about ISE 6.1i, go to *www.xilinx.com/xcell_ise/*. To order your copy of ISE 6.1i, contact your local sales support representative.

Put Hardware in the Loop with Xilinx System Generator for DSP

Co-simulating with hardware in the loop gives you faster simulations and eases hardware verification. System Generator for DSP Version 3.1 now lets you include FPGA hardware in Simulink simulations.

by Nabeel Shirazi
Senior Staff Software Engineer
Xilinx, Inc.
nabeel.shirazi@xilinx.com

Jonathan Ballagh
Senior Software Engineer
Xilinx, Inc.
jonathan.ballagh@xilinx.com

FPGA hardware can now be included in simulations controlled by the MATLAB™ language and its Simulink™ design tools. At the push of a button, the Xilinx System Generator for DSP Version 3.1 software tool produces an implementation of your model that is ready to run in hardware.

System Generator uses a special co-simulation block to control the design hardware during Simulink simulations. The co-simulation block looks just like traditional System Generator blocks. Its ports have names and types that match the ports on the original System Generator subsystem. The co-simulation flow for this architecture is shown in Figure 1. The simulation behavior of the co-simulation block is bit- and cycle-accurate when compared to the behavior of the original subsystem.

System Generator's ability to incorporate hardware in the loop enables significant simulation speedups, provides incremental hardware verification capabilities, and removes many of the hurdles to getting a design up and running in an FPGA.

Hardware Co-Simulation Flow

Together, System Generator and hardware-in-the-loop co-simulation solve many problems normally associated with simulating and verifying a Xilinx FPGA design.

- You don't have to know a hardware description language. System Generator performs automatic code generation when your model is translated into hardware.
- You don't need in-depth knowledge of the board that hosts the FPGA. System Generator makes sure the correct hardware platform is targeted and implements the design accordingly.
- You need not create a separate test bench application. The co-simulation block can be used with the same Simulink test bench apparatuses that were used to test the original System Generator model.

The starting point is the System Generator model – or subsystem – that you want to co-simulate. You identify the subsystem by dragging and dropping a

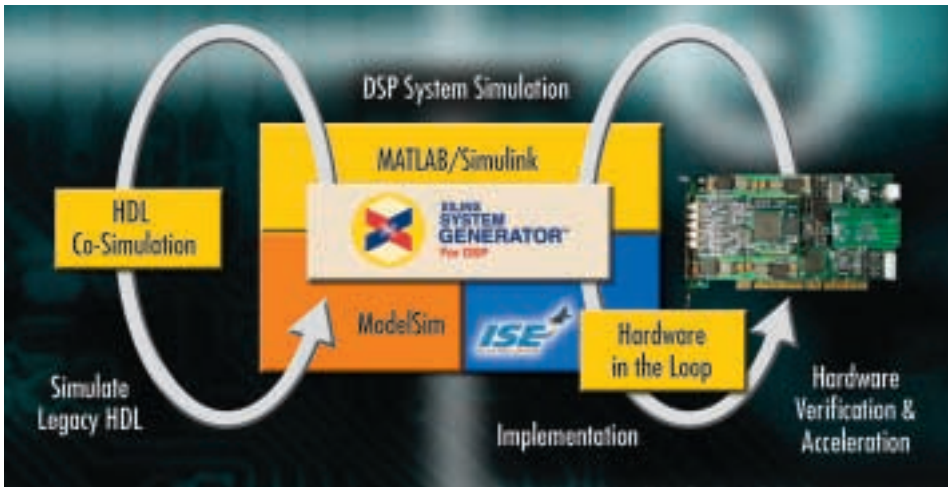

Figure 1 – MATLAB/Simulink and hardware co-simulation

compilation block from the XtremeDSP™ development kit library, for example, into the design, as shown in Figure 2.

A subsystem that contains a compilation block should also contain a System Generator block. System Generator uses the compilation block to provide information about the underlying hardware that the product, device, speed, and package fields have been grayed-out and are inactive. These parameters are used by System Generator when you press the "generate" button and the model is translated into hardware. This ensures that the correct device information is preserved when the FPGA configuration bit file is generated.

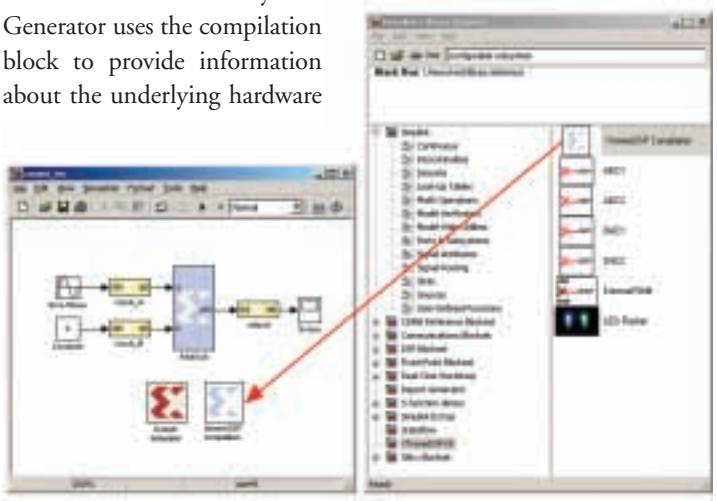

Figure 2 – System Generator block and an XtremeDSP compilation block

Generating a model for hardware co-simulation includes operations such as producing HDL code and netlists and invoking the Xilinx CORE Generator™ tool. After System Generator has generated a model, it invokes a compilation block script. Using the output from System Generator, the compilation block script invokes the tools necessary to produce a configuration file for the FPGA. Figure 4 illustrates the XtremeDSP compilation block script running ngdbuild.

The script creates a new library and adds a co-simulation block that is parameterized with information from the original subsystem. This information includes subsystem interface information such as port names, port data rates, port data types, and port directions. Figure 5 shows the library and co-simulation block created by the compilation block script for the cosim_ex model.

Figure 6 shows the co-simulation block first illustrated in Figure 5 as it is brought back into the cosim_ex model for simulation. Note that the ports on the co-simulation block match the port names on the input and output gateway blocks. The waveform in the scope shows equivalent output results for the AddSub block and the co-simulation block. The block's inputs must be the same signal type, as shown in Figure 7.

Figure 3 – System Generator dialog box with inactive product family, device, speed, and package fields

and how to prepare the model for co-simulation. Pressing the "generate" button on a System Generator block that has an accompanying compilation block automatically invokes the hardware co-simulation flow.

Some of the information obtained from a compilation block overrides values on the System Generator block parameters dialog box. The parameters that are overridden appear as grayed-out fields in the System Generator dialog box and cannot be modified. For example, the System Generator dialog box is shown for a subsystem that also contains an XtremeDSP compilation block. Notice in Figure 3

Using Co-Simulation Blocks in Simulink

A co-simulation block is a Simulink block that interacts with an underlying co-simulation platform. This block behaves as a normal Simulink block and can be simulated transparently with other blocks. Each co-simulation block has a port interface that is parameterized by the compilation block script to match the ports on the original subsystem.

System Generator co-simulation blocks can be driven by either Xilinx fixed-point data types or by double-signal data types. The hardware co-simulation block must be of the same type of signals (bit width, binary point position, signed/unsigned) that were used in the original model. If there are no inputs to the block, the block allows you to choose between Xilinx fixed-point or double-signal output types.

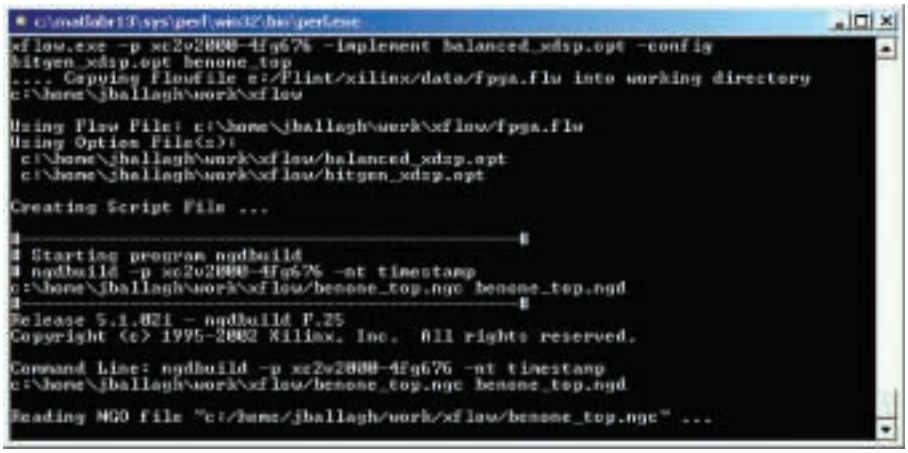

Figure 4 – Compilation block script for the XtremeDSP development kit

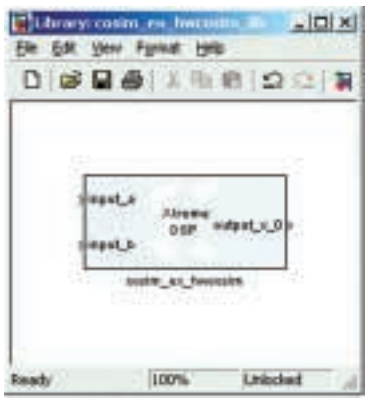

Figure 5 – Co-simulation library – cosim_ex_hwcosim_lib – created by the post-generation script.

When System Generator implements a model in hardware, it uses a single clock source to drive the design. Multi-rate blocks are supported by using derived clock enables. These clock enables have periods that are integer multiples of the period of the system clock. Data traveling in and out of the hardware also has a period that is a multiple of the system clock period. These multiples, or ratios, between the port data rates and the system clock rate are stored in the co-simulation block by the compilation block script when it is instantiated.

A co-simulation block must be driven by ports with data rates equivalent to the gateway data rates of the original subsystem. The block automatically adjusts its output ports to run at the appropriate rates. If the block detects an incorrect input port rate, it will generate an error message.

The easiest way to ensure that an input port has the appropriate rate is to drive the co-simulation port with the same signal that drove the original subsystem port or gateway.

Hardware Clocking

There are many possible synchronization techniques for the interface between FPGA hardware and Simulink. One technique uses a single-step clock to keep the hardware in lockstep with the software simulation. This is achieved by providing a single clock pulse to the hardware for each simulation cycle. Using this technique enables you to perform incremental design and verification.

On the other hand, the single-step clock technique has its disadvantages. When it is used for clocking the FPGA design, the communication overhead between hardware and Simulink can severely limit the effective processing rate in some designs. This condition is exacerbated by bus latency.

Simulation speed can be greatly increased by allowing the hardware to process more than one set of input samples at a time. One way to accomplish this is to provide a free-running clock to the design under test. Use an explicit synchronization mechanism such as a flag implemented as a memory-mapped register to coordinate data transfers between Simulink and the hardware. The inputs and outputs of the design are written to and sampled asynchronously.

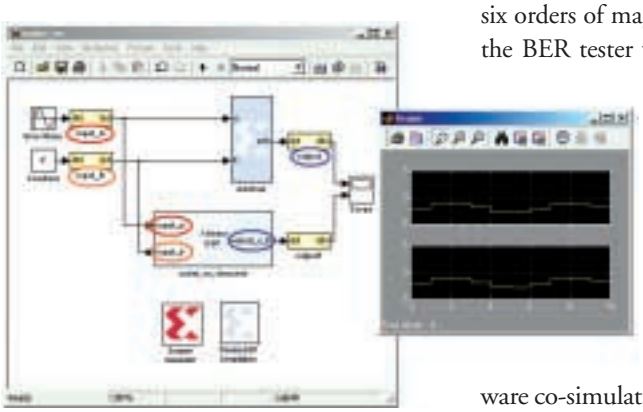

Figure 6 – Co-simulation block parameterized with port names that match the original model.

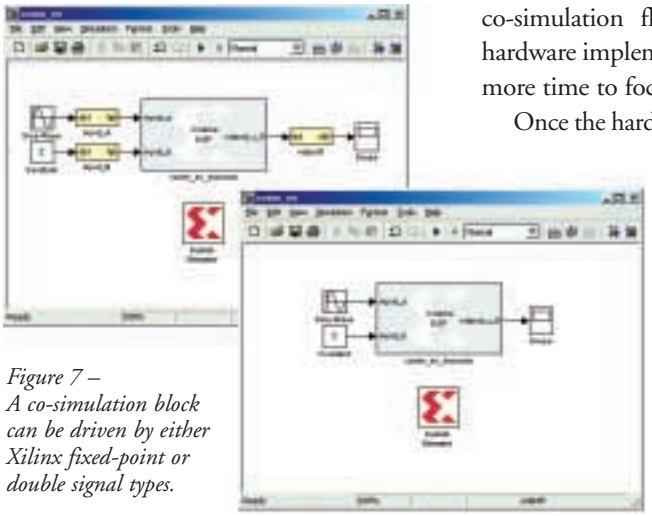

Figure 7 – A co-simulation block can be driven by either Xilinx fixed-point or double signal types.

Benchmarks

When using a single-step clock, speedups of seven to 50 times are typically achieved. The demos that come with System Generator were used as benchmarks. The results are shown in Table 1.

Referring to the table, the bit error rate (BER) tester was created to test forward error correction (FEC) blocks. A free-running clock source achieved a speedup of six orders of magnitude. The data source for the BER tester was an LFSR in the FPGA. The test is started, and after some time when a "done" flag is set, you can read the results from the FPGA and display them in Simulink.

Conclusion

System Generator hardware co-simulation provides you with a simple but powerful method of incorporating hardware co-simulation in Simulink and System Generator models. The hardware co-simulation flow automates the entire hardware implementation process. You have more time to focus on the design itself.

Once the hardware co-simulation software has produced a co-simulation block, it can be used just like any other System Generator block in a model.

To learn more about System Generator for DSP, go to *www.xilinx.com/system_generator/* and click on "System Generator for DSP."

Application	Software Simulation Time	Hardware Execution Time	Speedup
5x5 Image Filter	170 sec	4 sec	43x
Cordic Arc Tangent	187 sec	27 sec	7x
Additive White Gaussian Noise Channel (AWGN)	600 sec	80 sec	7.5x

Table 1 – Benchmark results for hardware-in-the-loop co-simulation

Get Physical with the PALACE Synthesis Solution

Aplus Design Technologies' PALACE physical synthesis tool can drive down your costs by helping you quickly meet or exceed your design requirements with a slower speed grade or a smaller device.

by Sheldon D'Paiva
Applications Engineer
Aplus Design Technologies
sdpaiva@aplus-dt.com

Xilinx Spartan™-3 FPGAs give designers advanced features at previously unattainable costs. Moreover, Spartan-3 densities reaching five million system gates provide cost-competitive solutions in markets once only addressable by ASICs. But as FPGAs penetrate the ASIC space, FPGA designers must deal with the same complexities faced by ASIC designers.

In particular, ASIC designers know that to meet design requirements and reduce iterations within the design flow, synthesis must be coupled with physical design considerations.

Synthesis design flows that do not take physical design into account can erode your time to market and cost advantages. With the large densities and complexities of modern low-cost FPGA architectures such as Spartan-3 FPGAs, you cannot afford to ignore interconnect delay and architecture-specific considerations during synthesis.

Poorly synthesized designs cause greatly increased runtimes for the place-and-route tools, as they struggle to meet timing constraints. Bad designs also create an additional burden on the designer by requiring expert manual intervention and time-consuming iterations between synthesis and implementation tools.

Furthermore, additional costs are incurred if the design requirements are not met in the targeted device. For high-volume and low-density devices, moving to a faster speed grade typically translates to a 13% increase in cost; moving to a larger part translates to a 16% increase in cost. This increase in cost is even greater for lower volume designs or when using higher density devices.

Physical synthesis couples synthesis with physical design. It has become a key enabling technology that can help you maximize the potential of FPGAs, especially low-cost FPGAs such as Spartan-3 devices.

Aplus Design Technologies, a leader in physical synthesis technology, has recently added Spartan-3 support to its fully automated physical synthesis product called PALACE™ (Physical And Logic Automatic Compilation Engine). The PALACE physical synthesis tool combines logic optimization, architecture-specific mapping, and placement-driven synthesis with detailed device modeling.

Using the PALACE tool, you can meet your design requirements without manual intervention, floorplanning, or numerous lengthy iterations between synthesis and place-and-route.

In fact, by fully exploiting the capabilities of the FPGA architecture, you can often meet your design requirements in a slower speed grade or a smaller part and significantly reduce your costs.

PALACE Physical Synthesis

PALACE physical synthesis employs a suite of architecture-specific optimization techniques in both the logic and physical domains. These techniques include logic restructuring, flip-flop repositioning, constraint-driven mapping with delay-area tradeoffs, logic duplication, placement, and placement-driven optimization.

PALACE technology highlights include:

- Unified timing models
  - Detailed device + interconnect modeling
  - Consistent model throughout physical synthesis
- Physical planning
  - Global interconnect planning
  - Timing-driven routability-aware placement
- Optimization in both logic and physical domains
  - Timing and area optimization
  - Multi-clock constraint-driven retiming
  - Logic restructuring and replication
- Technology mapping
  - Architecture-specific mapping
  - Constraint-driven with delay-area trade-off.

The benefits of the PALACE solution include:

- A fully automated solution with no expert user intervention required
- Faster timing closure by reducing iterations within the design flow
- Performance improvement by at least one speed grade
- Reduction in area utilization
- Reduction in place-and-route runtime.

PALACE Design Flow

There are two design flows that can be used with PALACE technology:

1. Standard flow: This is the initial flow that should be used for your design.
2. Guided flow: If necessary, this flow should be used to achieve incremental improvements for a fully placed design (for Virtex™-II FPGAs).

Standard Flow

When you use the PALACE tool in the standard flow, you can use any synthesis tool that generates an EDIF netlist. This dramatically reduces your learning curve and allows you to choose the design environment that you prefer. Your design flow is essentially preserved, with only a simple extra step required after synthesis.

Once you have synthesized your design, the resulting netlists and constraints are passed as inputs to the PALACE program. The PALACE tool processes the netlists together with any design constraints. Upon successful completion, the PALACE engine will generate an optimized netlist and a constraint file. You then use the optimized netlist and constraint file as you normally would in the rest of your design flow. Figure 1 shows how the PALACE solution fits seamlessly into your design flow.

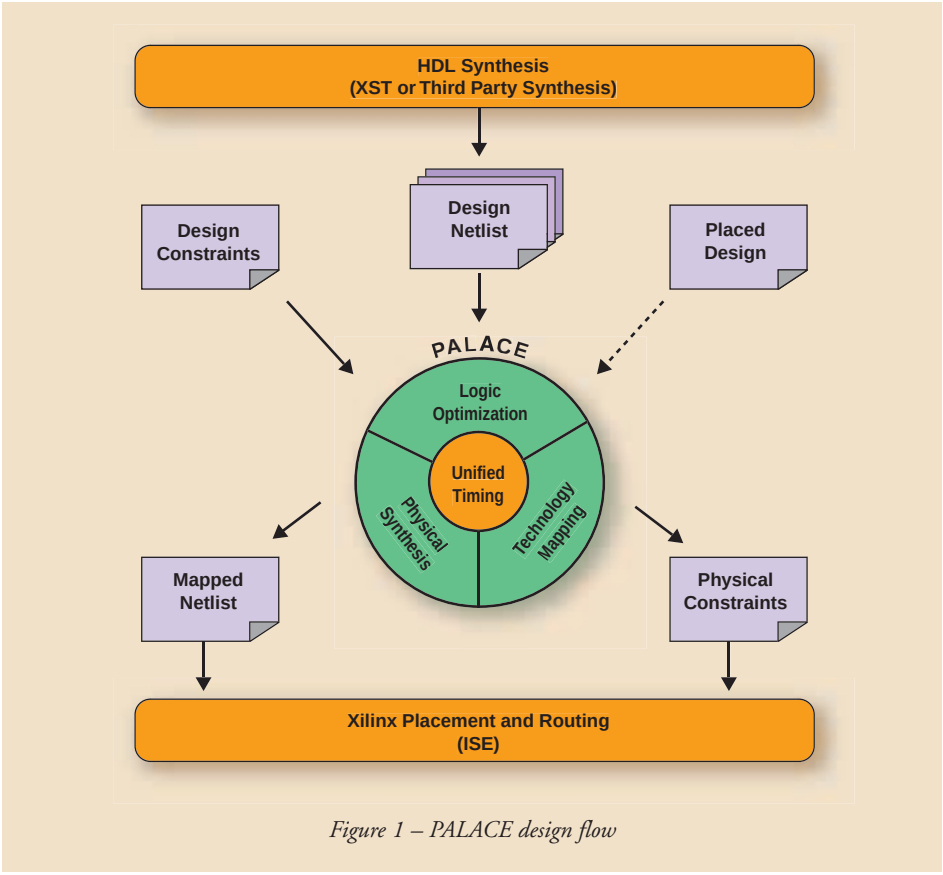

Figure 1 – PALACE design flow

You can control PALACE operations in the standard flow by a simple four-level effort option.

The effort level determines how aggressively the PALACE tool will try to optimize your design, with the first level optimizing for minimum area utilization and the remaining levels optimizing for maximum performance. The default level is the third level, and it is the initial recommended level when performance is an issue.

After specifying the effort level, no further user intervention is required. Using PALACE technology in this flow provides you with the advantages of physical synthesis in a simple push-button automated flow, and can meet most realistic design requirements.

Guided Flow

You can use the PALACE tool in the guided flow for Virtex-II devices to achieve the incremental improvements needed to meet your design requirements. In this flow, you first generate an NCD (native circuit description) file by placing and routing your design. This NCD file is then converted to XDL (Xilinx Design Language) format with the Xilinx XDL utility and passed as an input to the PALACE program, along with the design constraint file. The PALACE tool optimizes the paths that do not meet timing requirements to help you achieve timing closure. As with the standard flow, an optimized netlist and constraint file is produced that you use in the rest of your design flow.

You can control PALACE operation in the guided flow by a simple two-level effort option, which determines how aggressively the PALACE program will try to optimize your design. As with the standard flow, no further user intervention is required. Using PALACE physical synthesis in the guided flow provides you with the capabilities of an advanced physical synthesis solution in a simple push-button automated flow. Thus, you can meet the "last mile" performance requirements of your design.

Achieving Best Results with PALACE

Effective use of the PALACE physical synthesis solution can help you meet your design requirements quickly. With the standard and guided flows, you have two simple but powerful solutions that should be used together in order to achieve the best results.

When you use the PALACE engine in either flow, you should always include timing constraints in the input constraint file. Meaningful and accurate timing constraints are important because they help the PALACE program to focus on the problem areas of your design, while allowing trade-offs with other non-critical areas.

If you are particularly concerned about area utilization, you should use the first effort level on your initial run. Otherwise, you should use the default settings, which will automatically run PALACE physical synthesis in standard flow at a high-optimization effort level.

After you have obtained the optimized results, you can analyze either the PALACE report file or the report from an implementation run to determine if the result meets your requirements. If you miss your timing targets by a wide margin (more than a few nanoseconds in Spartan-3 devices), you should try running the PALACE engine in the standard flow with a higher effort level. If you are using a Virtex-II device and have still not met timing requirements, you should try the guided flow to achieve the last nanosecond or so of required performance.

The performance gains that you obtain with PALACE physical synthesis will vary depending on the type and complexity of the design. Table 1 shows performance improvements obtained when the PALACE solution was used in the standard flow for a few sample designs from a variety of categories. In all of these cases, the design flow was identical for two runs, except that the PALACE solution was used before the implementation stage in the second run. As you can see, PALACE physical synthesis provides a substantial increase in performance that allows you to move to a slower speed grade in many cases.

Category (Design)	Spartan-3 Device	Performance Gain (% Max Frequency)
DSP (DES)	XC3S400FT256-4	37%
Microcontrollers (uP1232a)	XC3S200FT256-4	28%
Communications (Reed-Solomon decoder)	XC3S200PQ408-4	8%
Bus Interfaces (I^2C Master)	XC3S50PQ208-4	12%
State Machine and Control Logic (Arbiter)	XC3S400FT256-4	93%

Table 1 – Examples of PALACE performance improvements

Conclusion

As FPGAs continue to increase in density and complexity, you need to ensure that your tools extract the maximum potential of the device's architecture in the minimum amount of time. This is especially true for cost-sensitive design cycles that involve the low-cost Spartan-3 FPGAs. Without a physical synthesis solution that can effectively exploit the architecture of your target device, you risk overrunning your forecasted costs by having to move to a more expensive part, or by spending too many engineering hours trying to achieve timing closure.

The Aplus PALACE tool is an advanced physical synthesis solution that fits seamlessly within your existing design flow for all Xilinx Spartan-II, Spartan-IIE, Spartan-3, Virtex, Virtex-E, Virtex-II, and Virtex-II Pro™ FPGAs, providing you with a fully automated solution to help achieve your design requirements in the minimum amount of time. In fact, you may even be able to move to a smaller part or a slower speed grade and dramatically reduce your costs.

To learn more about the PALACE physical synthesis tool, e-mail *info@aplus-dt.com* or visit *www.aplus-dt.com.* Σ

Seamless Co-Verification Accelerates Time to Market

The Seamless co-verification tool from Mentor Graphics has been customized to provide an efficient debug methodology for Virtex-II Pro Platform FPGAs.

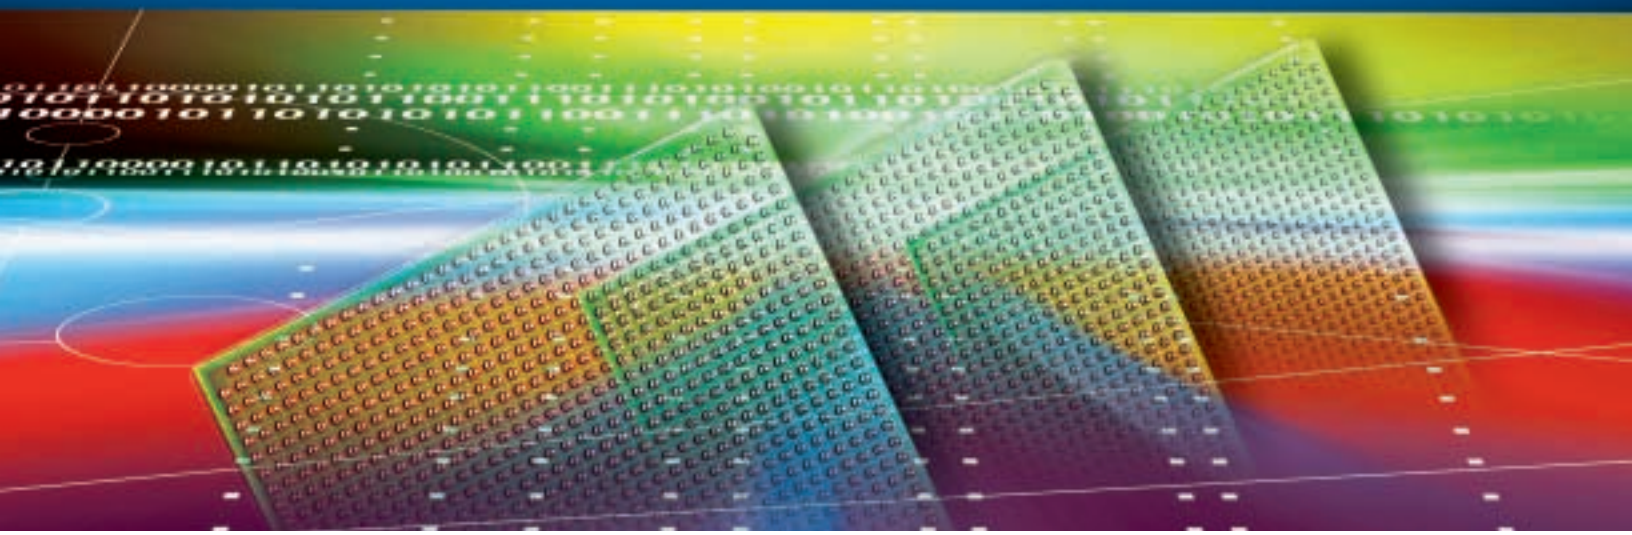

by Milan Saini
Technical Marketing Manager
Xilinx, Inc.
milan.saini@xilinx.com

The growth in the system-on-chip (SoC) market is fueled by the need to increase performance and reliability while reducing overall system costs.

With the introduction of the latest Xilinx Platform FPGAs, ASIC-style integration and performance is now available in programmable logic. This means that FPGA designers need access to newer electronic design automation (EDA) tools and methodologies to maintain design efficiency and productivity.

Hardware/software (HW/SW) co-verification is an example of one ASIC methodology that has recently gained relevance for platform FPGAs. In this article, I will introduce the concept of co-verification, and will describe the role, relevance, and realizable benefits in the context of programmable systems. I will also explain the actual development work we've done in cooperation with Mentor Graphics to extend existing HW/SW co-verification solutions to effectively target Xilinx devices.

The Debug Challenge

With as many as four IBM™ PowerPC™ 405 processors, ultra high-speed serial I/O technology, and up to 10 million system gates, Xilinx Platform FPGAs represent a paradigm shift in programmable platforms, as illustrated in Figure 1. These devices offer a unique and unprecedented combination of flexibility, performance, and integration to widely expand the size, scope, and range of applications that can now be deployed on FPGAs.

Just as leading-edge process technologies have helped silicon to achieve amazing capabilities, design software must keep pace by delivering productivity-enhancing tools that make it easier to design and debug. By various accounts, design verification is the most serious bottleneck engineers face today in delivering multimillion-gate SoCs. In the case of ASICs, it is not uncommon for verification teams to spend as much as 50% to 70% of their time and resources on the functional verification effort.

In the particular instance where a processor is part of a design, the interface between hardware and software becomes an area of increased focus and attention. Validation that the hardware and software will function correctly together becomes a

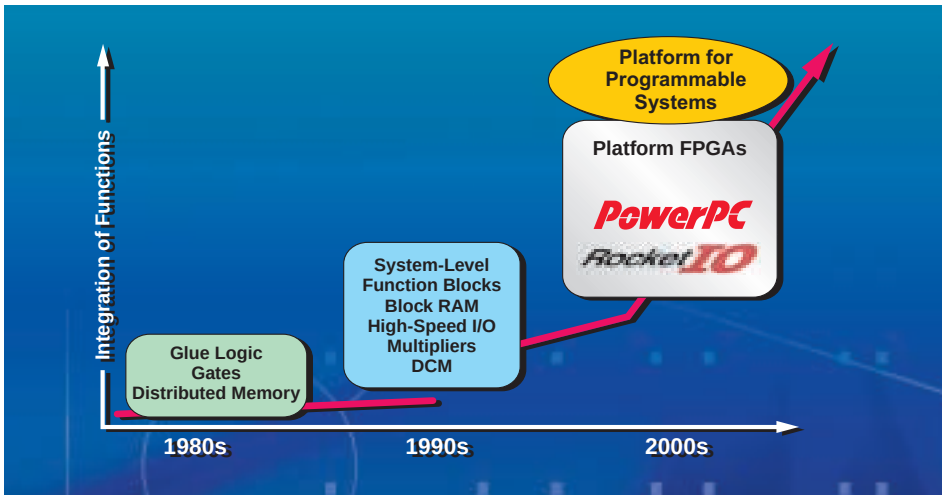

Figure 1 – Paradigm shift in programmable platforms

critical part of the overall debug process. It is therefore important that specialized tools and methodologies be developed that promote debug efficiency and provide a streamlined approach to verification.

For a while now, HW/SW co-verification has been quite commonly used to debug ASIC SoC designs. Now, with an increasing number of FPGA design starts involving embedded processors, this technology is becoming meaningful and important for FPGA designers as well.

Co-Verification for Platform FPGAs

Mentor Graphics' Seamless™ Co-Verification Environment (CVE) is the industry's leading HW/SW co-verification tool – and now it supports the Xilinx Virtex-II Pro™ family of FPGAs. The basic concept behind co-verification is to merge the respective debug environments used by hardware and software teams into a single framework.

For example, a logic simulator is made to communicate with a software debugger, enabling you to get simultaneous control and visibility into the internals of the processor, as well as the hardware peripheral logic that surrounds it, as shown in Figure 2.

An efficient co-verification tool can help uncover a range of HW/SW interface problems, including:

- Initial startup and boot sequence errors (including RTOS boot)
- Processor and peripheral initialization and configuration problems
- Memory accessing and initialization problems
- Memory map and register map discrepancies
- Interrupt service routine errors.

To realize the many benefits that co-verification has to offer, there are three prerequisites for the design under test. A co-verification design environment typically involves the following scenario:

1. The system has a processor executing software code that interacts with surrounding custom peripheral logic.
2. There is extensive interaction between software and hardware parts of the design during execution.
3. Both hardware and software engineering teams agree on using co-simulation early in the design stage.

These prerequisites guarantee a smooth methodology flow and a common communication medium between the two teams.

Co-Verification vs. Simulation

Seamless CVE advances the concept of "functional simulation" in traditional logic-only FPGA designs to "co-verification" in processor-based Virtex-II Pro Platform FPGAs. This methodology establishes value for multiple design teams including hardware engineers (peripheral logic debug), embedded software engineers (application and firmware debug), and system designers (performance analysis and tuning). Let's discuss some of the many advantages Seamless co-verification offers over simulation.

Faster Performance

Pure logic simulation can be used to simulate a design with a processor component. This is accomplished by including an RTL model of the processor to simulate the software code.

This approach, however, is painfully slow and not adequate to address all but the most basic debug requirements. The overall simulation speed is generally in the sub-100 Hz range.

Co-verification, on the other hand, is able to run simulation orders of magnitude faster. This speedup is achieved primarily through the use of clever tool optimizations and faster processor models known as instruction set simulators (ISSs).

To understand the concept of optimization, note that the real bottlenecks in simulation are due to the accurate but slow logic simulators. Every time the software needs to

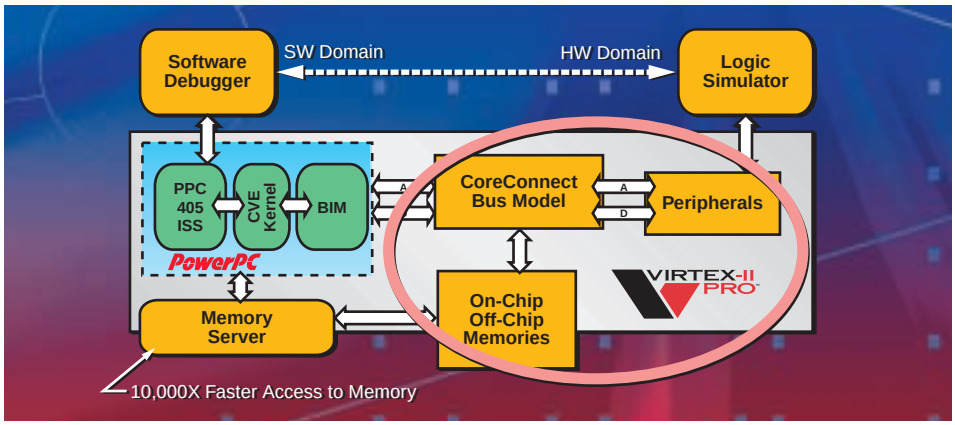

Figure 2 – The Seamless Co-Verification Environment connects hardware and software verification environments for Virtex-II Pro FPGAs.

communicate with hardware, the transaction must go through a logic simulator. Because software runs on the ISS and it runs much faster, the software and ISS are always waiting for the hardware and logic simulator to catch up.

The Seamless tool bypasses this fundamental limitation by introducing the concept of a coherent memory server (CMS), shown in Figure 3. Using the CMS, the ISS is able to read and write to memory about

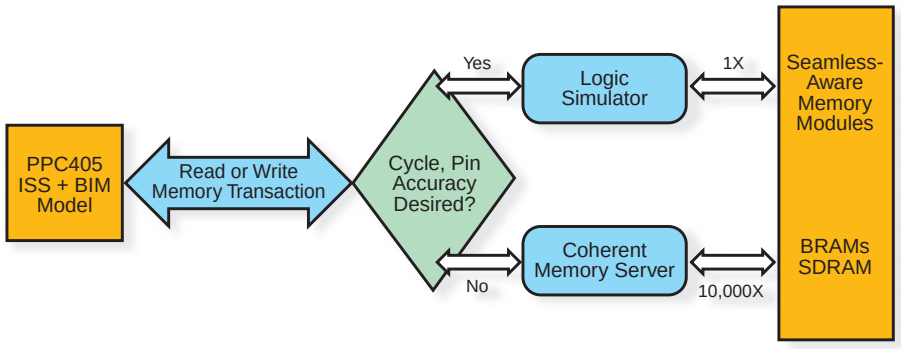

Figure 3 – The Seamless coherent memory server accelerates co-simulation with faster access to memory.

10,000 times faster than if it had to go through a logic simulator. Given that processor-to-logic interaction is mostly through read-write cycles to memory – fetching instructions, accessing peripheral registers, and such – the overall simulation speed can be dramatically increased by diverting most CPU-to-memory transactions to run through the faster CMS instead of through the logic simulator.

Only portions of code or regions of memory that are under active debug are run through the pin and cycle-accurate logic simulator. This means that the simulator bottleneck only factors in less than 1% of the software-hardware transactions, providing a significant overall throughput advantage over pure RTL simulation.

Increased Comprehension

To efficiently address debugging problems that span multiple teams, you need tools and methodologies that each team can relate to. For example, debugging processor code on an RTL model of the CPU is inherently inefficient and impractical.

In Seamless CVE, however, because a cycle-accurate ISS model replaces the RTL processor model, a symbolic debugger can be attached that enables interactive and graphical debugging capability. The standard features of a software debugger include the ability to step through source code (C and assembly), set breakpoints, and observe register and memory contents. This symbolic debugger gives you greater control and comprehension than what you would get by trying to achieve similar goals using the combination of a logic simulator and an HDL processor model.

Support for Abstract Models

Oftentimes, when very high data throughput is required to validate certain design functions, RTL models have to be replaced with faster, more abstract behavioral models. These high-speed models, usually written in C or C++, can interact with the ISS at very high speeds, allowing for complex protocols to be rapidly and comprehensively tested.

The Seamless co-verification tool has been extended to be able to plug in these behavioral models through its "C-Bridge" interface technology. By working less in the logic simulator and more with higher-level models, verification speeds can deliver significant performance gains. With increased simulation throughput, the virtual platform now gives you visibility into system performance and architectural trade-off issues at a very early stage in the design process. You can not only quickly validate functionality, but also analyze and tune important system attributes, such as bus bandwidth, latency, and contention – all leading to increased system performance.

Additional Co-Verification Benefits

With access to co-verification technology, processor-based designs are not only easier to debug, but the process starts much earlier in the design phase.

Find Problems Earlier

Design teams are highly motivated to identify and fix problems at an early stage in the design cycle. A well-known design axiom states: "The earlier a problem can be identified, the easier and cheaper it is to fix it."

Typically, designers cannot initiate software verification until a hardware prototype is available. As a consequence, when software verification occurs in a serial manner, HW/SW interaction problems may not be detected until much later in the design stage.

A virtual prototyping and debug environment removes this restriction by enabling product integration ahead of board and device availability, or even before the final design is committed. With the Seamless CVE, software teams do not have to wait for silicon before they can start developing and testing their portions of the design. As a result, problems can be found earlier and the time to working silicon is dramatically reduced.

Simplified Test Benches

To verify design functions, hardware engineers often write elaborate HDL test bench routines. These test benches can become very complex, and it is not uncommon for the test bench code size to approach that of the design itself. With co-verification, the ISS processor model allows test benches to be greatly simplified.

For hardware verification engineers testing protocols and device drivers, test benches are simplified, because actual embedded software code – and not contrived test bench code – is driving the hardware circuits.

Similarly, software engineers do not have to resort to writing stub code. Actual hardware devices provide real-life responses to calls made to hardware. Overall, this leads to fuller and more comprehensive test coverage, leading in turn to increased confidence in the working of the design in silicon in the first instance.

Greater Runtime Control

An important attribute of debugging in the virtual domain is the ability to "stop time."

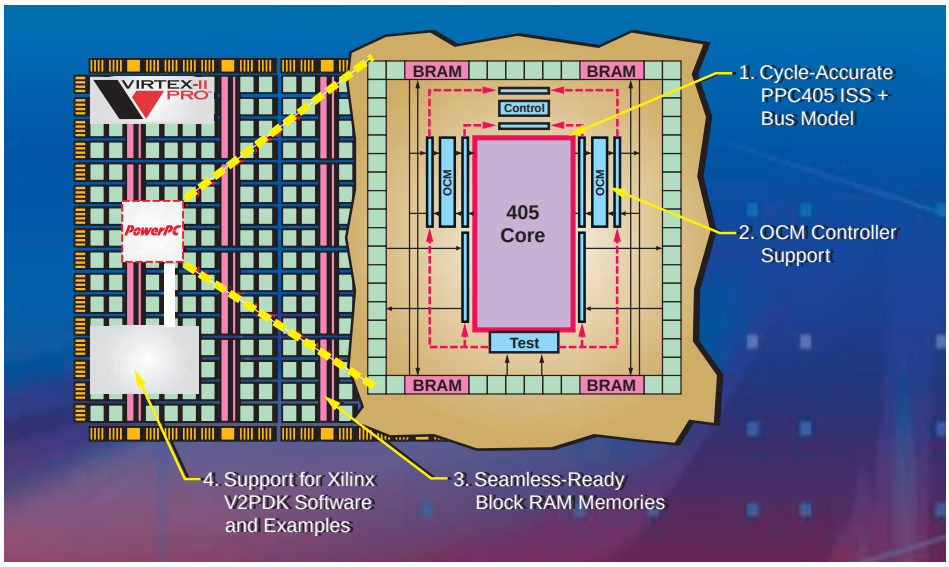

Figure 4 – Customized Seamless enhancements to support Virtex-II Pro devices

As a result, it is easily possible to simultaneously observe and modify the internal values of the CPU registers, as well as those of the hardware device registers with which the processor is communicating. This ability to freeze and synchronize the hardware and software domains offers the ultimate in control and observability – and it is invaluable in efficiently helping debug complex and intricate transactions.

Virtex-II Pro Enhancements

In supporting Xilinx Virtex-II Pro FPGAs, the Seamless tool is able to largely leverage off its existing support for the IBM PowerPC 405 core. However, subtle differences in Virtex-II Pro devices have required customized enhancements to the IBM 405 Processor Support Package (PSP) and the Seamless CVE product to provide you with an efficient and easy out-of-the-box experience, as shown in Figure 4.

First, the PowerPC 405 ISS model from Mentor Graphics had to be matched with the Xilinx version of the PowerPC 405 core, which resolved some of the pin differences.

Additionally, Virtex-II Pro devices have dedicated on-chip memory (OCM) controller circuitry that, in the actual silicon, is tied to the OCM port of the PowerPC 405 core. The same connection had to be stitched in the Seamless environment to the PowerPC 405 ISS.

Xilinx embedded memory blocks (block RAMs) also had to be specially coded to be compatible with the Seamless coherent memory server. This was accomplished by inserting special hooks into the HDL memory models for the block RAMs.

Finally, a conscious attempt was made to ease the learning curve for engineers already familiar with Virtex-II Pro Development Kit (V2PDK) flows. Three of the reference designs included in V2PDK were ported to the Seamless environment. This allows you to not only get a jump-start in understanding co-verification flows, but also the ability to compare and contrast co-verification with pure logic simulation.

All these enhancements are rolled into a customized PSP that represents the final integrated module, which can be used with the standard Seamless kernel.

In addition, a special Xilinx-only version of the Seamless CVE product has been introduced. This bundle includes the Seamless kernel and the PSP, and it is designed specifically for the Virtex-II Pro environment.

To prepare a design to run in the Seamless co-verification environment, the following steps must be taken:

1. Instantiate the Seamless PowerPC bus interface models into your Verilog™/VHDL design (replacing the PPC Swift model).

2. Provide the Seamless tool with access to the "Seamless-ready" block RAM memory models. These models mirror the HDL models but are enhanced with API hooks, which allows them to communicate with the coherent memory server.

3. Input the systems memory map into the Seamless graphical user interface and define startup options for the hardware simulator.

4. Load the compiled software executable (.elf with -gdwarf) into the software debugger and start the session.

Programmable Logic Complements Co-Verification

The chances for first-time success with your design are greatly increased by early integration and testing in the virtual prototype domain.

However, there are classes of problems involving behavior that can only be captured when the processor runs at full speed. In this regard, platform FPGAs serve as a perfect complement to virtual platform debug techniques. Designs can be downloaded into FPGA silicon for validation at full system speeds.

If problems escaped earlier attention, you can debug in-system with the Xilinx ChipScope™ Pro interactive logic analyzer, or you can go back to the co-verification environment for a more controlled analysis. Design errors can be fixed and re-implemented in silicon without incurring the huge delays and costly mask re-spins common with ASIC design flows.

Conclusion

The current generation of Xilinx Platform FPGAs with powerful RISC processors and multimillion-gate capacities requires powerful and matching co-verification methodologies. With the availability of Mentor Graphics' Seamless CVE, you now have access to an ASIC-strength, best-in-class debug solution.

Seamless CVE provides an efficient and easy-to-use methodology that can integrate, verify, and debug hardware and software interactions very early in the design cycle – thus preserving and enhancing the critical time-to-market advantage of FPGAs. To learn more about Seamless co-verification, go to *www.mentor.com/seamless/fpga/*. Σ

In FPGA Synthesis, Timing Is Everything

By optimizing its timing-driven synthesis tool for Spartan-3 devices, the Synplicity Synplify Pro FPGA synthesis tool makes it possible to meet timing constraints and deliver significant area savings.

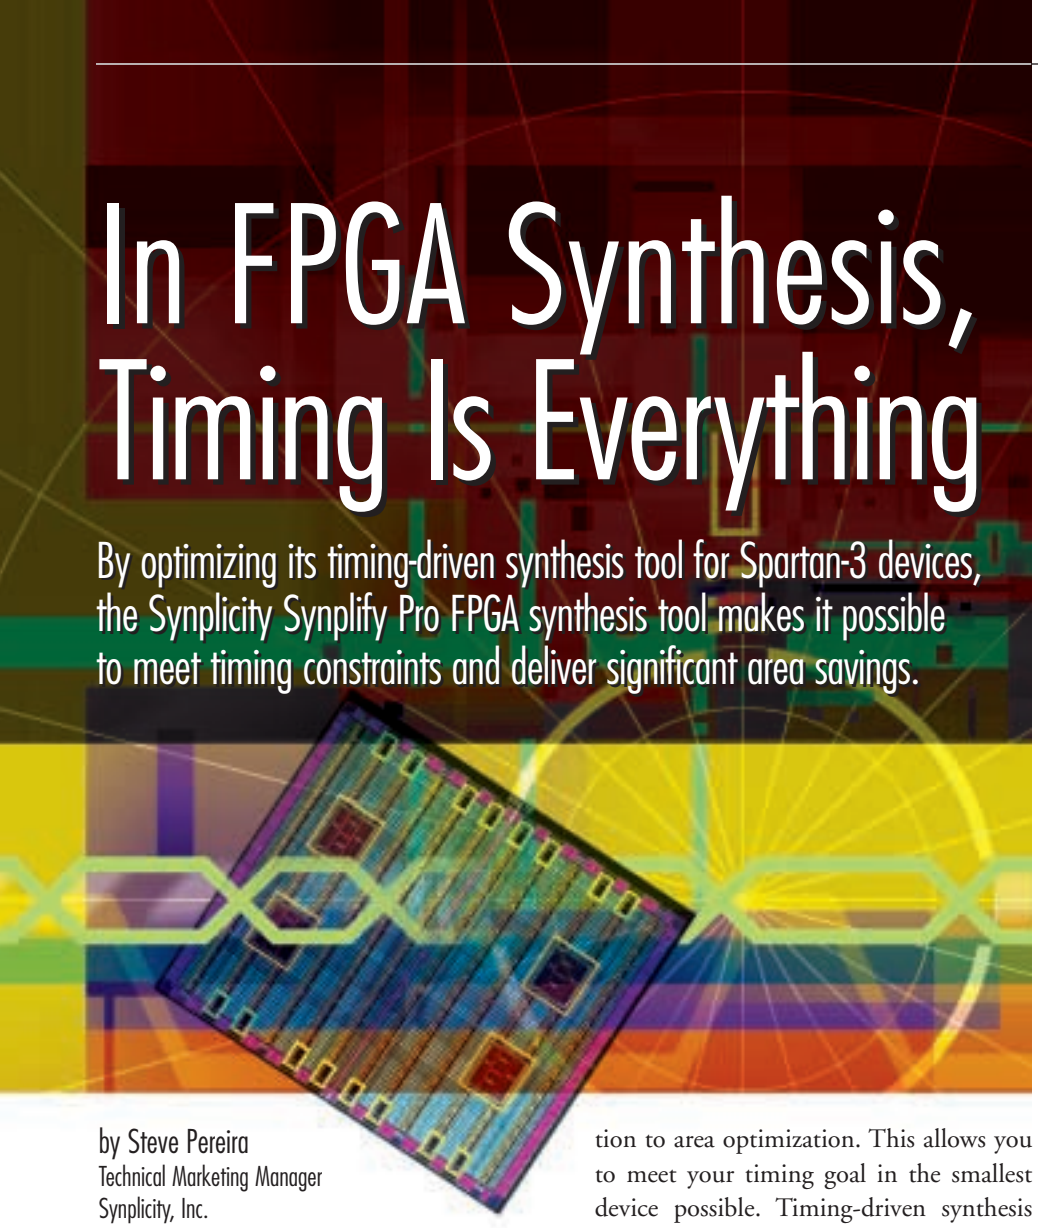

by Steve Pereira
Technical Marketing Manager
Synplicity, Inc.
stevep@synplicity.com

Xilinx Spartan™-3 devices let you take on a new set of applications once reserved for the ASIC domain – and complete those applications quickly and cost effectively. Closing timing is a challenge on large devices; however, Synplicity® has made the leap to the next performance level.

The Synplicity timing-driven synthesis solution has been optimized for the Spartan-3 family. It provides optimal implementations of your HDL to increase performance, reduce area, and drive down the solution cost.

Best of Both Worlds

The Synplify Pro® FPGA synthesis tool is truly timing driven. Critical paths are speed-optimized to meet timing constraints. Non-critical paths are optimized for area. Once your timing constraints are met, the Synplify Pro tool turns its attention to area optimization. This allows you to meet your timing goal in the smallest device possible. Timing-driven synthesis delivers huge area savings and still meets timing constraints.

Figure 1 shows a graph of requested and actual maximum operating frequency (Fmax) for the place-and-route (PAR) results of a design run through the Synplify Pro synthesis tool. The tool's estimates follow the requested frequency line until other critical paths start appearing that the tool cannot optimize. At that point, Fmax starts to degrade.

Figure 2 shows the LUT (look-up table) count for the same design. The results underline the importance of true timing-driven synthesis. When the requested frequency is low, the LUT count is low. The LUT count only starts to increase when the requested frequency is raised. This performance-on-demand feature can drastically reduce area while maintaining performance.

The Synplify Pro tool can also reduce area by automatically mapping logic to dedicated resources within the Spartan-3 architecture. The tool extracts ROMs, RAMs, SRL™ (scan ring linker) modules, and global control resources, among others. Extracting these logic modules increases performance and in most cases, decreases the number of logic gates in slices, which in turn reduces the LUT count.

Three Ways to Win

The Synplify Pro synthesis solution reduces Spartan-3-based design costs for three different design scenarios:

1. For designs with challenging performance goals, the tool offers advanced logic optimizations to meet your requested frequency – and still allows you to choose a slower speed-grade device. Switching to a lower speed-grade device can drastically reduce costs on high production runs.
2. For designs with both performance and device-size goals, the Synplify Pro engine provides performance on demand. As Figures 1 and 2 illustrate, when requested performance is non-taxing, the LUT count remains very low, which lets you choose a smaller device.
3. When modifications are required to an existing FPGA in a system, the cost can be high. In particular, when modifications fail to place or route properly, the most likely option is a complete redesign at enormous cost. The Synplify Pro tool reduces the probability of modification failure with its performance benefits and performance-on-demand area savings. It is much easier to route a less full device. Because the Synplify Pro engine can meet the performance by utilizing a small device, there is a higher probability of the modified design being successfully implemented with the low utilization.

Mapping to Spartan-3 without TBUF

The Synplify Pro tool provides simple conversion of tri-state bus drivers (TBUFs) to multiplexers (MUXs) for mapping to Spartan-3 devices that do not have TBUFs. For each net with multiple tri-state

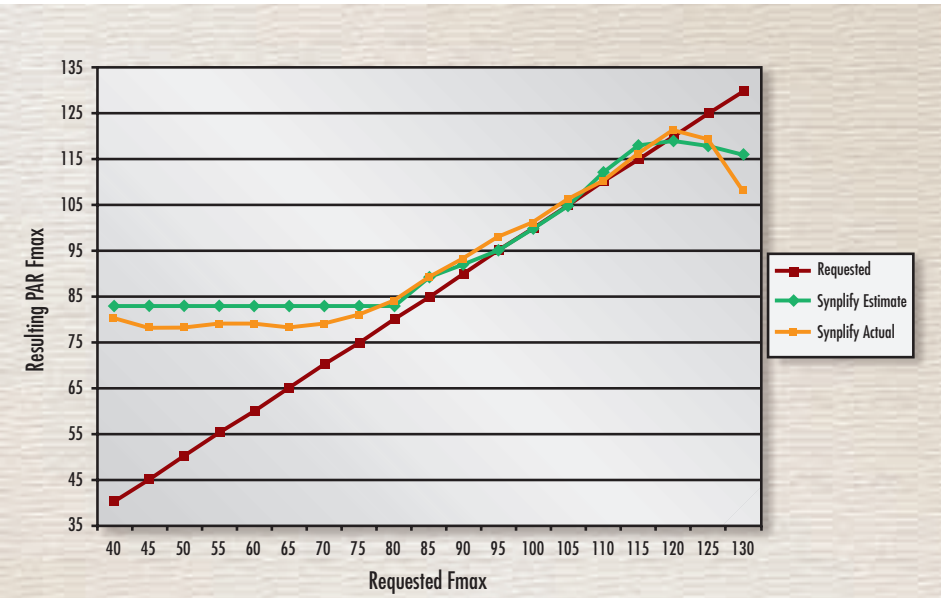

Figure 1 – Place-and-route frequency results meet design goals.

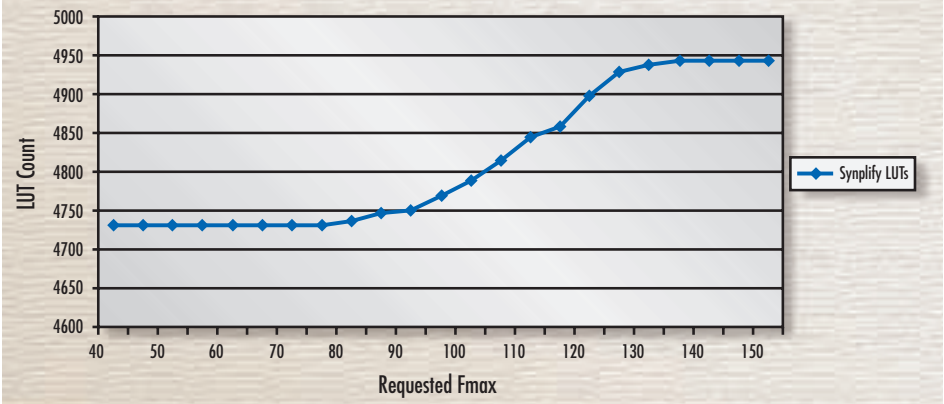

Figure 2 – LUT count can be traded for frequency.

drivers, the tool creates a parallel MUX. For example, if the Verilog™ design code is shown as follows:

```
n1 = e1 ? d1 : 1'bz;
n1 = e2 ? d2 : 1'bz;
n1 = e3 ? d3 : 1'bz;
n1 = e4 ? d4 : 1'bz;
```

then the conversion procedure in the Synplify Pro tool would create the following Verilog code:

```
n1 = e1 & d1 | e2 & d2 | e3 & d3 | e4 & d4
```

– this is the tool's pmux structure.

If the conversion is done very early in the flow, when most MUX-based optimizations are necessary, this logic and the Synplify Pro tool deliver an optimal implementation.

For example, if the following code is true:

```
e1 = !s0 & !s1;
e2 = s0 & !s1;
e3 = !s0 & s1;
e4 = s0 & s1;
```

then the logic driving `n1` will be implemented using 4x1 MUX with optimal use of MUXs. The Synplify Pro tool also supports more complex conditions. If `n1` is driving a primary output, appropriate logic would be generated so that the signal outside the chip shows a high impedance when `e1 = e2 = e3 = e4 = 0`. Single tri-states driving primary outputs are sucked into the pads. Internal single tri-states are converted to logic.

ASIC to FPGA Migration

Because Spartan-3 devices have a very high gates-per-dollar ratio, ASIC designers are increasingly adapting to FPGA solutions. In doing so, they also reap the benefits of reprogrammability and zero NRE (non-recurring engineering) costs.

Among the drawbacks of the ASIC-to-FPGA flow are coding structure and instantiated components. Synplicity has been providing a solution (the Certify® product) for ASIC prototyping in an FPGA that addresses both of these issues. But for generic ASIC code, the biggest pitfall is gated clocks.

With the release of Version 7.3, the Synplify Pro FPGA synthesis tool can automatically convert the gated clock structures commonly used in ASICs so that they map to clock enables in Xilinx devices while making use of global routing resources. This feature has two primary benefits for ASIC/FPGA designers:

1. Gated clock conversion eases the pain of migrating ASIC code.
2. The tool automatically provides higher performance and better utilization, because global routing resources are used instead of internal high fanout nets.

Save Money on Your Next Spartan-3 Design

The performance of the Synplify Pro solution, coupled with industry-leading FPGAs from Xilinx, gives you the ability to meet aggressive performance goals on time and on budget.

To test the Synplify Pro tool and see the cost savings for yourself, download the Synplify Pro tool at *www.synplicity.com/downloads/download1.html*.

When informed that you do not have a license, follow the subsequent instructions and send the information to *license@synplicity.com*. Synplicity will send you a temporary license immediately.

Also, to be sure you get the most from the Synplify Pro FPGA synthesis solution, download and read "Benchmarking Synplify and Synplify Pro Software" at *www.synplicity.com/literature/pdf/benchmarking_synplify.pdf*.

Advanced Switching Extends PCI Express

Advanced Switching supports flexible interconnects, distributed computing, and multicast options. Teamed with Xilinx FPGAs and RocketIO transceivers, it will dominate next-generation embedded computing and communications platforms.

by Kiran S. Puranik
Staff Design Engineer
Xilinx, Inc.
kiran.puranik@xilinx.com

The PCI Special Interest Group's (PCISIG) ratification of the PCI Express™ Base specification in July 2002 was a major milestone in the widespread adoption of serial system interconnect architectures. Now, the combination of Xilinx Virtex-II Pro™ Platform FPGAs, RocketIO™ multi-gigabit transceivers, and embedded IBM PowerPC™ 405 processors is making high-performance PCI Express designs easier and faster.

PCI Express is suitable for a wide range of computing and communications applications because it is a low-cost, high-performance, interoperable, and scalable solution. An important adjunct to the base standard is the Advanced Switching (AS) technology. It enhances the base standard to address the needs of the communications infrastructure for the next decade.

Single-chip AS solutions for server, storage, and communications applications can be configured with a Virtex-II Pro Platform FPGA, on-chip RocketIO multi-gigabit transceivers, and embedded IBM PowerPC 405 processors. The extensive Xilinx intellectual property (IP) offerings – including the world's first Real-PCI™ Express solution – will accelerate product design cycles and time to market.

PCI Express Architecture

PCI Express is a layered architecture consisting of physical, data link, and transaction layers. Device drivers and application software constitute upper layer protocols (ULPs).

PCI Express Base features include:

- LVDS, 2.5 Gbps, serial links scalable up to 32 lanes in each direction
- Embedded clock and 8b/10b transmission code
- Native hot attach/detach capability for high-availability applications
- Robust and efficient link protocol
- Classes of traffic and support for isochronous applications
- Packet-based protocol and credit-based flow control per virtual channel
- Link level and end-to-end data integrity (CRC-32)
- Completely in-band signaling support.

Advanced Switching

Advanced Switching leverages the physical and data link layers of the PCI Express Base specification. Communication and embedded computing enhancements are added at the transaction layer to address chip-to-chip, backplane, and inter-chassis data communication requirements. The AS fabric architecture is designed for both control and data-plane switching applications.

AS components can be broadly classified as being either switches or end systems. Figure 1 shows the AS protocol stack on an end system. The transaction layer creates an interface between ULPs and the data link layer. It serves as the tunnel for ULP encapsulation and extraction at an end system.

The AS transaction layer packet is shown in Figure 2. It contains an AS route header followed by a payload section. Ingress end systems provide route and encapsulation payload, while egress end systems extract payload. In order to perform packet switching, the AS route header is completely agnostic to the contents of the payload section.

An entire AS transaction layer frame is shown in Figure 3. The physical and link layer portions of the frame are added and removed at each hop between fabric components.

Important Features of AS

Legacy PCI platforms require strong parent-child relationships between connected components. Tree topologies are necessary for hardware and software compatibility. In contrast, AS fabric system topologies can be described as a graph of connected switches and end systems. Switches constitute internal nodes of this graph, providing interconnects with other switches and end systems. End systems, on the other hand, are the edge nodes, representing data ingress and egress

points. The ability to support many topologies gives platform architects enormous flexibility regarding placement of critical resources.

Supports Distributed Computing Architecture

AS enables distributed processing systems, which results in multiple memory-address domains within a platform. This is in sharp contrast to the single flat address domain found in legacy PCI platforms. AS platforms allow load-store and messaging protocol interactions between concurrent hardware and software processes on end systems. This facilitates peer-to-peer-based applications commonly found in server, storage, and communications arenas. Distributed processing offers better scalability and ultimately a lower cost of operation.

Path-Based Unicast, Multicast, and Broadcast Packet Routing

Unicast packet switching is based on path information embedded in the AS route header and takes the form of a turn pool. Path routing requires no up-front programming of switches. Switches simply look at a packet's route header to determine the egress port. This simplifies switch design and platform configuration enormously. A packet's forward route turn pool can also serve as a backward route to the source. Completions for read requests and event notifications to a packet's source end system are examples of backward-routed packets.

The multicasting feature allows an end system to target a packet to multiple end systems. A multicast group index is carried on each multicast packet's route header. A multicast group uniquely identifies a set of switch egress ports for each switch on a multicast packet's path. At a switch, a multicast group table is found in a lookup table using the packet's multicast group index.

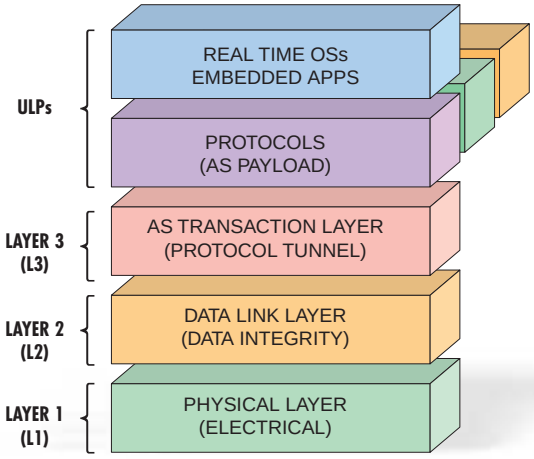

Figure 1 – PCI Express AS protocol architecture

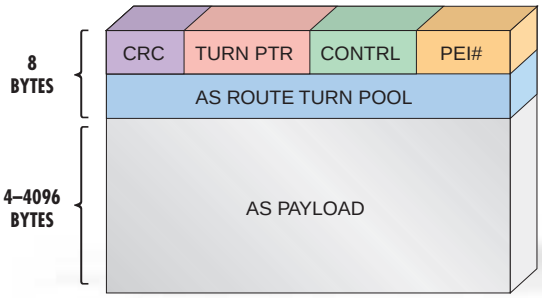

Figure 2 – AS transaction layer packet

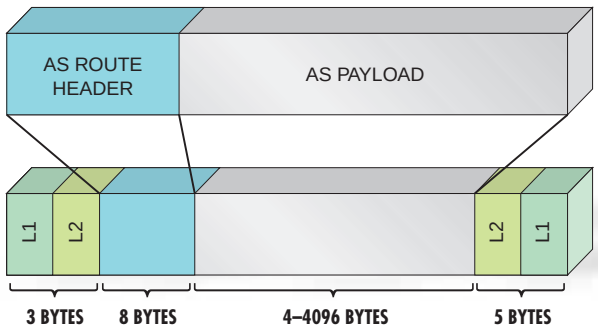

Figure 3 – AS transaction layer frame

The packet is then replicated on each port contained in the multicast group.

When a switch detects a broadcast route header on a packet, it is required to replicate it on all ports except the packet's ingress port.

Congestion Management and Quality of Service

Bandwidth provisioning is an important aspect of AS fabric management because it helps guarantee Quality of Service (QoS) levels to applications. End systems must operate within prescribed bandwidth limits by metering the rate of data transfer.

Congestion in AS fabric can be caused by unexpected transient events such as component failure, accidental removal, or errant end-system behavior. Congestion causes packets to suffer excessive latencies, resulting in a loss of expected levels of service.

AS congestion management mitigates fabric congestion and maintains QoS levels. Congestion is detected at switches; when detected, Backward Explicit Congestion Notification (BECN) messages are sent upstream to end systems contributing to congestion. End systems respond to BECN messages, reducing the rate of data injection by a specified amount.

Congested packets may also be marked with a Forward Explicit Congestion Notification (FECN) bit. FECN notifies downstream end systems of congestion on a certain path. Some AS applications may also set a discard bit on AS route headers, allowing downstream switches to selectively discard packets to alleviate local congestion. In the absence of BECN messages, end systems restore normal traffic flow in specified increments. Fabric management may use congestion notifications to initiate corrective action and activate built-in fail-over using alternate paths.

Differentiated Classes of Service

Class of Service (CoS) mechanisms reduce the complexity of maintaining QoS by mapping multiple traffic flows into a few service levels. AS fabric resources are allocated based on as many as eight service levels called Traffic Classes (TCs). Traffic flows are aggregated and forwarded by fabric components based on the TC of these packets.

Within a TC, AS fabric preserves the ordering of packets end-to-end, with the exception of those marked as bypassable at

the source. There is no such ordering requirement across TCs. AS components map TCs to Virtual Channels (VCs) corresponding to hardware channels within the components. AS components must implement at least two VCs. All AS link partners must share the same TC to VC mapping, downshifting if necessary to the smallest common number of VCs supported between the two link partners.

Each AS VC contains two independent queues – the main queue and a bypass queue. AS link flow control manages flow credits for each queue independently. A packet marked bypassable must enter the bypass queue if it causes VC head-of-line (HOL) blocking. The bypass queue is serviced as soon as bypass credits become available. This capability is compatible with legacy bus protocols that require write transactions to pass HOL-blocked read transactions, to avoid possible deadlocks.

Multi-Protocol Support

Protocol Encapsulation Interfaces (PIs) represent fabric management and application-level interfaces to the AS fabric. Table 1 details a list of currently supported PIs. PIs 0-7 represent fabric management interfaces, while PIs 8-254 are application-level interfaces. As shown in Figure 4, AS supports the tunneling of virtually any protocol. This makes AS platforms modular and cost-effective, as well as easy to deploy and support.

Application PI implementations mandate an efficient encapsulation, without losing a tunneled protocol's semantics, which enables effective extraction at the fabric's egress. A PI may be custom-tailored to suit a specific requirement or can tunnel standard protocols to support a broad range of applications, such as SPI, ATM, Ethernet, or TDM.

Support for Segmentation and Reassembly

The Maximum Payload Size (MPS) of the AS platform is the least common denominator of MPS supported by all components within the platform. All PIs must restrict AS Transaction Layer Packet payload size to

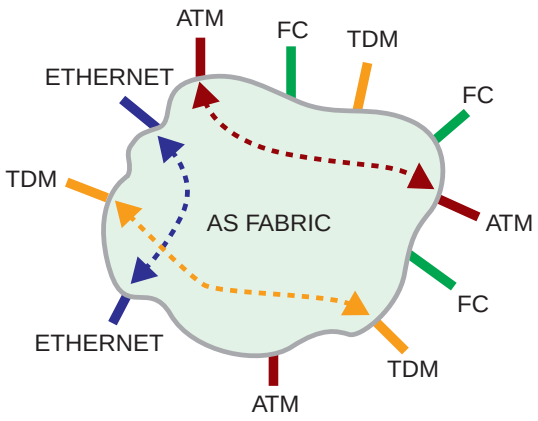

Figure 4 – AS tunneled protocols

platform MPS. End systems that need to encapsulate larger-than-MPS protocol packet sizes must split these into sub-MPS-sized segments. This also requires keeping track of multiple segments of the original packet and reassembling them at the fabric egress end. AS supports a standard mechanism to perform segmentation and reassembly (SAR) via the SAR PI. Individual PIs may choose other SAR implementations.

PI Index	Protocol Encapsulation Identity
0	Fabric Discovery
1	Multicasting
2	Congestion Management
3	Segmentation and Reassembly
4	Configuration Management
5	Fabric Event
6	Reserved
7	Reserved
8	PCI-Express Base
9-223	Unassigned, Future PEIs
224-254	Vendor-Defined PEIs
255	Invalid

Table 1 – AS protocol encapsulation interfaces

Xilinx Real-PCI Express

The Xilinx Real-PCI Express solution is a combination of two leading technologies: the PCI Express specification and Virtex-II Pro FPGAs. It provides user-configurable options, excellent flexibility, and Xilinx Smart-IP™ technology guarantees critical timing. Key benefits include:

Availability

The world's first PCI Express solution is available for download today. It enables your compute and communication systems to achieve the highest level of performance using serial I/O technology.

Performance

The RocketIO 3.125 Gbps-capable transceivers on Virtex-II Pro FPGAs enable multiple 2.5 Gbps lane implementations on a single chip.

Flexibility

The inherently programmable nature of FPGAs allows you to continually tune your design to changing platform performance and functional requirements, reducing your risk in adopting the standard.

Faster Time to Market

Today, there is simply no easier way to develop PCI Express applications with minimal impact to your overall system development cycle. Xilinx will provide updates for any specification changes.

Conclusion

Advanced Switching architecture addresses all of the major requirements for next-generation system interconnect solutions, such as scalability, expandability, modularity, high availability, and peer-to-peer capability, with built-in QoS and CoS support.

Xilinx Virtex-II Pro Platform FPGAs, together with the Real-PCI Express solution, are well-positioned to provide compliant Advanced Switching solutions, enabling rapid product deployments. More information about the Xilinx Real-PCI Express solution may be found at *www.xilinx.com/pciexpress.* Σ

Get the Platform Flash PROMise

Reduce your costs and minimize board space for any Xilinx FPGA design.

by Rob Schreck
Senior Marketing Manager, Product Solutions
Xilinx, Inc.
rob.schreck@xilinx.com

Designing with FPGAs often requires two different types of configuration PROMs. In the design and test phase, in-system programmable (ISP) PROMs offer maximum flexibility for design changes, while one-time programmable (OTP) PROMs are used for manufacturing once the design is complete.

You may, however, have to change your board design to accommodate higher density, lower cost OTP PROMs. This problem becomes more acute if you are using high-density FPGAs that need as many as 10 ISP PROMs and five OTP PROMs.

Xilinx has a solution with the new Platform Flash PROM family, which can configure any Xilinx FPGA – most with just one PROM. You get low-cost, in-system programmability, so you no longer have to change from one PROM to another to reduce production costs. See Table 1 for details.

Economical and Flexible

The December 2004 100K-unit price of Platform Flash PROMs is projected to be $1.05 for the 1 megabit (Mb) version (XCF01S) and $10.40 for the 32 Mb version (XCF32P). This not only represents a dramatic price decrease over previous generations of configuration ISP PROMs, but for OTP PROMs as well.

Engineers know that they need to reprogram FPGAs during logic design and testing, and Platform Flash PROMs offer flexible in-system programmability to ease development. However, don't forget that flexibility can be used at other times in the product life. For example, during manufacturing, you can program the PROM to test the board, and

then reprogram the PROM with the final FPGA design. You increase quality while decreasing manufacturing costs.

In-system programmability also lets you reprogram the FPGA remotely to add new features or fix problems. This added dimension of flexibility means you reduce maintenance and field repair costs, plus you keep your customers happier.

Big Functionality in Small Packages

The Platform Flash PROMs come in a variety of densities in two packages. The 1 Mb, 2 Mb, and 4 Mb density PROMs are offered in the very small (6.4 mm x 6.5 mm) VO20 TSSOP package, the smallest package area per megabit in the industry. These versions offer serial configuration for FPGAs and a very low-cost solution for Spartan™-IIE, Spartan-3, and Virtex-II Pro™ FPGAs.

The family is also available in 8 Mb, 16 Mb, and 32 Mb densities, which come in a small (8 mm x 9 mm) FS48 thin flat ball grid array package. These versions offer both serial and parallel configuration, and are well suited for higher density Spartan-3 and Virtex-II Pro FPGAs.

To reduce board space and costs, you can pack on average 50% more bits into your configuration PROMs using the compression capability in our higher density packages. You can use one low-cost PROM to configure multiple FPGAs. You can even put multiple programs on one PROM and change the FPGA program on the fly.

Conclusion

Now you can use one PROM family to configure all of your FPGAs, reducing your manufacturing and inventory costs. The Platform Flash PROM family offers you ultimate flexibility at a very low cost. For more information, visit *www.xilinx.com/product/platformflash/*. Σ

	XCF01S	XCF02S	XCF04S	XCF08P	XCF16P	XCF32P
Density	1 Mb	2 Mb	4 Mb	8 Mb	16 Mb	32 Mb
JTAG Prog	■	■	■	■	■	■
Serial Config	■	■	■	■	■	■
SelectMap Config				■	■	■
Compression				■	■	■
VCC (V)	3.3	3.3	3.3	1.8	1.8	1.8
VCCO (V)	1.8-3.3	1.8-3.3	1.8-3.3	1.5-3.3	1.5-3.3	1.5-3.3
Clock (MHz)	33	33	33	40	40	40
Package	VO20	VO20	VO20	FS48	FS48	FS48

Table 1 - Platform Flash PROM specifications

Performance + Time = Memory

By approaching FPGA designs as three-dimensional endeavors, you can radically reduce device size – and cost.

by Ken Chapman
Senior Staff Engineer
Xilinx, UK
ken.chapman@xilinx.com

"Performance + Time = Memory" may sound like an odd formula, but when you understand it, you can realize significantly lower implementation costs within Xilinx FPGAs. In this article, I'll show you how to use three-dimensional (3-D) design to accomplish a 15X reduction in the number of logic blocks in a sensing application.

Although vital for DSP applications, I really like the way the formula can be applied to so many designs. It is particularly useful for applications that are suited to the range of Spartan™ devices, where cost savings are always welcome for high-volume applications.

But let's understand the formula first.

2-D Parallel Design

In most hardware designs, we treat the Xilinx FPGA as a two-dimensional (2-D) fabric, as shown in Figure 1. Complex logic blocks (CLBs) provide the logical functions and blocks of RAM are used for buffers, such as first-in first-out (FIFO) memories.

The tendency is for a design to become larger as more functionality is required. Therefore, it must use larger devices. The clock speed can often be well below 100 MHz, and many of the functions are clock-enabled at even lower rates.

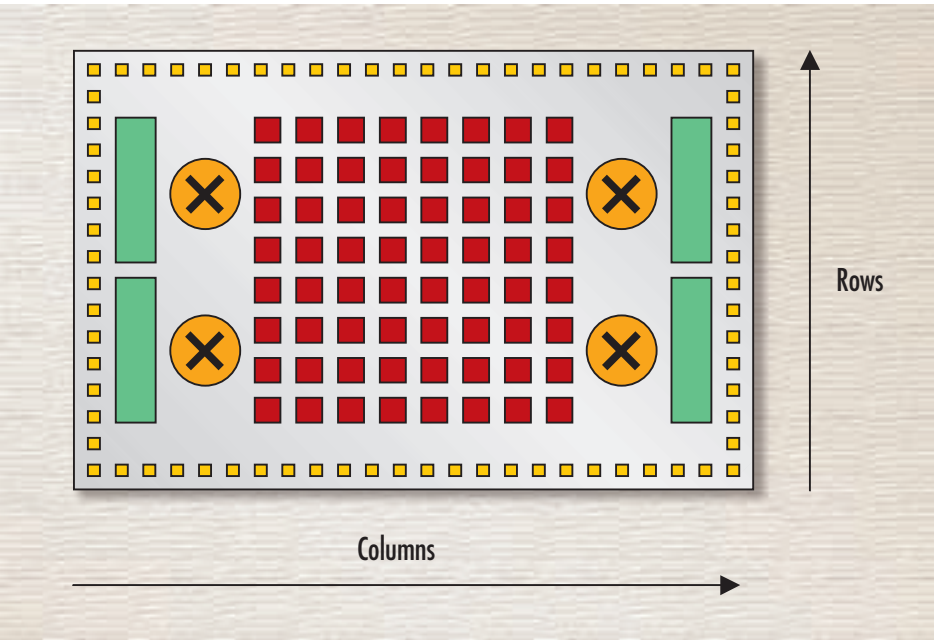

Figure 1 – Two-dimensional designs consider only logic and memory.

Because the cost of a design is proportional to the size of the device, parallel implementations, even if well optimized, will be relatively expensive. They cannot be avoided where maximum performance is required.

Applications such as bus interfaces that need a predefined number of pins and clock rates are also fundamentally constrained in the way they can be implemented. However, when processing functions need only be completed in a relatively long time period, such 2-D design is wasteful and unnecessarily expensive.

A parallel design provides logic for each and every function that must be

implemented. This means that there is actually a zero requirement for memory, because a signal (wire) exits for every value to be calculated. The addition tree example in Figure 2 shows how the value "A+B+C+D" is created. Because the value is immediately applied to the final adder, however, the value does not need to be stored.

Of course, the parallel implementation offers the very highest performance. The adder tree can easily exceed 100 MHz in a Spartan-II device, which is equivalent to more than 700 million additions per second. However, such a structure cannot benefit from having more time to complete the required operation, other than consuming less power if it is clocked slower.

If there is 1 ms available to perform the addition tree, then it can be clocked at 1 KHz. It will work, but it really is a waste of the Spartan-II silicon performance potential. Even worse, the more values that need to be added, the larger the circuit becomes – and this increases the cost of your product.

Processors Obey 3-D Formula

Now, take a closer look at the familiar world of processors. A processor is a very good time-sharing engine. The ALU is directed to perform many different operations over many clock cycles to (it is hoped) complete the desired process in the required time period. The higher the performance of the processor, the faster the ALU will be clocked, and hence, the more that ALU can be time-shared to achieve the algorithmic process, as illustrated in Figure 3.

For example, given that a particular process must be completed in a maximum time of 1 ms, the number of clock cycles available for the processor to exploit depends on the performance:

- A clock speed of 1 MHz provides 1,000 clock cycles per 1 ms.
- A clock speed of 100 MHz provides 100,000 clock cycles per 1 ms.
- A clock speed of 200 MHz provides 200,000 clock cycles per 1 ms.

This is all very obvious, but less obvious is the direct link this has to memory.

Suppose the available clock cycles are used by the ALU to perform the trivial task of summing data values. In a 1 ms time period, a 1 MHz clock rate means that the processor has the ability to sum 1,000 data values. It will have to get these values from somewhere, and that place will be memory. As the clock increases to 200 MHz, it can then use the same amount of logic to sum 200,000 data values – and it now needs a memory to hold 200,000 words.

In a more realistic case, a process tends to apply multiple instructions to each data set, so the memory requirement to store data is not so high; all the same, there is a very strong relationship.

3-D Sequential Design

Making the decision to operate the logic functions at a higher rate than the processing rate allows operations to be achieved sequentially. As with a processor, logic is time-shared over multiple clock cycles. Because "Performance + Time = Memory," we also need to use memory to hold all the values not being used in a given clock cycle, as well as partial/temporary results created during the processing. See Figure 4 for a 3-D rendering.

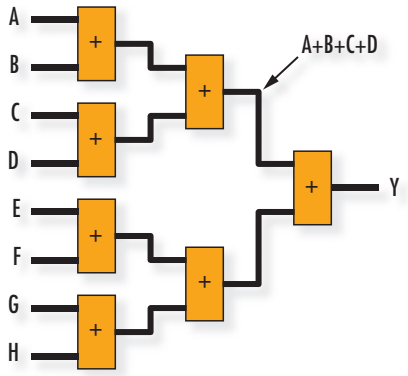

Figure 2 – Addition tree example of parallel design

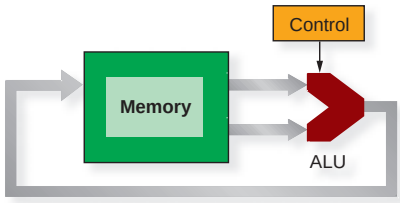

Figure 3 – ALUs can be time-shared.

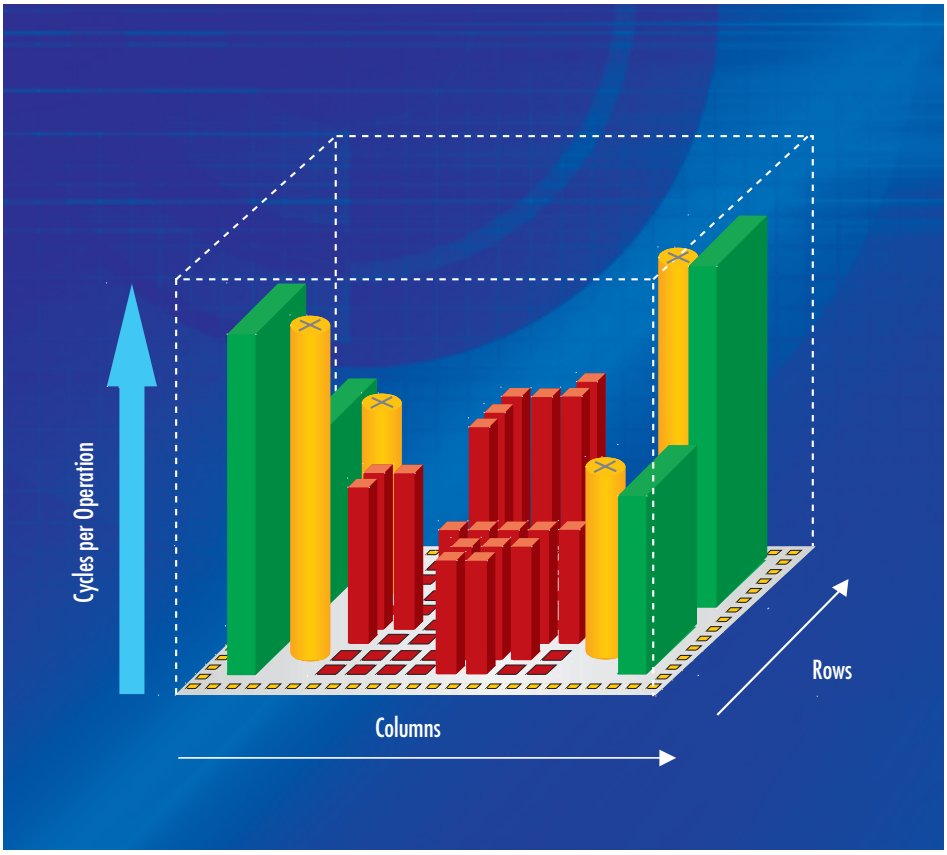

Figure 4 – Logic can be time-shared over multiple clock cycles.

The FPGA can now be thought of as a 3-D volume to be filled. The best part is that you just pay for the 2-D fabric being occupied. The only limits to "building" upwards are the maximum clock rate of the device and the amount of RAM available in a given block. In addition to the dedicated blocks of RAM, each CLB can be used to provide distributed RAM, allowing the correct amount of memory to be allocated in each position. This prevents memory access bottlenecks from forming in your design.

3-D Approach to Design

When any function is implemented, two basic questions should be asked:

1. How much time is available to complete the process?
2. Given the performance of the selected Xilinx device, what clock rate will be used?

The answer to the first question comes from the design specification. The way you partition the design into functions can have quite an impact, so consider some alternatives. As to the performance of Xilinx FPGA devices, this has more to do with "design comfort" than the actual peak performance of the devices.

Regarding the second question, I personally like to see devices clocked above 75 MHz, and I find this relatively easy to achieve. However, the higher the clock rate is, the more challenging the design is. Anything lower than a 50 MHz clock rate is very slow and wasteful of the performance potential offered by Xilinx logic devices. Remember that the embedded DLL (delay locked loop) and DCM (digital clock manager) blocks can be used to create internal clocks of a higher rate than those available on the PCB.

The answers to the first two questions will let you know if there is any potential for time-sharing of logic resources. This leads to a third and final question:

3. How can the memory resources of the device be utilized to reduce the size of implementation?

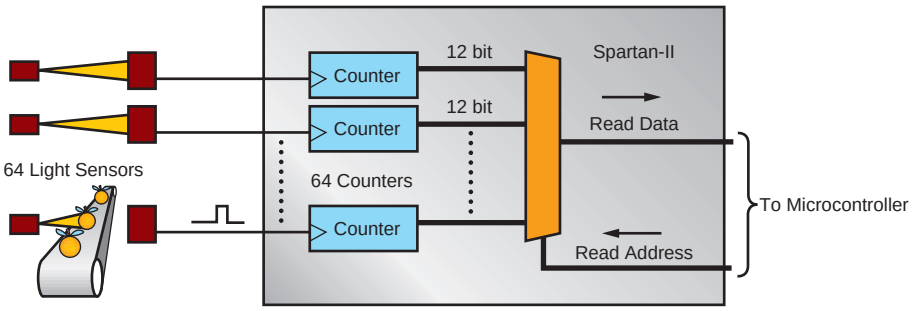

Figure 5 – Light sensors collect data for processing.

Now the engineering starts. It does take some practice, so what follows is a design specification for you to consider.

Design Challenge

The challenge is to design a small box to be used in factories processing items such as fruits and vegetables. The card is used to collect data from light sensors located on various conveyor belts, along which the fruits and vegetables pass as they are sorted for type, quality, and size.

The initial design concept is to employ a microcontroller (or similar small processor) to collate the information and communicate it via serial (RS-232) links to a PC in the factory control room.

An FPGA is being considered to interface the processor to the sensors. The product is required in high volume (50K to 100K units), so a Xilinx Spartan-II FPGA is the target for a cost-effective solution.

The card supports 64 sensors. A logic "1" signal is generated when the light beam is broken by a passing object. The maximum speed of the conveyor belt is 1 meter/sec. The minimum width of a single item is 3 cm, and there is a minimum 10 cm between items on the belts.

Each pulse is recorded by a separate counter, which can support a maximum value of 4,095 (12 bits). A simple interface to the microcontroller is then able to read the value of any of the 64 counters in the card by supplying a 6-bit address.

Initial Observations

Taking a very direct approach to the design, we could simply identify the need to implement 64 counters of 12 bits followed by a 64:1 data multiplexer. In fact, this is a direct representation of the block diagram shown in Figure 5.

However, we need to apply some fundamentally good engineering here, because we certainly wouldn't want to have 64 independent clocks in a design. Such a design would lead to very poor utilization of the device and have a high probability of unreliable operation. The signal inputs really should be synchronized to a single internal clock, and then clock enables should be used with the counters.

First Estimate

Given a basic understanding of the device architecture, you can easily make an estimate of the device resources used.

- Counters – Because each slice of an FPGA can implement a 2-bit counter, six slices are required to implement a 12-bit counter. Therefore, a total of 384 slices are required for all 64 counters. (Two slices form a CLB within the Virtex™ and Spartan-II FPGA families.)
- Multiplexer – Each slice contains two lookup tables and a dedicated multiplexer (MUXF5), enabling a 4:1 multiplexer to be implemented. However, each pair of slices within a CLB share an additional dedicated multiplexer (MUXF6), enabling a complete 8:1 multiplexer to be implemented in two slices. Nine of these 8:1 multiplexers are required to construct a 64:1 multiplexer, which then must be replicated 12 times to support the data width of the counters. The total size of the multiplexer is then 2 x 9 x 12 = 216 slices.

- Synchronizing Logic – At this stage, we have not designed the logic to capture the input signals and synchronize them to the internal clock. For now, we will allow a slice per input (two flip-flops and some gates). This gives us a total of 64 slices.

Based on these major building blocks of the design, our estimate is for 664 slices. Thus, a Spartan-II XC2S50 device is suitable with its 768 slices, providing a surplus of 104 slices to complete the processor interface.

There are many ways to implement the "Performance + Time = Memory" formula – and we will look at just one. But as long as you can significantly lower the cost, you are well on your way to improving the profit margins on your own designs in the future.

Remember, the target to beat is 664 slices in a Spartan XC2S50 device, which was the result of a full parallel 2-D design.

Implementing a 3-D Design

We must begin our 3-D design process by asking the right questions that relate to the "Performance + Time = Memory" formula.

How Much Time Is Available?

Taking the minimum fruit size and minimum spacing between fruit passing on a belt at the maximum speed of 1 meter/sec, we derive the timing of the fastest pulses from a light sensor.

We discover that the pulses are of a long duration and that the pulse rate is very low. In fact, the maximum pulse rate is less than 8 Hz, which is very slow indeed. However, we must consider that there are 64 sensors to be monitored; we could be unlucky enough to have them all triggered at the same time. So, all 64 sensors must be serviced in a maximum of 30 ms, and the aggregate data rate is more like 500 Hz.

What Performance Is Available?

We know that a Spartan-II FPGA is the target architecture. This device is capable of operation above 100 MHz, so device performance should not limit us at all in this case. Although we want to get the most out of the silicon, there is no point overdoing it and burning power unnecessarily. In this case, it is better to work out the minimum clock rate required to process all 64 channels, and then tie this rate in with a suitable clock source on the PCB.

Looking at the timing waveform shown in Figure 6, the pulse width caused by the smallest fruit breaking the light beam is the most demanding. We must guarantee that we observe each sensor at least once every 30 ms.

If, however, the 64 sensors are observed and processed sequentially, rather than in parallel, then 30 ms divided by 64 is the maximum time that can be allocated to each sensor. This means that the minimum processing rate is 2,133 Hz. Obviously, this is still desperately slow, but it only emphasizes that "Performance + Time = Memory" must be a valid formula to be applied in this case.

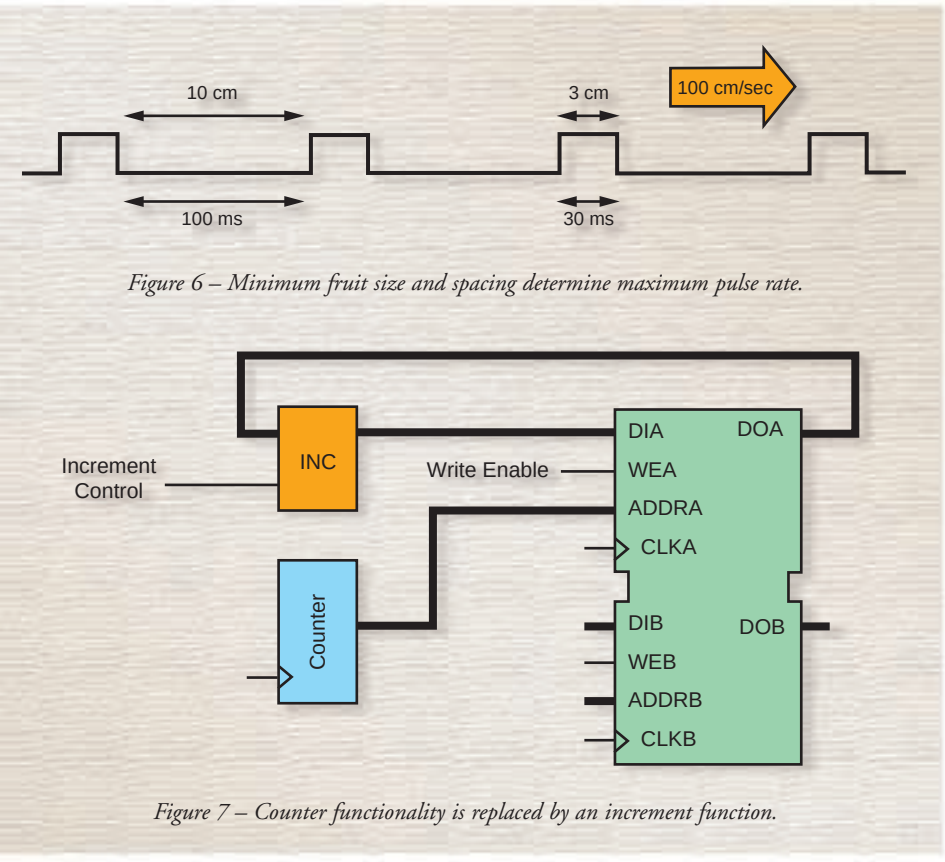

Figure 6 – Minimum fruit size and spacing determine maximum pulse rate.

Figure 7 – Counter functionality is replaced by an increment function.

Replacing Counters with Memory

We have "Time" and we have ample "Performance," so now it is a case of working out how to make the whole thing a sequential 3-D design. How can the memory resources of the device be utilized to reduce the size of implementation?

Because memory is used to hold data values, we must identify where the data is in the system. These may be complete values or partial values, so we must have a good look through the block diagram and identify where the data values are. In this system, they are fairly obvious in that the counters each hold a value. In the parallel implementation, they are distributed across the 384 slices, forming the 64 counters, but we want to consolidate them into a single memory.

We can choose between distributed (CLB) memory and dedicated (block) memory, and we could really use either to form storage for 64 values of 12 bits. However, as the dedicated block RAM isn't required for anything else, let's take that option. Configured as 256 words of 16 bits, a single block provides more than adequate storage.

The counter functionality is then replaced with a single increment function, as shown in Figure 7. A "count value" is read from the RAM, passed through the increment block, and then written back into the RAM at the same location. This is

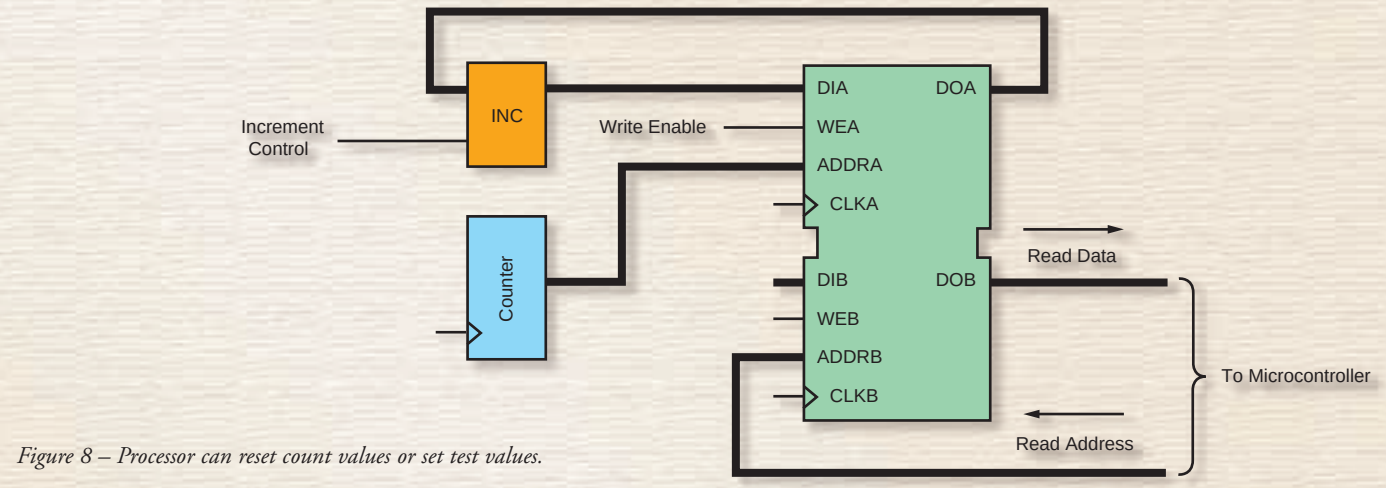

Figure 8 – Processor can reset count values or set test values.

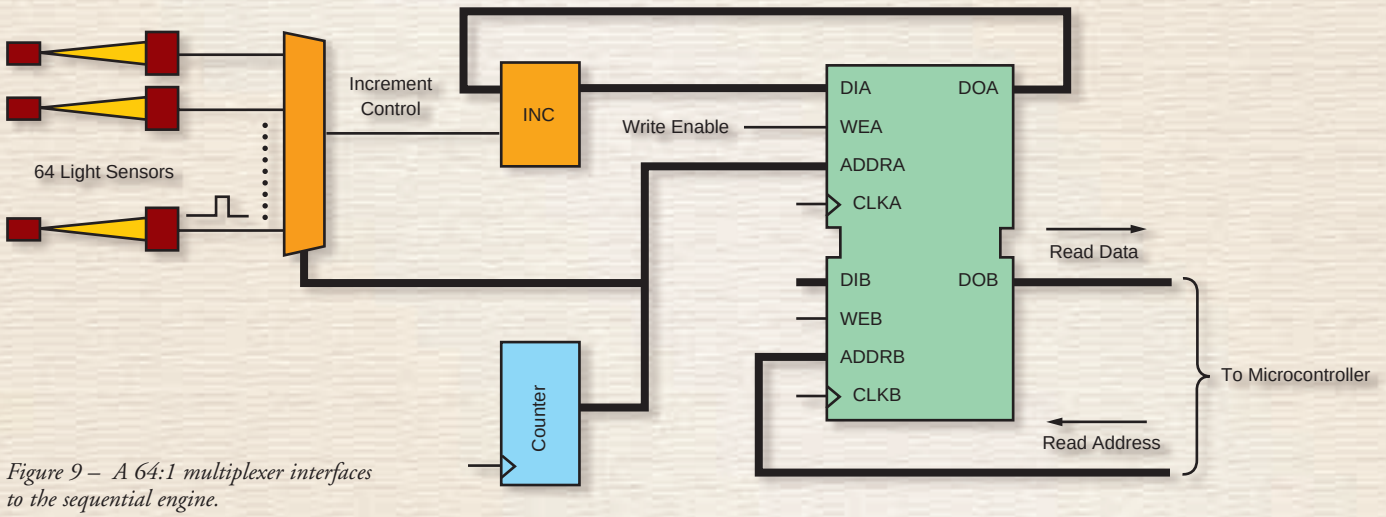

Figure 9 – A 64:1 multiplexer interfaces to the sequential engine.

best organized as a two-cycle process, but is no issue given the "Performance + Time" that is available.

Although we could selectively access the count values to be recorded as a corresponding light beam is broken, it is much easier to scan sequentially through all 64 count values and record only those which must be increased before the value is written back into the RAM. This reduces the address generation to a simple 6-bit counter.

At this stage, we have replaced 384 slices with one block RAM and just nine slices of logic (six for the increment and three for the address counter). This is a huge savings. Now, however, we must find ways to connect the inputs and outputs to this 3-D processing engine.

Eliminate the Data Multiplexer

The parallel data multiplexer is simply not required in this design. We save 216 slices instantly because the count values are now held in one consolidated memory. The dual-port nature of the block RAM really makes it very easy to connect the external processor.

As illustrated in Figure 8, the memory also offers the opportunity for the processor to have a write mode to reset count values or set test values. As with the parallel implementation, there is a risk that the processor will try to read a count value that is in the process of being modified. However, it's very easy to allocate time for the recording process and time for the processor to read values.

Although a clock rate of a few KHz is adequate for the processing, a clock of 2 MHz (or similar clock rate associated with the microcontroller) would achieve a count value update scan in 64 µs, leaving nearly all of the 30 ms processing period available for the microcontroller to read or write count values.

Connecting the Sensors

At some point in all 3-D designs, the parallel world must be interfaced to the sequential processing engine. This does not have to be difficult, and often a simple method is adequate, as seen in Figure 9.

The counter used to access each count value from the RAM can be used to select the associated sensor via a 64:1 multiplexer.

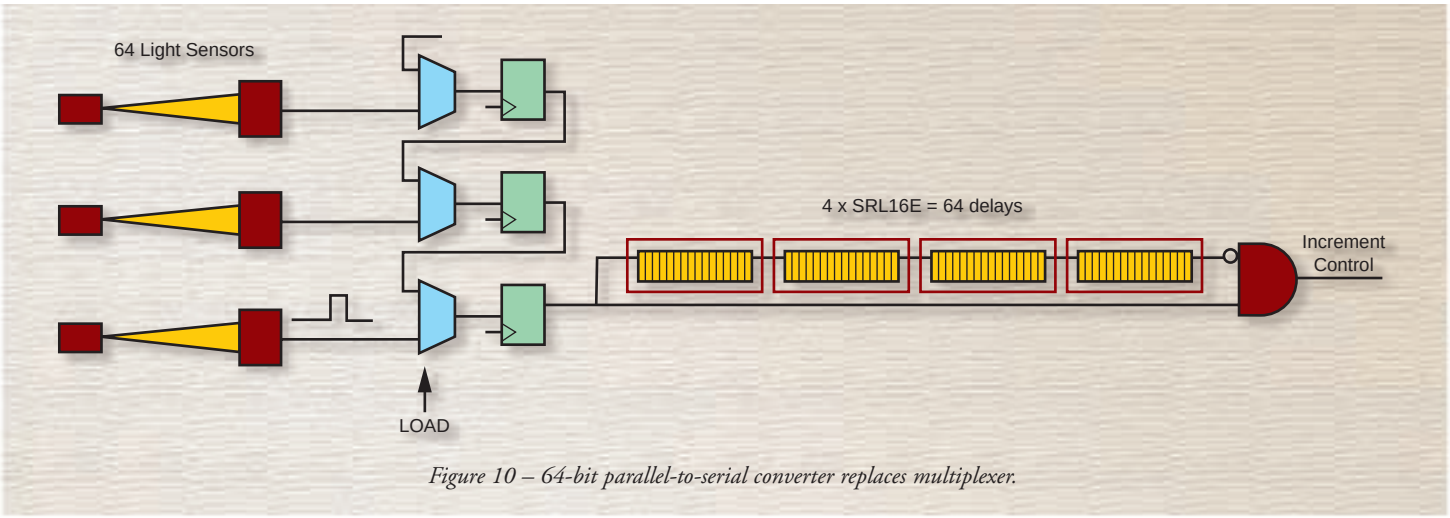

Figure 10 – 64-bit parallel-to-serial converter replaces multiplexer.

Although this requires multiplexer logic, it is just for one bit and therefore only requires nine slices.

Each sensor still requires its own logic. This is partly to synchronize the input signals, but is also required to ensure that each "beam broken" pulse is only used to record a count value once. For this reason, the one slice per sensor is unlikely to be reduced.

When you see that the logic size is increasing because the function is becoming more parallel, it is worth looking to see if anything else can be time-shared and moved into memory. In this case, we can indeed improve things.

We can replace the multiplexer with a 64-bit parallel-to-serial converter (32 slices), which converts the parallel domain into a serial sequential process, as demonstrated in Figure 10. To detect only the start of a new pulse, a memory is used to remember the last state of each of the 64 sensors. Because the operation is so predictable, we can use the SRL16E memory mode, which requires just two slices.

Dramatic Cost Reduction

So was it worth it? I think the diagrams in Figure 11 speak for themselves.

To reduce the function from 332 CLBs to just 22 CLBs is a dramatic change: 15 times smaller. Our design now fits in the smallest Spartan-II device (XC2S15) – and actually only uses 25% of that.

This reduction in size and cost is not just specific to this particular design. For example, much of 3G wireless processing is involved with "chip rates" of 1.2288 MHz and 3.84 MHz. This provides the time to allow the performance and memory of Virtex devices to process at least 32 channels sequentially, in just the same way as our simple fruit counter.

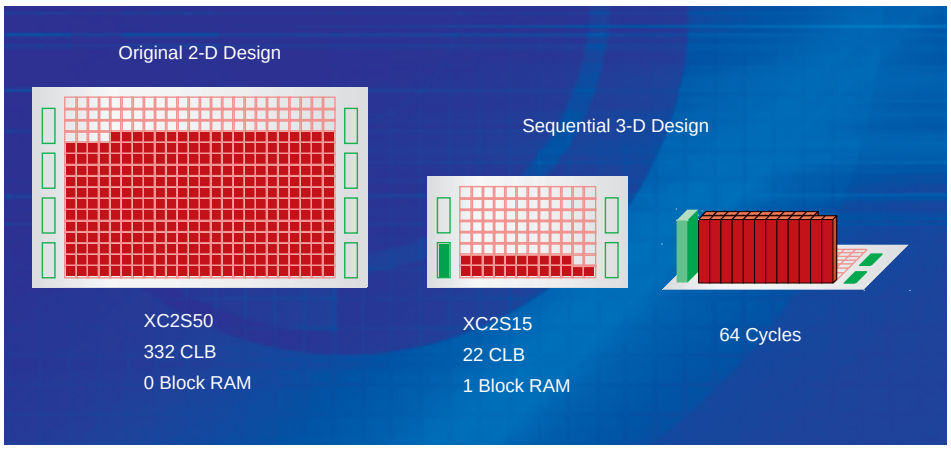

Figure 11 – 3-D design yields compelling results.

Final Considerations

The Spartan-II XC2S15 has only 86 user I/Os, and our design has high I/O demands. Having used 64 for sensor inputs and applied a clock, only 21 I/Os are left for the microcontroller interface. Given an 8-bit data bus, it is possible to connect to the microcontroller, but it does illustrate how I/O can limit a design once these highly efficient techniques are employed.

Of course, it would be a pity for 75% of the XC2S15 to be completely wasted. It would be nice to embed the microcontroller and the UART in the same device. This is also possible, but it's a topic for another article.

Meanwhile, once you discover that 3-D designs are possible, you are well on your way to improving the profit margins on your own designs.

[Editor's note: This article was derived from a two-part TechXclusive on the *support.xilinx.com* website. To see the original TechXclusive, go to *support.xilinx.com/support/techxclusives/3-D-techX22.htm* and *support.xilinx.com/support/techxclusives/3-D2-techX23.htm.*

To see more TechXclusives, go to *support.xilinx.com* and search for *"TechXclusives,"* then click on *"Xilinx TechXclusives Home."*]

How to Implement a DDR SDRAM Controller

The DDR SDRAM specification can make designing a controller tricky, but using an innovative approach with Xilinx Virtex-II devices lets you avoid the time and costs associated with ASIC implementations.

by Tim Murphy
Senior Design Engineer
Plexus
tim.murphy@plexus.com

Our customer's application required buffering a large amount of data, which prevented us from using internal memory and also made external SRAM cost prohibitive. Product volumes were too low to justify spending the non-recurring engineering costs associated with an ASIC, especially for initial prototypes. The data transfer rate through the buffer ruled out a single data rate SDRAM. We needed a double data rate (DDR) SDRAM controller running at 200 MHz.

Several built-in Virtex™-II features enable a design at DDR SDRAM clock rates. Although implementing the controller was still a challenge, we found the ideal solution – from both a cost and bandwidth perspective – was a DDR SDRAM with the controller implemented in a Virtex-II XC2V2000-5FG676 device.

DDR SDRAM Basics

DDR SDRAM is very similar to single data rate SDRAM, except that data is clocked into or out of the device on both the rising and falling edges of the clock. All control signals, however, still change only on the rising clock edge. During any data access, the controller provides the DDR SDRAM with a clock, inverted clock, address, and control signals. During a write cycle, the controller also provides data and a data strobe signal (DQS). During a read cycle, the DDR SDRAM provides data and the DQS signal.

The DQS signal is used to clock data into the DDR SDRAM on a write and into the controller on a read. Bank precharging, refreshes, and so forth are handled in a DDR SDRAM controller in much the same way they are handled in a standard SDRAM.

Design Challenges

Keys to a successful DDR SDRAM controller design include a thorough understanding of the clocking specifications and the clever use of the Virtex-II digital clock manager (DCM).

Clocking

The DDR SDRAM specification requires that the clock and inverted clock received from the controller cross within a very tight window, as shown in Figure 1. The crossing point of these clocks is considered the clock edge in the DDR SDRAM specification. To ensure our design met this window, we used an LVDS signal pair configured using the SSTL signaling standard. Even though LVDS wasn't used, the fact that this pair of outputs could be used to drive 622 Mbps differential data ensured that the skew between the outputs was small enough to meet the DDR SDRAM requirement.

Data Capture

As mentioned earlier, data is clocked into and out of the DDR SDRAM using the DQS signal. The DQS signal is bidirectional. During a write, it's driven by the controller. To maximize setup and hold time windows, the controller must drive DQS 90° out of phase with the data. Data is clocked by the DDR SDRAM on both edges of DQS.

During a read, the DDR SDRAM provides both data and DQS. However, the DDR SDRAM provides data and DQS coincident with each other. This means the controller must either provide the 90° phase shift internally or find another way to clock in data. In addition, DQS is a strobed signal. It is driven while there is a transaction in progress, but tri-stated otherwise. These two items complicate the controller design a bit.

Because DQS is strobed, a DCM can't be used to provide the shift because the DCM needs a continuously running clock.

The DQS net could potentially be delayed inside the FPGA, but there is no way to specify a minimum delay to the tools. FPGA Editor could be used to add delay by hand, but relying on a minimum delay is bad design practice. The DQS net could be delayed on the board. This is a bit more straightforward. Assuming the PC board is made of FR4 material, the net will see a delay of ~150 ps/inch. This helped our read timing, but unfortunately hurt our write timing by the same amount.

The Solution

Our design solved the problem by not using DQS at all – except as a data enable. Instead, we started with the internal clock sent to the DDR SDRAM (referred to as the write clock) and ran it through a Virtex-II DCM. The DCM phase-shifted the clock by an amount carefully determined through a detailed timing analysis of the system. This clock is referred to as the read clock.

In addition, we treated the read and write clocks as asynchronous and used the DQS as a data enable, as shown in Figure 2. As a result, our design can tolerate a skew between the read and write clocks of anywhere from 0 ns to 10 ns (one clock period). Therefore, if the SDRAM timing is affected by a new SDRAM part or board change, the FPGA does not have to be re-done.

Other Useful Virtex-II Features

We also took advantage of several other Virtex-II features in our design:

- The Virtex-II has DDR registers built into the input/output blocks (IOBs). This feature allowed us to isolate the 200 MHz DDR signaling in the IOBs, and allowed the rest of our design to run at 100 MHz.
- XCITE – Xilinx digitally controlled impedance (DCI) technology allowed us to terminate the DDR SDRAM signals without using any external resistors.
- The DCMs allowed us to synthesize clocks, limit jitter, and reduce clock skews.

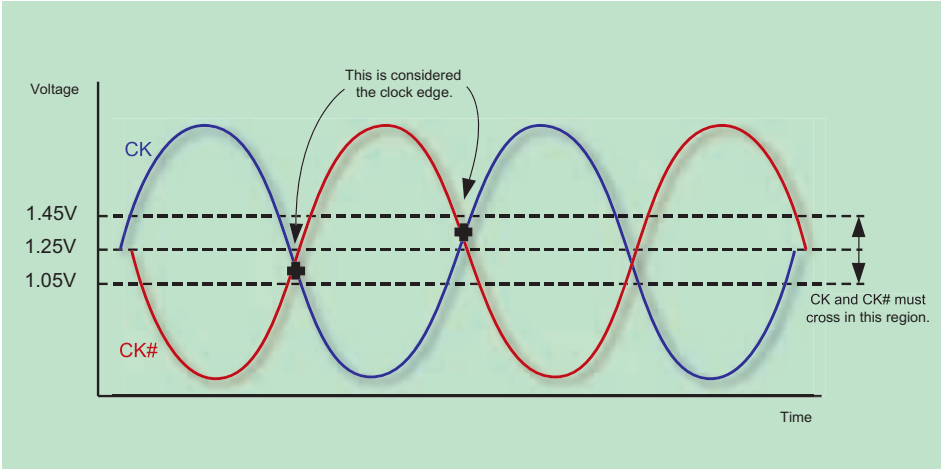

Figure 1 – Clock crossing

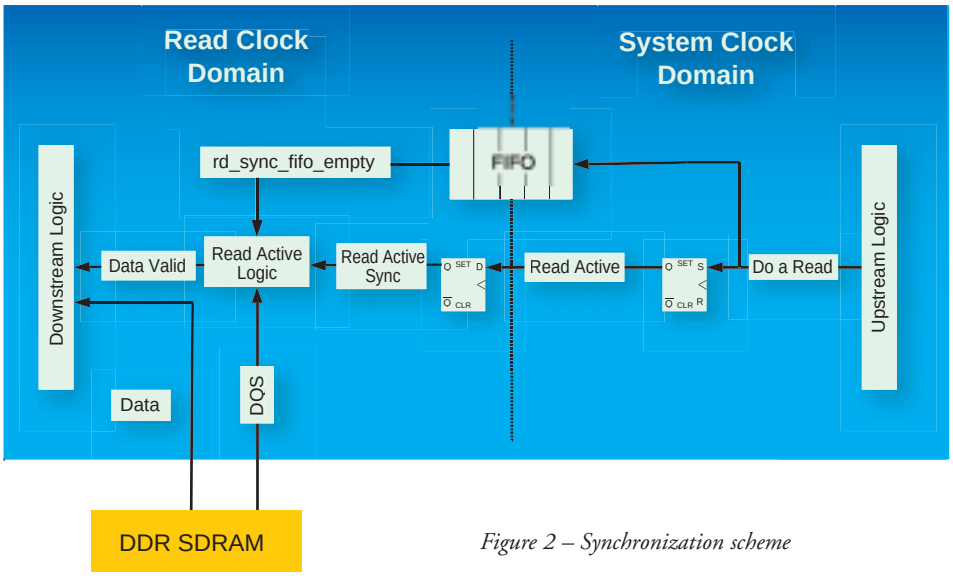

Figure 2 – Synchronization scheme

Process

The architecture for the entire design was created before any RTL code was written. This practice allowed us to explore various alternatives quickly and select the best one. When we started our RTL coding, we ran into very few problems.

A DDR SDRAM has many corner cases. It is extremely difficult to find problems in hardware. However, using Mentor Graphics' ModelSim™ simulator, Xilinx UNISIM libraries, and Micron's DDR SDRAM memory model, we successfully simulated and debugged all corner cases.

DDR SDRAM corner cases can be difficult to control because they are independent of the design's intent. Here's one example: A test writes a series of words into the SDRAM, but the SDRAM needs a refresh in the middle of the burst. These corner cases are very hard to target. Generating random tests and running code coverage helped determine which corner cases were hit and which needed more work.

Conclusion

Designing a DDR SDRAM controller can be tricky. By utilizing the Virtex-II features along with a solid process and a little creativity, however, we were able to implement a 200 MHz DDR SDRAM controller in a –5 speed grade Virtex-II FPGA.

To design this controller in an ASIC would have cost a great deal more, and it could not have been possible in previous FPGA architectures. Visit Plexus at *www.plexus.com* for more information.

Is Your FPGA Design Secure?

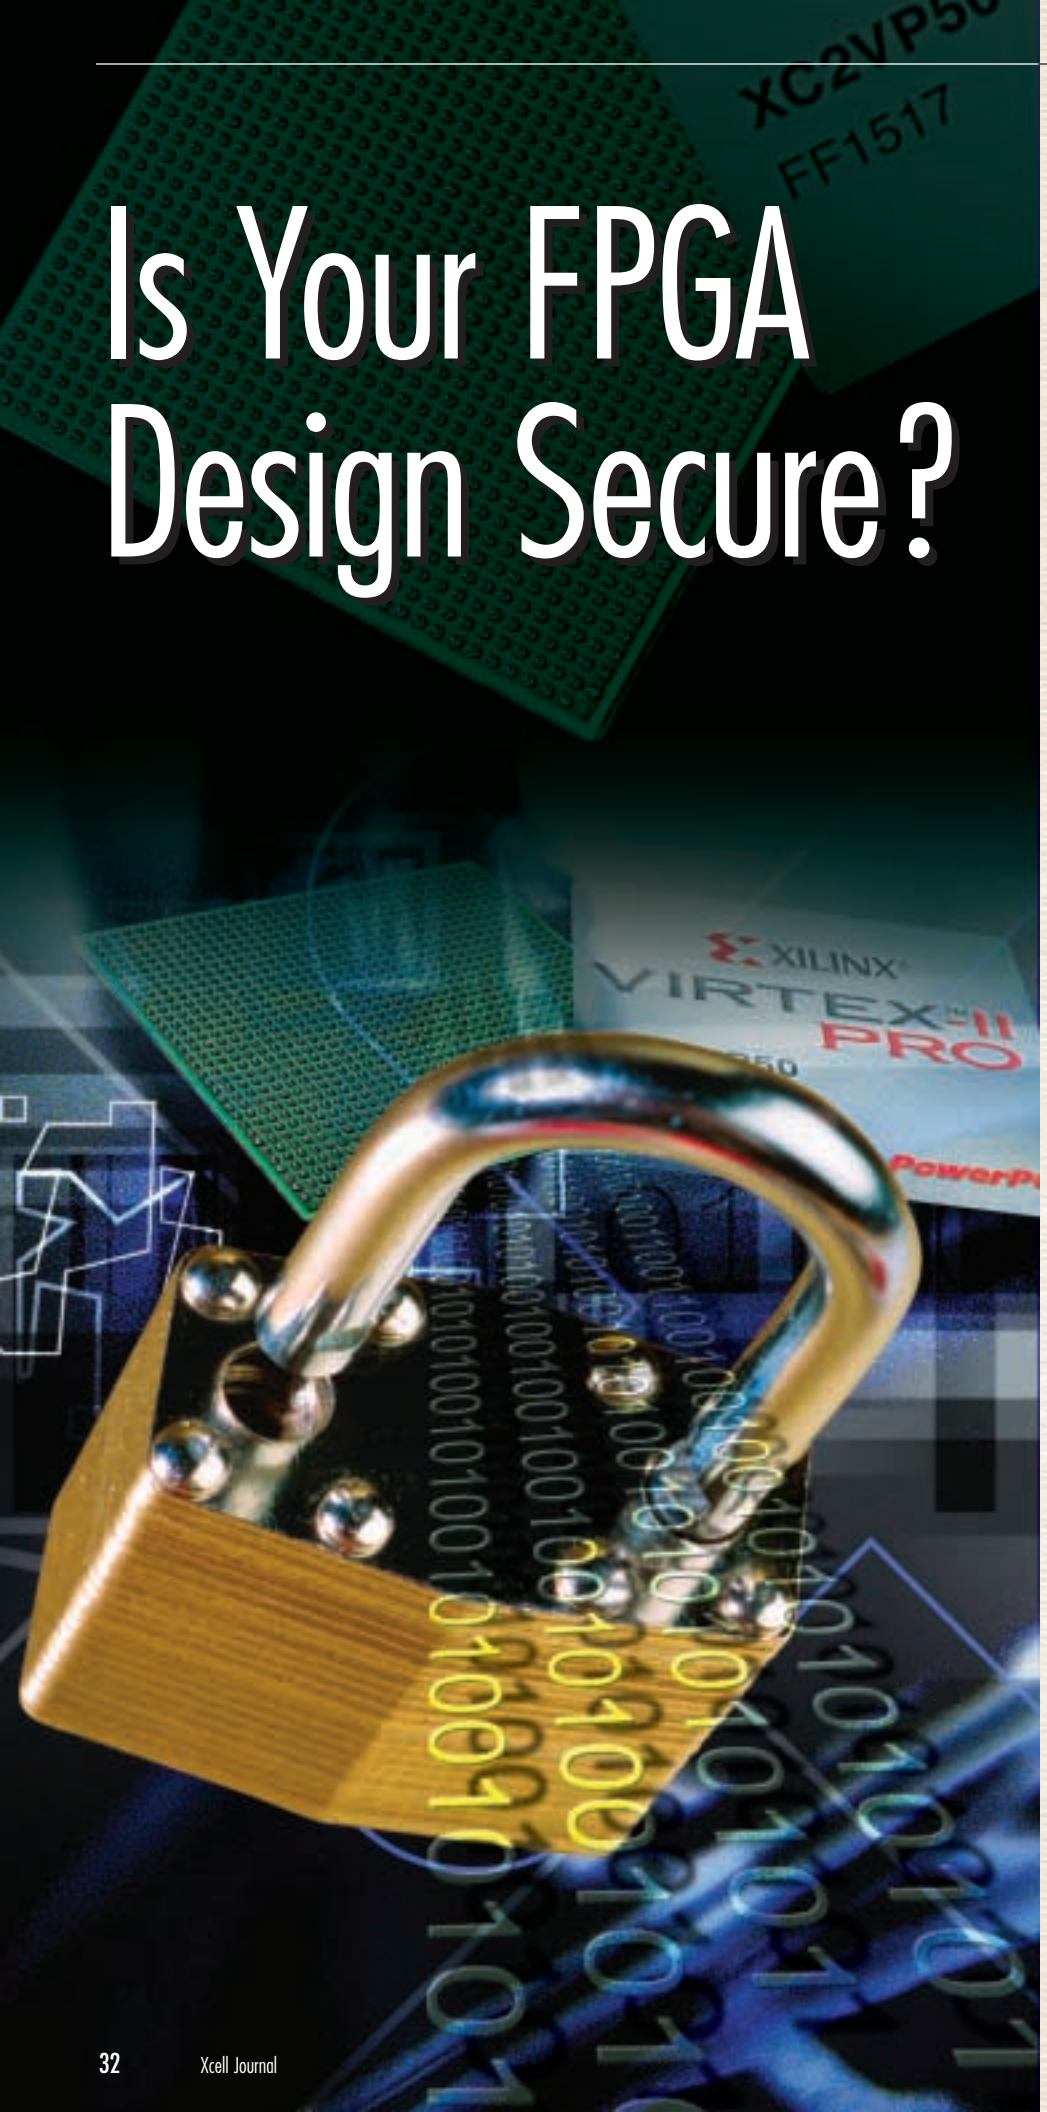

After spending months on your design, the last thing you want is to find your design has been stolen. Say goodbye to "locks," "fuses," "antifuses," and other contraptions. You can sleep peacefully when you design with Xilinx.

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

With an increased focus on security and terrorism in the world, it is natural for you (and me) to worry about the security of our designs and products. Are your ICs really foolproof? Can someone steal your design inside the chip? What about cloning? Reverse engineering?

Here at Xilinx, we want you to rest at ease. All Xilinx devices (FPGAs and CPLDs) have robust security mechanisms that make it nearly impossible to steal designs. In particular, Virtex-II Pro™ Platform FPGAs have advanced Triple DES (Data Encryption Standard) security embedded in them. We are not talking about my grandmother's "secure" peppermint box with its heavy lock that you or I could crack open with a toothpick. Triple DES algorithms provide the same security that the world's financial institutions rely on every day for transactions involving trillions of dollars. With security mechanisms you can "bank" on, let's get into the details of design security and how Xilinx protects your valuable proprietary designs.

Design Security

Consider the philosophy of theft. There is nothing such as absolute security in this world. A determined thief can break any barrier if given enough time and resources.

Our goal is to make it extremely difficult to break that barrier. Hence, we've raised that barrier to the state of the art. Breaking the Xilinx barrier is virtually impossible for the vast majority of pirates. To help you understand just how we do this, let's look at the three basic levels of attack on a security barrier.

Categories of Security Attacks

IBM™ defined three categories of security attacks in a paper published in 1991. Although the paper is somewhat dated, it still serves as a pretty good reference (Abraham, D., G. Dolan, G. Double, and J. Stevens, 1991. Transaction Security System. *IBM Systems Journal* 30(2): 206-229.).

1. Class I: This is a clever outsider and a curious person with negligible resources. Representative of the majority of the population of hackers, he is not interested in wholesale piracy. He simply wants a capability for personal use. Seemingly harmless, this class of thief can pose a significant, worldwide threat if he shares the ill-gotten information on the Internet.
2. Class II: This person is a knowledgeable insider with access to some sophisticated resources. Examples include university students and unscrupulous corporate employees with access to the Internet. However, the information provided by this class tends to be esoteric and only usable by the same class or higher.
3. Class III: This is a funded organization with a determined team of experts – people who can crack open literally anything. Examples include the FBI, CIA, NSA, and any other large commercial or governmental operation with unlimited funds for wholesale reverse engineering operations.

Class I and Class II attacks can be deterred; a good IC helps the designer achieve this objective. Xilinx Virtex-II Pro FPGAs provide a redundant key structure to thwart even most Class III attacks.

Programmable Device Security

There are two major types of design theft – cloning and reverse engineering. Cloning is copying the design as-is. Reverse engineering is more sophisticated stealing, where the thief extracts the design implemented in a device and then improves, modifies, and disguises it so that it no longer looks like the original implementation.

ASICs and ASSPs implement a fixed function and are vulnerable to attack, because every via/connection represents real logic. An attacker can de-cap such chips to observe the vias and hack the design. Sometimes bogus logic is included to confuse the attacker, but the chips are still at risk. In general, the more fixed an IC is, the more the risk is. Plus, in a fixed IC, if you suspect that the security of one chip is compromised, then all the chips are compromised. With programmable ICs, however, if one chip is compromised, all the other chips can be reprogrammed.

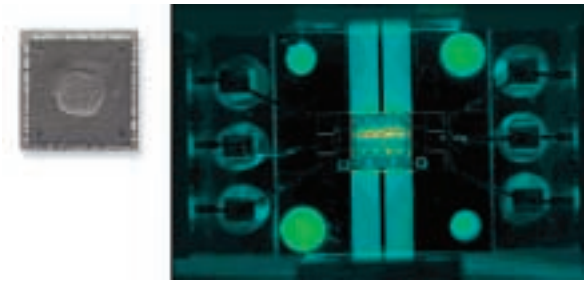

Figure 1 – An attacker can extract nonvolatile device die from a chip package and find security bits using thermal imaging.

Programmable devices (FPGAs and CPLDs) come in two types – nonvolatile (antifuse/flash) and volatile (SRAM) devices. There are many differences between them, each with unique advantages.

For our discussion on security, we will only focus on the security of the programming/configuration bits. Nonvolatile devices retain the configuration bits and do not need to load the bitstream from an external PROM every time they are powered up. In the past, the external PROM was a cause for security concerns and nonvolatile devices had a slightly upper hand.

But that was before the debut of the industry's first secure SRAM FPGAs in 2000 – the Xilinx Virtex™-II Platform FPGAs with Triple DES encryption.

Nonvolatile Programmable Devices

Let's take a look at today's nonvolatile programmable logic devices (PLDs) that include both FPGAs and CPLDs from competing vendors. These devices provide a test mode to observe the internal circuit to improve yields as well as a read-back mode for customers to inspect the programmed bits.

Both these modes expose the devices to attacks. To prevent such attacks, fixed security bits are built with antifuse or flash gates to prevent read/write from the device after it is ready. Here is the weakness: An attacker can implement a simple circuit in two identical devices – one secured and the other not. The security bits have fixed locations and can be distinguished by observing both the devices using thermal imaging (Figure 1). Once distinguished, the security bits can be disabled with a high-intensity light source impinged on the die to allow read-back. Any circuit implemented in that particular device is now available to the thief.

As device size increases in programmable devices, the configuration bitstream size increases. When key size also increases with the device size, it indicates a fixed location pattern of security bits. We can conclude that if one such device is cracked, the security of devices of all densities is cracked. Sounds scary? You bet.

Xilinx CoolRunner™-II CPLDs use flash technology, but they are different from other nonvolatile devices. CoolRunner-II CPLDs add multiple layers of security, with read/write-protect security bits that are hard to find. The bits are placed among the programming bits deep under several layers of metal.

The security bits are also placed in such a way that requires a specific sequencing of signals to set and clear

Figure 2 – Several metal layers in CoolRunner-II CPLDs prevent the theft of your design.

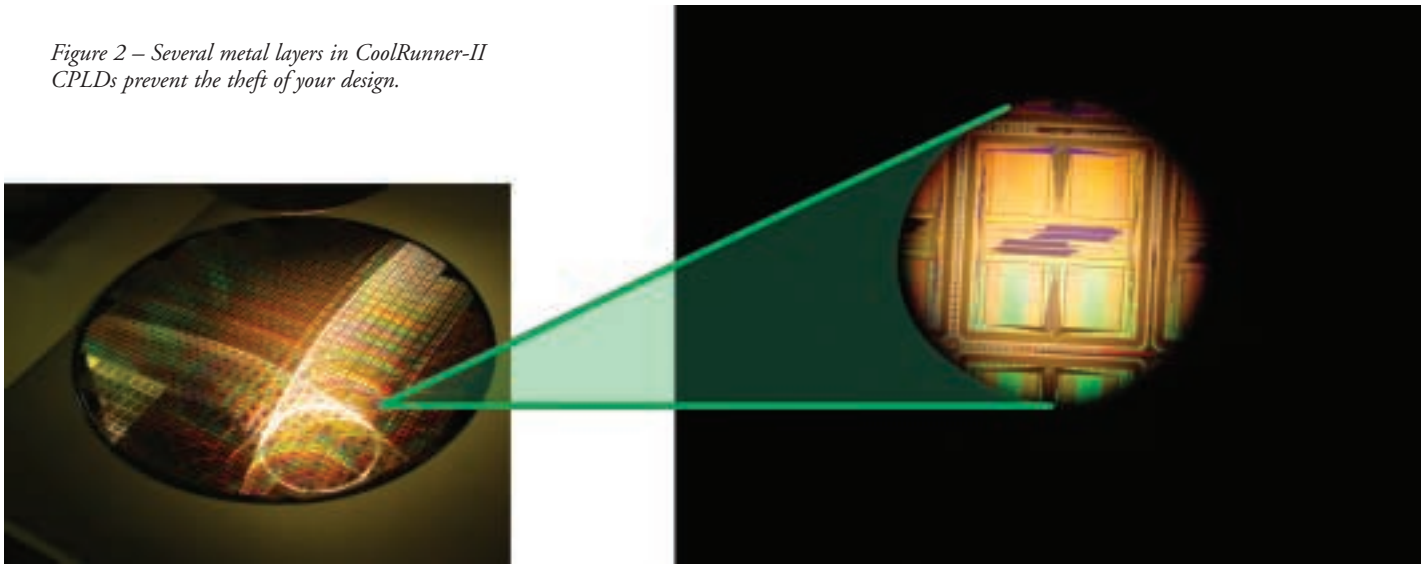

them, as well as charge pumping. In addition, four or five layers of metal (Figure 2) rule out direct exposure of top die to laser/electrical tampering.

Two CoolRunner-II CPLD technologies – DualEDGE and DataGATE – are schemes that also confuse attackers with double data operation and input signal locking under internal macrocell control. Furthermore, the legendary low power operation of CoolRunner-II CPLDs makes it difficult to see anything under thermal imaging. For more information on CoolRunner-II security, see the white paper "CoolRunner-II CPLDs in Secure Applications" at *www.xilinx.com/publications/whitepapers/wp_pdf/wp170.pdf.*

Volatile Programmable Devices

Volatile SRAM FPGAs lose the configuration bitstream whenever the power goes off, and thus they need an external memory to hold the bitstream. A typical application includes an external PROM next to the FPGA. The bitstream between the PROM and the FPGA is a cause for security concern for some designers.

To alleviate these concerns, Xilinx embedded Triple DES encryption technology into Virtex-II FPGAs three years ago.

The second generation of this technology is now embedded in the latest Virtex-II Pro FPGAs. With this scheme, the PROM contains the encrypted bitstream, which is decrypted inside the FPGA. No read-back is allowed. Any form of attack on the FPGA erases the design. The best an attacker can do is to intercept the design between the PROM and the FPGA – but all he will obtain is the triple DES encrypted bitstream.

Security You Can Bank On

Triple DES is the pre-eminent encryption standard that is used by financial institutions for millions of transactions every day involving trillions of dollars.

This same technology is used to provide security for your designs in Xilinx FPGAs. Although there have been reports of the older DES standard being cracked, Triple DES is yet to be cracked.

Triple DES is an official NIST (National Institute of Standards and Technology) and American National Standards Institute (ANSI) X9.52 standard. AES (Advanced Encryption Standard) is the next generation of encryption that will be adopted by the industry if and when Triple DES is cracked. Encryption gurus are already working on developing future versions of AES in a constant effort to keep hackers at bay. With trillions of dollars at stake, millions of financial transactions, and perhaps thousands of hackers at work on various types of attacks, Triple DES has held its ground to date. This is why Xilinx chose Triple DES as the best encryption capability to use in its flagship Virtex-II Pro product line.

Triple DES is a symmetric encryption algorithm, which means the keys used to encrypt and decrypt are the same. The security of the data lies in the key – in contrast to public key systems such as RSA or PGP. Triple DES uses three keys and the encryption algorithm is repeated for each key for added security. Each key is 56 bits wide and encrypts 64-bit blocks at a time. For more information on the Triple DES standard, go to: *www.itl.nist.gov/.*

Triple DES in Virtex-II Pro FPGAs

Virtex-II Pro devices have an on-chip decryption engine that can be enabled to secure the configuration bitstream, and hence the FPGA. You can encrypt the bitstream in the Xilinx software with a set of keys, and the Virtex-II Pro device decrypts the incoming bitstream internally using the same set of keys (Figure 3).

Once the design is placed and routed in ISE tools, the encrypted configuration bitstream is generated using the ISE's BitGen program with user-selected keys (Figure 4). The same keys are loaded into the FPGA through the JTAG port using the ISE iMPACT tools (Figure 5).

To program the keys into the device, the device has to be put into a key access mode, which automatically erases the FPGA – including any old keys and the actual bitstream. This also provides protection from attacks.

After the device is programmed with

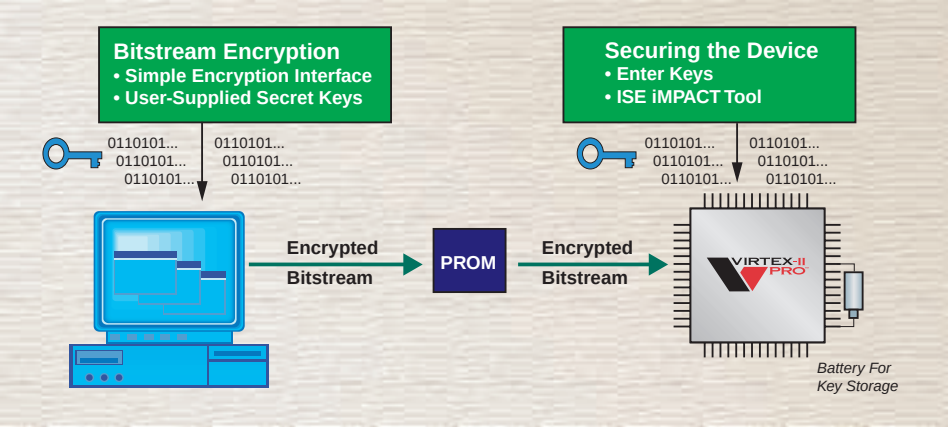

Figure 3 – Design security in Virtex-II Pro FPGAs using Triple DES encryption/decryption

keys, it can be configured with the encrypted bitstream. Once an encrypted bitstream has been programmed into the FPGA, it cannot be reconfigured, or partially configured, or read-back by unintended or intended tampering or snooping. Any attempt to steal a design automatically erases the FPGA completely. (Note that non-encrypted bitstreams may be programmed into an FPGA loaded with keys to facilitate test and debug.)

The keys reside in a special Triple DES block at the corner of the device. An external battery is used to hold the keys when the board is powered off. Any standard battery between 1V and 4V can be used, yielding a 15-year life.

Although nonvolatile PLD security bits disable read/write into the device, Virtex-II Pro Triple DES provides real security using encryption. Even though the location of the Triple DES block in Virtex-II Pro FPGAs is known, the keys cannot be observed without erasing the chip.

Nine layers of metal pose yet another barrier to would-be thieves. Any attempt to use thermal imaging to find the security bits buried inside Virtex-II Pro FPGAs is destined to failure, because there isn't even hard wiring involved to store the security bits.

Each Virtex-II Pro device provides six separate keys, thus allowing two sets of Triple DES keys. You can program the devices up-front with both sets of keys. This provides security even against Class III attacks. For example, if the first key set is compromised, you can remotely reprogram devices in the field with a bitstream using the second set of keys.

In addition, some of our customers use this technique for new business strategies and pricing models – one key set bitstream enables a low-price, low-feature service, and the other set activates a high-price, high-feature service.

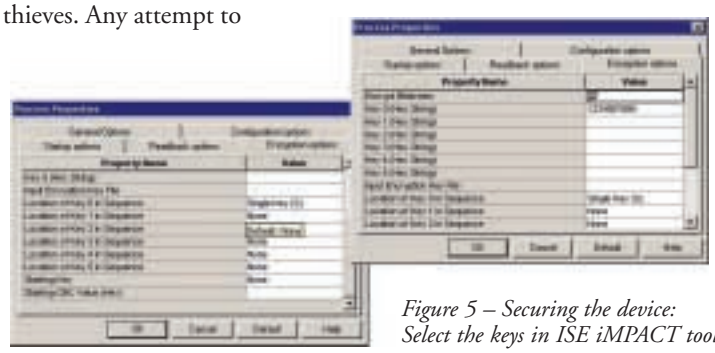

Figure 5 – Securing the device: Select the keys in ISE iMPACT tools.

Export Considerations

Encryption standards are tightly controlled in the U.S. by the Department of Commerce. The decryptor on the chip can only decrypt the incoming bitstream and is not available as a design block; hence, Virtex-II Pro FPGAs have been classified as field programmable logic devices (3A001.a.7), the same as any other FPGA. The software has been classified under ECCN# 5D002 and can be exported globally – with the exception of countries banned by the Department of Commerce. The bottom line is that no changes to your current export practices are necessary to include Virtex-II Pro FPGAs in your system.

Figure 4 – ISE design tools flow

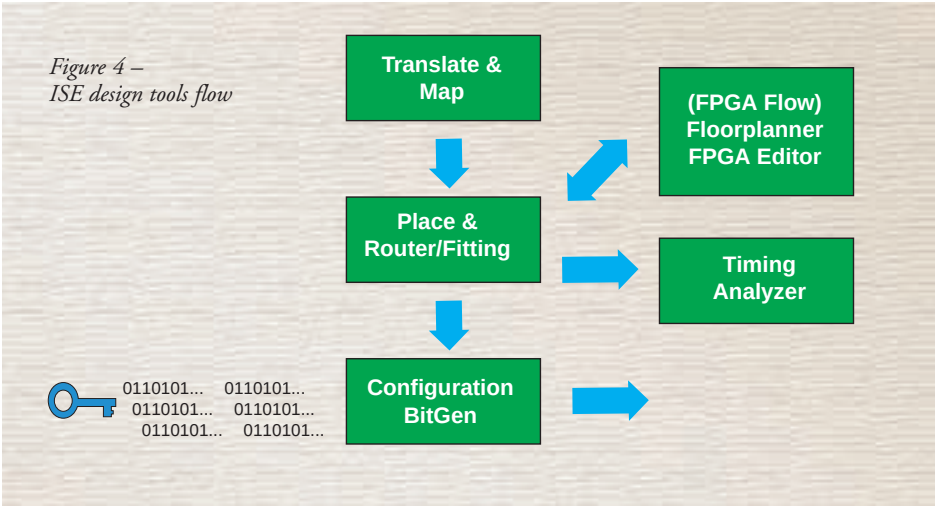

Conclusion

Xilinx FPGAs and CPLDs provide unparalleled security for your designs, helping you achieve your design goals while maintaining the highest design security you can get.

This dedication to the protection of your designs is one of the reasons why Xilinx has become the largest vendor of programmable logic devices in the world. We offer the highest performance, lowest cost solutions with maximum design security so you can rest easy.

Virtex FPGAs Tame DSP-Hungry Broadband Wireless Designs

Second-generation Broadband Fixed Wireless Access standards are still coalescing – but all require high-performance signal processing power. You can meet your processing goals while maintaining standards flexibility by using Virtex FPGAs.

by David Nicklin
Senior Manager, Strategic Solutions Marketing
Xilinx, Inc.
dave.nicklin@xilinx.com

David Eden
Systems Specialist, Europe
Xilinx, Inc.
david.eden@xilinx.com

Broadband Fixed Wireless Access (BFWA) technology standards are consolidating behind the IEEE 802.16 specification in the U.S. and the ETSI BRAN (European Telecommunications Standards Institute Broadband Radio Access Networks) specification in Europe. This convergence makes BFWA technology commercially attractive, particularly as a flexible, cost-effective, 3G (third-generation) cellular backhaul solution.

Both standards take broadly similar approaches to solving the Quality of Service (QoS) and bandwidth-utilization limitations experienced by first-generation BFWA services. However, both standards demand greater signal processing performance than conventional DSPs can deliver today, which leaves the door open for developing custom hardware to boost performance.

The BFWA market remains difficult to predict, especially in the domestic and SME (small and medium enterprise) sectors. Some observers, for example, believe standards must harmonize further to make the technology viable. Therefore, developers must find ways to achieve hardware acceleration without compromising flexibility, time to market, or cost-effectiveness.

System Challenges

The differing requirements of 3G backhaul and residential/SME applications mean that, as a designer, you must create access solutions that are capable of handling a wide range of services. These include legacy services such as time division multiplexing (TDM), IP (Internet Protocol), and VoIP (Voice over Internet Protocol). A variety of backhaul requirements must also be accommodated, including ATM and packet-based protocols. In other words, your design must have enough flexibility to efficiently carry any traffic type.

Maintaining QoS and delivering 99.999% availability, or "five 9s," is made more difficult by environmental effects such as rain, obstructions, or other non-line-of-sight conditions. Both the IEEE and ETSI standards have introduced adaptive coding and modulation algorithms that make the best use of the available bandwidth under favorable link conditions, and they invoke a more reliable alternative to maintain availability under unfavorable link conditions. Co-channel interference also acts to degrade link quality, and both standards address this issue as well.

A wireless access system may be presented with multiple connections per terminal, multiple QoS levels per terminal, and a large number of statistically multiplexed users. As a result, second-generation systems demand an extremely high-performance, scalable, signal processing platform designed for wireless processing and dataflow.

Physical Layer Design

While the Media Access Control (MAC) layer handles algorithms relevant to the various traffic classes, complex adaptive coding and modulation are performed at the PHY (physical) layer.

The PHY specifications for IEEE 802.16 and ETSI BRAN standards establish a burst specification that allows both time-division duplexing (TDD) and frequency-division duplexing (FDD). Both TDD and FDD support adaptive burst profiles in which transmission parameters relevant to modulation and coding may be assigned dynamically on a burst-by-burst basis. Further complexity is added by including support for half-duplex FDD subscriber stations (which may help to reduce the cost of subscriber equipment because they do not simultaneously transmit and receive).

Variable burst profiles require extensive processing resources, including Reed-Solomon forward error correction (FEC) with variable block size and error correction capabilities. Moreover, FEC is combined with 16-state quadrature amplitude modulation (16-QAM) or 64-QAM in both ETSI BRAN and IEEE 802.16. ETSI BRAN also defines 4-QAM.

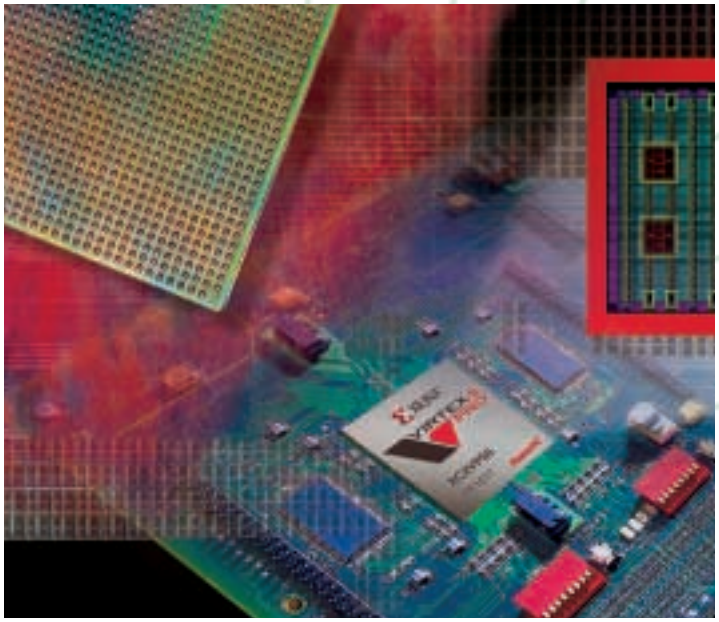

A PHY implementation for either standard integrates transceiver and CODEC (coder/decoder) functions as well as backplane interfaces and data queuing. This implementation calls for complex mixers, frequency synthesizers, mapping capability, automatic gain control, and access point synchronization in the transceiver. Additional CODEC blocks include: Viterbi and convolutional coding; interface buffering; PDU header controls and scramble/descramble; and MUX/DEMUX (multiplexer/demultiplexer) functions. Additional interfacing and data-queuing functions include frame formatting, cell queuing, and lookup or PDU overhead formatting.

As a result, ETSI- or IEEE-compliant base stations must integrate extensive processing capabilities, many of which you can perform in hardware.

Hardware Acceleration

In both second-generation BFWA standards, the algorithmic complexity of the PHY layer in particular has accelerated faster than Moore's law can empower DSPs to keep pace. Still, a software-reconfigurable solution is desirable given the current uncertain market conditions. Market projections are under constant review, and some manufacturers believe further harmonization of standards must take place before the true potential of BFWA will be realized.

All of these factors make it extremely difficult to plan an ASIC development with confidence. These factors also make creating the ICs for standards-compliant equipment a significant challenge, especially if you want to avoid the fixed engineering development costs.

On the other hand, high-speed FPGAs allow you to create custom hardware and exploit the massive parallelism needed to meet performance goals without sacrificing flexibility. You can also achieve high integration by implementing DSP functions alongside protocol translation, glue logic, and other system functions. The Virtex™-II and Spartan™ FPGA families also provide a wealth of processing resources, including Xilinx MicroBlaze™ soft processor cores, as many as four embedded IBM PowerPC™ cores (in Virtex-II Pro™ Platform FPGAs), and user-selectable I/Os.

These resources combine well with large IP (intellectual property) libraries, through which Xilinx offers functions such as Viterbi, turbo-product coding, and other off-the-shelf functions.

FPGA Implementation

Figure 1 shows a block diagram for an ETSI-compliant BRAN PHY layer. You can create the design represented here with either two Virtex-II XC2V3000 FPGAs or one XC2V6000.

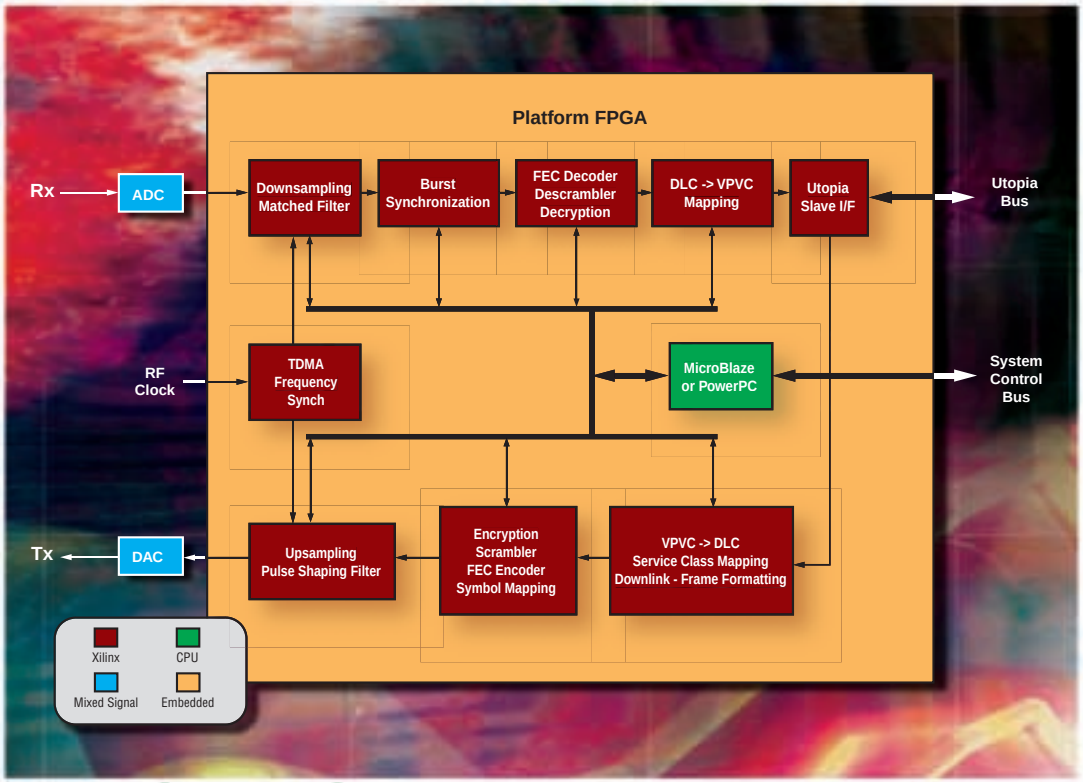

Figure 1 – Access-point transceiver for ETSI-compliant BRAN PHY layer

Transceiver Design

The access point transceiver shown in the diagram includes a Hilbert Transform decimator and interpolator and complex mixers running at around 100 Mbps, or double the symbol rate. You can implement these using the Xilinx System Generator tool, exploiting block RAM and embedded multipliers available in the Virtex architecture as well as its extensive logic resources and look-up tables (LUTs). The mapper function uses block RAM, multipliers, and LUTs to perform LUT-based mapping.

Additional major functional blocks include complex variable decimation and interpolation (CVD/I), linear interpolated digital frequency synthesis, and access point synchronization. You can easily build all of these functions using the XC2V3000/6000's plentiful block RAM, multiplier, and LUT functions.

CODEC

When designing the CODEC, the combination of Virtex high-speed processing performance and off-the-shelf IP is especially powerful. You can implement a standard Viterbi decoder directly in the Virtex-II architecture using approximately 1,000 slices and two block RAMs. Supporting OC3 (Optical Carrier 3 – 155 Mbps) data rates and higher, the Viterbi decoder implementation offers fully synchronous two-clock or one-clock versions to reduce latency, size, and power dissipation. Because the system-code rate for ETSI BRAN and IEEE 802.16 varies, the decoder also allows you to change puncturing on the fly by using control bits generated from data formatting. The decoder is available as VHDL source code or as a fixed netlist from the Xilinx LogiCORE™ IP program.

Both BFWA standards call for Reed-Solomon forward error correction with variable block size and error correction capabilities. We have modified the standard Xilinx Reed-Solomon IP core to meet the required frame duration and symbol rate, and to reduce latency below 92 byte periods. You can also implement convolutional decoding and encoding using standard Xilinx IP cores. The ETSI BRAN specification includes headroom for optional Turbo product codes, which are also readily available as complete IP cores from Xilinx.

At the Backplane

The Virtex-II XC2V3000/6000 devices can also implement backplane interface and data queuing functions by using the standard Utopia2 (Universal Test and Operations PHY Interface for ATM Level 2) backplane interface IP core. Furthermore, you can custom-define a state machine to perform functions such as scheduling and timing in about 200 slices of the XC2V3000/6000 devices. With all of this functionality, you still have plenty of block RAM and logic resources for queuing, overhead formatting, and header processing.

Conclusion

New-generation standards for broadband fixed wireless access have helped make the technology commercially attractive. But considerable uncertainty remains regarding future standards and markets. This uncertainty demands flexible, adaptable solutions.

At the same time, to answer the QoS and availability limitations of first-generation systems, second-generation standards require processing power beyond the capabilities of conventional reconfigurable DSPs.

Xilinx FPGAs deliver the flexibility and raw DSP performance you need as well as the extensive IP cores and engineering support to solve these diverse challenges.

For more information, visit these websites:

- Xilinx Wireless Networks Solutions – *www.xilinx.com/esp/wireless_networks/index.htm*
- IP cores suitable for BFWA – *www.xilinx.com/ipcenter/.*

nVisage captures it all

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 one application.

Only
$2,995

nVisage DXP™

Multi-dimensional design capture for every engineer

1225-nVsg-AD-XJ-US

nVisage™
Multi-Dimensional design capture from Altium™

nVisage, nVisage DXP, Design Explorer and Altium and their respective logos are trademarks or registered trademarks of Altium Limited or its subsidiaries.

Altium
Making Electronics Design Easier™

Call Altium Sales and Support now:
Free call: 800-470-2686
Fax: 858-485-4610
Email: nVisage.sales.na@altium.com
Web: www.nVisage.com

Refer to VALUE CODE 1225

Biology Goes Digital

An array of 5,700 Spartan FPGAs brings the BioWall to "life."

by Gianluca Tempesti, Ph.D.
Assistant Professor
Swiss Federal Institute of Technology
gianluca.tempesti@epfl.ch

Christof Teuscher
Research Assistant
Swiss Federal Institute of Technology
christof@teuscher.ch

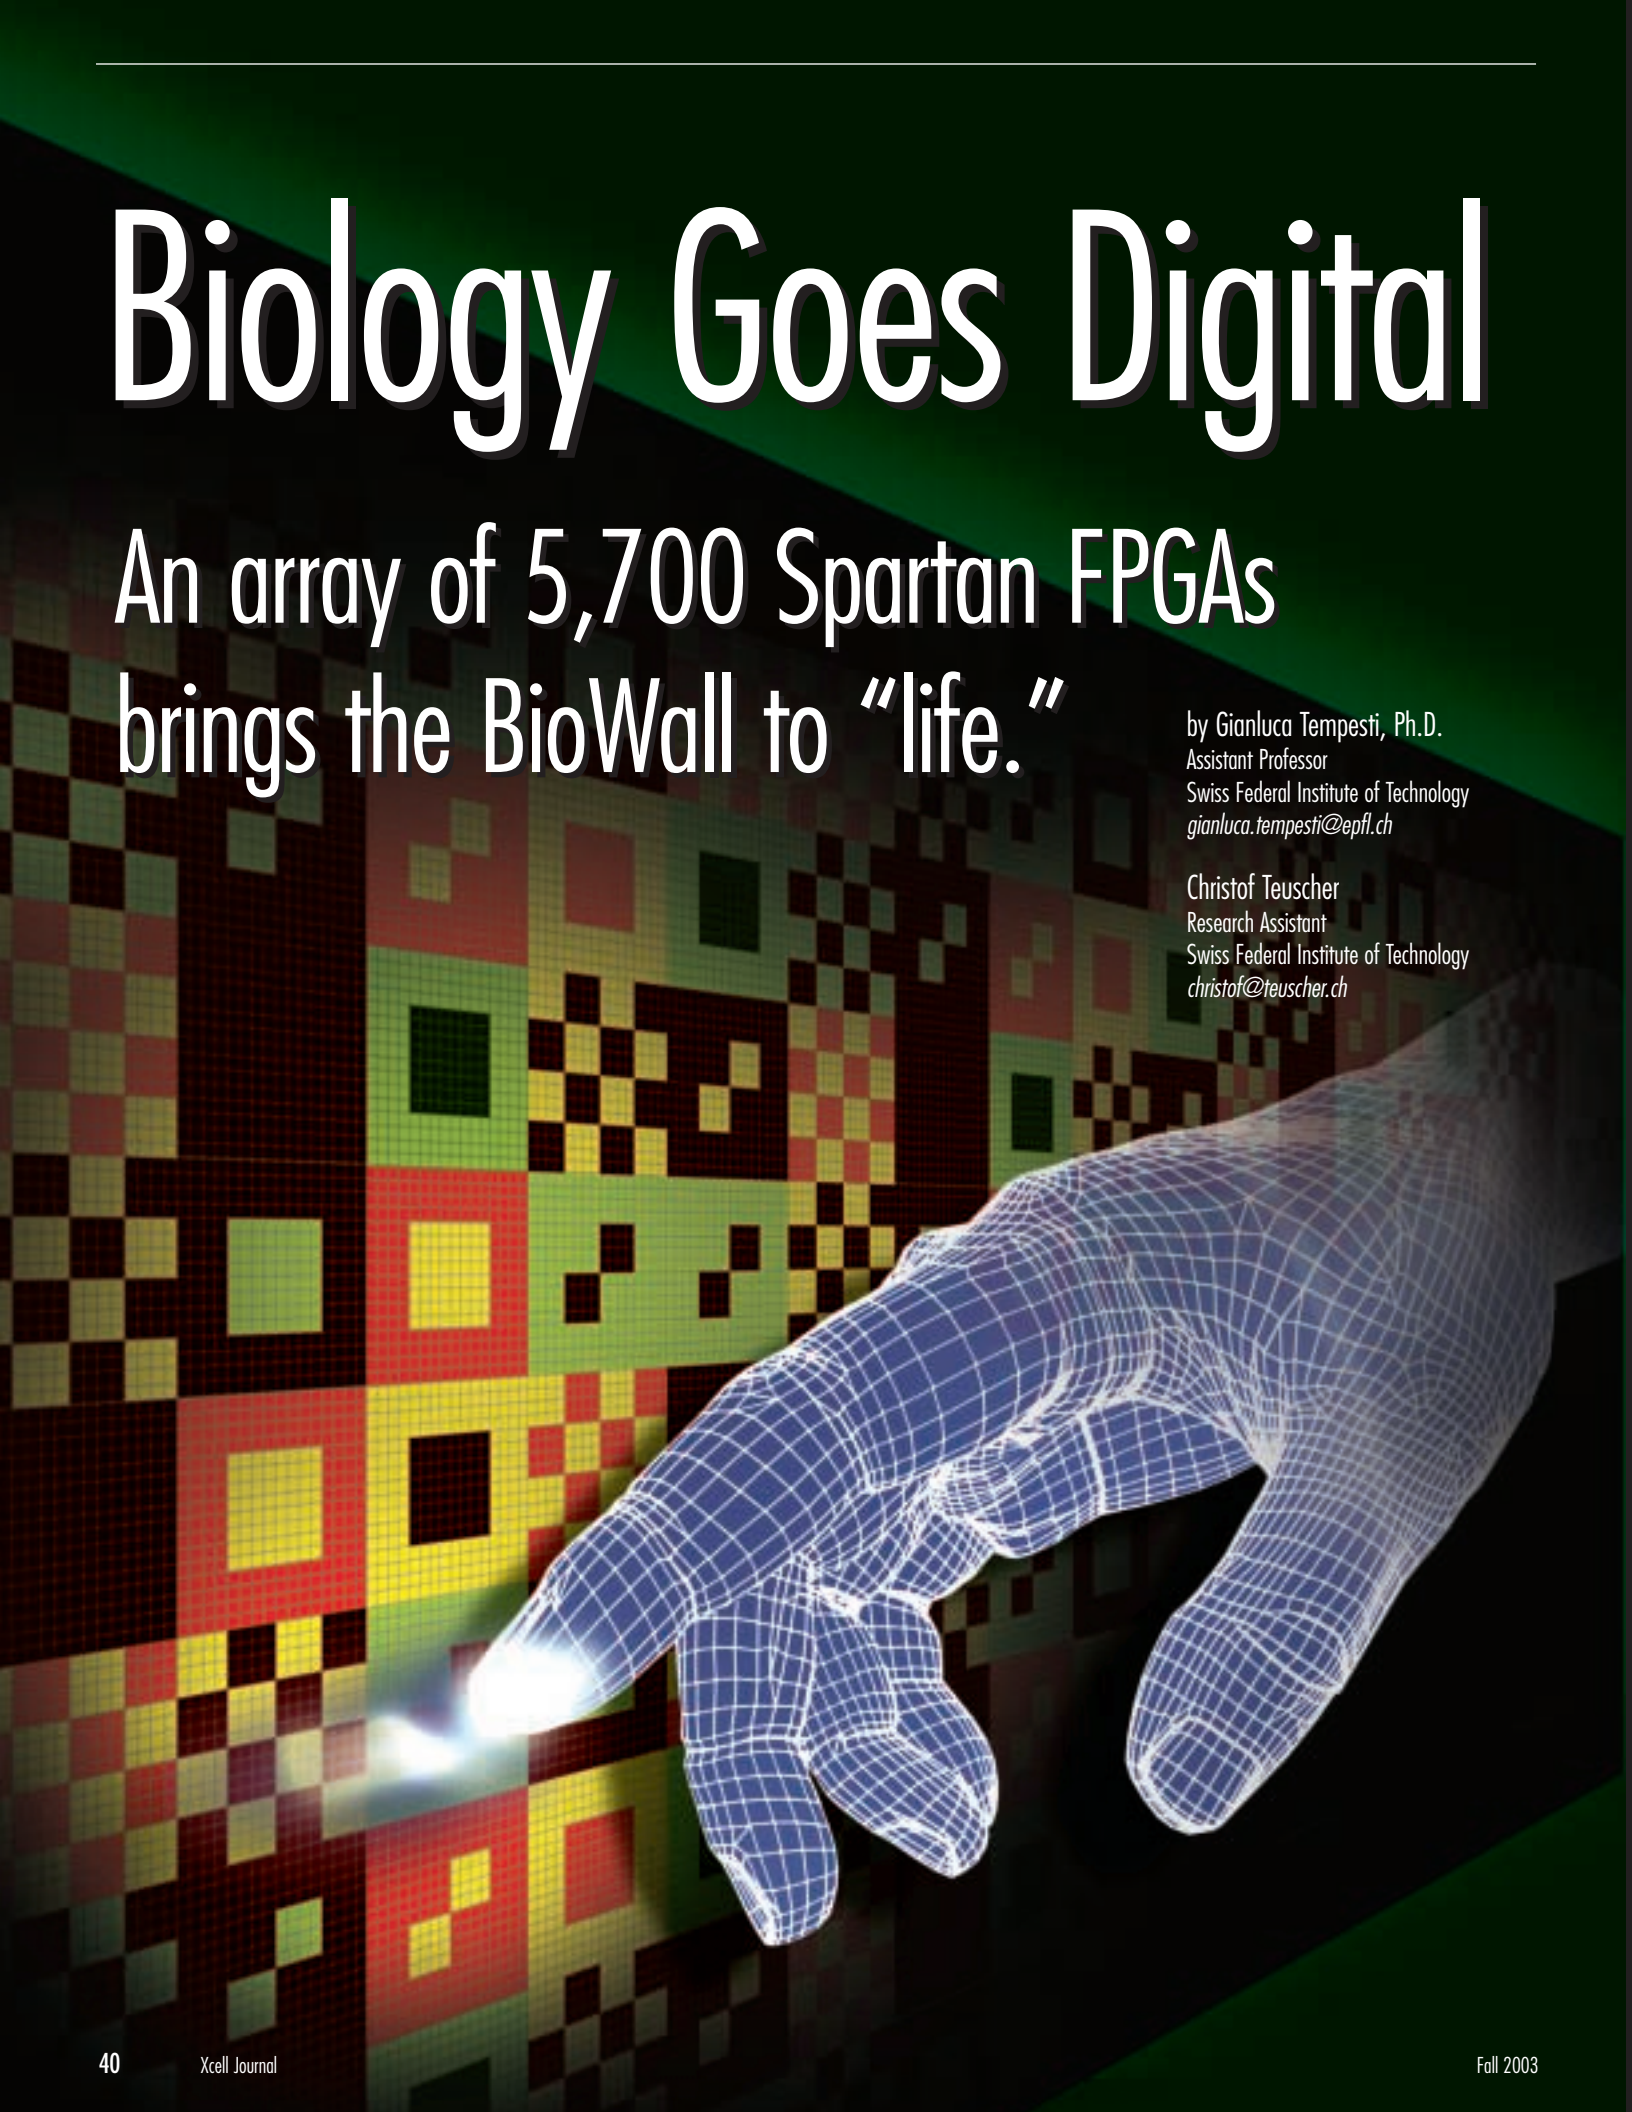

A unique, scientific research instrument and art piece, the BioWall models bio-inspired electronic tissues capable of evolution, self-repair, self-replication – and learning.

Much pure and applied scientific research has focused on replicating biological functions in digital hardware. Here, at the Logic Systems Laboratory of the Swiss Federal Institute of Technology in Lausanne (EPFL), we have utilized 5,700 Xilinx Spartan™ FPGAs, in multiple configurations, to build bio-inspired computing machines that exploit three essential biological models:

- **Phylogenesis** – the history of the evolution of the species
- **Ontogenesis** – the development of an individual as directed by his genetic code
- **Epigenesis** – the development of an individual through learning processes (nervous system, immune system), influenced both by genetic code (the innate) and environment (the acquired).

Although we have individually and jointly investigated all three models, we have concentrated on the ontogenetic model through the Embryonics (embryonic electronics) Project. This project studies the development of multi-cellular organisms for the purpose of obtaining in digital hardware some of the features of biological organisms, notably growth and fault tolerance.

Our work has attracted a flattering amount of interest in the most varied and sometimes unexpected milieus. Among the most unexpected sources of funding and support came from Mrs. Jacqueline Reuge. Mrs. Reuge decided to fund the construction of the BioWall to display the principles of embryonics within a museum built to honor the memory of her late husband.

Her generous support has allowed us to maintain our tradition of verifying our theoretical concepts in hardware. Without Mrs. Reuge's support, we could not have constructed a computing machine of such magnificent proportions.

We named this machine BioWall because of its biological inspiration, as well as its size

Figure 1 – The BioWall reacts to touch.

(5.3m x 0.6m x 0.5m = 3.68m^3, or 130 cubic feet). The main purpose of creating the BioWall is to demonstrate the features of our embryonics systems to the public through a visual and tactile interaction (Figure 1).

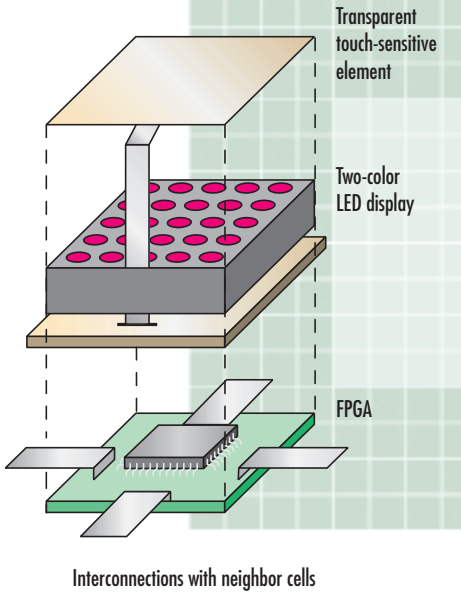

Figure 2 – The BioWall's basic building block

Bio-Inspired Machines

We implemented – for the first time in actual hardware – an organism endowed with all of the features of an embryonics machine, as it has often been defined in the literature. One of the functions of our organism is the BioWatch, which counts hours, minutes, and seconds. It demonstrates the growth and self-repair capabilities of our systems.

The implementation of the BioWatch would have been sufficient to justify the effort that went into the construction of our embryonic BioWall. However, in developing our machine, we quickly realized that the capabilities of such a platform were not limited to a single application. In fact, it is an ideal platform to prototype many different kinds of two-dimensional cellular systems, which are systems comprising arrays of small, locally connected elements.

For example, "cellular automata" (CA) are very common environments in bio-inspired research. The BioWall is ideally suited to the implementation of CAs, but it is by no means limited to it. We have only begun to explore the possibilities of the BioWall as a research tool.

Xilinx Behind the Scenes

We built the BioWall to demonstrate an embryonic machine. The structure of such machines is hierarchical: Organisms (application-specific systems) are created by the parallel operation of a number of cells (small processors). Each cell is implemented as an array of molecules (programmable logic elements).

To implement this kind of machine, the BioWall is structured as a two-dimensional tissue comprising units (each unit corresponds to a molecule). As shown in Figure 2, each unit consists of:

- An input element (a touch-sensitive membrane)
- An output element (an array of 64 two-color LEDs)
- A programmable computing element (a Xilinx Spartan XCS10XL FPGA).

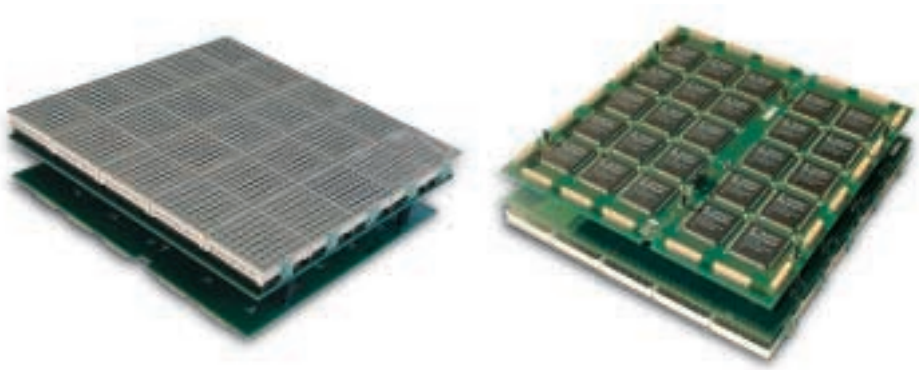

Figure 3 – Front and back view of 25 FPGAs assembled on a BioWall panel

The circuits are mounted on double boards. The logic board hosts 25 FPGAs, and the display board holds the displays and membranes (Figure 3). The two boards are rigidly bound together and connected by a bus to allow two-way communication between the logic and the display (a dedicated circuit on the logic board automatically distributes the signals to and from the displays).

On the logic board, the Spartan devices are placed in a regular two-dimensional grid. A subset of the pins of each FPGA (approximately 20 per side) are used to make a direct pin-to-pin connection between each circuit and its four cardinal neighbors. The pins of the FPGAs placed along the edges of the board are brought to a set of connectors to allow the pin-to-pin association to continue across boards (Figure 4), thus creating perfectly uniform surfaces of FPGAs spanning as many boards as required.

The remaining pins are connected to a centralized circuit that handles the distribution of the global signals (the clocks, resets, and FPGA configurations) arriving from the outside.

We have built 228 such boards (not including spare materials) for a total of 5,700 units. The architecture of the boards implies that they can be seamlessly connected with each other to form a uniform surface of any shape and size. Throughout the development phase and lifetime of the machine, we have so far constructed several independent machines:

- A 3,200-unit machine (Figure 5) displayed at the Villa Reuge museum
- A 2,000-unit machine kept in our laboratory to develop and test new applications

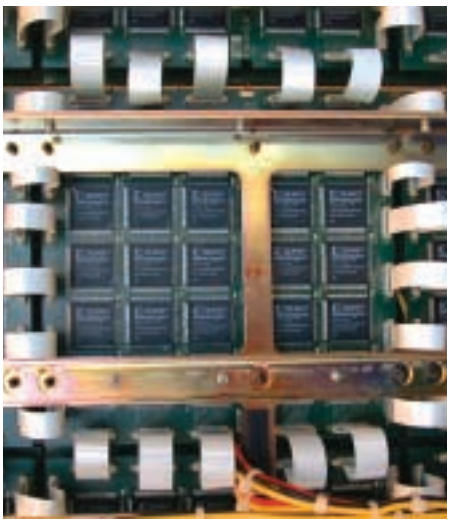

Figure 4 – Connectors allow pin-to-pin association to join multiple boards.

- A 150-unit machine embedded together with the necessary control logic to charge applications into a suitcase for portability
- A 4,000-unit machine that will be on display at the Telecom '03 conference in Geneva in October 2003.

This tissue of 5,700 FPGAs represents an impressive amount of computational power, coupled with I/O interfaces (comprising the membranes and LED arrays) that allow for large-scale visual and tactile interaction. The advantage of this solution is the size of the display, which enables an immediate interaction with applications normally limited to software simulation on a computer screen. Furthermore, the computing power and programmability of the Xilinx FPGAs enables the prototyping of new bio-inspired systems.

In the current version of the BioWall, the Xilinx FPGAs can only be programmed with the same configuration, which limits the functionality of the units to the 10,000 equivalent logic gates of Spartan devices. The considerable delays inherent in propagating a global signal over distances measured in meters seriously limit the clock speed. Considering the role of the BioWall as a demonstration tool, we have not tried to push the clock to its limits, as the current frequency of 1 MHz is more than adequate if coupled with the massive parallelism of the machine.

Besides the I/O capabilities of the membranes and LED displays, a set of modules placed on the borders of the machine allows the tissue to be interfaced with standard logic either via a PC or directly with user-defined modules. The modules allow access only to the borders of the array, but, if necessary, signal propagation logic can be programmed in the FPGAs.

The software tools developed for the BioWall are rudimentary but complete. A simple interface on a PC allows users to define a set of files to configure the tissue. Four kinds of files are currently defined: the configuration file for the Xilinx FPGAs, and three different formats used to send user-defined data on the input pins at the borders of the tissue (used, for example, to provide an initial configuration for a cellular automation). The values on the output pins at the borders of the tissue can be read by the PC and either stored on disk or used as required.

Applications

The cellular structure of the BioWall is well suited to the implementation of all sorts of bio-inspired applications. The BioWall can exploit the versatility inherent in its programmable logic and in its architecture to implement hardware inspired by all the three models of biological inspiration: phylogenesis, ontogenesis, and epigenesis.

BioWatch

To illustrate the implementation of the BioWatch application on the BioWall, we will introduce a slightly simplified example. Whereas the complete BioWatch is an organism capable of counting hours, min-

Figure 5 – The BioWall on display in the Villa Reuge Museum

utes, and seconds, the "counter application" we describe here only counts seconds. Otherwise, the principles of operation of the two machines are identical.

The counter counts seconds, from 00 to 59. From left to right, the display shows tens of seconds (from 0 to 5), units of seconds (0 to 9), and a spare zone, which remains inactive during normal operation (Figure 6a). The counter is divided into four cells: two active (indicating tens and units, respectively) and two spare. Each unit of the BioWall is a molecule of the

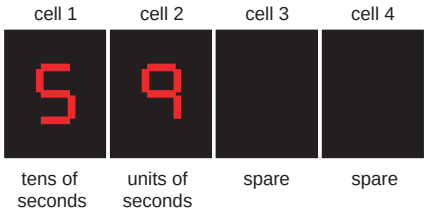

Figure 6a – Counter implemented on BioWall

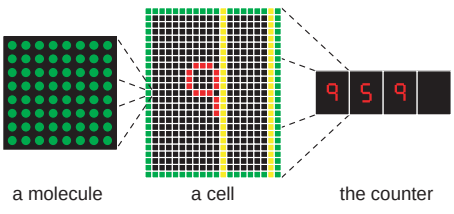

Figure 6b – Hierarchical structure of the counter application

embryonics hierarchy. A cell is then a mosaic of (20 x 25) 500 molecules (Figure 6b), and contains two repair columns for a total of (2 x 25) 50 molecules.

You have control over the "life" of each molecule. A stuck-at-fault can be inserted in any molecule simply by pressing on the corresponding unit's membrane. The fault detection mechanism included in the embryonics molecular layer (embedded into the programmable logic of the Spartan FPGAs) automatically detects the error and activates the molecular self-repair mechanism. A "dead" molecule is instantly replaced by the neighbor immediately to its right, and so on, until the nearest repair column (Figure 7a).

The limits of this kind of self-repair imply that only a single molecule per line, between two repair columns, can be killed. If this constraint is respected, the cell survives any amount of faults, although the figure displayed is distorted. Each cell can thus tolerate up to two faults per line (one fault between each pair of repair columns), equaling (2 x 25) 50 faults in total.

If the above rule is not respected, and several faults are inserted on the same line of the same cell between two repair columns, the molecules can no longer repair themselves and the cell dies. However, the death of a cell does not imply

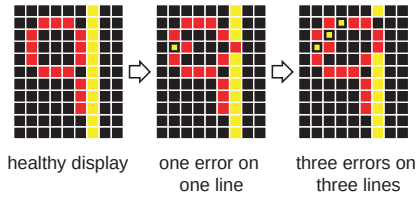

Figure 7a – Molecular self-repair of the counter application

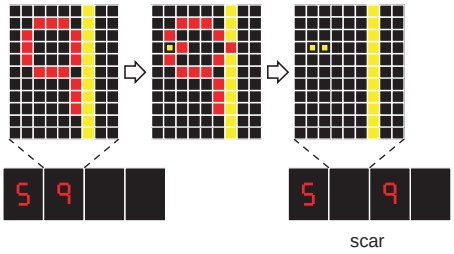

Figure 7b – Cellular self-repair of the counter application

the death of the organism. It is instantly replaced by a spare cell to its right (Figure7b), while the dead cell is switched off and becomes a scar.

Throughout this self-repair process, the counter continues to work without fault. The tissue remembers its state and recovers the correct time after repair. Moreover, we are currently implementing an "unkill" mechanism to address the issue of transient faults. If a sufficient number of faults are removed (by pressing the membranes of dead molecules), this mechanism will automatically re-activate a dead cell, which will recover its functionality (and its state) within the organism.

The self-repair capabilities of the embryonics machines are based on a general principle of life – cell differentiation. Each organism is a collection of cells, each containing a full copy of the genetic program, the genome. This structure makes the whole organism extremely robust, because each cell contains the complete plan and can therefore replace any other defective cell.

Nevertheless, like all artificial and natural organisms, the death of a sufficiently large number of cells cannot be repaired, causing the death of the organism. The advantage of the controlled environment in which the machine operates is that the death of the organism causes a general reset

of the system, the obliteration of all injected faults, and the "birth" of a new, perfectly functioning machine.

The complete implementation of the BioWatch on the BioWall uses eight cells of (20 x 20) 400 molecules each, with two spare columns of molecules in each cell. Six of the eight cells are active during normal operation, while two are spares, ready to replace a dead cell. All the theories of the Embryonics Project have been tested and verified in hardware through this implementation.

Self-Replicating Loops

Initiated by von Neumann more than 50 years ago, the study of self-replicating computing machines has produced a plethora of results. Much of this work is motivated by a desire to understand the fundamental information processing principles and algorithms involved in self-replication independently of their physical manifestation. The construction of artificial self-replicating machines can have diverse applications ranging from nanotechnology to space exploration to reconfigurable computing tissues.

To render the self-replication process more interactive and visible, we implemented self-replicating loops on the BioWall, initially of size 2 x 2 and then of variable size (Figure 8). In this implementation, every unit of the BioWall is one cell of the CA. Pressing on the membrane of a unit belonging to a loop causes the unit to replicate in one of four cardinal directions (Figure 9).

Turing Neural Networks

In 1948, Alan Turing wrote a little-known report entitled "Intelligent Machinery." Turing never had great interest in publicizing his ideas, so the paper went unpublished until 1968, 14 years after his death.

Few people know that Turing's "Intelligent Machinery" paper contains a fascinating investigation of different connectionist models that would today be called neural networks. In describing randomly connected networks of artificial neurons, Turing wrote one of the first manifests of the field of artificial intelligence (although he did not use this term).

Recently, we implemented Turing's neural networks on the BioWall's reconfigurable tissue (Figure 10). Each of the 3,200 units of the machine can be interactively configured by choosing one out of five possible functions:

1. Empty cell
2. Neuron
3. Connection
4. Synapse
5. Input cell.

You can discover and affect the behavior of this "unorganized" machine by opening and closing synapses, "organizing" the machine and modifying the network's inputs. All modifications occur by simply pressing on the touch-sensitive membranes. This application is first and foremost a demonstration of Turing's neural networks on reconfigurable hardware (and to the best of our knowledge, the first one). However, it also exemplifies the fusion of the ontogenetic and epigenetic models in a single artificial tissue.

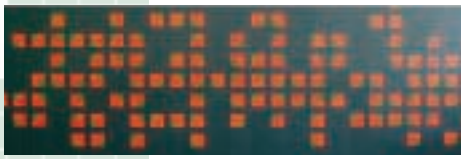

Figure 8 – Self-replicating loops on the BioWall

Figure 9 – Self-replication of a set of loops

Figure 10 – Turing neural networks on the BioWall

Figure 11 – Two-dimensional Firefly

Firefly

In 1997 the Logic Systems Laboratory presented an evolving hardware system called Firefly, based on a cellular programming approach, in which parallel cellular machines evolve to solve computational tasks. The computational task studied – and successfully solved – is known as "synchronization": Given any initial configuration, the nonuniform CA must reach, within M time steps and using only local connections, a final configuration in which all cells oscillate synchronously between all 0s and all 1s on successive time steps.

The novelty of Firefly is that it operates with no reference to an external device, such as a computer that carries out genetic operators. Thus, the Firefly demonstrates online autonomous evolution.

The original Firefly machine was able to find a solution for a one-dimensional CA. Subsequently, we have been able to evolve, on the BioWall's 3,200 FPGAs, a CA that solves the synchronization task in two dimensions.

The implementation on the BioWall (Figure 11) consists of a two-state, nonuniform CA, in which each cell (FPGA) may contain a different rule. The cells' rule tables are encoded as a bit-string, known as the genome. This genome has a length of 25 = 32 bits for our two-dimensional CA (the binary CA has a neighborhood of 5).

Rather than employ a population of evolving CAs, our algorithm evolves a single, nonuniform CA the size of the entire BioWall (one cell of the CA in each unit of the BioWall, or 3,200 cells) whose rules are initialized at random. Initial configurations are then randomly generated and for each configuration the CA is run for M time steps.

Each cell's fitness is accumulated over C initial configurations: a single run's score is 1 if the cell is in the correct state after M + 4 iterations, and 0 otherwise. The local fitness score for the synchronization task is assigned to each cell by considering the last four time steps (M + 1 to M + 4). If the sequence of states over these steps is pre-

cisely 0-1-0-1, the cell's fitness score is 1; otherwise this score is 0.

After every C configuration, the rules are evolved through crossover and mutation. This evolutionary process is performed in a completely local manner; that is, genetic operators are applied only between directly connected cells.

Unlike standard genetic algorithms, where a population of independent problem solutions globally evolves, our approach involves a grid of rules that coevolves locally. The CA implemented on the BioWall performs computations in a completely local manner, each cell having access only to its immediate neighbors' states. In addition, the evolutionary process is also completely local, because the application of genetic operators as well as the fitness assignment occurs locally.

Using the above-described cellular programming approach on the BioWall, we have shown that a nonuniform CA of radius 1 can be evolved to successfully solve the synchronization task. In addition, after having found a set of successful rules, our machine allows the state of each CA cell to be changed by pressing on its membrane. You can then observe how the machine resynchronizes the 3,200 cells.

DNA Sequence Comparison

The comparison and alignment of characters taken from a finite alphabet is a fundamental task in many applications, ranging from full-text search to computational biology. In particular, string comparison is a critical issue in the field of molecular biology.

In fact, both DNA fragments and proteins can be represented as sequences of characters (taken from alphabets of four and 20 symbols, respectively). Sequence similarities provide useful information about the functional, structural, and evolutionary relationships between the corresponding molecules.

Biological sequences may differ because of local substitutions, insertions, and deletions of one or more characters. The complexity of string comparison comes from the large number of possible combinations of these three basic mutations.

The similarity between two strings can

		G	A	A	T	T	C	A	G	T	T	A
	0	0	0	0	0	0	0	0	0	0	0	
G	0	1	1	1	1	1	1	1	1	1		
G	0	1	1	1	1	1	1	1	2			
A	0	1	2	2	2	2	2	2				
T	0	1	2	2	3	3	3					
C	0	1	2	2	3	3						
G	0	1	2	2	3							
A	0	1	2	3								

Figure 12 – The Needleman-Wunsch algorithms in two dimensions

Figure 13 – The Needleman-Wunsch algorithm implemented on the BioWall

be evaluated either in terms of edit distance or similarity score. The edit distance is a measure of the minimum number of edit operations (mutations) required to make the two strings equal to each other.

The similarity score is a measure of the maximum number of residual matches between the two strings. Both metrics can be evaluated in polynomial time by means of dynamic programming techniques.

The key algorithm for evaluating the similarity between two strings of length N and M was developed by Needleman and Wunsch and takes O (N x M) steps to complete execution. The two-dimensional structure of the algorithm (Figure 12) makes it suitable for a parallel implementation on a systolic array.

In particular, hardware parallelism can be exploited to perform string comparison in O (N + M) steps. In this experiment, we present a parallel implementation of the Needleman-Wunsch algorithm on the BioWall (Figure 13).

The BioWall cannot compete in performance with existing parallel implementations of the Needleman-Wunsch algorithm, because it suffers from the typical performance limitations of a large prototyping platform. Nevertheless, the implementation of the Needleman-Wunsch algorithm on the BioWall is a significant design experiment in the field of reconfigurable computing, because of the peculiarities of the target architecture.

Conclusion

The current configuration of the BioWall is a mosaic of more than 3,000 transparent electronic modules. Each module enables visitors to communicate with the surface simply by touching it with their fingers. The BioWall calculates its new status and indicates it immediately on an electronic display. The usefulness of this approach has been demonstrated through a number of experiments.

In fact, the applications presented here are just a small sample of the capabilities of the BioWall – capabilities that we are still discovering. The cellular structure of the machine makes it an ideal platform for prototyping bio-inspired systems.

The size and structure of the BioWall impose a certain number of limitations, such as clock speed. But its complete programmability provides outstanding versatility, and the visual and interactive components of the system are invaluable tools both for the dissemination of ideas as well as the verification of research concepts often limited to software simulations.

Some of the other bio-inspired applications we have implemented or plan to implement on our machine include L-systems, ant simulations, predator-prey environments, other kinds of CAs, and more conventional artificial neural networks.

We invite you to come and "play" with the machine at one of the events at which it will be on public display, or even at our laboratory.

Finally, we are extremely interested in putting our machine at the disposal of other research groups interested in hardware realizations of their ideas and concepts. For more information, please visit us at *lslwww.epfl.ch/biowall/*.

[Daniel Mange, André Stauffer, Fabien Vannel, André Badertscher, and Enrico Petraglio also contributed to this article.]

Photos: © André Badertscher, Alain Herzog, EPFL.

Formula 1 Racing: The Xilinx Advantage

The BMW WilliamsF1 team lifts the hood on its fifth-generation vehicle control system.

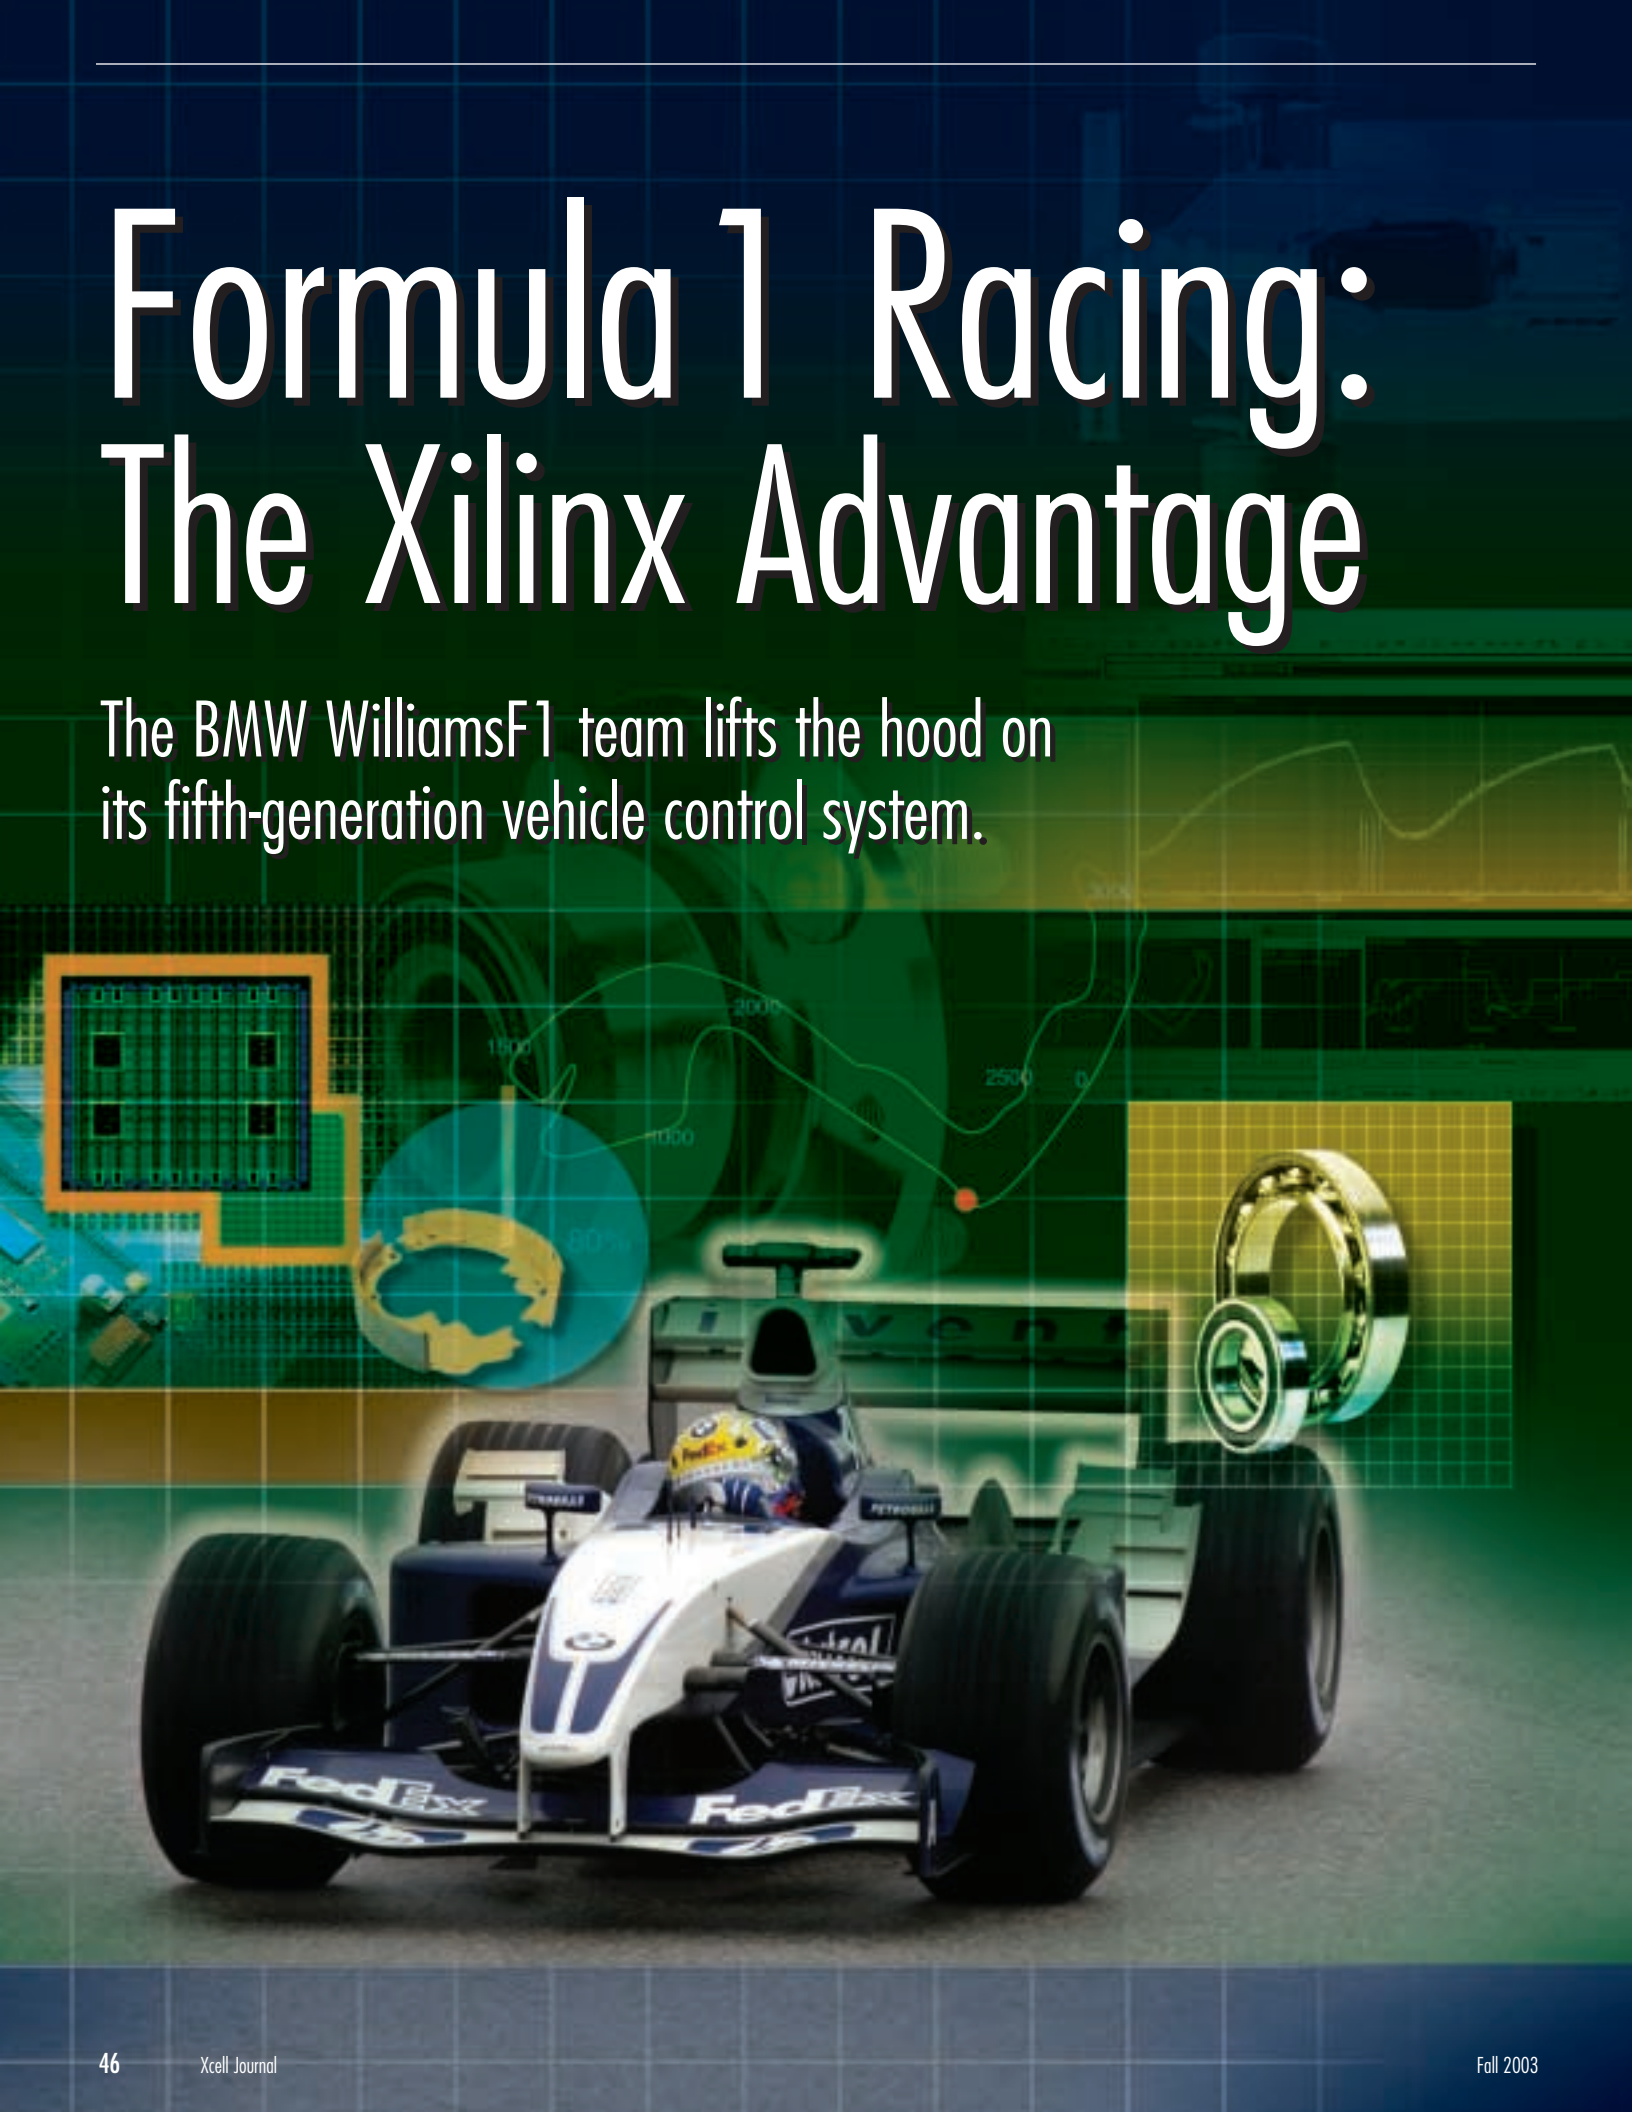

by Liza Boland
System Architect, Consumer & Automotive Technologies
Xilinx, Inc.
liza.boland@xilinx.com

A modern Formula 1 (F1) car communicates an enormous amount of data to the engineering teams that support it. How the teams collect, analyze, and react to that data is vital to competitiveness. The collect-analyze-and-respond process is also subject to change, as the sport's rule-makers aim to reduce onboard electronics to control costs and achieve closer racing competition.

Throughout the last decade, the electronics department at BMW WilliamsF1, led by Steve Wise along with engineers Dave Walker and Ian Cartwright, can claim to have been in the driver's seat as onboard electronics have spread to every extremity of a Formula 1 chassis.

The team is now running its fifth-generation VCM (vehicle control and monitoring) unit; the current Stage 5 VCM uses a Xilinx Virtex™-E XCV600-E FPGA alongside a Texas Instruments™ DSP. Shown in Figure 1, the VCM controls all aspects of the 2003 FW25's chassis (other than the engine). And the team is continuing to extend and refine the unit's capabilities.

Figure 1 – The VCM controls all of the race car's chassis functions except the engine.

Adapting to Rule Changes

Rules change quickly in F1 racing. Teams must be able to respond with a competitive solution that meets every new detail. The regulatory landscape has been especially volatile over the last decade. Driver aids such as active suspension, anti-lock braking, traction control, and launch control all have been banned at one time. Some driver-assist features have been granted a reprieve for this year – only to be completely outlawed for the 2004 season.

The ability to quickly add and remove functional blocks is therefore critical to each team's ability to compete and to demonstrate compliance.

Successive generations of the BMW WilliamsF1 VCM have made great use of FPGAs to achieve rapid compliance. This flexibility has allowed BMW WilliamsF1 to migrate more and more sophisticated functions into hardware as new generations of FPGA silicon offer extra capacity and complexity.

Controlling an F1 Car

The VCM is a data logger and processor on a grand scale. Its vehicle control functions include overseeing the hydraulically actuated gear change, which calls for the VCM to initiate a gear shift either in response to a driver request or automatically.

Shifting Gears

The gear change sequence in a F1 gearbox requires precise control over the positions of the gearbox actuators, as well as controlling the clutch and coordinating engine revolutions per second with the BMW engine controller. A typical gear change takes less than 50 ms. The system also prevents the driver from damaging or over-revving the engine.

For the 2003 racing season, the VCM is allowed to initiate all gear changes with no input from the driver. But automatic gear changes will be outlawed for 2004, requiring the driver to initiate each gear change manually. Once the driver has initiated the change, the VCM will be allowed to handle the rest of the gear change sequence.

Traction

The VCM also handles traction control by performing calculations based on complex tire models to predict the amount of wheel slip required to achieve maximum traction and minimum tire wear. The VCM calculates control targets and sends this data to the engine controller.

Generating control signals for the hydraulically actuated differential also comes under the VCM's domain – with the aim of providing maximum traction from each of the rear wheels to optimize stability of the car when cornering.

Launch Control

The use of "launch control" – which optimizes the car's standing start as it leaves the grid – has received much press coverage in recent seasons. Incoming rules will outlaw this feature, but for the 2003 season, the launch start controller is a functional block within the VCM. Clutch and engine targets communicated by the VCM help the car to accelerate from the grid to 100 m/hr in less than three seconds.

Real-Time Feedback

The VCM displays driver information signals and warning indicators on the steering wheel. Even though a modern F1 steering wheel can cost upwards of $33,000, there is relatively little intelligence onboard – that is, until next year's BMW WilliamsF1 team adds a Xilinx FPGA to perform some of the instrumentation processing functions locally.

Logging and Telemetry

The VCM logs around 220 channels of data at rates as high as 1 KHz. As many as 90 of these channels are signals from sensors used to analyze the car's performance over an entire race or test distance. Some channels monitor the actions of the control software, while others monitor the driver's inputs. An example of the logged data is shown in Figure 2.

Each of the Stage 5 VCMs fitted to the two FW25 cars of drivers Juan Montoya and Ralf Schumacher can log 256 MB of data generated by the cars during a race. This data is stored on a standard, commercial CompactFlash™ (CF) card mounted permanently inside the VCM, which provides convenient and cost-effective non-volatile storage. The CF card contents are downloaded by wire link after the race for analysis.

A subsection of the logged data is also transmitted via serial link to a telemetry transmitter on the car. This transmits the data in real time back to the garage where it is displayed on a PC, enabling technicians to monitor a car's performance from the pit.

Failure Analysis and Reporting

Embedded in the VCM are functions to detect failures of sensors on board the car. The output of these routines prevents failed sensors from being used by the control algorithms, and allows problems to be identified quickly.

The failure of a small sensor cannot be allowed to ruin a driver's race. To prevent this, the VCM monitors inputs from multiple sensors and acts on the majority decision in each case.

The level of redundancy is determined according to the importance of the function, the vulnerability of the sensor, and the potential weight penalty. The VCM is capable of monitoring as many as 100 input sensors, while controlling numerous hydraulic systems.

This failure detection capability is augmented by pit-to-car telemetry. The race engineer can arbitrate against a sensor or set of sensors believed to be providing inaccurate signals.

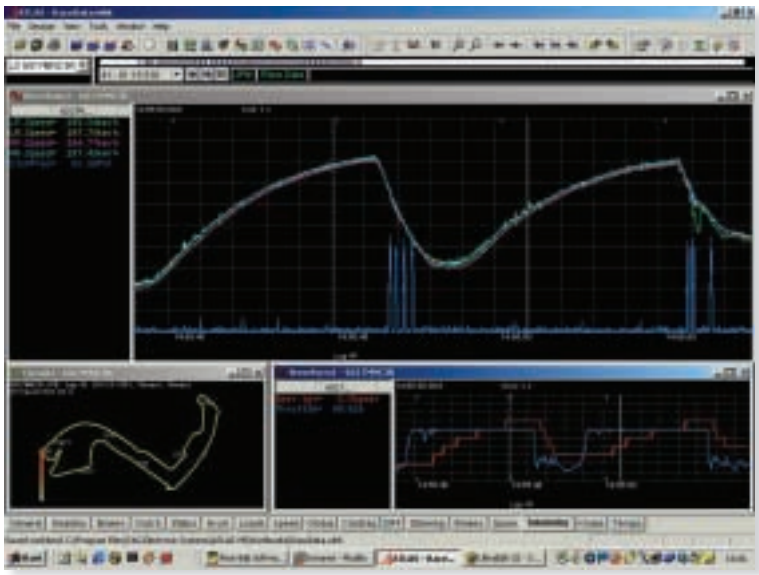

Figure 2 – Data from more than 90 sensors monitor the car's performance.

Extra Processing Horsepower

BMW's engine designers would say that a racing car can never have too much horsepower. BMW WilliamsF1's electronic designers have a similar opinion about processing capability: More MIPS (million instructions per second) allow them to consolidate functions within the VCM, control more aspects of the car, and use more sophisticated models to generate accurate control signals.

Despite the rule-makers' attempts to simplify and reduce the cost of F1 technology, the processing demands on the VCM continue on a steeply upward trend.

Paradoxically, rule changes for 2004 prohibiting pit-to-car telemetry, which are intended to reduce the complexity of electronics, will drive the processing load higher in future years. The decision-making performed by human operators in the garage must now migrate to the car itself.

Dave Walker and his team are faced with an ever-increasing demand for processing power. But they do not enjoy a corresponding increase in thermal budget or physical size limits. The growing list of duties for the VCM, originally used only as a pure data logger, has increased processing requirements.

Hardware Solution

With 220 channels to monitor, as well as numerous control algorithms to execute, Dave Walker says processor cycles are far too precious to waste on basic input functions, such as filtering or controlling A-to-D conversions.

Among the potential solutions to reducing processor load, the BMW WilliamsF1 team did not favor a microcontroller because of the need to write software and maintain the code base. In addition, the software inspection demands imposed by the FIA make the use of a microcontroller even less desirable.

A pure hardware solution is more optimal, according to Walker. Before the advent of modern multi-MFLOP (mega floating-point operations per second) DSPs, it was only possible to meet the timing demands for data logging and controlling high-speed hydraulic actuators by using hardware to accelerate certain functions.

The team has used field programmable gate arrays from a variety of vendors since the first-generation design, but BMW WilliamsF1 has now been working exclusively with Xilinx for a number of years. The current Stage 5 VCM uses the Xilinx Virtex™-E XCV600-E FPGA alongside a Texas Instruments™ DSP.

Although processing capabilities of DSPs have advanced dramatically since the first versions of VCM entered service, the team finds the hardware processing capabilities of the companion FPGA a valuable resource, which allows it to focus on using the DSP for pure number crunching.

There will never be a surfeit of processing power, as this can always be consumed by the increasingly sophisticated models provided to the VCM. These models demand more accurate control outputs for gearbox, differential management, and communications with the engine. Today's XCV600-E is around 85% utilized. The

team plans to migrate to a more complex, next-generation Virtex-II Pro™ FPGA in the VCM incarnation for 2004.

The team also plans to use Xilinx Titanium and Platinum technical support packages, which will include on-site design and training by expert Xilinx engineers.

Accordingly, the sixth-generation BMW WilliamsF1 VCM will perform even more processing, with a more powerful DSP chip alongside a more complex Xilinx FPGA. The Virtex-II Pro™ platform FPGA will also be able to offload additional DSP functions by virtue of its onboard multiply-and-accumulate (MAC) block to deliver a greater boost in effective system MIPS.

For 2004, the Virtex-II Pro device will host new programmable filtering and communication protocol control functions. The platform FPGAs will also manage UARTs and PWMs (pulse-width modulators) that were formerly implemented in separate components on the printed circuit board.

This ability to progressively soak extra discrete logic into the FPGA – at the same time as assimilating software routines – has enabled the VCM's footprint to shrink significantly over its lifetime.

Conclusion

For those determined to win, Formula 1 has always demanded the utmost speed and nimbleness – not only on the racecourse, but also in terms of technical innovation. Rule changes introduced at the start of 2003 have taxed all teams' abilities to respond quickly and effectively.

At the vehicle control level, engineers must to be able to add and remove functional blocks from the VCM within a short timeframe to meet FIA regulations. The teams must also be able to incorporate performance modifications as the current season progresses.

Fast feature swapping – combined with the ability to consolidate more of the low-level hardware and discrete logic into an FPGA to save weight, space, and power dissipation – is crucial to BMW WilliamsF1's ability to demonstrate compliance with FIA rules. Moreover, Xilinx FPGA reprogramability enables the team to innovate quickly to maintain a competitive edge over rival teams.

To learn more about the Xilinx partnership with the BMW WilliamsF1 team, go to *www.bmw.williamsf1.com* and click on "Xilinx Joins the Team" under Partner News. To get the latest Formula 1 standings, go to *www.formula1.com*.

As of this writing (July 8), the BMW WilliamsF1 team had back-to-back wins, taking first and second place at both the European Grand Prix at Nurburgring and the French Grand Prix at Magny-Cours. Earlier this year, the team posted its first win of the season in the Monte-Carlo Grand Prix at Monaco.

According to a July 7 Formula 1 press release: "Williams has suddenly vaulted past both McLaren and Ferrari to establish itself as the team most in form ... If proof were needed of the growing Williams menace, in the last five races McLaren, initially the 2003 series leader, has scored 34 points; Ferrari has scored 55; and Williams has scored 70."

In overall standings, BMW WilliamsF1 was trailing the series leader Ferrari by just three points.

CoolRunner-II Has Hidden Talents

Girsberger Elektronik AG's avalanche rescue training transmitters make efficient use of board space and power, thanks to Xilinx CoolRunner-II CPLDs.

by Felix Meier
Senior Engineer
Girsberger Elektronik AG
felix.meier@smile.ch

Every year in Switzerland, approximately 27 winter sports enthusiasts lose their lives in avalanches. If caught in an avalanche, your life depends upon the speed of your companions as rescuers to locate you and dig you out. After only 30 minutes, your chance of survival is just 40%. Speed is of the essence.

One rescue tool is an avalanche beacon, a small portable unit that transmits a beep about every second. Using one significantly increases the chances of being quickly and successfully rescued if buried by an avalanche. Your companions need only to switch their units to "receive" mode and locate your beeps.

As with any rescue and survival skill, the correct use of avalanche beacons requires education and practice. However, for use in classroom situations, the signals from regular avalanche beacons are too strong to allow for a meaningful demonstration of the basic search principles.

At Girsberger Elektronik AG, we developed a special reduced-strength transmitter for search training (Figure 1). It meets our key requirements: low cost, battery efficient, and compatible with the various timing signals (that is, the interval and duration of beeps) of all beacon

brands on the market. We use a 73.768 KHz watch oscillator followed by a Xilinx CoolRunner™-II device to generate the timing signals (Figure 2). This provides an optimal solution, saving on both power consumption and PCB real estate.

Figure 1 – Training transmitter

I/O Port Expanders

The best practical training application for beacons is, of course, to bury the transmitters into mountain snow. However, changing the location of the transmitters to build new search scenarios cost valuable training time. By designing a control console (Figure 3) for as many as 16 remote transmitters activated via radio link, we enabled trainers to quickly and easily simulate new accident situations.

In order to handle 16 remote transmitters we required more I/O lines than what is available on a microprocessor. We chose to use two Xilinx CoolRunner-II devices, which provide a compact solution with low console battery consumption.

PCM Codec to RAM Interface

We also found CoolRunner-II CPLDs effective in reducing transmission loss. To initiate transmission over a private mobile radio, you must press a PTT (push to talk) button and wait before the transmitter is ready for you to start speaking. If you're under stress, there's a strong tendency to start talking before the transmitter is ready, and the beginning of your transmission will be lost. You may have to repeat the whole transmission in order for the receiver to make sense of it. This is inconvenient and, especially in emergency situations, can cause additional stress and frustration.

Figure 2 – Internal view of training transmitter

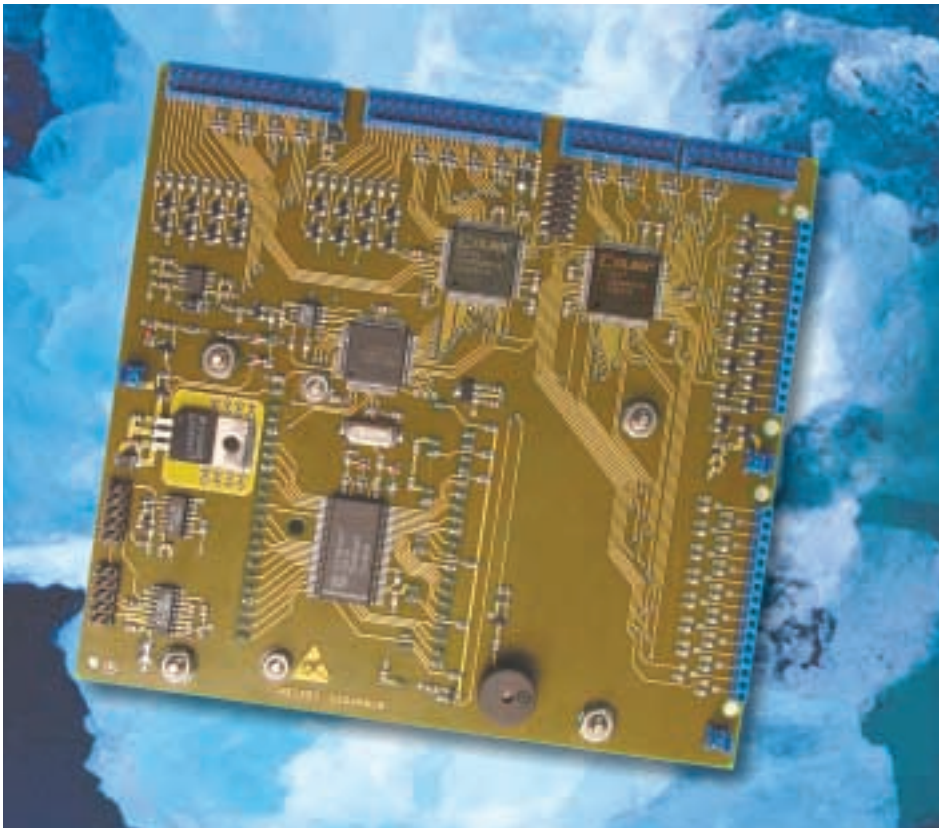

Figure 3 – Transmitter control console

Figure 4 – Digital audio delay line

In some cases the PTT button is replaced by a voice activity detector, which has an inherent turn-on delay, meaning that the beginning of your transmission will always be lost.

To eliminate loss at the beginning of transmissions, we combined a PCM codec with a static RAM to implement a digital audio delay line, which holds up the voice signal (Figure 4).

The PCM data transfers in to and out of the codec over a high-speed 2.048 MHz serial bus (Figure 5). Encoded PCM samples are converted to parallel data and written into a ring buffer in the RAM. By reading parallel data from the RAM location with a fixed offset (modulo the size of the ring buffer) to the write address, the output is delayed. The parallel data is then reconverted and fed into the codec over a serial bus. One PCM sample reads/writes every 125 microseconds. The delay can be adjusted by modifying the read-to-write address offset.

These operations, in particular the emulation of the serial bus, cannot be performed by a standard microprocessor, which would be too slow. We chose a Xilinx CoolRunner-II device with 64 macrocells for the interface between the codec and the RAM. This solves our speed problem and also saves PCB real estate that would have otherwise been taken by standard logic components.

The Development Process

Excellent support from Impact Memec, the local Xilinx representative, kept the development start-up time extremely short. We used ISE WebPACK™ software for all of the designs, and found it especially helpful for the following reasons:

- Access to error messages on the Xilinx website provided instant in-depth information about the particular message, often with hints on how to solve the problem.
- The designs were simulated using ModelSim XE Starter, leading to some design modifications and allowing flaws to be fixed early in the development process. Some problems were due to signal spikes on clock lines, which were properly indicated by the simulator.

All designs loaded and operated correctly on the first attempt when moved to the target hardware. We have no doubt that the design environment of the Xilinx CoolRunner-II family of CPLDs is mature and sound.

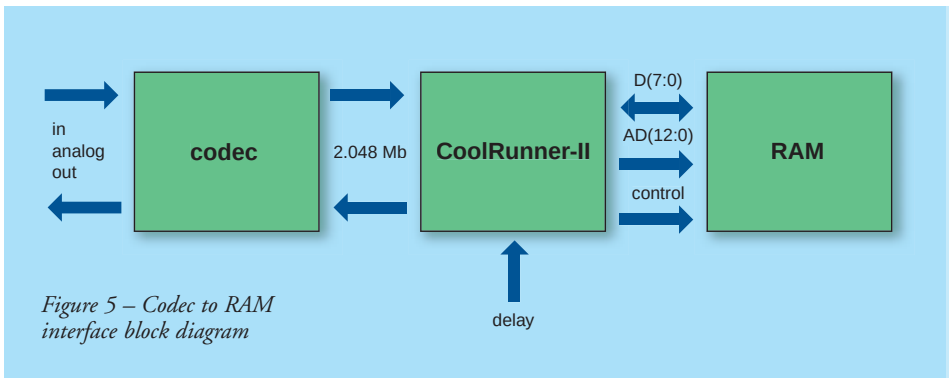

Figure 5 – Codec to RAM interface block diagram

Conclusion

Xilinx CoolRunner-II devices are well suited for our applications. They provide a reliable technical solution combined with a short development time and a low start-up cost.

Based on this positive experience, Girsberger Elektronik AG is planning to incorporate Xilinx CoolRunner-II devices in such products as:

- The upgrade and replacement of an old avalanche beacon ASIC to eliminate the high NRE cost.
- A beacon to locate stolen cars. By taking over part of the training transmitter CPLD design, we will save significant development time.
- A data over voice interface for use on ISDN B-channels. The CPLD beats the speed of any microprocessor in manipulating the bitstream on an industry-standard IOM bus.

Although CPLDs are often associated with high speed or high density, Xilinx CoolRunner-II devices are an interesting alternative for small-scale applications in which top speed and high density are not key concerns. For more information about Girsberger Elektronik AG, visit *www.girsberger-elektronik.ch/*.

Boeing Harnesses Virtex-II Pro Platform for Image Processing

Designing a state-of-the-art airborne imaging system requires massive parallel processing, high-speed I/Os, and heavy-duty signal processing. Boeing's design team found all three plus configurability in the Virtex-II Pro Platform FPGA.

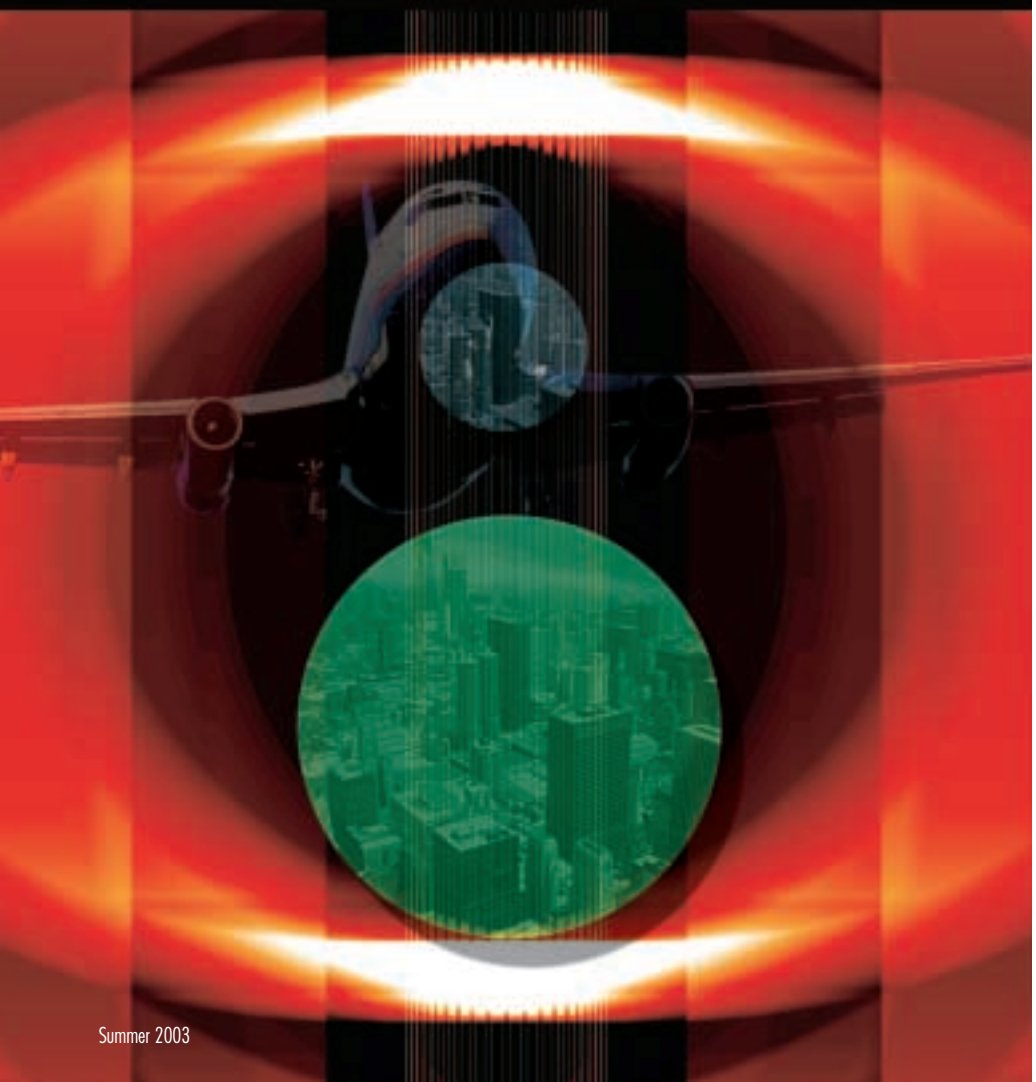

by Christopher Musial
Real Time Systems Engineer
The Boeing Co.
christopher.j.musial@boeing.com

Katie DaCosta
Solutions Marketing Manager, DSP Solutions
Xilinx, Inc.
katie.dacosta@xilinx.com

Over the past 10 years, the aerospace industry has adopted the use of commercial off-the-shelf (COTS) electronics in favor of developing custom systems. This change has been fueled in part by the high development costs associated with ASICs or specialized software processors.

As a result, commercial suppliers seeking the best available technology have focused on the design and fabrication of multi-processor hardware. Today, many of these systems feature DSP devices or RISC-based CPUs. Although Boeing has successfully integrated these solutions into many systems, the application of COTS to high-end signal- and image-processing designs has remained very challenging.

Even "high-end" DSP and RISC architectures offer limited parallelism, memory flexibility, interface capability, and determinism. The integration of dozens of processors in a hard real-time architecture is non-trivial. The complications of packaging, weight, volume, thermal management, and power consumption to system design are obvious. The challenges often appear to grow exponentially with the number of processors required to solve the problem. The advent of Virtex-II Pro™ Platform FPGAs and the promise of more advanced, future Platform FPGAs are completely changing the game of programmable system design.

Airborne Imaging

Boeing is designing a state-of-the-art airborne, electro-optical imaging system. The end product will provide exceptional surveillance and engagement capability for aircraft in the 21st century. To achieve these objectives, Boeing designers are integrating a broad array of technologies into a complete programmable system. Key design elements include:

- A high-resolution gimbaled telescope is driven by line-of-sight control electron-

ics. This system implements control algorithms that accept operator inputs to point the telescope to ground-based objects of interest.

- Precision optics direct the light collected by the telescope into a series of imaging cameras. Aircraft engine vibration, structural resonance, and other disturbances can wreak havoc on the alignment of these devices. To compensate for these effects, an automatic alignment system is incorporated. The goal is to stabilize the imaging system so it can yield a clear, jitter-free image for the system operators.
- Sensitive imaging cameras act as the "eyes" of the electro-optical payload. These devices provide high-resolution daytime and nighttime video to real-time image processing hardware.
- Sophisticated image-processing algorithms that perform video enhancement, segmentation, and feature extraction.

Together, these subsystems allow an operator to identify and interrogate ground-based objects with unprecedented capability. Figure 1 highlights some of the essential functions of this system.

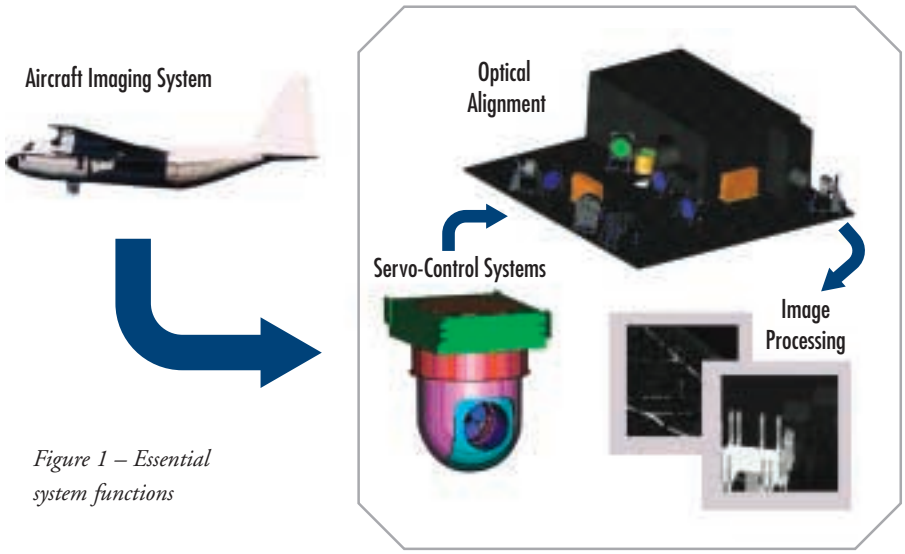

Figure 1 – Essential system functions

Boeing Subsystems

Boeing has chosen the Virtex-II Pro Platform as a critical element in this programmable system. The FPGA device and the supporting chips connected to it execute the following functions:

- High-speed image processing
- Digital-video scene generation and storage
- Fiber-optic media conversion of video data
- DVI display generation
- NTSC video generation
- Servo-control functions
- High-bandwidth data transfer interfaces.

Embedded DSP

One reason Boeing chose the Virtex-II Pro platform is that it incorporates three of the most important elements for a modern, embedded DSP device.

1. **A high-performance, high-density, programmable logic fabric** – this capability is the heart and soul of true parallel processing. FPGA technology lends itself very well to the repetitive, systolic nature of our algorithms. Embedded multipliers, SRL16 capabilities, true dual-port Block RAM, and a segmented routing architecture give us key enabling technologies to produce powerful designs.
2. **A wealth of interface pins that can be configured for different signaling standards and easily integrated with devices from other vendors** – using a single Xilinx chip, Boeing's design team can arrange a device with a PCI bus interface, video encoder/decoder connections, digital-video display outputs, multiple SRAM interfaces, and a high-speed digital-video input port. These functions run in parallel, concurrently, without the hassles of bus contention and multi-processor scatter/gather issues.

   The Virtex-II Pro multi-gigabit transceivers (MGTs) raise the bar to another level. MGTs are extremely useful for the imaging camera interfaces. Boeing's design team can use a low-end Virtex-II Pro device to serialize several parallel digital video outputs. These serial outputs are easily converted to fiber-optic media. And because the data transfer is point-to-point, designers can program the serial transceivers with a user-friendly data protocol.
3. **An embedded IBM PowerPC™ 405 CPU** – the point of using an FPGA in many image-processing applications is to unburden the host CPU from simple, brute-force calculations. With the embedded PowerPC CPU, the host is left unimpeded to handle the complex decision-making logic that forms the brains of a system.

PowerPC Advantages

Many people might argue that it's easy enough to hook an FPGA to a PCI bus interface – so why bother with the embedded PowerPC CPU in the Virtex-II Pro platform? For some applications, a PCI interface is sufficient, but Boeing's design team has found many compelling reasons to choose an embedded CPU:

- By using logic to accelerate the embedded CPU, most of the residual back-end processing for our applications requires only a modest amount of compute power. We all like having multi-GHz CPU power, but it is not always necessary.
- The embedded PowerPC CPU consumes a fraction of the power of a high-end CPU, while remaining directly coupled to the logic fabric. In Boeing's business, it seems that Moore's law does not always apply. RISC CPUs for embedded applications have not generally kept pace with desktop systems in terms of clock speed and performance. Packaging and thermal dissipation con-

straints make it very hard for embedded hardware developers to offer multi-GHz-class processor subsystems in a small form factor. DSP processors address many power consumption issues, but lack the logic interface flexibility. For many applications, the embedded PowerPC provides the right balance of performance, power, interface, and general-purpose computing capability.

- Those who think PCI is easy should try using the CoreConnect™ bus with dual-port RAMs. In a matter of minutes, you can configure an FPGA interface to a 405 processor via high-speed BRAM. It's like having your data put right into the cache RAM. Now consider that multiple, independent FPGA processes can stuff these RAMs in parallel – and with completely different clocks than those used by the PowerPC CPU. You get high-bandwidth, extremely simple data transfer without the vices or overhead of PCI. Put simply, there is far more room for creative interfaces, application-specific tailoring, and flexibility in the Virtex-II Pro platform.
- Consider your application model for your microprocessor. Boeing's design team gets plenty of mileage out of using the embedded CPU as a "microcontroller." If you dispose of a full-blown operating system and run the PowerPC CPU with a lightweight kernel, the internal Block RAM can hold a significant amount of user code. Even better, the FPGA is now both hardware and software reconfigurable – in real time. A single unified bitstream defines the operating characteristics for both the CPU and logic.

Major Features

In addition to the device's embedded DSP capability, Boeing's design team makes use of almost every major feature of the Virtex-II Pro logic fabric:

- The embedded multipliers are very useful for DSP-based algorithms. They provide high silicon efficiency and effortlessly support the clock rates our designs require.

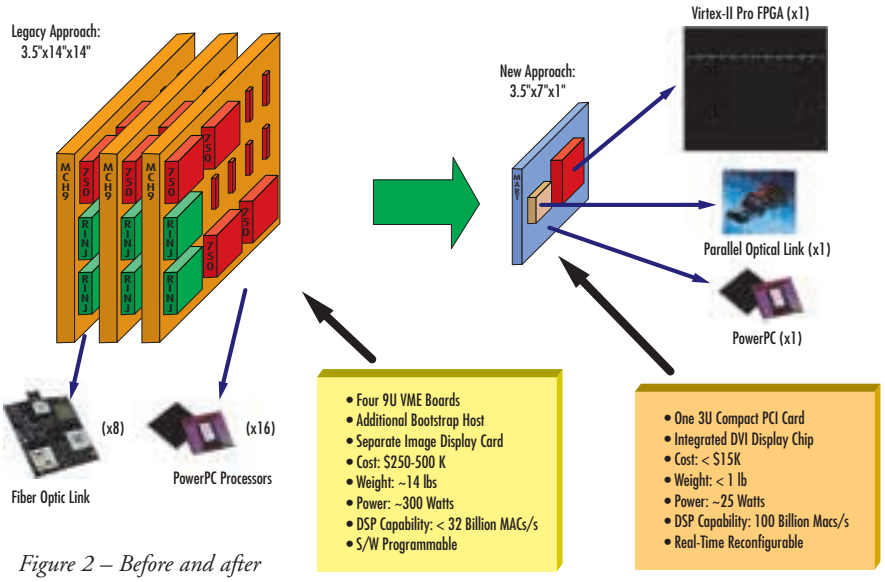

Figure 2 – Before and after Virtex-II Pro implementation

- The dual-port Block RAM is a veritable Swiss Army knife in terms of functionality. It can be used for tapped delay lines, command and status data buffers, lookup tables, bus-width conversion, FIFOs, and so forth.
- The SRL16 capability with dedicated fast-carry logic gives the Virtex-II Pro platform a key advantage in DSP applications. Both features are critical to the design of high-performance digital filters and the SRL16 mode helps to significantly reduce the logic footprint of DSP algorithms.
- DCI (digitally controlled impedance) technology is a significant aid when implementing functions such as DDR RAM interfaces. Boeing's PCB designer no longer struggles with finding a way to place hundreds of termination resistors.
- The DCMs (digital clock managers) provide an easy-to-use, flexible clock management scheme. Our design team has leveraged both the frequency synthesis and precision phase shifting capabilities in many designs. The new ISE architecture wizards make instantiating them easy.
- The MGTs allow us to create high-bandwidth interfaces to a high-speed serial link that operates as if it were a parallel data bus. Ninety-five percent of the design problems in this application are simply about sending pixel data from A to B. Rapid IO, 3GIO, InfiniBand, or Fibre Channel are generally overkill for this type of data transfer requirement. The MGTs provide serial interfaces in a way that has reduced board size tremendously. The 8b/10b and SerDes (serializer/deserializer) chips that we used in a prior design were actually bigger than an entire Virtex-II Pro device.

Figure 2 shows how the Virtex-II Pro platform enabled Boeing's design team to dramatically reduce system size and cost.

Boeing's high-performance imaging system is taking advantage of the complete Xilinx solution. We are using most of the features in the silicon as well as Xilinx handcrafted IP cores, System Generator for DSP, and CORE Generator™ tools. These tools enable us to focus on the specific elements of the design and spend less time constructing the architecture that surrounds it.

Conclusion

Virtex-II Pro Platform FPGAs raise the bar of both capability and complexity for programmable system design. Boeing's design team chose Xilinx solutions because the Virtex-II Pro Platform FPGA offers unmatched "system" capability along with a broad range of devices from which to choose. For more information on Virtex-II Pro Platform FPGAs and design resources, go to *www.xilinx.com/virtex2pro*.

Support Acr

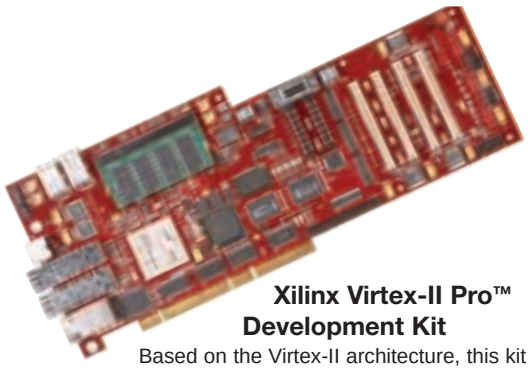

Xilinx Virtex-II Pro™ Development Kit

Based on the Virtex-II architecture, this kit provides a unique environment for developing high-performance microprocessor and I/O intensive applications.

Part Number	Description	Resale Price
ADS-XLX-V2PRO-DEVP7-5	Xilinx Virtex-II Pro Development Kit, with XC2VP7, -5 speed grade	$1,995.00 USD
ADS-XLX-V2PRO-DEVP7-6	Xilinx Virtex-II Pro Development Kit, with XC2VP7, -6 speed grade	$2,495.00 USD
ADS-XLX-V2PRO-DEVP20-5	Xilinx Virtex-II Pro Development Kit, with XC2VP20, -5 speed grade	$2,995.00 USD

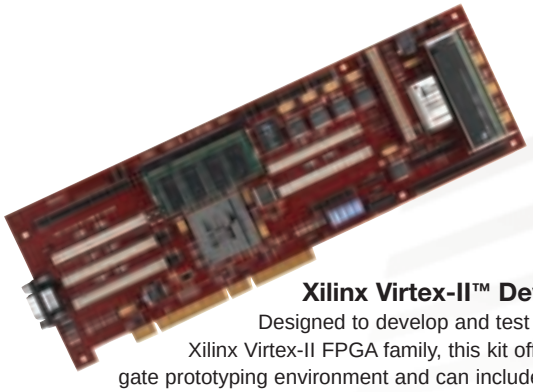

Xilinx Virtex-II™ Development Kit

Designed to develop and test designs targeted to the Xilinx Virtex-II FPGA family, this kit offers up to a 6 million system gate prototyping environment and can include a MicroBlaze Core License.

Part Number	Description	Resale Price
ADS-XLX-V2-DEV1500	Xilinx Virtex-II Development Kit populated with an XCV1500 device	$1,000.00 USD
ADS-XLX-MB-DEV1500	Bundled with Communications/Memory Module and MicroBlaze Core License	$1,400.00 USD
ADS-V2-MB-DEV1500XP	Bundled with Communications/Memory Module, MicroBlaze Core License and high-current power supply	$1,900.00 USD
ADS-XLX-V2-DEV4000	Xilinx Virtex-II Development Kit populated with an XC2V4000 device	$2,500.00 USD
ADS-XLX-MB-DEV4000	Bundled with Communications/Memory Module and MicroBlaze Core License	$2,900.00 USD
ADS-XLX-V2-DEV4000XP	XC2V4000 and high-current power supply	$3,000.00 USD
ADS-V2-MB-DEV4000XP	Bundled with Communications/Memory Module, MicroBlaze Core License and high-current power supply	$3,400.00 USD
ADS-XLX-V2-DEV6000XP	Xilinx Virtex-II Development Kit populated with an XC2V6000 and high-current power supply	$6,000.00 USD
ADS-V2-MB-DEV6000XP	Bundled with Communications/Memory Module, MicroBlaze Core License and high-current power supply	$6,400.00 USD

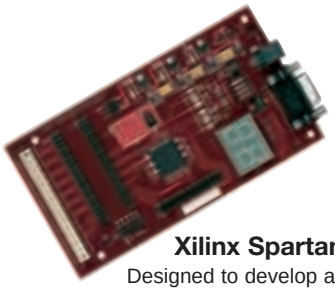

Xilinx Spartan-IIE™ Evaluation Kit

Designed to develop and test designs targeted to the Xilinx Spartan-IIE FPGA family, this kit is optimized for the low-cost, consumer digital convergence market and can include a MicroBlaze Core License.

Part Number	Description	Resale Price
ADS-XLX-SP2E-EVL	Xilinx Spartan-II Evaluation Kit	$249.00 USD
ADS-SP2E-MB-EVL	Xilinx Spartan-IIE Evaluation Kit bundled with Communications/Memory Module and MicroBlaze Core License	$650.00 USD

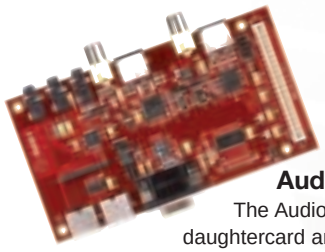

Audio/Video Module

The Audio/Video Module is an expansion daughtercard and offers a platform to develop and test products that require audio and/or video.

Part Number	Description	Resale Price
ADS-AV-DAU	Audio/Video Module	$300.00 USD

Other Evaluation Kits Available from ADS:

Denali RLDRAM IP Core Demonstration Kit

Part Number	Description	Resale Price
ADS-DLI-RLDRAM	Denali RLDRAM IP Core Demo Kit	Contact your local Avnet FAE for more information.

Communications/Memory Module

Part Number	Description	Resale Price
ADS-EMU-DAU	Communications/Memory Module	$250.00 USD
ADS-XLX-MB-DAU	Communications/Memory Module bundled with MicroBlaze Core License	$499.00 USD

Motorola ColdFire® Evaluation Kit

Part Number	Description	Resale Price
ADS-MOT-5282-EVL	Motorola ColdFire Evaluation Kit	$250.00 USD
ADS-MOT-5282PE-EVL	Motorola ColdFire Evaluation Kit with P&E BDM interface	$300.00 USD

oss The Board.™

Avnet Design Services Accelerates Design Cycles

Avnet Design Services (ADS), the technical arm of Avnet Cilicon, provides a variety of technical resources and capabilities to speed your design cycles. From concept to reality – ADS can help get your products to market faster with:

- Reference designs
- Development & evaluation kits
- Fee-for-service FPGA & ASIC designs
- Support of Field Applications Engineers (FAEs)
- Product & system-level architecture selection assistance
- IP evaluation

Partnering with Avnet Cilicon means working with a company that is dedicated to giving you Support Across The Board.™

For more information or to purchase an evaluation kit, visit www.avnetavenue.com/xcell

Ready.
Set.
Go to market.™

AVNET® CILICON

1 800 332 8638
www.avnetcilicon.com

© Avnet, Inc. 2003. All rights reserved. AVNET is a registered trademark of Avnet, Inc.

Explore the Possibilities of the CoolRunner-II Design Kit

A new low-cost development kit for Xilinx CoolRunner-II CPLDs helps experts and novices alike get the most from their designs.

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

The new CoolRunner™-II RealDigital CPLD Design Kit contains everything you need to explore all the new capabilities of the CoolRunner-II RealDigital CPLD. With this kit, you will be able to do everything necessary to test your designs in the industry's best 1.8V CPLD.

For low-power/high-speed designs, a 256 macrocell CoolRunner-II device in a 144-pin thin quad flat pack (TQ144) package is included with every kit. If you need more programmable logic or want to try a different Xilinx CPLD, an additional socket accepts a CoolRunner-II 64 or 32 macrocell device for other low-power/high-speed designs.

Want to test run a 2.5V or 3.3V Xilinx CPLD? This socket also accepts Xilinx best-selling XC9500XL or XC9500XV devices in 36 or 72 macrocell sizes.

The CoolRunner-II CPLD design kit board (Figure 1) can be powered either by a standard power supply or by two AA batteries (battery holder included). The kit also includes a parallel cable for programming your CPLD. The cable simply goes from your PC parallel port to the parallel port connection on the board.

To help you get started, we have included a Programmable Logic Quick Start Book that contains installation instructions for the ISE WebPACK™ software, as well as step-by-step instructions on how to use, enter, and synthesize designs, and how to target Xilinx programmable logic devices.

To give you a good jump-start on designing your own applications, we give you design examples in both schematic and high-level language, such as VHDL.

The kit comes with a resource CD that contains a great deal of useful information:

- Reference designs in VHDL, user constraint file (UCF), and HDL design files
- Instructional video demonstrating the entire process, from design to programming of the device
- ISE WebPACK overview and software options for different Xilinx devices
- Links to related websites
- Application notes, data sheets, power consumption comparison, product pages, articles, and other product presentations
- Technical presentations and packaging information
- Top 10 FAQs and where to find other answers to questions
- Product brochures, reference manuals, and product briefs
- Video featuring Xilinx CEO Wim Roelandts.

Simple Tutorial, Advanced Features

Whether you already design with programmable logic or would like to begin, this kit can help you. The accomplished designer will appreciate the many advanced features and new interfaces on CoolRunner-II, including high-speed transceiver logic (HSTL) and stub-series-terminated transceiver logic (SSTL).

Or you might experiment with the other new memory device interfaces not supported on your current microprocessor. A breadboard area allows you to add components for use in future designs, and you can also connect a daughtercard to the CoolRunner-II design kit board via the dual inline headers.

Designers new to programmable logic can use the demo board with its tutorial to quickly get up to speed on tool usage. The demo covers all aspects of designing a CPLD, from design entry with WebPACK, through simulation, to production of a JEDEC file used to program the device. This straightforward tutorial should get you up and designing within hours, and can also be used to encourage lab technicians to explore their own design aspirations.

The demo board is handy to have in your lab to test new design ideas for implementation on current products. If your system already includes a CPLD, it serves as a nice test bed for running upgrade tests. If power is your concern, the CoolRunner-II development board is especially useful for running different designs to test power consumption. If you want some ideas about reducing power consumption, check out the low-power tips and tricks application note XAPP346 on the CLPD application notes webpage (*www.xilinx.com/apps/epld.htm* – see Table 1 for a list of all of the application notes described in this article).

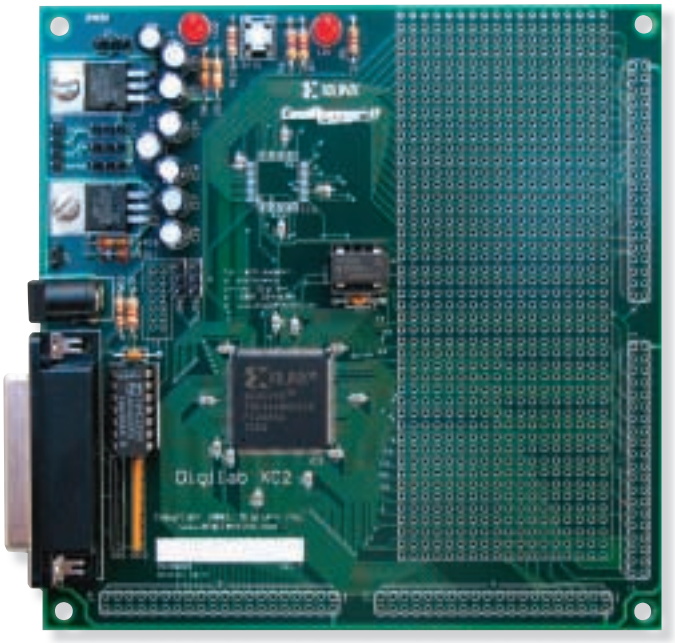

Figure 1 – CoolRunner-II design kit board

The 256-macrocell (XC2C256) CPLD in a TQ144 – with room for other devices – gives you a good sense of design capacity. With this device configuration you get 118 I/Os, two I/O banks, clock doubling/division, and voltage-referenced I/O. If you need voltage-level translation, this device can handle 1.5V, 1.8V, 2.5V, or 3.3V I/Os.

If your design is relatively simple, try the 32 or 64 macrocell device. If your design is complex, try the 256 macrocell configuration. If you think partitioning your design offers the best solution, use both devices. This board is designed to give you the maximum amount of flexibility.

What Can You Do With It?

With both a breadboard area and dual inline header sockets, the design kit board is suitable for developing on-board applications, attaching daughtercard boards, or communicating with another board via headers. It is large enough to accept many useful designs, including standard bus protocols (SPI and I^2C), microcontroller interfaces, and serial communications, such as Infrared Data Association and universal asynchronous receiver transmitter (UART) designs, which have already been created by Xilinx application engineers. These last items are available free by downloading application note XAPP345.

PicoBlaze Soft Processor

The CoolRunner-II design board is extremely useful for evaluating the PicoBlaze™ soft processor in your design. Available for downloading free from application note XAPP387, the PicoBlaze soft processor offers a constant K-coded programmable state machine, and is written in VHDL and C programming languages. Extremely customizable, both its size and functionality can easily be changed, making PicoBlaze an ideal application for content-sensitive designs. The PicoBlaze soft processor provides 49 different instructions, eight 8-bit registers, 256 directly addressable ports, and a maskable interrupt. If you simply want an

8-bit microcontroller interface, take a look at application note XAPP349.

Memory Interface

The CoolRunner-II design board may well be the best solution for testing new memory interfaces. With multiple memory interfaces – including HSTL, SSTL, and LVCMOS – you can test SDRAM, SSDRAM, MSDRAM (mobile SDRAM), UtRAM (UniTransistor RAM), cellular RAM, and flash memories. You can even design your own memory controller for new flavors of portable or mobile RAM as they become available.

The CPLD application notes webpage contains many memory design examples, including beginning designs for flash and DDR SDRAM memory. Check this website often for new memory interfaces and updates or to download an HDL design and get email notification of new application notes.

Crosspoint Switch

The non-blocking architecture allows each output to be independently connected to any input, and any input to be connected to any or all outputs. The double-row latch architecture utilized in this design allows switch reprogramming to occur in the background during operation. Activation of the new configuration occurs with a single configuration pulse. Each output can be individually disabled and set to a high-impedance state, allowing easy expansion to larger switch array sizes. This is covered in detail in application note XAPP380.

OTF Reconfiguration

On the fly (OTF) reconfiguration permits the CPLD to operate with an initial design pattern while simultaneously acquiring a second pattern. For example, your first design pattern can be a power-up application, such as a built-in self-test (BIST), and the second pattern can be one of normal device operation. If you want to run diagnostics on your system on power-up, you would load a functional pattern to perform additional bus interface support. The second pattern can be configured into the device with minimal disturbance to its operation.

Another important consideration is that OTF reconfiguration enhances security, because switching keys on the fly permits robust protocols for data communication designs that can be tough to crack. OTF reconfiguration takes in-system programming (ISP) to a new level and will undoubtedly spawn many new applications that take advantage of this capability. More information about OTF can be found in application note XAPP388.

FPGA Downloader

Another common use for CPLDs is as an FPGA downloader. At power-up the CPLD sequentially loads each FPGA with a pattern, which is stored in external memory. By assigning a CPLD to accomplish this task, you can get double duty from the CPLD. On power-up, it loads the FPGA; after this power-up task is done, you can OTF reconfigure the CPLD and use it as a functioning device in your design. It's like getting two devices for the price of one. To learn more about this application, see application note XAPP137.

If you do experience problems, a website has been set up for you at *www.digilentinc.com* to help troubleshoot the board.

Conclusion

CoolRunner-II RealDigital CPLDs can be the perfect solution for any high-speed/low-power design and cover many application areas. In addition to their features, these devices also come in packages that suit high volume and small form factor packaging.

To order your CoolRunner-II RealDigital Design Kit, just go to *www.xilinx.com/cpld/* and click on "CoolRunner-II Design Kit for $49.99."

XAPP137	Configuring Virtex FPGAs from Parallel EPROMs with a CPLD
XAPP345	IrDA and UART Design in a CoolRunner CPLD
XAPP346	Low Power Tips for CoolRunner Design
XAPP349	CoolRunner CPLD 8051 Microcontroller Interface
XAPP380	Building Crosspoint Switches with CoolRunner-II CPLDs
XAPP387	PicoBlaze 8-Bit Microcontroller for CPLD Devices
XAPP388	"On the Fly" Reconfigurations

Table 1 – All CoolRunner-II application notes can be found at www.xilinx.com/apps/epld.htm.

Develop Applications Right Out of the Box

Avnet's Virtex-II Pro development kit is a full-featured hardware development environment that includes Linux OS, 10 Gigabit Ethernet modules, extensive memory, and interface support.

by Warren Miller
Vice President of Marketing
Avnet, Inc.
warren.miller@avnet.com

One of the best ways to accelerate your design time is to use standard hardware development platforms to prototype, evaluate, and test your complex designs. With the right platform, you can quickly create and test complex hardware – a task that is much too difficult to do in software.

Complex networking applications such as Gigabit Ethernet switches and routers require hardware platforms so various design alternatives can be evaluated. Compliance and field tests are also only possible with hardware platforms. Acceleration of these tasks can make all the difference in the time-to-market race.

Avnet's development kit with a Xilinx Virtex-II Pro™ Platform FPGA provides the right mix of hardware, support, third-party tools, and application examples to allow you to create a platform for real-life designs. The kit includes up to a Virtex-II Pro XC2VP20 device and enough memory, peripheral devices, and interfaces to create a complete hardware development environment right out of the box.

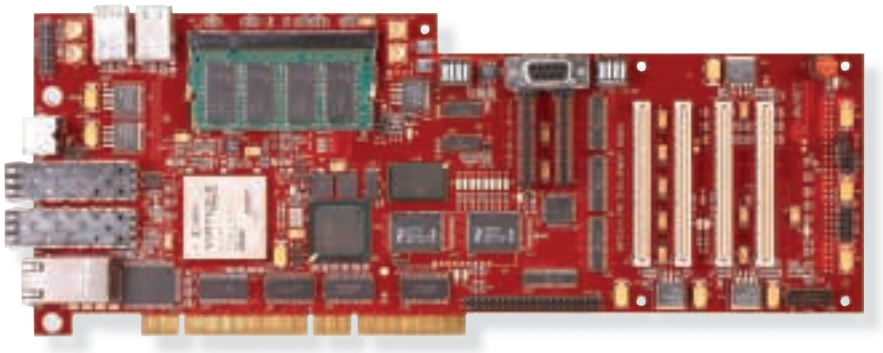

Figure 1 – Avnet-designed Virtex-II Pro XC2VP20 development board

10 Gigabit Ethernet Designs

Many features of the Avnet Virtex-II Pro development board, shown in Figure 1, are especially useful to support 10 Gigabit Ethernet designs. The Virtex-II Pro device, an XC2VP20-FF896, provides the logic capacity needed for 10 Gigabit MAC and other supporting IP cores. Extensive on-board memory provides sufficient storage for both code and data space for the embedded IBM™ PowerPC™ processors found on the Virtex-II Pro device. High-speed interface functions using the RocketIO™ features of the Virtex-II Pro Platform FPGA are supported on the board with a variety of standard connectors. Xilinx IP cores provide the logic functions required to implement the physical and logical layers of common serial communications standards. A detailed block diagram of the board is shown in Figure 2.

The development board has enough memory devices to support even a complex design. Memory support includes:

- Micron DDR SDRAM SODIMM (128 MB expandable to 1 GB)
- Micron Mobile SDRAM (two 8 Mb x 16 Mb devices, 32 MB total)
- Cypress Asynchronous SRAM (512 Kb x 32 Kb, 2 MB total)
- Intel StrataFlash (16 MB total)
- CompactFlash card
- Configuration memory using Xilinx PROM XC18V02-VQ44 and the Xilinx System ACE™ CF solution.

With 1 GB of high-speed DDR SDRAM memory, bandwidth will not be a bottleneck for most 10 Gigabit Ethernet applications.

The board supports a variety of high-speed interface standards. These include:

- A standard XPAK module (located on the back of the board) that can implement either a 10 Gigabit Ethernet port or an OC-192 serial interface

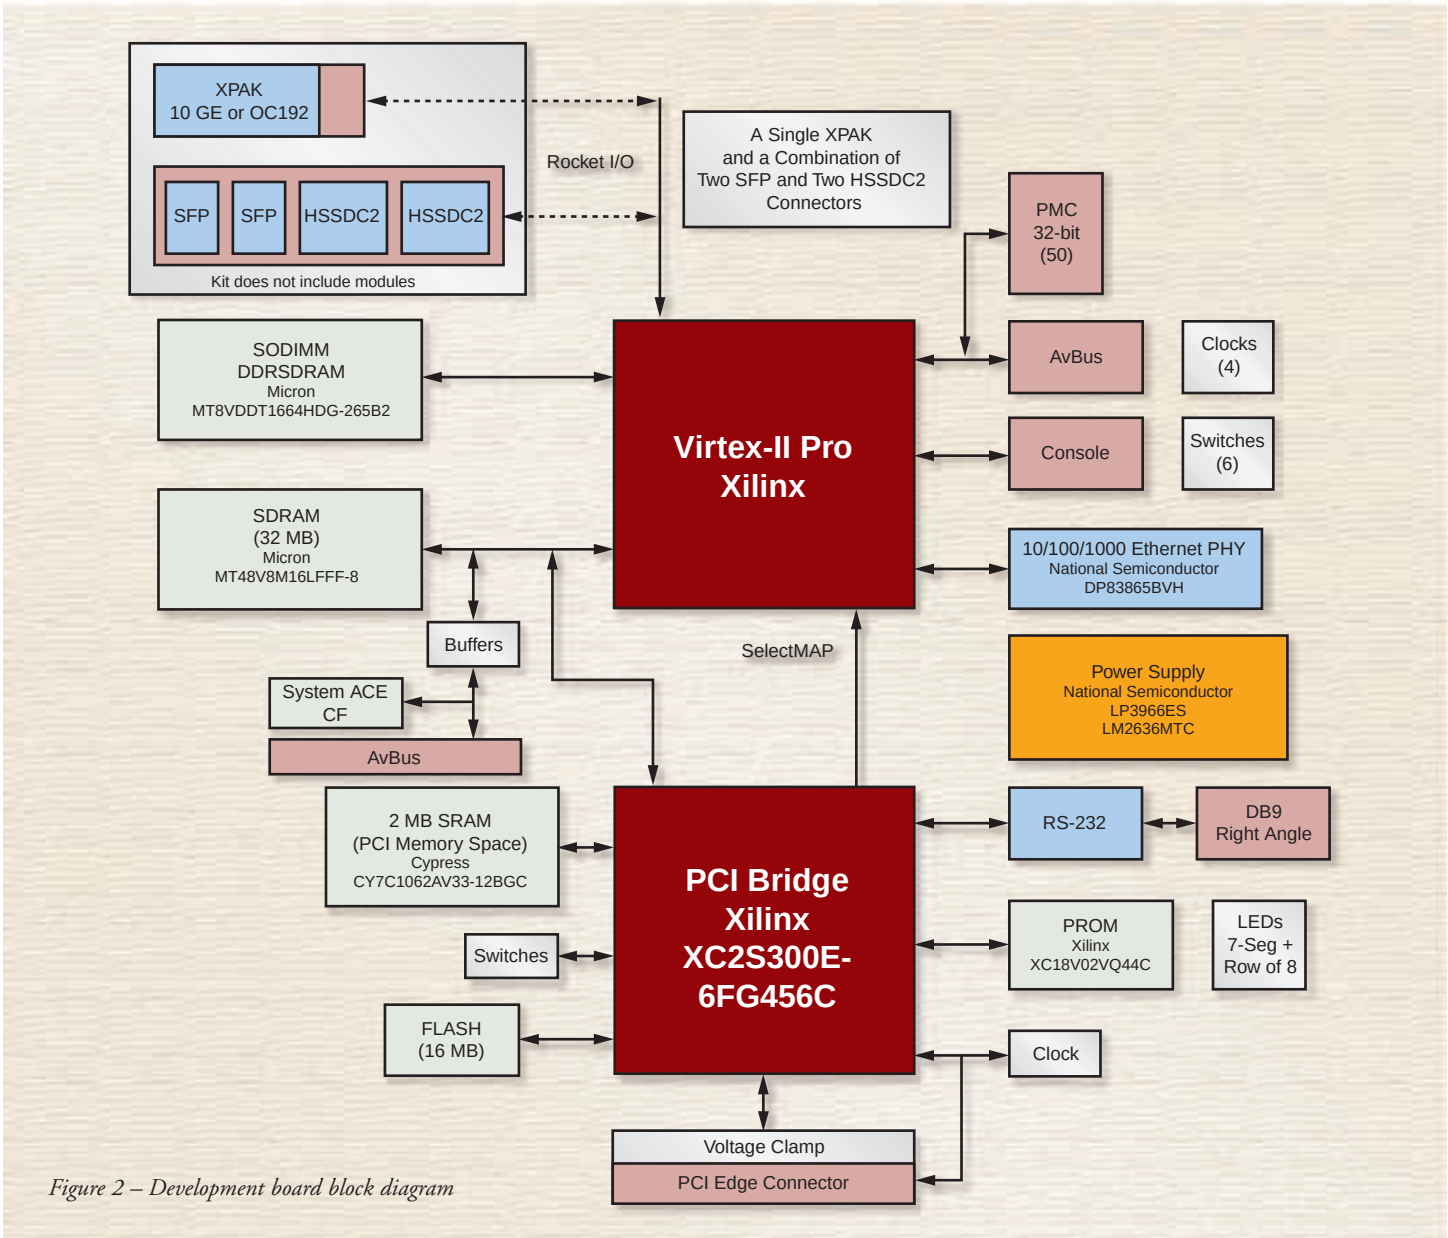

Figure 2 – Development board block diagram

- Two HSSDC2 connectors for implementing InfiniBand™ or Fibre Channel interfaces
- Two small form-factor pluggable (SFP) modules for implementing Gigabit Ethernet, Fibre Channel, or InfiniBand protocols.

These interfaces can be used in conjunction with RocketIO transceivers on the Virtex-II Pro device and the appropriate Xilinx IP cores to implement the desired serial interface standard. A bridging or switch application could use one or more of these ports as the other end to a 10 Gigabit Ethernet interface.

Expansion of the basic features of the board is possible by using the PCI bus or the four 140-pin general purpose connectors (AvBus connectors). A variety of AvBus-compatible boards, from Avnet and third parties, provide even more hardware options.

A Speedy Start

In addition to the extensive hardware features described above, the kit comes with a variety of supporting documentation and design files to make it easy to get started. The kit ships with a detailed user manual, quick start guide, and an extensive board support package that includes a full-featured Linux kernel, PCI bus, and memory drivers. Several demonstration examples show how to best use the features of the Virtex-II Pro device and interface it to additional devices on the development board. Two of the projects are PowerPC-based and give examples of memory interfaces – one for an OPB interface and one for a PLB interface.

Another example project is RocketIO-based and illustrates the use of a custom peripheral to allow PowerPC access to the multi-gigabit transceiver (MGT) at 3.125 Gbps.

Implementing 10 Gigbit Ethernet

Several of the design considerations in the development kit will be common to any Virtex-II Pro-based design and are shared here. On our design, the MGTs of the Virtex-II Pro Platform FPGA are connected to two HSSDC2 connectors, two SFP connectors with EMI cages, and pads for an XPAK host connector with mounting holes for the mid-board module holder. In the case of the two HSSDC2 and two SFP connections, the MGTs were treated individually. This means that the lengths of the transmit and receive signals were matched per MGT but not matched to any other MGT. However, in the case of the XPAK interface, four MGT channels were bonded together to create the 10 Gigabit Attachment Unit Interface (XAUI). The four transmit pairs and the four receive pairs have matched lengths.

ADS-XLX-V2PRO-DEVP7-5 (Populated with XC2VP7, –5 speed grade)	$1,995
ADS-XLX-V2PRO-DEVP7-6 (Populated with XC2VP7, –6 speed grade)	$2,495
ADS-XLX-V2PRO-DEVP20-5 (Populated with XC2VP20, –5 speed grade)	$2,995
ADS-XLX-V2PRO-DEVP20-6 (Populated with XC2VP20, –6 speed grade)	$3,495

Table 1 – Part numbers and prices of the Avnet Virtex-II Pro development kit

The XPAK is a SFP transceiver module for 10 Gbps serial data transmission. The XPAK interface on the board was designed to run at the IEEE 10GBASE-R optical rate of 10.3125 Gbps with a four-lane electrical interface at 3.125 Gbps (XAUI interface). This interface requires the –6 speed grade Virtex-II Pro device and the use of the Xilinx 10 Gigabit Ethernet MAC core. The XPAK interface on the Virtex-II Pro development board complies with the XPAK MSA (multi-source agreement) Revision 2.1 (except for the programmable supply, which is implemented via jumpers). The XPAK MSA closely resembles the XENPAK MSA and makes frequent references to it. The XPAK form factor was used because it is half the size of a XENPAK, does not require a large cutout in the PCB, and has a mid-board mounting option, allowing the module to be placed anywhere on the board instead of on the faceplate.

The 156.25 MHz differential clock input is used for the reference clock to the MGT macro. Because the phase-locked loop of the MGT always multiplies by a factor of 20, using the 156.25 MHz clock results in a transmission rate of 3.125 Gbps. The transmit and receive signals are directly connected to the XPAK host connector or are DC-coupled. However, all XPAK-compliant modules have AC coupling on both the transmit and receive signals inside the module itself. The transmit differential pairs are routed on the solder side of the board, while the receive pairs are routed on the component side. This keeps the signals from crossing on the way to the XPAK connector. The analog ground planes in the layers directly adjacent to the outer layers provide the return paths. The analog ground planes are separated from the digital ground used for the rest of the components on the board, but are referenced to digital ground in several locations through ferrite beads.

More design details like these are documented in the user guide, design files, and example designs included with the development kit.

Conclusion

The Avnet Virtex-II Pro development kit is designed with real applications in mind. It has all of the features demanded by this class of application – a large FPGA, extensive memory, and standards-based flexible high-speed serial ports, as well as considerable expansion capability. The kit should enable a faster time to market for your next 10 Gigabit or other high-speed embedded design.

The Avnet Virtex-II Pro development kit is available now. See Table 1 for available Virtex-Pro devices and board prices. Contact your local Avnet sales office to get detailed ordering information and to talk to an Avnet FAE about your specific design need. Σ

Memec Design Simplifies Bluetooth Development

A new development kit from Memec Design lets you implement Bluetooth wireless technology on Spartan-IIE FPGAs with MicroBlaze soft processors.

Embedded Bluetooth Development Kit

by Jim Beneke
Director of Technical Marketing
Memec
jim_beneke@mei.memec.com

Dennis Schlaht
Director of Technical Marketing
Insight Memec
dennis_schlaht@ins.memec.com

After several false starts, Bluetooth™ short-haul wireless connectivity technology for embedded systems finally seems ready for prime time. Bluetooth technology is just now beginning to gain momentum, and we expect it to grow significantly in the coming years.

For that reason, Insight Memec and Memec Design have developed an Embedded Bluetooth Development Kit featuring a Xilinx Spartan™-IIE FPGA coupled with a Xilinx MicroBlaze™ soft-processor core.

Bluetooth Basics

Bluetooth technology is a frequency hopping spread spectrum (FHSS) system that operates in radio frequencies in the 2.4 GHz to 2.5 GHz ISM band, and it has a maximum data throughput of 723.2 Kbps.

As with most standards, interoperability is key to success. The Bluetooth Special Interest Group's goals include maintaining an open specification, delivering voice and data capability, and providing worldwide usability for short-range wireless solutions.

The specification defines the protocols and profiles used by Bluetooth-certified products. Usage models define the real-world applications, and these usage models result in profiles as defined in the Bluetooth specification.

Profiles are basically instructions for implementing usage models. The profiles assure interoperability by providing a well-defined set of higher layer procedures and uniform ways of using the lower layers of the Bluetooth protocol. The serial port profile (SPP), for example, provides basic RS-232 serial cable emulation for Bluetooth devices. Legacy applications do not have to be modified to use Bluetooth technology; they can simply treat a Bluetooth link as a serial cable connection.

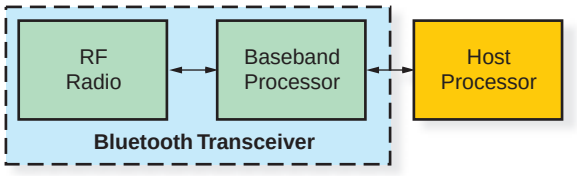

Figure 1 – Key elements of a typical Bluetooth application

Adding Bluetooth capability requires the key elements shown in Figure 1. The RF radio and the baseband processor are typically available as an integrated Bluetooth transceiver solution.

The transceiver chip provides a serial interface to the host processor via USB or UART ports. The host processor typically executes the upper protocol-specific functions defined by the Bluetooth standard.

These protocol functions are known as the Bluetooth protocol stack and can be licensed in processor-specific binary or independent source code formats from stack providers, such as Stonestreet One.

Integrated Solutions

In a Windows™ or PocketPC-based environment, a Pentium™ or ARM processor actually executes the upper stack functions, which are usually included as part of some Bluetooth software applications.

In an embedded application, the Bluetooth upper stack and host processor function must be accounted for in some other way, typically with an embedded core.

The host processor performance requirements for supporting the Bluetooth protocol are minimal, usually less than 1 MIPS. Thus, it is very easy to add Bluetooth support to an embedded application if the host processor has processing bandwidth and a Bluetooth stack is available for the targeted processor.

The 150 D-MIPS performance of the MicroBlaze soft-processor core means the additional 1 MIPS of processing overhead can often be accommodated with relative ease.

The availability of the Bluetooth protocol stack presents a bigger obstacle, because the stack must be ported to the targeted embedded processor. However, with the introduction of the Memec Design Embedded Bluetooth Development Kit, this issue has been addressed for any MicroBlaze-based system.

The Embedded Bluetooth Kit

The Memec Design kit provides all the necessary components to build and test a Bluetooth system:

- An evaluation version of the Bluetooth stack ported to the MicroBlaze core
- A Spartan-IIE development board
- Two P160 Bluetooth modules
- An evaluation version of the BTExplorer™ Windows application
- Xilinx EDK software
- Cables
- Power supplies
- Documentation.

Figure 2 shows a typical setup of the Bluetooth kit.

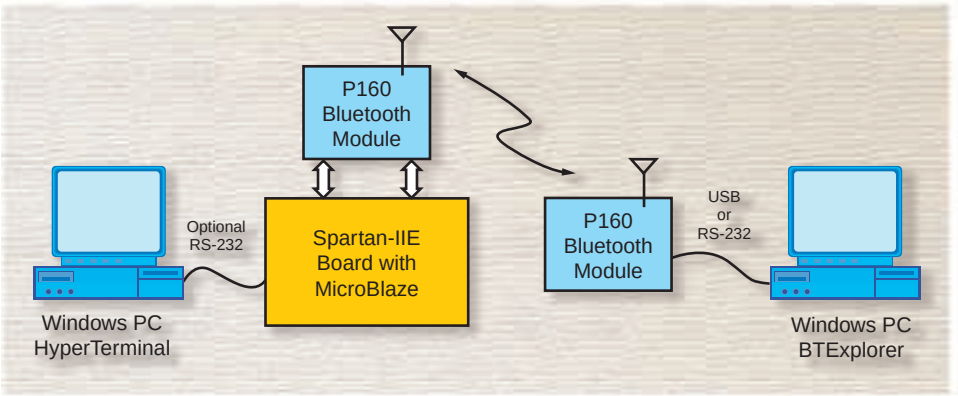

Figure 2 – Bluetooth Development Kit with wireless link

A simple point-to-point radio system can be created with the kit components and a Windows-based PC. One end of the communication system is created with a standalone P160 Bluetooth module connected via USB or RS232 to a Windows PC. This PC runs the evaluation version of the BTExplorer application, which executes the Bluetooth stack within the Windows environment.

The other radio system uses the Spartan-IIE board and second P160 Bluetooth module. The Spartan-IIE platform implements a MicroBlaze design, combining the evaluation Bluetooth binary stack and serial port profile with your application code. An optional second PC connects to the Spartan-IIE board for user feedback and control.

Figure 3 shows the functionality of the P160 Bluetooth module based on Broadcom's BCM2035 – a monolithic, single-chip, baseband processor with an integrated 2.4 GHz transceiver for Bluetooth v1.1 and 1.2 applications. It minimizes the footprint and system cost of implementing a Bluetooth link by integrating all critical components into the device.

The BCM2035 is an ideal solution for any voice or data application that requires the Bluetooth standard Host Controller Interface (HCI) via either UART or USB, and PCM (pulse code modulation) audio interfaces. The integrated microprocessor unit stores the lower level protocol stack in ROM plus patch RAM to provide the maximum flexibility while eliminating the need for external flash memory for the lower stack.

The BCM2035 radio transceiver provides enhanced radio performance to meet the most stringent industrial temperature applications or the tightest integration into portable devices. It provides the highest available radio performance of any single-chip device with -90 dBm sensitivity and +7 dBm programmable output power.

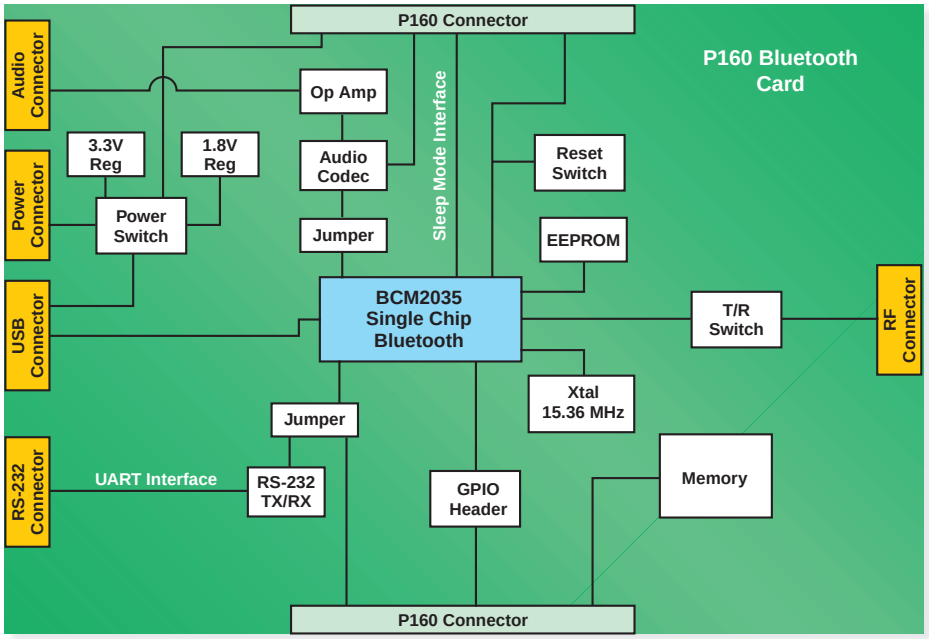

Figure 3 – P160 Bluetooth Module functional diagram

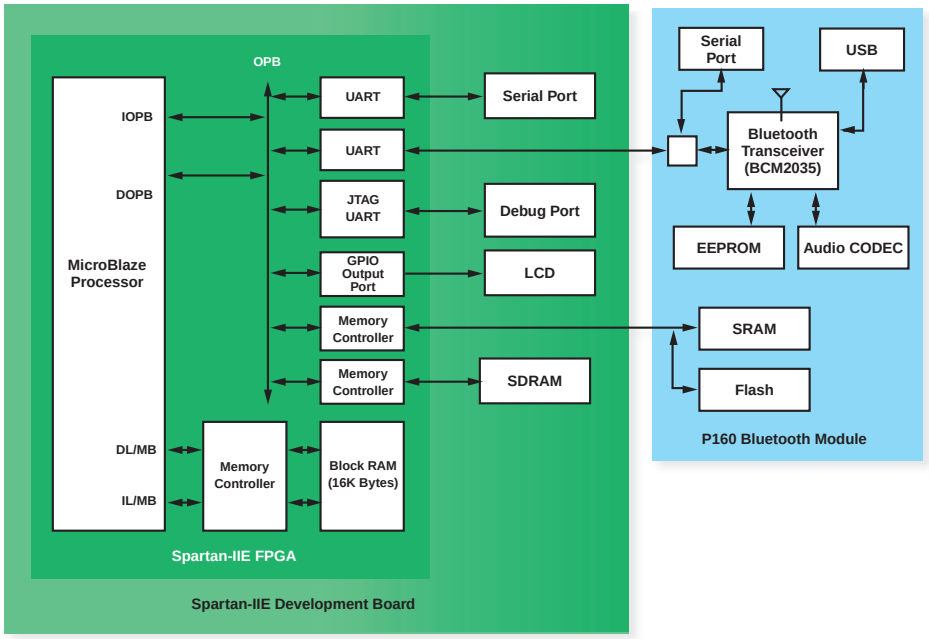

Figure 4 – Typical MicroBlaze-based system

Development System Capabilities

The Spartan-IIE board from Memec Design provides all the necessary features for implementing a simple MicroBlaze design. The board includes the P160 expansion module slot for connection to the P160 Bluetooth module, SDRAM, serial port, and miscellaneous support circuits.

Figure 4 shows an example MicroBlaze system that incorporates the Bluetooth interface. Because the Bluetooth module looks like a standard serial port to the MicroBlaze system, the only hardware modification required to the system architecture is the addition of a UART block.

The Memec Design Embedded Bluetooth Development Kit includes a fully functional evaluation version of Stonestreet One's Bluetopia™ protocol stack. An implementation of the upper Bluetooth protocol stack, the Bluetopia software eases application development by providing a robust yet easy-to-use development tool that implements the Bluetooth protocols above the HCI.

The Bluetopia application programming interface (API) provides access to the upper-layer protocols, including Logical

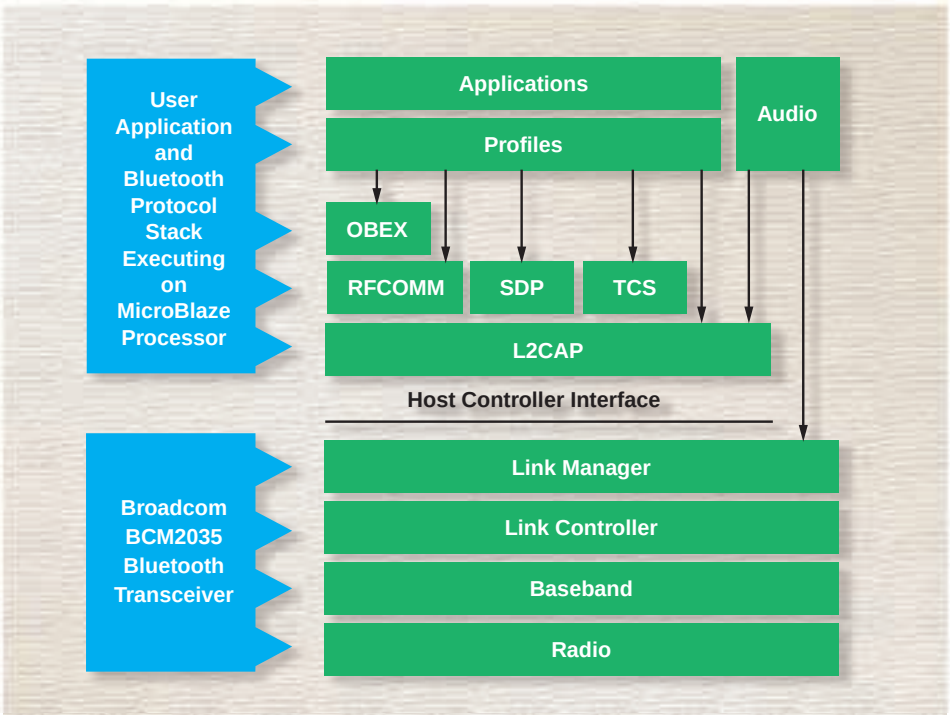

Figure 5 – Bluetooth protocol stack

Link and Adaptation Protocol (L2CAP), Service Discovery Protocol (SDP), RFCOMM protocol, and OBEX protocol. Bluetopia software also provides APIs for all mandatory and most optional profiles.

In addition, Stonestreet One provides lower level HCI transport drivers to interface to various Bluetooth devices such as the BCM2035. These lower level HCI drivers have been written to support the addition of new HCI transport drivers with little or no change to existing applications.

The evaluation version of the Bluetopia protocol stack provides support for the following Bluetooth protocols: HCI, L2CAP, SDP, and RFCOMM. The Bluetopia stack also provides support for generic access profile (GAP) and serial port profile (SPP) Bluetooth profiles. Figure 5 illustrates an example Bluetooth protocol stack.

In addition to the above listed protocols and profiles, Stonestreet One will make available an extensive suite of existing and future Bluetooth protocols/profiles as required. Full development or production versions are available through license from Stonestreet One along with additional profiles if required.

BTExplorer is a user-friendly Windows application that is used to connect and manage Bluetooth devices. BTExplorer discovers Bluetooth devices in the vicinity and presents them in an easy-to-understand format similar to Windows Explorer.

As used in the development kit, BTExplorer facilitates a simple peer-to-peer system. When connected to the standalone P160 Bluetooth Module, BTExplorer provides the upper stack capability to the module, thus creating a complete Bluetooth node.

Conclusion

The Memec Design Embedded Bluetooth Development Kit allows designers to enhance any new or existing MicroBlaze-based designs with the addition of a Bluetooth interface. By providing a MicroBlaze-specific port of the Bluetooth stack along with the necessary prototyping hardware, the kit eases the development process, shrinks the design cycle, and speeds your time to market.

The Embedded Bluetooth Kit is available for $1,495 from Insight Memec at *www.memec.com/insight-memec/*. Plans are under way to develop a similar development kit utilizing the new Xilinx Spartan-3 FPGA.

Save Time, Money, and Space

Xilinx leads the industry in 300 mm wafers and 90 nm process technologies.

by Steve Sharp
Senior Manager, Programmable Logic Solutions Marketing
Xilinx, Inc.
steve.sharp@xilinx.com

Xilinx is aggressively driving down costs for its customers, especially in the area of 90 nm process technology and 300 mm wafer manufacturing.

The industry leader in PLD (programmable logic device) process technology since the introduction of its Virtex™-E and Virtex-EM FPGA families in 1999, Xilinx was already well ahead of the Semiconductor Industry Association's process roadmap. And we've maintained that lead by more than a year ever since.

Our successful strategy includes partnerships with two world-class fabrication partners: IBM™ in the United States and UMC™ in Taiwan. Both companies are experts in process technology and manufacturing. They have allowed Xilinx to be the first PLD manufacturer to deliver products based on 300 mm wafers and, more recently, 90 nm process technology.

Advantages Beyond Wafer-Thin

Larger 300 mm wafers are imperative for lowering costs. With almost twice the usable wafer area and approximately 2.5 times the number of die per wafer, cost savings of 30% to 40% over 200 mm wafers are typical.

According to Glen Yeung of the investment firm Salomon Smith Barney, "[These] 300 mm facilities are the wave of the future. For both chip and equipment companies, the move is an economic imperative."

The semiconductor industry has been successful in the use of "copy exact," which makes it effective to transition a product

from one fabrication line to another. However, when moving from 200 mm wafers to 300 mm wafers, copy exact does not exist. This leaves the door open for process and equipment variabilities that can affect yield.

"We would expect a significant investment to convert an existing product from 200 mm to 300 mm wafers," said Jim Feldhan, president of Semico Research Corp.

Re-Qualification Headaches

Designers must always consider a re-qualification when a device moves from 200 mm to 300 mm wafer manufacturing. These re-qualification efforts can include:

- Review of the process change notice (PCN) by the designer
- Assembling test boards
- Re-characterization of boards
- Full systems review, including evaluation of thermal stress and high/low voltage testing.

Re-qualification can take as long as three to nine months and consume valuable engineering resources.

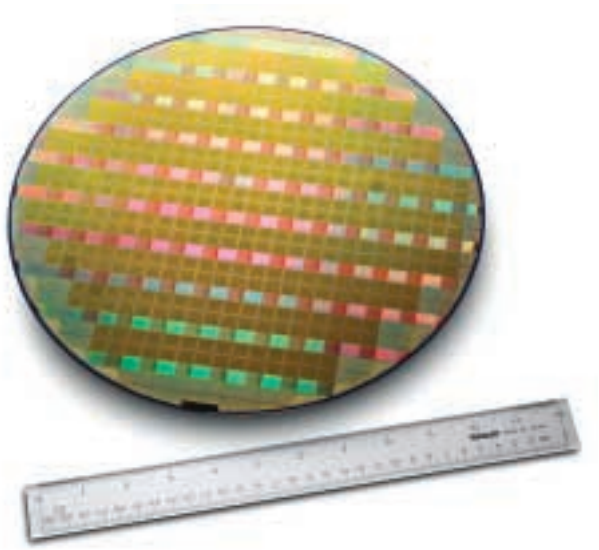

Xilinx: At the Cutting Edge

Xilinx moved established products, such as the Spartan™-IIE and Virtex-II FPGA families, to 300 mm wafers last year. Xilinx has also chosen to release its Virtex-II Pro™ and Spartan-3 FPGA families directly on 300 mm wafers, saving customers from performing a system re-qualification later on.

We realized that to deliver the maximum cost benefit to our customers, it would be necessary to avoid requiring companies to re-qualify a product in their systems when moving from 200 mm to 300 mm wafers. Therefore, we have done the re-qualification for them.

Cheaper Solutions in 90 nm

The use of the latest 90 nm process technology can also have a profound impact on cost. With half the die size for an equivalent circuit on 130 nm technology, 90 nm technology is rapidly becoming the choice of industry leaders as they strive to reduce costs for their customers.

As Dan Hutcheson of VLSI Research Inc points out, "Companies that get into 90 nm production first will get a tremendous advantage in lower cost. Rivals who are late in [adopting] 90 nm process technology will fall behind and may not be able to catch up."

When the die size advantage of 90 nm process technology is combined with the manufacturing efficiency of 300 mm wafers, the savings of a combined 90 nm/300 mm solution is a 5X increase in gross die per wafer when compared to 130 nm technology on 200 mm wafers.

These cost savings – the largest in recent semiconductor history – could result in pricing of under $20 for one million system gates in 2004, which would be a seven-fold cost decrease from the year 2000.

Conclusion

Not all companies have the resources to develop such advanced technology, but as leaders in their respective industries, Intel™, IBM™, Texas Instruments™, and Xilinx are all adopting – and adapting to – the latest 90 nm/300 mm process technologies. Our customers will be the first to realize the tremendous cost benefits this new technology can deliver. Σ

Xilinx Events and Tradeshows

Xilinx participates in numerous trade shows and events throughout the year. This is a perfect opportunity to meet our silicon and software experts, ask questions, see demonstrations of new products and technologies, and hear other customers' success stories with Xilinx products.

For more information and the most up-to-date schedule, visit *www.xilinx.com/events/*.

Worldwide Events Schedule

September 9-11
Military and Aerospace Programmable Logic Devices (MAPLD) International Conference
Washington, DC

September 15-18
Intel Developer Forum
San Jose, CA

September 16-17
Embedded Systems Conference
Boston, MA

November 17-19
Software Defined Radio Technical Conference and Product Exposition
Orlando, FL

March 8-10, 2004
Wireless Systems Design
San Diego, CA

March 30-April 1, 2004
electronicaUSA with Embedded Systems Conference
San Francisco, CA

PCI

AWESOME ASIC Prototyping Solutions

PCI-X

Cool Emulation Platforms

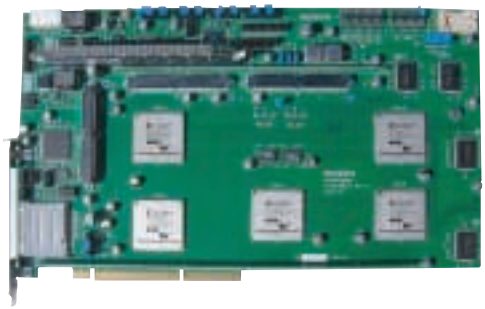

DN3000K10

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

- Up to 5 2V6000/2V8000'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

DN6000K10S

Single VirtexII - Pro PWB

- 2VP70/2VP100
- 4 Separate DDR SDRAM's (16Mb x 16)
- 4 Separate SSRAM's (512k x 36)
- 2 Flash's (4M x 16)
- Misc. serial I/O's
- 2 PowerPC's
- PCI / PCI-X or standalone

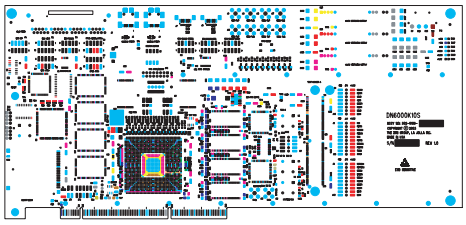

DN3000K10S

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

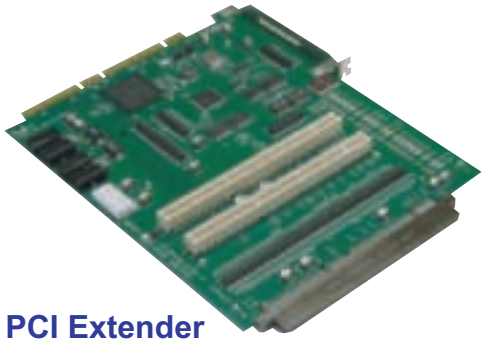

PCI Extender

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

A Sample of Our Services:

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

Mr. Freaky Says:

"This stuff is WAY cool!"

- White-Backed Mousebird

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

600M+ ASIC Gates SOLD!

Take the Titanium Solution

Xilinx Titanium Technical Service engineers coach clients through design issues.

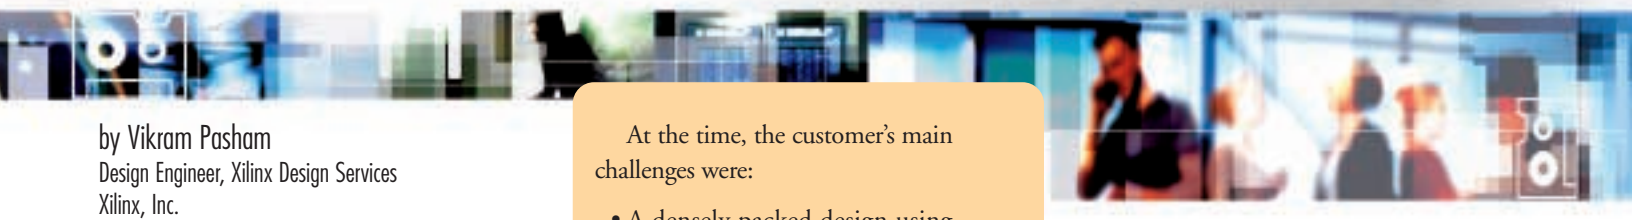

by Vikram Pasham
Design Engineer, Xilinx Design Services
Xilinx, Inc.
vikram.pasham@xilinx.com

Need to enhance your design productivity? Decrease your design costs? Accelerate time to market? With the complexities and possibilities of today's designs, it's easy to get sidetracked from your main goals.

Xilinx Titanium Technical Service provides on-site or off-site support for clients on a contract basis. Titanium application engineers are adept at ensuring that you start and finish your designs the right way. Our engineers provide design methodology coaching to make sure you take the most efficient approach. And our engineers' skill at tracing debug issues back to the design is one of our most powerful services. Here's an example from one of our clients.

Customer Challenges

During the development of a high-definition television (HDTV) video application, the engineering team of our client company (which prefers to remain anonymous for competitive reasons) had developed proprietary DSP algorithms for uniform video correction. These were implemented on a custom ASIC developed by a third party, which had several silicon bugs.

Fearing that the bugs would cost them the chance to be first to market with their technology, the engineers ported part of the ASIC's functionality to a Xilinx FPGA. Yet they still noticed that prototypes, when tested in the lab, had power-up and image distortion issues. Thus, they decided to bring in a dedicated Xilinx Titanium Technical Service application engineer.

At the time, the customer's main challenges were:

- A densely packed design using 99% slices, 70% block RAMs, and multipliers in a Virtex™-II device. They wanted to add additional features into the FPGA, but did not have the budget to use a larger one.
- They couldn't meet timing in their Virtex-II –4 speed grade part and were using –5 speed grade. This increased their bill of materials.
- This design had to support six different pin configurations for different RGB connectors on different boards.
- The prototypes were behaving inconsistently on power-up and exhibiting image distortion.

The Titanium Solution

Over the course of one week, a Xilinx Titanium engineer gave the client's engineering team a crash course on FPGA design techniques and constraints. He suggested replacing the existing design's clocking structure with Virtex-II digital clock managers (DCMs) for multiple clock generations. After learning about the features of DCMs, the engineering team discovered they could use DCM status bits to determine if the digital video interface (DVI) clock was disconnected. Based on this discovery, the team designed a robust recovery mechanism. This solved the power-up and image distortion problems.

The existing design was already using all of the resources of the FPGA, yet there were further requirements to add functionality. The Titanium engineer identified portions of the design that weren't suited for FPGA architecture. He recommended other changes for optimal efficiency, while still preserving design functionality. Knowing that Virtex-II silicon can easily meet 60 MHz DVI clock rates, he suggested using tight packing options in map tools, and using the multipass place-and-route and floorplanning features in Xilinx ISE tools. This freed up enough resources for the engineers to add the additional features in the FPGA.

To meet timing in the Virtex-II –4 speed grade part for the six different pin configurations, the Titanium engineer placed timing constraints to cover all the clocks, cross clock domain paths, and multicycle paths. Together, they put together working prototypes in the lab supporting the six different pin configurations.

After these adjustments to the design, the company is now able to use a lower-speed grade part, reducing their system cost. They are on schedule to meet their time-to-market goals.

Conclusion

A Titanium Technical Service application engineer can work at Xilinx, on-site, or a mix of both. This flexibility allows the engineers to fully understand the needs and requirements of our clients, as well as leverage Xilinx factory resources to resolve problems and accelerate production.

For more information about Titanium Technical Services, please call 800-888-FPGA (3742) or visit *http://support.xilinx.com/support/services/titanium.htm.* Σ

Maximize Your Success with Spartan-3 Designs

Xilinx Global Services helps you get the most from your Spartan-3 designs.

by Bill Okubo
Solutions Marketing Manager, Global Services
Xilinx, Inc.
bill.okubo@xilinx.com

Increasingly, FPGAs are being used as systems platforms because they are a more strategic purchase for customers who require a comprehensive solution that includes software, silicon, intellectual property (IP), and specialized services.

At the same time, FPGAs are expanding into new application areas such as digital signal processing (DSP), automotive, medical imaging, and video broadcast, where engineers new to FPGA technology need to learn quickly. Xilinx Global Services enables customers to reduce the learning curve and development costs associated with creating programmable system solutions using the Spartan™-3 family.

Xilinx Global Services Portfolio

Xilinx 90 nm Spartan-3 Platform FPGAs are reshaping the semiconductor landscape, and "on-demand" application engineers, global support centers, and personalized online support are some examples of how Xilinx is expanding its services to meet this evolving market trend. For Spartan-3 designs, the portfolio of support and services to meet these advanced requirements includes Education Services, Support and Titanium Technical Service, Xilinx Design Services, and the Xilinx Productivity Advantage Program.

Education Services

On any given business day, as many as six Xilinx Education Services classes may be taking place simultaneously at any number of cities around the world. Xilinx Education Services has new locations that expand our presence to 50 locations in North America, Europe, and Asia Pacific. In addition to public classes at Xilinx training facilities and regional training centers, we also offer private classes that can be taught either at your location or a nearby facility.

For Spartan-3 and ASIC designers, Xilinx offers several new and free technical lecture modules to help get you up to speed quickly:

- Overview of the Spartan-3 Family – Learn about the Spartan-3 architecture, its underlying technology and target markets, and software support features.
- ASIC User to FPGA Design – Choose from three modules that include a technology comparison between FPGAs and ASICs, design flow differences, and how to optimize ASIC code for implementation in Xilinx FPGAs.

For an in-depth curriculum on design methodology for the Spartan-3 family, Education Services has recently updated several key courses, including:

- Fundamentals of FPGA Design
- Designing for Performance for the FPGA User
- Designing for Performance for the ASIC User
- Advanced FPGA Design.

To help reduce your overall development and education costs, FPGA Essentials is a new promotional package that delivers both the Fundamentals of FPGA Design and the Designing for Performance courses at a savings of $200.

For course schedules or more information about Education Services for Spartan-3 FPGAs, visit *www.xilinx.com/education* and select the region where you would like to take a course.

Support and Titanium Technical Service

Support Tools and *MySupport.xilinx.com*

Support for software and IP tools can dramatically improve the success of your design projects. *Support.xilinx.com* reduces downtime with 24/7 online support tools such as:

- Answer Browser – Browse for information using the Answer Browser or use the powerful Advanced Search to pinpoint your exact search location by category and product name.
- Tech Tips – Learn tips and techniques from the experts through access to a categorical listing of tools, reference materials, and other information.
- Forums – A new Spartan-3 forum lets you solve design issues by collaborating online with other designers.
- Problem Solvers – Find answers to your design questions in a logical manner with this interactive tool that can save you hours of work.

For personalized support on the Spartan-3 family, you can enroll in *MySupport.xilinx.com* to receive automatic alerts on specific information you choose, such as updated information on data sheets, product change notices, and product discontinuation notices.

Visit *support.xilinx.com* to find out more about MySupport and the support tools available for Spartan-3 designs.

Titanium Technical Service

Titanium Technical Service provides a dedicated applications engineer who will supply expert advice on the best strategies to achieve timing closure and design fit in the shortest amount of time. The Titanium team of engineers has extensive knowledge of the new Spartan-3 architecture, 90 nm technology, and design techniques that will help you achieve your design goals for developing a low-cost FPGA. Titanium engineers are available on a contract basis and can work remotely or at your site. For more information about Titanium Technical Service, visit *support.xilinx.com/titanium/*.

Xilinx Design Services

Xilinx Design Services (XDS) combines skills and experience in system, logic, and embedded software designs to provide a unique development partner. Leveraging these skill sets will help you optimize your budget, schedule, and performance requirements for your project.

FPGA Design from Specification

XDS will develop designs from specification for projects such as ASIC conversions or driver development and integration. XDS has the expertise in optimizing Xilinx technology to provide the best solution, and excellent project management to ensure on-time and correct deliveries.

FPGA System Design

XDS has broad applications expertise – including an immense depth of experience with Xilinx embedded processing tools and products – for system architecture consulting, FPGA logic, and embedded software design projects.

Embedded Software Design

XDS also has tremendous experience with FPGA platform design, including processors and gates, MicroBlaze™ soft processors, and hardware/software co-design techniques. Projects such as complex embedded software designs with real-time constraints, driver development, and integration with hardware are areas where XDS can help.

For more information about Xilinx Design Services, visit *www.xilinx.com/xds/*.

The XPA Program

To help you use Spartan-3 FPGAs to their maximum potential, we offer everything you need – software, education and support services, and IP cores – delivered in one comprehensive package called the Xilinx Productivity Advantage (XPA) Program. It allows you to get the tools and services you need, when you need them, without the worry of ordering each piece of the solution separately.

The XPA Program helps to reduce the overall cost of developing your designs by bundling Xilinx tools and services at the best value. Because the XPA Program is set up with a single purchase order, it cuts down on the paperwork required to equip each designer and saves valuable time. Our packaged solution ensures that your engineering team will have everything they need to begin designing immediately. For more information about the XPA Program, visit *support.xilinx.com/xpa/*.

Conclusion

Xilinx Global Services is committed to helping you succeed with Spartan-3 devices, the world's lowest cost FPGA family. Our complete portfolio of support and services provides you with ready access to the knowledge and expertise that you need. By taking advantage of Xilinx Global Services, you will learn faster and be able to get better performance from your Spartan-3 designs. For more information about Xilinx Global Services, visit *support.xilinx.com*, where you can learn about our support tools and MySupport, or select the "Education" or "Services" menus for more information on those topics.

It's a Programmable World

Programmable World 2003 workshops offer detailed instruction in tools, technologies, and methodologies of system designs.

by David Vornholt
AllianceEDA Program Manager
Xilinx, Inc.
david.vornholt@xilinx.com

Why spend your evenings and weekends wading through product datasheets and application notes? Time to knowledge is just as important as time to market.

In the international series of forums and workshops that make up Programmable World 2003, Xilinx and its partners have created dozens of free technical seminars designed to stir your imagination, accelerate your design cycle, and reduce the overall design cost.

An Ecosystem of Leaders

The driving concept behind Programmable World 2003 has been the recognition that we do not work in a vacuum. Recognizing this, Xilinx has partnered with such industry leaders as IBM™, Agilent Technologies™, Cadence™ Design Systems, Wind River™ Systems, and Mentor Graphics®. The synergy of Xilinx and our partners plays an inherent role in the success of your designs.

The definition and implementation of your design involves support from not just your IC supplier but also from your EDA, embedded tool, IP, and design service partners. This is our "ecosystem" of companies that want to support your Xilinx-based design.

Cost-Sensitive Designs

The accessibility and usability of cutting-edge technologies will be demonstrated at Programmable World.

For example, the Xilinx Spartan™-3 family is based on IBM and UMC's advanced 90 nm, 8-layer metal process technology. Xilinx is using 90 nm technology to drive high-volume pricing down to under $20 per unit by the end of 2004 for one-million-gate FPGAs. This represents a cost savings of up to 80 percent compared to competitive offerings.

Programmable World Forums

The Programmable World forums, which were held earlier this year, comprised high-level discussions and demonstrations from Xilinx and the partners mentioned above. Other forum participants were Boeing™, Intel™, and Texas Instruments™, among others.

Attendees included hardware engineers, ASIC designers, system engineers, system architects, executives, and researchers from more than 27 industries and government agencies, ranging from home networking to the military.

In the forums, panelists and presenters explored the end-use possibilities and real-world design challenges of using multi-gigabit I/Os, embedded processors, low-cost platforms, and DSP solutions:

- IBM discussed the design methods associated with the integration of new processor architectures in FPGA-based systems.
- Cadence Design Systems and Xilinx described the challenges – and solutions – when upgrading existing systems for higher bandwidth.
- Agilent Technologies addressed the growing difficulties of system verification and hardware/software debugging.
- Xilinx presented strategies for solving the connectivity challenges of integrating ASSPs, network processors, and memory devices.
- Xilinx described the use of processor cores to combine the strengths of both the processor and logic for optimum performance and cost.
- Xilinx examined technical trends in DSP and presented a case for new architectures and system-level design methodologies.
- Xilinx also demonstrated PCI Express™, RocketPHY™, MicroBlaze™/Spartan-3, and DSP system products. Agilent, Cadence,



Celoxica, IBM, Mentor Graphics, Nallatech, Synplicity®, and The MathWorks demonstrated their respective technologies.

- In a panel discussion entitled "What's Required to Make High-Speed Serial Design a Reality," panelists from Agilent, Cadence, IBM, Intel, Mentor Graphics, Texas Instruments, and Xilinx confirmed the impending proliferation of serial technology and systems. The panelists also highlighted the technologies they're creating in anticipation of high-speed serial design demands.

User Applications

The Programmable World forums also included presentations on actual programmable system designs:

- Christopher Musial, a technical fellow at Boeing-SVS, gave a presentation entitled "Programmable Systems Transform the Design and Implementation of Adaptive Optics Systems." This case study illustrated the design processes and issues related to the implementation of a Xilinx Platform FPGA in a real-time adaptive optics control system. As Musial explained, the project had stringent constraints for cost, performance, power, and maintainability – parameters challenging many designers today.
- "A Highly Scalable FPGA-Based Car Infotainment Platform Solution" was the topic of Dr. Karlheinz Weiss' presentation. Dr. Weiss, the director of electronic system development at Elektroniksystem und Logistik GmbH, described the planning, specification, and implementation phases of an automobile infotainment system using Virtex-II Pro™ and Spartan FPGAs for the processing platform.

This design integrated audio, navigation, telephone, Internet services, telematics services, and video.

For those of you who missed the forums component of Programmable World – as well as those who wish to relive the experience – many of the presentations are viewable through the video-on-demand feature at *www.xilinx.com/pw2003*.

Free "University Style" Workshops

Xilinx has also developed a second component to the Programmable World forums – a series of workshops from September to November 2003. After learning about the high-level concepts discussed at the forums, you now have the opportunity to drill down to the details of creating your particular design.

The workshops are organized into four major "tracks," each comprising six one-hour sessions. You can attend any of the 24 sessions.

- **Track A:** Designing Flexible and High-Performance DSP Systems. When does it make sense to use DSP processors or platform FPGAs individually, and when does it make sense to combine the unique signal processing capabilities of both? Example scenarios help illustrate how to best answer these design dilemmas. Also, industry-leading DSP solution providers including The MathWorks, Nallatech, Texas Instruments, and Xilinx explain how they can help designers outsmart Moore's law and dramatically reduce the development time and cost of high-performance DSP systems.
- **Track B:** Integrating High-Performance Embedded Processing Solutions. This track focuses on technologies of interest to embedded systems designers from companies such as Altium™, IBM, Wind River Systems, and Xilinx.
- **Track C:** Implementing Cost-Effective Connectivity Solutions. Seeking low-cost connectivity applications? Cadence, Intel, and Xilinx present examples of gigabit Ethernet applications; implementation of high-speed parallel and multi-gigabit serial interfaces for chip-to-chip and backplane communications; and designing with serial technology.
- **Track D:** Effective System Design for Programmable Systems. FPGAs provide a variety of flexible and powerful capabilities and have evolved into true system platforms. This track offers engineers new and innovative information from Agilent, Celoxica, Mentor Graphics, Synopsys™, Synplicity, and Xilinx on how to extract the most performance from your system designs.

Locations and Dates

The schedule for Programmable World 2003 workshops this year is:

September 2003	
San Jose, CA:	September 3
Woodland Hills, CA:	September 4
San Diego, CA:	September 5
Stuttgart, Germany:	September 9
Munich, Germany:	September 10
Stockholm, Sweden:	September 11
Helsinki, Finland:	September 12
Madrid, Spain:	September 15
Paris, France:	September 16
Eindhoven, Netherlands:	September 17
Leuven, Belgium:	September 18
Boston, MA:	September 19
Ottawa, Canada:	September 22
Toronto, Canada:	September 24
Dallas, TX:	September 25
Chicago, IL:	September 26
October 2003	
Shanghai, China:	October 20
Hsinchu, Taiwan:	October 22
Seoul, Korea:	October 24
November 2003	
Tokyo, Japan:	November 14

Conclusion

The Programmable World 2003 series of tradeshows, forums, and workshops can solve your time-to-knowledge headaches by giving you access to critical information from industry leaders in international venues.

The topics are focused on providing cost-effective solutions to meet your time-to-market and performance goals.

For more information or to register for workshops, visit *www.xilinx.com/pw2003/workshop/*.

Like the forum presentations, the first Programmable World workshop in San Jose will be available in the last quarter of 2003 as a video-on-demand.

A Brand New Creation!

SUNDANCE

Xilinx & FPGA TI DSP

Adding Flexibility to power

SMT365E DSP & FPGA Module
- 64 bits 600MHz TMS320C6416 DSP
- Xilinx Virtex II FPGA XC2V6000
- 256MB of SDRAM @100MHz and 8MB of FLASH
- Six 400Mbytes/s Sundance High-speed Bus
- Six ComPorts for InterDSP communications

SMT374 DSP & FPGA Module
- Two 225MHz TMS320C6713 Floating Point DSP
- Xilinx Virtex II FPGA XC2V2000-4
- 256MB of SDRAM @100MHz and 8MB of FLASH
- Six ComPorts and two 400MB/s SHB ports

SMT398 FPGA Module
- Virtex II range of FPGA devices
- Up to 8 million gates as a re-configurable resource
- 4MB of high speed ZBT SRAM,
- 2MB of high speed QDR-SRAM,
- External Clock & LVDS Interface
- Multiple FPGA/DSP support

SMT406 FPGA PMC Module
- Virtex-II FPGA from V4000 to V8000
- 128MB of 133Mhz SDRAM, accessible from FPGA and PCI bus
- 16MB, 200MHz DDR SRAM or QDR SRAM
- 32MB, 250MHz DDR SDRAM and 16MB of Flash
- 64 bits, 66MHz PCI interface
- 4 programmable clocks generating up to 500MHz

SUNDANCE DIGITAL SIGNAL PROCESSING INC. Tel: (775) 827-3103 USA

SUNDANCE MULTIPROCESSOR TECHNOLOGY LTD. Tel: +44 (0) 1494-793167 UK

SUNDANCE ITALIA S.R.L. Tel: +39 0185 91055 ITALY

Email: sales@sundance.com http://www.sundance.com

XILINX PARTNER YELLOW PAGES

Xilinx Alliance EDA Members and Products

Alatek
www.alatek.com
Hardware acceleration (HES)

Altium
www.altium.com/products/nvisage.htm
Complete PCB/FPGA design flow (nVisage) targeting Spartan™ Series FPGAs

Apache Design
www.apache-da.com/Products/nspice.htm
HSPICE®-compatible simulator using actual S-parameter data (NSPICE) for Virtex-II Pro™ MGT-based PCB design

Aplus
www.aplus.com
Architecture-specific synthesis and layout optimization for Virtex™-II series FPGA (PALACE™)

Aptix
www.aptix.com
ASIC emulation (Prototype Studio™)

Atrenta
www.atrenta.com
RTL Linter technology for Virtex-II series FPGAs (SpyGlass®)

Cadence Design Systems
www.cadence.com/feature/fpga_design.html
Virtex-II Pro MGT support for high-speed PCB design (SPECCTRAQuest®); NCSim family for FPGA design simulation; SPW for DSP data path development on FPGAs

Celoxica
www.celoxica.com
C-level design creation, compilation, and verification for Virtex-II and Virtex-II Pro FPGAs

Endeavor
www.endeav.com
Co-verification for Virtex-II Pro (CoSimple™)

EVE
www.eve-team.com
RTL-level Simulator Accelerator (ZEBU)

Forte Design
www.forteds.com
HLL compiler to RTL (Cynlib Tool Suite); timing specification and analysis tool (TimingDesigner)

Future Design Automation
www.future-da.com
ANSI-C algorithms to RTL models (SystemCenter)

Hierarchical Design Inc. (HDI)
www.hierdesign.com
Hierarchical floorplanning and analysis software for Virtex-II series FPGAs

Mentor Graphics
www.mentor.com
Virtex-II Pro MGT support for high-speed PCB design (HyperLynx™, ICX™, Tau™); complete FPGA design flow using the ModelSim® family of simulators and Leonardo-Spectrum™/Precision® synthesis; Seamless® co-verification support for Virtex-II Pro

Novilit
www.novilit.com
Design partitioning using ISE EDK (AnyWare) for Virtex-II series FPGAs

Product Acceleration Inc.
www.prodacc.com
FPGA symbol generation (LiveComponent™); automated pin assignment optimization; PCB layer count control (DesignF/X)

Simucad
www.simucad.com
Verilog™ simulator for FPGAs (Silos)

Synopsys
www.synopsys.com
Virtex-II Pro multigigabit transceiver support for high-speed PCB design (HSPICE®); FPGA Compiler II™ RTL synthesis for Virtex series FPGAs; VCS™(MX)/Scirocco™(MX) simulators; Formality® formal verification, LEDA® RTL linter, and PrimeTime™ static timing analysis; high-level design using CoCentric™ System C Compiler and CoCentric System Studio

Synplicity
www.synplicity.com
RTL synthesis (Synplify®), debugger technology (Identify®), and physical synthesis (Amplify®) for all Xilinx FPGAs

Translogic
www.translogiccorp.com
FPGA symbol generation for fast PCB design iterations (FPGA Connector™) for Virtex-II series FPGAs

Verplex
www.verplex.com
Formal verification technology for Virtex-II series FPGAs (Conformal™ LEC)

Xilinx AllianceCORE Third-Party IP Providers and Products

Alcatel Internetworking, Inc.
www.ind.alcatel.com/emac/
Ethernet media access controller cores

Amphion Semiconductor Ltd.
www.amphion.com
DSP, wireless communications, and multimedia cores

Barco-Silex
www.barco-silex.com
DSP cores, with a specialization in image-audio applications

CAST, Inc.
www.cast-inc.com
General-purpose IP for processors, peripherals, multimedia, networking, encryption, serial communications, and bus interfaces

Crossbow Technologies, Inc.
www.crossbowip.com
System Interconnect IP for on-chip and chip-to-chip multiprocessing

Deltatec S.A.
www.deltatec.be/
IP and services for digital imaging, DSP, multimedia, PCI, and datacom

Derivation Systems, Inc.
www.derivation.com
IP cores for high assurance hardware/software applications and embedded systems products; 32-bit Java processor core

Digital Communications Technologie
www.dctl.com
Embedded Java solutions

Digital Core Design
www.dcd.pl/
Microprocessor, microcontroller, floating point, and serial communication controller cores

Dolphin Integration
www.dolphin.fr/
Bus interface, processor, peripheral, and DSP cores; EDA software

Duma Video, Inc.
www.dumavideo.com
Digital video, AES/EBU audio, MPEG and DCT cores; real-time MPEG-2 HD and SD compression

eInfochips Pvt. Ltd.
www.einfochips.com
PCI, USB, 80186, multimedia, and wireless cores

Eureka Technology
www.eurekatech.com
Silicon-proven IP for CPU interface, PCI, PCMCIA, memory control, and other peripheral functions for SoC designs

GDA Technologies, Inc.
www.gdatech.com
Hypertransport, PCI, Ethernet, design services

GV & Associates, Inc.
www.gvassociates.com
Boards and expertise in DSP, processor, software, analog, and RF design

Helion Technology Limited
www.heliontech.com
Encryption and DSP algorithms; design services

iCoding Technology, Inc.
www.icoding.com
Turbo Code-related error correction products including system design

Intrinsyc, Inc.
www.intrinsyc.com
Single board embedded computers (FPGA + external processor); peripheral IP

Loarant Corporation
www.loarant.com
Proprietary 16-bit RISC CPU; embedded system design

MEET Ltd.
www.meet-electronics.com
IP cores dedicated to industrial servomotor control

Memec Core
www.memecdesign.com/xilinx/
Peripheral, bus interfaces, encryption, and communications cores

ModelWare, Inc.
www.modelware.com
Datacom/telecom cores including ATM, HDLC, POS, IMA, UTOPIA, and bridges

NewLogic GmbH
www.newlogic.com
Cores for wireless applications including Bluetooth and 802.11

Paxonet Communications
www.paxonet.com
Solutions for interworking metro networking technologies to SONET

Perigee, LLC
www.perigeellc.com
DSP and image processing cores; embedded firmware, GUI, and application software development

Pinpoint Solutions, Inc.
www.asic-design.com
Standards-based telecom and datacom IP

QualCore Logic, Inc.
www.virtualipgroup.com
Cores for microcontrollers, telecommunications, datacom, PC peripherals, and bus interface standards

RealFast Operating Systems AB
www.realfast.se/
Real-time systems and RTOS; operating system cores

Robert Bosch GmbH
www.can.bosch.com
Automotive, CAN

SOC Solutions, L.L.C.
www.socsolutions.com
Embedded µP and software; networking, wireless, audio, GPS, and handheld; peripheral cores

SysOnChip, Inc.
www.sysonchip.co.kr/
IP and products for CDMA Cellular/PCS/WLL modem, FEC, and Bluetooth

Telecom Italia Lab S.p.A.
www.idosoc.com/soc/viplibrary/
Cores for ATM, optical networking, and switched digital video broadcast equipment

Tensilica, Inc.
www.tensilica.com
Configurable processor cores and development tools

TurboConcept
www.turboconcept.com
Turbo Code forward error correction solutions

Wipro, Ltd.
www.wipro.com
IP building blocks for technologies like IEEE 1394, USB, Bluetooth and IEEE 802.11

Xylon d.o.o.
www.logicbricks.com
Human-machine interfaces and industrial communication controllers

Zuken, Inc.
www.zuken.co.jp/soc/
PCI2.2, 10/100 Ethernet, and Gigabit Ethernet cores

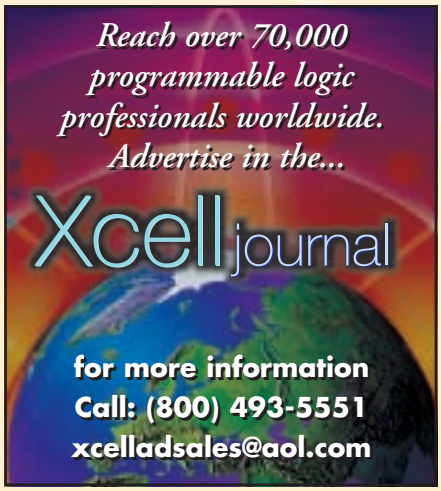

Reference Design Partners

The Xilinx Reference Design Alliance Program builds partnerships with industry-leading semiconductor manufacturers to develop reference designs for accelerating our customer's product and system time to market. The goal is to build a library of high-quality, multicomponent system-level reference designs in the areas of networking, communications, image and video processing, DSP, and emerging technologies and markets.

Agere Systems
www.agere.com

AMCC
www.amcc.com

Analog Devices
www.analog.com

Bay MicroSystems
www.baymicrosystems.com

BitBlitz Communications
www.bitblitz.com

Broadcom Corp.
www.broadcom.com

Centillium
www.centillium.com

Congnigine Corporation
www.cognigine.com

Cypress Semiconductor
www.cypress.com

IBM Microelectronics
www.ibm.com/us/

IDT
www.idt.com

Ignis Optics
www.ignis-optics.com

Infineon Technologies AG
www.infineontechnologies.com

Intel Corporation
developer.intel.com/design/network/

Intrinsity
www.intrinsity.com/home.htm

Micron Technology
www.micron.com

Motorola
www.motorola.com

Octasic Semiconductor
www.octasic.com

PMC-Sierra Inc.
www.pmc-sierra.com/index.html

SiberCore Technologies
www.sibercore.com

Teracross
www.teracross.com/~web/

Texas Instruments
www.ti.com

TransSwitch Coporation
www.transwitch.com/index_flash.jsp

Velio Communications, Inc.
www.velio.com

Vitesse
www.vitesse.com

West Bay Semiconductor Inc.
www.westbaysemi.com

ZettaCom
www.zettacom.com

Xanoptix
www.xanoptix.com

XPERTS Partners

3T BV
www.3T.nl/
Design services; measurement instruments; fail-safe controlling; medical equipment, communication, and computer peripherals

Accent S.r.l.
www.Accent.it/
Complex FPGAs; SoCs; PowerPCs; IP platforms; software, PCB, and IC foundry services, from concept to silicon

ADI Engineering
www.adiengineering.com
Reference designs; custom designs; design services focused on Intel IXA and XScale

Advanced Electronic Designs, Inc.
www.aedbozeman.com
Design services and manufacturing support; hardware and software design for embedded systems

Advanced Principles Group, Inc.
www.advancedprinciples.com
Design services and consulting; software algorithm acceleration and application-specific computing systems

Alpha Data
www.alpha-data.com
PCI/PMC/CPCI FPGA boards; embedded applications using HDL or System Generator methodologies

AMIRIX Systems Inc.
www.amirix.com
Design services for electronic product realization; high-performance digital boards; programmable logic; embedded software

Andraka Consulting Group, Inc.
www.andraka.com
Exclusively FPGAs for DSP since 1994; applications include radar, sonar, digital communications, imaging, and video

Array Electronics
www.array-electronics.de/
SoC/FPGA/PCB designs in telecom, industrial, and DSPs; concepts; verification; prototypes; cores

Avnet Design Services
www.ads.avnet.com
Embedded processing FPGA design services; Virtex-II Pro; high-speed SERDES; IP core integration; timing closure; RLDRAM; RapidIO™

Baranti Group Inc.
www.baranti.com
Award-winning design/manufacturing firm; expertise in digital video, FPGAs

Barco Silex
www.barco-silex.com
Total solutions consisting of IP, FPGA, and board design

Benchmark Electronics Inc.
www.bench.com
EMS/JDM team; product design, support, introduction, and launch into volume production

Birger Engineering, Inc.
www.birger.com
Development of algorithms, electronics, and systems for digital imaging

BitSim AB
www.bitsim.com
Design services; DSP, video, 3G, and WLAN designs; high-speed prototyping boards

Bottom Line Technologies Inc.
www.bltinc.com
FPGA, product, and system design services; networking; data communications; telecommunications; video; PCI; signal processing; military; aerospace

CES Design Services @ Siemens Program and System Engineering
www.ces-designservices.com
Platform-oriented SoC/FPGA/PCB and firmware design in telecom, industrial, and video processing; customized solutions; concept; verification; prototypes

CG-CoreEl Programmable Solutions P. Ltd.
www.cg-coreel.com
Design services; telecom, networking, and processor IP and designs; turnkey FPGA; RTL design/validation

Chess iT/Embedded
www.chess.nl/
Embedded hardware and software design of products and services at the chip, board, and system level

Colorado Electronic Product Design, Inc.
www.cepdinc.com
Designing embedded systems with FPGAs for DSP; interfaces for medical, consumer, aerospace, and military

Comit Systems, Inc.
www.comit.com
High-end FPGA/SoC design services, including the integration of customer-developed and third-party IP

ControlNet India Pvt Ltd.
www.controlnetindia.com
ASIC/SoC Analog/RF FPGA design services; IP development; domains: Ethernet, IEEE1394, USB, security, PCI, WLAN, Infiniband

Convergent Design LLC
www.convergent-design.com
Audio/video design for professional/consumer products

DATA Respons ASA
www.datarespons.com
Professional design and development services for embedded hardware/software solutions; FPGA/digital hardware application specialists

Deltatec International Inc.
www.deltatec.be/
Design services; digital imaging applications; broadcast video, image processing, and image synthesis

Digicom Answers Pty Ltd.
www.fpga.com.au/
Design services; experienced FPGA resources; design-on-demand; pre-design advice; reviews and problem solving

Digital Design Corporation (DDC)
www.digidescorp.com
Design services and products; image processing; video; communications; DSP; audio; automotive; controls; interfaces

Dillon Engineering, Inc.
www.dilloneng.com
Design services; FPGA-based DSP algorithms; high-bandwidth, real-time digital signal and image processing applications

DRS Tactical Systems West Inc. (formally Catalina Research)
www.catalinaresearch.com
Design services; FPGAs for DSP, including the manufacturer of Virtex reconfigurable computing boards

Eden Networks, Inc.
www.edennetworks.com
Turnkey design services; telecom and networking designs (Ethernet, SONET, ATM, VoIP)

Edgewood Technologies, LLC
www.edgewoodtechnologies.com
Design services; FPGAs for telecom, DSP, embedded processors

Enea Data AB
www.enea.se/
Telecommunication, data communication, DSP, RTOS, automotive, defense, high-reliability, and consumer electronics design

Enterpoint Ltd.
www.enterpoint.co.uk/
Design services; development boards; video, telecommunications, military, and high-speed design

ESD Inc.
www.esdnet.com
Design services; FPGAs for embedded systems, software/hardware design, PCB layout

Evatronix S.A.
www.evatronix.pl/
Design services; FPGA design (Virtex, Virtex-E, Spartan-II); embedded systems; SoC and SoPC designs

Evermore Systems Inc.
www.evermoresystems.com
Design services; FPGA designs for communications, audio/video, wireless, and home networking

ExaLinx
www.exalinx.com
Design services; ASIC emulation for high-speed communications and wireless applications; reference designs

Fidus Systems Inc.
www.fidus.ca/
Design services; FPGAs; communications; video; DSP; interface conversions; ASIC verification; hardware; PCB layout; software; SI/EMC

First Principles Technology, Inc.
www.fptek.com
Design/assembly/test services; FPGA + PCB + DSP/processor + system; imaging; video; test; satellite telecom; control

Flexibilis Oy
www.flexibilis.com
IP development; FPGA design services; Linux device drivers

Flextronics Design
www.flextronics.com/design/
Complete product design and development services for high-speed communications; reconfigurable computing; imaging; audio/video; military

GDA Technologies, Inc.
www.gdatech.com
A leading services company focused on designing systems, SoC, ASIC, FPGA, and IP

HCL Technologies
www.hcltech.com
High-speed board, ASIC, and FPGA design and verification services in avionics, medical, networking/telecom, and consumer electronics

Helion Technology Limited
www.heliontech.com
Design services and IP for FPGA with a focus on data security, wireless, and DSP

IDERS Inc.
www.iders.ca/
Contract electronic engineering and EMS resource, providing complete product cycles ranging from specifications to production

Image Technology Laboratory Corporation
www.gazogiken.co.jp/
Design services; FPGAs for image processing; design tools: EDK, System Generator, and Forge

INITEC AG
www.initec.ch/
Design services for system design; FPGA applications in telecommunication systems and medical electronics

Innovative Computer Technology
ictconn.com
Digital, analog, Xilinx FPGA, and software design solutions for customers in a wide variety of industries throughout the United States

iWave Systems Technologies Private Ltd.
www.iwavesystems.com
Embedded hardware and software design; board design; FPGA and ASIC development services

Liewenthal Electronics Ltd.
www.liewenthal.ee/
Embedded systems design services; development of software and hardware including Virtex and Spartan FPGAs

Logic Product Development
www.logicpd.com
Full-service product development, including industrial design, mechanical, electrical, and software engineering

M&M Consulting
hellman@mmcons.com
Design services; imaging; video and audio designs; high-density ASIC emulation

Memec Design

www.memecdesign.com/xilinx

Focused on FPGA design for Xilinx technology, we have completed more than 1,300 designs worldwide

Memondo Graphics
www.memondo.com
Implementation of digital signal processing, communications, and image processing solutions and algorithms over Xilinx FPGAs

Mikrokrets AS
www.mikrokrets.no/
FPGA development comprising telecommunication, network communications, and embedded systems; PCB development

Millogic Ltd.
www.millogic.com
Design services and synthesizable cores; expertise in PCI, video, imaging, DSP, compression, ASIC verification, and communications

Milstar Inc.
www.Milstar.co.il/
Design services and manufacturers (industrial and military specs); DSP; communication; video; audio; high-speed boards

Misarc srl - Agrate
www.misarc.com
Design services and manufacturers; telecommunications, SoC development, and test equipment boards using FPGA solutions

Multi Video Designs
www.mvd-fpga.com
Design services; DSP for video, telecommuications; SoC with Virtex-II/Pro and Spartan (PowerPC/MicroBlaze™); Xilinx training

Multiple Access Communications Ltd.
www.macltd.com
Design services; products and consultancy; DSP for wireless communications

Nallatech
www.nallatech.com
Embedded reconfigurable computers for high-performance applications in DSP, imaging, and defense; Xilinx Diamond XPERTS

NetQuest Corporation
www.netquestcorp.com
Turnkey engineering design and production services in advanced high-speed embedded communications

North Pole Engineering Inc.
www.northpoleengineering.com
Hardware and software design services for FPGA, ASIC, and embedded systems

Northwest Logic
www.nwlogic.com
PCI, PCI-X, and SDRAM controller IP; high-end FPGA design services

NovTech Engineering
www.novtech.com
Board-level and IP cores; turnkey design services; imaging; communication; PCI cores; encryption

NUVATION
www.nuvation.com
Design services; imaging and communications for defense/security, medical, consumer, and datacom/telecom markets

Plextek Ltd.
www.plextek.com
Design services; FPGA, DSP, radio, microwave, and microprocessor technologies for defense and communications

POLAR-Design
www.polar-design.de/
High-end FPGA design for telecom, networking, and DSP applications

Polybus Systems Corporation
www.polybus.com
Test patterns; customizable in-system FPGA test patterns; design services; networking and communications

Presco, Inc.
www.prescoinc.com
Twenty-five years of experience in custom electronics design/manufacturing for high-performance image processing, data communications, inkjet

Productivity Engineering GmbH
www.pe-gmbh.com
Design services; PCI design; microprocessor cores; obsolete component replacements; ASIC prototyping

Programall Technologies Inc.
www.ProgramallTechnologies.com
Design services; embedded systems targeting Virtex, Spartan, Virtex-II, and Virtex-II Pro FPGAs

R&D Consulting
972-9-7673074
FPGA development; video; imaging; graphics; DSP; communications

Rapid Prototypes, Inc.
www.FPGA.com
Design services; high-performance reconfigurable FPGA applications requiring maximum speed or density using physical design principles

RDLABS
www.rdlabs.com
FPGA/ASIC design; wireless communications; 802.11b/a IP cores; Matlab, Sysgen, ModelSim; instruments HP1661A, HP8594E, HP54520A

RightHand Technologies, Inc.
www.righthandtech.com
Design services; Virtex-II Pro FPGAs, boards, and software for video gaming, storage, medical, and wireless

Rising Edge Technology Inc.
www.risingedgetech.com
Complete design solutions from the ground up; FPGA interface designs

Riverside Machines Ltd.
www.riverside-machines.com
Design services and project management; DSP, wireless, network, and processor development; Verilog, VHDL, SystemC, Specman

Roke Manor Research
www.roke.co.uk/
Design services; IP, ATM, DSP, imaging, radar, and control solutions for FPGAs

Roman-Jones, Inc.
www.roman-jones.com
Design services; FPGAs for embedded microprocessors; PCI; DSP; low-cost 8051 softcore

Sapphire Computers, Inc.
drudolf@voyager.net
Engineering design consulting; high-speed digital and FPGA design

Siemens AG, Electronic Design House
www.eda-services.com
Design services; system-on-chip; ASICs; FPGAs; IPs; DSPs; boards; EMC; embedded software; prototypes

Silicon & Software Systems Ltd. (S3)
www.s3group.com
World-class electronics design company uniquely combining IC, FPGA, software, and hardware design expertise

Silicon Interfaces America Inc.
www.siliconinterfaces.com
Design services; develops IP in areas of networking, wireless, optical, datacom, interconnect, and microcontrollers

Silicon System Solutions P/L
www.silicon-systems.com
Design services; FPGAs for very high-speed applications including communications and DSP applications

Siscad S.p.A.
www.siscad.it/
Xilinx FPGA design services; IP development and integration

Smart Logic, Inc.
www.smart-logic.com
Design services; FPGAs for rapid prototyping including imaging, compression, and real-time control

SoleNet, Inc.
www.solenet.net
FPGA, board, and system designs for wireless, telecommunications, consumer, and reference design applications

so-logic electronic consulting
www.so-logic.net
Embedded systems (PowerPC, MicroBlaze); system-level design (Forge, System Generator, Handel-C); Xilinx training centers (Austria, Eastern Europe)

Styrex AB
www.styrex.se/
Design services for embedded systems; programmable logic, hardware development, and microprocessor programming

Synopsys Professional Services
www.synopsys.com
Design services to solve your design challenges in system-level design, RTL design, and physical design

Syntera AB
www.syntera.se/
Design services; FPGAs for radar, digital communications, telecom automotive, aerospace, and DSP

Tachyon Semiconductor
www.tachyonsemi.com
Design services for portable computing, networking, and embedded systems

Tao of Digital, Inc.
www.taoofdigital.com
High-speed, large-gate-count SoC designs using IP cores for maximum modularity and efficiency

Technolution B.V.
www.technolution.nl/
Innovative hardware and software solutions

Thales Airborne Systems
www.thalesgroup.com/airbornesystems
Provider and integrator with high expertise in FPGA design for radars, DSP, and communications

TietoEnator R&D Services AB
www.TietoEnator.com
European design services for telecom, industrial IT, and automotive; Xilinx exclusive training provider in Scandinavia

V-Integration
www.V-Integration.com
Providing first-pass success in complex FPGA designs; applications include imaging, communications, and interface design

Williams Consulting, Inc.
chipw@wciatl.com
Embedded systems hardware and software; video; MPEG; control

Zaiq Technologies, Inc.
www.zaiqtech.com
Design; partitioning and timing convergence; VIP; platform experts for embedded computer, networking, wireless, and instrumentation

How Would You Use 16 Million Gates?

Two Xilinx XC2V8000 FPGAs per Board
- Six 2Mx18-bit QDR SRAMs per FPGA
- Over 200 direct off-board digital I/O links
- ChipScope Pro headers
- 128Mbytes SDRAM per FPGA/DSP

Two TigerSHARC DSPs
- Floating-Point DSP
- Inter-DSP and FPGA data links
- 4Mbytes FLASH per DSP/FPGA

CompactPCI
- 32/64-bit PCI
- Hotswap support

Software & Tools
- FFT, radar receiver, FIR, QR, Floating-point IP cores
- Loader & diagnostic utilities
- VHDL examples
- Low level data communications libraries
- VxWorks drivers

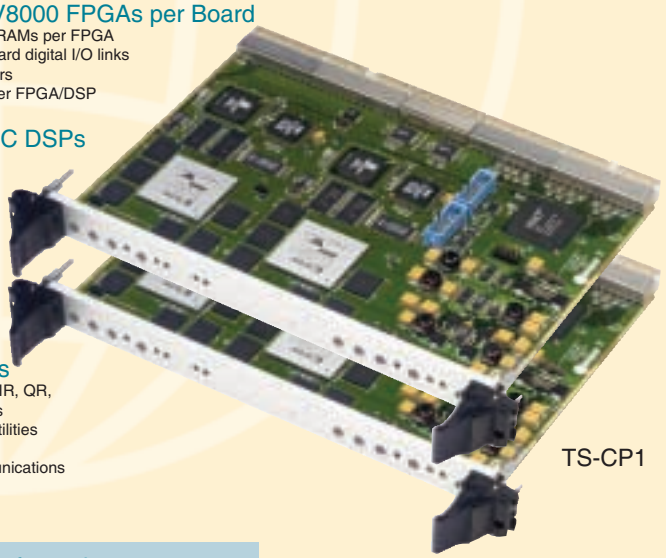

TS-CP1

Contact us for Information on:
PMC, PCI, VME and Virtex-II Pro Solutions

For More Information Call
USA 607 257 8678
Europe/Asia +44 (0)1494 464432
Israel +972 3 9518718

TRANSTECH DSP™
www.transtech-dsp.com/fpgas.htm
sales@transtech-dsp.com

Coming Soon
to a location near you!

X-fest 2003

From October through December

Learn how the latest Xilinx technology can help you design cost effective solutions faster.

Gain hands-on experience to speed your next development cycle.

What Are You Waiting For?
Register now for the event nearest you.

LOCATIONS

Albuquerque	Colorado Springs	Huntsville	Philadelphia	St. Louis
Atlanta	Connecticut	Irvine	Phoenix	Syracuse
Austin	Dallas	Kansas City	Portland	Tallahassee
Baltimore	Dayton	Los Angeles	Rochester	Toronto
Boston	Denver	Minneapolis	Salt Lake City	Vancouver
Calgary	Detroit	Montreal	San Diego	
Chicago	Ft. Lauderdale	Orlando	San Jose	
Cleveland	Houston	Ottawa	Seattle	

Insight MEMEC XILINX®

Call 888.488.4133 (x915) or register at
www.insight.na.memec.com/x-fest2003

Product Selection Matrix

		System Gates	Macrocells	Product terms per Macrocell	Input Voltage Compatible	Output Voltage Compatible	I/O Features: Maximum I/O	I/O Features: I/O Banking	Speed: Min. Pin-to-pin Logic Delay (ns)	Speed: Commercial Speed Grades (fastest to slowest)	Speed: Industrial Speed Grades (fastest to slowest)	Speed: IQ Speed Grade	Clocking: Global Clocks	Clocking: Product Term Clocks per Function Block
	CoolRunner-II Family – 1.8 Volt													
CoolRunner-II	XC2C32	750	32	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	33	1	3	-3 -4 -6	-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	-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	-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	-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	-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	-10	3	17
	CoolRunner XPLA3 Family – 3.3 Volt													
CoolRunner XPLA3	XCR3032XL	750	32	48	3.3/5	3.3	36		5	-5 -7 -10	-7 -10	-10	4	16
	XCR3064XL	1500	64	48	3.3/5	3.3	68		6	-6 -7 -10	-7 -10	-10	4	16
	XCR3128XL	3000	128	40	3.3/5	3.3	108		6	-6 -7 -10	-7 -10	-10	4	16
	XCR3256XL	6000	256	48	3.3/5	3.3	164		7.5	-7 -10 -12	-10 -12	-12	4	16
	XCR3384XL	9000	384	48	3.3/5	3.3	220		7.5	-7 -10 -12	-10 -12	-12	4	16
	XCR3512XL	12000	512	48	3.3/5	3.3	260		7.5	-7 -10 -12	-10 -12	-12	4	16
	XC9500XV Family – 2.5 Volt													
XC9500XV	XC9536XV	800	36	90	2.5/3.3	1.8/2.5/3.3	36	1	5	-5 -7	-7	NA	3	18
	XC9572XV	1600	72	90	2.5/3.3	1.8/2.5/3.3	72	1	5	-5 -7	-7	NA	3	18
	XC95144XV	3200	144	90	2.5/3.3	1.8/2.5/3.3	117	2	5	-5 -7	-7	NA	3	18
	XC95288XV	6400	288	90	2.5/3.3	1.8/2.5/3.3	192	4	6	-6 -7 -10	-7 -10	NA	3	18
	XC9500XL Family – 3.3 Volt													
XC9500XL	XC9536XL	800	36	90	2.5/3.3/5	2.5/3.3	36		5	-5 -7 -10	-7 -10	-10	3	18
	XC9572XL	1600	72	90	2.5/3.3/5	2.5/3.3	72		5	-5 -7 -10	-7 -10	-10	3	18
	XC95144XL	3200	144	90	2.5/3.3/5	2.5/3.3	117		5	-5 -7 -10	-7 -10	NA	3	18
	XC95288XL	6400	288	90	2.5/3.3/5	2.5/3.3	192		6	-6 -7 -10	-7 -10	NA	3	18

Package Options and User I/O

Pins	Body Size	CoolRunner-II: XC2C32	XC2C64	XC2C128	XC2C256	XC2C384	XC2C512	CoolRunner XPLA3: XCR3032XL	XCR3064XL	XCR3128XL	XCR3256XL	XCR3384XL	XCR3512XL	XC9500XV: XC9536XV	XC9572XV	XC95144XV	XC95288XV	XC9500XL: XC9536XL	XC9572XL	XC95144XL	XC95288XL
PLCC Packages (PC) – wire-bond plastic chip carrier (1.27mm lead spacing)																					
44	16.5 x 16.5 mm	33	33					36	36					34	34			34	34		
PQFP Packages (PQ) – wire-bond plastic QFP (0.5mm lead spacing)																					
208	28 x 28 mm				173	173	173				164	172	180				168				168
VQFP Packages (VQ) – very thin TQFP (0.5mm lead spacing)																					
44	10 x 10 mm	33	33					36	36					34	34			34	34		
64	10 x 10 mm																	36	52		
100	14 x 14 mm		64	80	80				68	84											
TQFP Packages (TQ) – thin QFP (0.5mm lead spacing)																					
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 7 mm							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 Fine-line 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 in this document with the device data sheet at: http://www.xilinx.com/partinfo/databook.htm.

Automotive products are highlighted: -40C to +125C ambient temperature for CPLDs

Xilinx IQ Solutions for Automotive Intelligence

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

Product Selection Matrix

			CLB Resources				Memory Resources			DSP	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	Maximum I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Platform Flash PROM	System ACE	Configuration Memory (Bits)	RocketIO™ Transceiver Blocks	PowerPC Processor Blocks	Virtex-II Series EasyPath Solution (see note 4)
Virtex-II Pro Family – 1.5 Volt																								
Platform FPGAs	VIRTEX-II PRO	XC2VP2	*	16 x 22	1,408	3,168	2816	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	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[3]	*	80 x 46	13,696	30,816	27,392	428	136	2448	136	24/420	8	YES	372	644		-5 -6 -7	-5 -6			11.36M	8	2	✔
		XC2VP40[3]	*	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[3]	*	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[3]	*	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[3]	*	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[3]	*	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	2	✔
Virtex-II Family – 1.5 Volt																								
	VIRTEX-II	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	ISP	0.4M	–	–	
		XC2V80	80K	16 x 8	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 x 16	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[3]	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[3]	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[3]	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[3]	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	-4 -5			29.1M	–	–	✔

* Logic cell counts are a more meaningful measurement of density for the Virtex-II Pro family since system gate count does not take into consideration the benefits of the immersed special blocks such as PowerPC processors and multi-gigabit transceivers.
** RocketIO unavailable in this package.

Notes: 1. System Gates include 20-30% of CLBs used as RAM
2. Logic cell = One 4-Input Look Up Table (LUT) + 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

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

Package Options and User I/O

		Virtex-II Pro (1.5V)									
		XC2VP2	XC2VP4	XC2VP7	XC2VP20	XC2VP30	XC2VP40	XC2VP50	XC2VP70	XC2VP100	XC2VP125
Pins	Body Size	204	348	396	564	692	804	852	996	1164	1200
Chip Scale Packages (CS) – wire-bond chip-scale BGA (0.8 mm ball spacing)											
144	12 x 12 mm										
BGA Packages (BG) – wire-bond standard BGA (1.27 mm ball spacing)											
575	31 x 31 mm										
728	35 x 35 mm										
FGA Packages (FG) – wire-bond fine-pitch BGA (1.0 mm ball spacing)											
256	17 x 17 mm	140	140								
456	23 x 23 mm	156	248	248							
676	27 x 27 mm				404	416	416				
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					
1152	35 x 35 mm				564	644	692	692			
1148*	35 x 35 mm						804	812			
1517	40 x 40 mm							852	964		
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										

		Virtex-II (1.5V)									
		XC2V40	XC2V80	XC2V250	XC2V500	XC2V1000	XC2V1500	XC2V2000	XC2V3000	XC2V4000	XC2V6000	XC2V8000
Pins	Body Size	88	120	200	264	432	528	624	720	912	1104	1296
Chip Scale Packages (CS) – 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	88	120	172	172	172						
456	23 x 23 mm			200	264	324						
676	27 x 27 mm						392	456	484			
FFA Packages (FF) – flip-chip fine-pitch BGA (1.0 mm ball spacing)												
672	27 x 27 mm											
896	31 x 31 mm					432	528	624				
1152	35 x 35 mm								720	824	824	824
1148*	35 x 35 mm											
1517	40 x 40 mm									912	1104	1108
1704	42.5 x 42.5 mm											
1696*	42.5 x 42.5 mm											
BFA Packages (BF) – flip-chip fine-pitch BGA (1.27 mm ball spacing)												
957	40 x 40 mm							624	684	684	684	

	Virtex-II Pro Packages 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				
FG672	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*

*RocketIO unavailable in this package.

Note: The FF package is a Fine Pitch Flip BGA Chip (1.00mm ball spacing).

Notes: 1. Numbers indicated in the Package Options and User I/O matrices indicate the maximum number of user I/Os for that package and device combination.
2. The number of I/Os for RocketIO MGTs are not included in this table.
3. Within the same family, all devices in a particular package are pin-out (footprint) compatible.
4. Virtex-II packages FG456 and FG676 are also footprint compatible.
5. Virtex-II packages FF896 and FF152 are also footprint compatible.

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

Product Selection Matrix

		CLB Resources					Memory Resources			DSP	CLK Resources				I/O Features				Speed		PROM	
		System Gates (see note 1)	CLB Array (Row x Col)	Number of Slices	Logic Cells (see note 2)	CLB Flip-Flops	Max. Distributed RAM Bits	# Block RAM	Block RAM (kbits)	Dedicated Multipliers	DCM Frequency (min/max)	# DCMs	Frequecny Synthesis	Phase Shift	Digitally Controlled Impedance	Number of Differential I/O Pairs	Maximum I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Platform Flash PROM	Configuration Memory (Bits)
	Spartan-3 Family – 1.2 Volt																					
SPARTAN-3 FPGA	**XC3S50**	50K	16 x 12	768	1,728	1,536	12K	4	72K	4	25/326	2	YES	YES	YES	56	124	Single-ended LVTTL, LVCMOS3.3/2.5/1.8/ 1.5/1.2, PCI 3.3V – 32/64-bit 33MHz, SSTL2 Class I & II, SSTL18 Class I, HSTL Class I, III, HSTL1.8 Class I, II & III, GTL, GTL+ Differential LVDS2.5, Bus LVDS2.5, Ultra LVDS2.5, LVDS_ext2.5, RSDS, LDT2.5	-4 -5	-4	XCF01S	.4M
	XC3S200	200K	24 x 20	1,920	4,320	3,840	30K	12	216K	12	25/326	4	YES	YES	YES	76	173		-4 -5	-4	XCF01S	1.0M
	XC3S400	400K	32 x 28	3,584	8,064	7,168	56K	16	288K	16	25/326	4	YES	YES	YES	116	264		-4 -5	-4	XCF02S	1.7M
	XC3S1000	1000K	48 x 40	7,680	17,280	15,360	120K	24	432K	24	25/326	4	YES	YES	YES	175	391		-4 -5	-4	XCF04S	3.2M
	XC3S1500	1500K	64 x 52	13,312	29,952	26,624	208K	32	576K	32	25/326	4	YES	YES	YES	221	487		-4 -5	-4	XCF08P	5.2M
	XC3S2000	2000K	80 x 64	20,480	46,080	40,960	320K	40	720K	40	25/326	4	YES	YES	YES	270	565		-4 -5	-4	XCF08P	7.7M
	XC3S4000	4000K	96 x 72	27,648	62,208	55,296	432K	96	1,728K	96	25/326	4	YES	YES	YES	312	712		-4 -5	-4	XCF16P	11.3M
	XC3S5000	5000K	104 x 80	33,280	74,480	66,560	520K	104	1,872K	104	25/326	4	YES	YES	YES	344	784		-4 -5	-4	XCF16P	13.3M

Note: 1. System Gates include 20-30% of CLBs used as RAMs
2. Logic cell is defined as a 4 input LUT + Flip-Flop + Carry Logic

Package Options and User I/O

		Spartan-3								
			XC3S50	XC3S200	XC3S400	XC3S1000	XC3S1500	XC3S2000	XC3S4000	XC3S5000
Pins	**Body Size**	**I/O's**	**124**	**173**	**264**	**391**	**487**	**565**	**712**	**784**
VQFP Packages (VQ) – very thin TQFP (0.5 mm lead spacing)										
100	**14 x 14 mm**		63	63						
TQFP Packages (TQ) – thin QFP (0.5 mm lead spacing)										
144	**20 x 20 mm**		97	97	97					
PQFP Packages (PQ) – wire-bond plastic QFP (0.5 mm lead spacing)										
208	**28 x 28 mm**		124	141	141					
FGA Packages (FT) – wire-bond thin BGA (1.0 mm ball spacing)										
256	**17 x 17 mm**			173	173	173				
FGA Packages (FG) – wire-bond fine-pitch BGA (1.0 mm ball spacing)										
456	**23 x 23 mm**				264	333	333			
676	**27 x 27 mm**					391	487	489		
900	**31 x 31 mm**							565	633	633
1156	**35 x 35 mm**								712	784

Note: Numbers indicated in the Package Options and User I/O matrix indicate the maximum number of user I/Os for that package and device combination.

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

XILINX®
The Programmable Logic Company℠

XILINX SPARTAN FPGAs

http://www.xilinx.com/devices/

Product Selection Matrix

	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)	DLL Frequency (min/max)	# DLLs	Frequecny Synthesis	Phase Shift	Number of Differential I/O Pairs	Maximum I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grade (slowest to fastest)	IQ Speed Grade	Platform Flash PROM	Configuration Memory (Bits)
	CLB Resources					Memory Resources			CLK Resources				I/O Features			Speed				
Spartan-IIE Family – 1.8 Volt																				
XC2S50E	50K	16 x 24	768	1,728	1,536	24K	8	32K	25/320	4	YES	YES	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	-6	ISP	0.6M
XC2S100E	100K	20 x 30	1,200	2,700	2,400	37K	10	40K	25/320	4	YES	YES	86	202		-6 -7	-6	-6		0.9M
XC2S150E	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	25/320	4	YES	YES	114	265		-6 -7	-6	-6		1.1M
XC2S200E	200K	28 x 42	2,352	5,292	4,704	73K	14	56K	25/320	4	YES	YES	120	289		-6 -7	-6	-6		1.4M
XC2S300E	300K	32 x 48	3,072	6,912	6,144	96K	16	64K	25/320	4	YES	YES	120	329		-6 -7	-6	-6		1.9M
XC2S400E	400K	40 x 60	4,800	10,800	9,600	150K	40	160K	25/320	4	YES	YES	172	410		-6 -7	-6	-6		2.7M
XC2S600E	600K	48 x 72	6,912	15,552	13,824	216K	72	288K	25/320	4	YES	YES	205	514		-6 -7	-6	-6		4.0M
Spartan-II Family – 2.5 Volt																				
XC2S15	15K	8 x 12	192	432	384	6K	4	16K	25/200	4	YES	YES	NA	86	LVTTL, LVCMOS2, PCI33 (3.3V & 5V), PC166 (3.3V), GTL, GTL+, HSTL I, HSTL III, HSTL IV, SSTL3 I, SSTL3 II, SSTL2 I, SSTL2 II, AGP-2X, CTT	-5 -6	-5	-5	ISP	0.2M
XC2S30	30K	12 x 18	432	972	864	13.5K	6	24K	25/200	4	YES	YES	NA	132		-5 -6	-5	-5		0.4M
XC2S50	50K	16 x 24	768	1,728	1,536	24K	8	32K	25/200	4	YES	YES	NA	176		-5 -6	-5	-5		0.6M
XC2S100	100K	20 x 30	1,200	2,700	2,400	37.5K	10	40K	25/200	4	YES	YES	NA	196		-5 -6	-5	-5		0.8M
XC2S150	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	25/200	4	YES	YES	NA	260		-5 -6	-5	-5		1.1M
XC2S200	200K	28 x 42	2,352	5,292	4,704	73.5K	14	56K	25/200	4	YES	YES	NA	284		-5 -6	-5	-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	77		-4 -5	-4	-4	ISP	0.05M
XCS10XL	10K	14 x 14	196	466	392	6.1K	NA	NA	NA	NA	NA	NA	NA	112		-4 -5	-4	-4		0.09M
XCS20XL	20K	20 x 20	400	950	800	12.5K	NA	NA	NA	NA	NA	NA	NA	160		-4 -5	-4	-4		0.18M
XCS30XL	30K	24 x 24	576	1,368	1,152	18.0K	NA	NA	NA	NA	NA	NA	NA	192		-4 -5	-4	-4		0.25M
XCS40XL	40K	28 x 28	784	1,862	1,568	24.5K	NA	NA	NA	NA	NA	NA	NA	224		-4 -5	-4	-4		0.33M

Note: 1. System Gates include 20-30% of CLBs used as RAM
2. Logic cell = (1) 4 Input (LUT) and a register

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 (PC) – wire-bond plastic chip carrier (1.27mm lead spacing)																				
84	29.3 x 29.3 mm															61	61			
PQFP Packages (PQ) – wire-bond plastic QFP (0.5mm lead spacing)																				
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) – very thin TQFP (0.5mm lead spacing)																				
100	14 x 14 mm									60	60					77	77	77	77	
TQFP Packages (TQ) – thin QFP (0.5mm lead spacing)																				
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 (BG) – wire-bond standard BGA (1.0 mm ball spacing)																				
256	27 x 27 mm																		192	205

Note: Numbers indicated in the Package Options and User I/O matrix indicate the maximum number of user I/O for that package and device combination

Automotive products are highlighted: -40C to +125C junction temperature for FPGAs. Xilinx IQ Solutions for Automotive Intelligence

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

QPRO

Device	SMD	Voltage	System Gates	Logic Cells	RAMbus	CFG Bits	MULTs	DLLs	Flip-Flops	Max I/Os	Packages
XC3020*	5962-89948	5	1500	256	–	14779	–	–	256	64	PG84, CB100
XC3030*	None	5	2000	360	–	22176	–	–	360	80	PG84
XC3042*	5962-89713	5	3000	480	–	30784	–	–	480	96	PG84, PG132, CB100
XC3064*	None	5	4500	688	–	46064	–	–	688	120	PG132
XC3090*	5962-89823	5	6000	928	–	64160	–	–	928	144	PG175, CB164
XC4005*	5962-92252	5	9K	466	6272	151960	–	–	616	112	PG156, CB164
XC4010*	5962-92305	5	20K	950	12800	283424	–	–	1120	160	PG191, CB196
XC4013*	5962-94730	5	30K	1368	18438	393632	–	–	1536	192	PG223, CB228
XQ4005E	5962-97522	5	9K	466	6272	95008	–	–	616	112	PG156, CB164
XQ4010E	5962-97523	5	20K	950	12800	178144	–	–	1120	160	HQ208, PG191, CB196
XQ4013E	5962-97524	5	30K	1368	18438	247968	–	–	1536	192	HQ240, PG223, CB228
XQ4025E	5962-97525	5	45K	2432	32768	422176	–	–	2560	256	PG299, CB228
XQ4028EX	5962-98509	5	18K-50K	2432	32768	668184	–	–	2560	256	HQ240, BG352, PG299, CB228
XQ4013XL	5962-98513	3.3	10-30K	1368	18K	393632	–	–	1536	192	PQ240, BG256, PG223, CB228
XQ4036XL	5962-98510	3.3	22-65K	3078	42K	832528	–	–	3168	288	HQ240, BG352, PG411, CB228
XQ4062XL	5962-98511	3.3	40-130K	5472	74K	1433864	–	–	5376	384	HQ240, BG432, PG475, CB228
XQ4085XL	None	3.3	55-180K	7448	100K	1924992	–	–	7185	448	HQ240, BG432, CB228
XQV100	None	2.5	108904	2700	40K	781248	–	4	2400	180	PQ240, BG256, CB228
XQV300	5962-99572	2.5	322970	6912	64K	1751840	–	4	6144	316	PQ240, BG352, BG432, CB228
XQV600	5962-99573	2.5	661111	15552	96K	3608000	–	4	13824	512	HQ240, BG432, CB228
XQV1000	5962-99574	2.5	1124022	27648	128K	6127776	–	4	24576	512	BG560, CG560
XQV600E	None	1.8	986K	15552	288K	3961632	–	8	13824	512	BG432, CB228
XQV1000E	None	1.8	1569K	27648	384K	6587520	–	8	24576	660	BG560, CG560
XQV2000E	None	1.8	2542K	43200	640K	10159648	–	8	38400	804	BG560, FG1156
XQ2V1000	TBD	1.5	1000K	11520	720	4082592	40	8	10240	432	FG456, BG575
XQ2V3000	TBD	1.5	3000K	32256	1728	10494368	96	12	28672	720	BG728, CG717
XQ2V6000	TBD	1.5	6000K	76032	2592	21849504	144	12	67584	1104	FC1144

* Not for new designs

QPRO Radiation Tolerant

Device	Voltage	System Gates	Logic Cells	RAMbus	CFG Bits	MULTs	DLLs	Flip-Flops	Max I/Os	Packages	Total Ionizing Dose (TID) in KRADs
XQR4013XL	3.3	10-30K	1368	18K	393632	–	–	1536	192	CB228	60
XQR4036XL	3.3	22-65K	3078	42K	832528	–	–	3168	288	CB228	60
XQR4062XL	3.3	40-130K	5472	74K	1433864	–	–	5376	384	CB228	60
XQVR300	2.5	322970	6912	64K	1751840	–	4	6144	316	CB228	100
XQVR600	2.5	661111	15552	96K	3608000	–	4	13824	512	CB228	100
XQVR1000	2.5	1124022	27648	128K	6127776	–	4	24576	512	CG560	100
XQR2V1000	1.5	1000K	11520	720	4082592	40	8	10240	432	BG575	200
XQR2V3000	1.5	3000K	32256	1728	10494368	96	12	28672	720	CG717	200
XQR2V6000	1.5	6000K	76032	2592	21849504	144	12	67584	1104	FC1144	200

QPRO Configuration PROMs

Device	SMD	Voltage	System Gates	Logic Cells	RAMbus	CFG Bits	MULTs	DLLs	Flip-Flops	Max I/Os	Packages	Total Ionizing Dose (TID) in KRADs
XC1736D	None	5	–	–	–	32768	–	–	–	–	DD8	–
XC1765D	5962-94717	5	–	–	–	65536	–	–	–	–	DD8	–
XC17128D	None	5	–	–	–	131072	–	–	–	–	DD8	–
XC17256D	5962-95617	5	–	–	–	262144	–	–	–	–	DD8	–
XQ1701L	5962-99514	5	–	–	–	1048576	–	–	–	–	SO20, CC44	–
XQ17V16	TBD	3.3	–	–	–	16777216	–	–	–	–	VQ44, CC44	–
XQR1701L	–	3.3	–	–	–	1048576	–	–	–	–	SO20, CC44	60
XQR18V04	–	3.3	–	–	–	4194304	–	–	–	–	VQ44, CC44	40
XQR17V16**	–	3.3	–	–	–	16777216	–	–	–	–	VQ44, CC44	TBD

** Under development

Device Nomenclatures
XC = Qualified Prior to QML Certification

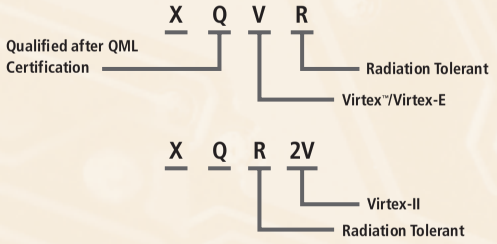

XILINX®
The Programmable Logic Company℠

XILINX CONFIGURATION STORAGE SOLUTIONS

http://www.xilinx.com/configsolns

Product Selection and Package Option Matrix

Platform Flash	PROM Solution	V020	FS48	Core Voltage	VCCI (V)	VCCO (V)
PlatformFlash PROMs						
XC3S50	XCF01S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC3S200	XCF01S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC3S400	XCF02S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC3S1000	XCF04S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC3S1500	XCF08P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC3S2000	XCF08P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC3S4000	XCF16P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC3S5000	XCF16P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2S50E	XCF01S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2S100E	XCF01S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2S150E	XCF02S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2S200E	XCF02S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2S300E	XCF02S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2S400E	XCF04S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2S600E	XCF04S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2VP2	XCF02S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2VP4	XCF04S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2VP7	XCF08P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2VP20	XCF16P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2VP30	XCF16P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2VP40	XCF16P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2VP50	XCF32P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2VP70	XCF32P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2VP100*	XCF32P + XCF01S	Y Y	Y Y	1.8 & 3.3	1.5 - 3.3 & 3.3	1.8 - 3.3 1.5 - 3.3
XC2VP125*	XCF32P + XCF16P		Y	1.8	1.5 - 3.3	1.8 - 3.3 1.8 - 3.3
XC2V40	XCF01S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2V80	XCF01S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2V250	XCF02S	Y		3.3	1.8 - 3.3	1.8 - 3.3
XC2V500	XCF04S	Y		3.3	1.8 - 3.3	1.5 - 3.3
XC2V1000	XCF04S	Y		3.3	1.8 - 3.3	1.5 - 3.3
XC2V1500	XCF08P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2V2000	XCF08P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2V3000	XCF16P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2V4000	XCF16P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2V6000	XCF32P		Y	1.8	1.5 - 3.3	1.5 - 3.3
XC2V8000	XCF32P		Y	1.8	1.5 - 3.3	1.5 - 3.3

*A single XCF32P can be used when compression is enabled. The compression ratio is bitstream dependent. A typical application will achieve 50% more bits.

FPGA	PROM Solution	PD8	V08	S020	PC20	PC44	VQ44	Core Voltage	I/O Voltage 3.3V	I/O Voltage 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

FPGA	PROM Solution	PD8	V08	S020	PC20	PC44	VQ44	Core Voltage	I/O Voltage 2.5V	I/O Voltage 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

Notes
1. All Platform Flash and PROM devices are temperature marked "C" indicating they operate from -40C to +85C.
2. Platform Flash should be quoted for all new designs.
3. PROM information provided as legacy cross-reference only.

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

Note: For multiple FPGA Configuration and for designs utilizing system level features, use SystemACE.

PROM Package Table

PROM	Density	PD8	V08	S020	PC20	PC44	VQ44	VO20	FS48*	Core Voltage	I/O Voltage 2.5V	I/O Voltage 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
Platform Flash												
XCF01S								Y		3.3V		
XCF02S								Y		3.3V		
XCF04S								Y		3.3V		
XCF08P									Y	3.3V		
XCF016P									Y	3.3V		
XCF032P									Y	3.3V		

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

XILINX SOFTWARE

http://www.xilinx.com/ise

	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 - V600 **Virtex-E:** V50E - V600E **Virtex-II:** 2V40 - 2V500 **Virtex-II Pro:** 2VP2, 2VP4, 2VP7	ALL	ALL
	Spartan™ Series	**SpartanII/IIE:** ALL (except XC2S400E and XC2S600E) **Spartan-3:** 3S50, 3S200, 3S400	**Spartan-II/IIE:** ALL **Spartan-3:** 3S50, 3S200, 3S400	**Spartan-II/IIE:** ALL **Spartan-3:** ALL	**Spartan-II/IIE:** ALL **Spartan-3:** 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	MS Windows Only	MS Windows Only
	HDL Editor	Yes	Yes	Yes	Yes
	State Diagram Editor	Yes	Yes	MS Windows	MS Windows
	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	Yes	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	No	Yes (Available with optional EDK)	Yes (Available with optional EDK)	Yes (Available with optional EDK)
	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 (MS Windows)	Integration Interface (MS Windows)
	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 (MS Windows)	CPLD (MS Windows)
Implementation	FloorPlanner	Yes	Yes	Yes	Yes
	Constraints Editor	Yes	Yes	Yes	Yes
	Timing Driven Place & Route	Yes	Yes	Yes	Yes
	Modular Design	Yes	Yes	Yes	Yes
	Timing Improvement Wizard	Yes	Yes	Yes	Yes
Programming	iMPACT	Yes	Yes	Yes	Yes
	System ACE Configuration Manager	Yes	Yes	Yes	Yes
Board Level Integration	IBIS Models	Yes	Yes	Yes	Yes
	STAMP Models	Yes	Yes	Yes	Yes
	LMG SmartModels	Yes	Yes	Yes	Yes
	HSPICE Models*	Yes	Yes	Yes	Yes
Verification	HDL Bencher™	Yes	Yes	MS Windows	MS Windows
	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), Linux	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®
The Programmable Logic Company℠

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-3	Spartan-IIE	Spartan-II	Spartan
Communication & Networking										
1.25 Gb/s GFP Framer	Paxonet Communications	AllianceCORE	V-IIP	V-II			S-3	S-IIE		
2.5 Gb/s GFP Framer	Paxonet Communications	AllianceCORE	V-IIP	V-II						
3G FEC Package	Xilinx	LogiCORE		V-II	V-E	V				
8b/10b Decoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
8b/10b Encoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
8-Bit Multilchannel GFP Framer	Paxonet Communications	AllianceCORE	V-IIP	V-II						
8-Bit Transparent GFP Framer	Paxonet Communications	AllianceCORE	V-IIP	V-II						
ADPCM, 1024 Channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II						
ADPCM, 256 Channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E					
ADPCM, 512 Channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II		V				
ADPCM, 768 Channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II		V				
AES Decryption family	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V			S-II	
AES Encryption	CAST, Inc.	AllianceCORE	V-IIP	V-II	V-E		S-3	S-IIE		
AES Encryption Family	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V			S-II	
AES Standard Encryptor/Decryptor	Helion Technology Limited	AllianceCORE	V-IIP	V-II	V-E		S-3	S-IIE		
AES Tiny Encryptor/Decryptor	Helion Technology Limited	AllianceCORE	V-IIP	V-II	V-E		S-3	S-IIE		
ATM Cell Processor	Paxonet Communications	AllianceCORE	V-IIP	V-II						
AWGN - Additive White Gaussian Noise	Xilinx	LogiCORE	V-IIP	V-II						
Bit Stream Analyzer and Data Extractor	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II				S-IIE		
Bluetooth Baseband Processor	NewLogic GmbH	AllianceCORE		V-II		V				
Bluetooth Hardware Baseband Controller	Wipro, Ltd.	AllianceCORE		V-II		V				
CDMA2000 Turbo Convolutional Codec	Xilinx	LogiCORE	V-IIP	V-II			S-3			
Convolutional Encoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Convolutional Encoder	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II				S-IIE		
CRC10 Generator and Verifier	Paxonet Communications	AllianceCORE				V				S
CRC32 Generator and Verifier	Paxonet Communications	AllianceCORE				V				S
DES and DES3 encryption engine	Memec Core	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
DES Encryption	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
DES3 Encryption	CAST, Inc.	AllianceCORE	V-IIP	V-II	V-E		S-3	S-IIE		
Distributed Sample Descrambler	Telecom Italia Lab S.p.A.	AllianceCORE				V			S-II	S
Distributed Sample Scrambler	Telecom Italia Lab S.p.A.	AllianceCORE				V			S-II	S
DVB Satellite Modulator	Memec Core	AllianceCORE		V-II	V-E	V			S-II	
E1 Framer	Paxonet Communications	AllianceCORE	V-IIP	V-II				S-IIE		
E1 Mapper	Paxonet Communications	AllianceCORE	V-IIP	V-II				S-IIE		
Ethernet MAC, 1 Gigabit	Wipro, Ltd.	AllianceCORE		V-II						
Ethernet MAC, 1 Gigabit Full Duplex	Alcatel Internetworking, Inc.	AllianceCORE		V-II						
Ethernet MAC, 1 Gigabit Half/Full duplex with GMII or Gigabit serial PHY	Xilinx	LogiCORE	V-IIP	V-II	V-E			S-IIE		
Ethernet 1000BASE-X PCS/PMA	Xilinx	LogiCORE	V-IIP							
Ethernet MAC, 10 Gigabit Full Duplex with XGMII or XAUI	Xilinx	LogiCORE	V-IIP	V-II						
XAUI	Xilinx	LogiCORE	V-IIP							
Ethernet MAC, 10/100	Alcatel Internetworking, Inc.	AllianceCORE		V-II					S-II	
Ethernet MAC, 10/100	CAST, Inc.	AllianceCORE	V-IIP	V-II	V-E			S-IIE	S-II	
Ethernet MAC, 10/100	Paxonet Communications	AllianceCORE				V			S-II	S
Ethernet MAC, 10/100	Zuken, Inc.	AllianceCORE		V-II	V-E					
Ethernet MAC, 10G	Paxonet Communications	AllianceCORE		V-II						
Ethernet PCS, 10G	Paxonet Communications	AllianceCORE	V-IIP	V-II						
Ethernet PCS, 10G	Memec Core	AllianceCORE		V-II						
G.709 Compliant FEC Core	Paxonet Communications	AllianceCORE	V-IIP	V-II						
HDLC, Single-Channel	Memec Core	AllianceCORE	V-IIP	V-II			S-3			
HDLC, Single-Channel	Memec Core	Submitted	V-IIP	V-II			S-3			
HDLC-PPP for Packet over SONET	Paxonet Communications	AllianceCORE				V			S-II	
HyperTransport Cave	GDA Technologies, Inc.	AllianceCORE	V-IIP	V-II						
HyperTransport Cave	GDA Technologies, Inc.	Submitted	V-IIP	V-II						
Interleaver/De-interleaver	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Interleaver Deinterleaver	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II				S-IIE		
Interleaver/Deinterleaver	Telecom Italia Lab S.p.A.	AllianceCORE				V			S-II	S
Internet Appliance	SOC Solutions, L.L.C.	AllianceCORE		V-II	V-E					
Inverse Multiplexer for ATM	ModelWare, Inc.	AllianceCORE		V-II	V-E					
ITU-T J.83 Annex B Modulator	Xilinx	LogiCORE	V-IIP	V-II			S-3			
MD5 Message Digest Algorithm	CAST, Inc.	AllianceCORE	V-IIP	V-II			S-3	S-IIE		
Noisy Transmission Channel Model	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II				S-IIE		
OC12 Framer	Paxonet Communications	AllianceCORE	V-IIP	V-II						
OC12c Path Processor	Paxonet Communications	AllianceCORE	V-IIP	V-II						
OTU2 Framer	Paxonet Communications	AllianceCORE	V-IIP	V-II						
Reed Solomon Decoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	S
Reed Solomon Decoder	Memec Core	AllianceCORE		V-II		V		S-IIE	S-II	
Reed Solomon Encoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	S
Reed Solomon Encoder	Memec Core	AllianceCORE		V-II		V		S-IIE	S-II	
Reed-Solomon Decoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II		V				
SDLC Controller	CAST, Inc.	AllianceCORE		V-II		V				
SHA-1 Encryption processor	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
SPI-3 (POS-PHY L3) Link Layer Interface, 1-256Ch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
SPI-3 (POS-PHY L3) Link Layer Interface, 1-Ch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
SPI-3 (POS-PHY L3) Link Layer Interface, 2-Ch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
SPI-3 (POS-PHY L3) Link Layer Interface, 4-Ch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
SPI-3 (POS-PHY L3) Physical Layer Interface	Xilinx	LogiCORE			V-E					
SPI-4.1 (Flexbus 4) Interface Core, 1-Channel	Xilinx	LogiCORE		V-II						
SPI-4.1 (Flexbus 4) Interface Core, 4-Channel	Xilinx	LogiCORE		V-II						
SPI-4.2 (POS-PHY L4) Multi-Channel Interface	Xilinx	LogiCORE	V-IIP	V-II						
SPI-4.2 (POS-PHY L4) to SPI-4.1 (Flexbus 4) Bridge	Xilinx	LogiCORE		V-II						
SPI-4.2 (POS-PHY L4) to XGMII (10GE MAC) Bridge	Xilinx	LogiCORE		V-II						
SPI-4.2 Lite (POS_PHY L4)	Xilinx	LogiCORE		V-II			S-3			

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-3	Spartan-IIE	Spartan-II	Spartan
STS192/STM64 Framer	Paxonet Communications	AllianceCORE	V-IIP	V-II						
STS192/STM64 Path Processor	Paxonet Communications	AllianceCORE	V-IIP	V-II						
STS48 OTN Framer/Digital Wrapper	Paxonet Communications	AllianceCORE	V-IIP	V-II						
T1 Framer	Paxonet Communications	AllianceCORE	V-IIP	V-II				S-IIE		
T1 Framer	QualCore Logic, Inc.	AllianceCORE								
Turbo Decoder, Convolutional, 3GPP Compliant	Xilinx	LogiCORE		V-II	V-E	V				
Turbo Decoder, Product Code	Xilinx	LogiCORE	V-IIP	V-II						
Turbo Decoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II		V				
Turbo Decoder, 3GPP	iCoding Technology, Inc.	AllianceCORE		V-II	V-E			S-IIE		
Turbo Decoder, 3GPP	SysOnChip, Inc.	AllianceCORE		V-II	V-E					
Turbo Decoder, DVB-RCS	iCoding Technology, Inc.	AllianceCORE		V-II	V-E	V				
Turbo Decoder, DVB-RCS	TurboConcept	AllianceCORE		V-II	V-E					
Turbo Encoder, Convolutional, 3GPP Compliant	Xilinx	LogiCORE		V-II	V-E	V				
Turbo Encoder, Product Code	Xilinx	LogiCORE	V-IIP	V-II						
Turbo Encoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II		V			S-II	
Turbo Encoder, DVB-RCS	iCoding Technology, Inc.	AllianceCORE		V-II	V-E	V				
Turbo Product Code Decoder, 25 Mbps	TurboConcept	AllianceCORE		V-II	V-E					
UMTS ADDRESS Gen	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II	V-E			S-IIE		
UMTS Turbo Decoder	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II	V-E					
Utopia 2 Master	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II				S-IIE		
Utopia 2 Master	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II				S-IIE		
UTOPIA Level-2 PHY Side RX Interface	Telecom Italia Lab S.p.A.	AllianceCORE				V			S-II	S
UTOPIA Level-2 PHY Side TX Interface	Telecom Italia Lab S.p.A.	AllianceCORE				V			S-II	S
UTOPIA Slave (CC143S)	Paxonet Communications	AllianceCORE				V			S-II	S
Viterbi Decoder, General Purpose	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Viterbi Decoder, IEEE 802-compatible	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
Viterbi Decoder	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II	V-E			S-IIE	S-II	
XAPP 289: Common Switch Interface CSIX-L1 Reference Design	Xilinx	Reference Design	V-IIP	V-II						
Digital Signal Processing										
Bit Correlator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
C320C25 Digital Signal Processing	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
Cascaded Integrator Comb (CIC) Filter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
Comb Filter	Xilinx	LogiCORE								S
CORDIC	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Digital Down Converter (DDC)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
Direct Digital Synthesizer (DDS)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Fast Fourier Transform	Xilinx	LogiCORE	V-IIP	V-II			S-3			
FFT/IFFT for Virtex-II, 1024-Point Complex	Xilinx	LogiCORE		V-II						
FFT/IFFT for Virtex-II, 16-Point Complex	Xilinx	LogiCORE		V-II						
FFT/IFFT for Virtex-II, 256-Point Complex	Xilinx	LogiCORE		V-II						
FFT/IFFT for Virtex-II, 64-Point Complex	Xilinx	LogiCORE		V-II						
FFT/IFFT, 1024-Point Complex	Xilinx	LogiCORE			V-E	V				
FFT/IFFT, 16-Point Complex	Xilinx	LogiCORE		V-II	V-E	V				
FFT/IFFT, 256-Point Complex	Xilinx	LogiCORE		V-II	V-E	V				
FFT/IFFT, 32-Point Complex	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
FFT/IFFT, 64-, 256-, 1024-Point Complex	Xilinx	LogiCORE	V-IIP	V-II			S-3			
FFT/IFFT, 64-Point Complex	Xilinx	LogiCORE			V-E	V				
FIR Filter using DPRAM	eInfochips Pvt. Ltd.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
FIR Filter, Distributed Arithmetic (DA)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
FIR Filter, Distributed Arithmetic Parallel	Xilinx	LogiCORE								S
FIR Filter, Distributed Arithmetic Serial	Xilinx	LogiCORE								S
FIR Filter, MAC	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
LFSR, Linear Feedback Shift Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Oscillator, Dual-Channel Numerically Controlled	Xilinx	LogiCORE			V-E	V		S-IIE	S-II	S
Oscillator, Numerically Controlled	Xilinx	LogiCORE			V-E	V		S-IIE	S-II	S
Parallel Distributed Arithmetic FIR Filter	eInfochips Pvt. Ltd.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
Time-Skew Buffer, Nonsymmetric 16-Deep	Xilinx	LogiCORE								S
Time-Skew Buffer, Nonsymmetric 32-Deep	Xilinx	LogiCORE								S
Time-Skew Buffer, Symmetric 16-Deep	Xilinx	LogiCORE								S
Virtex-E DSP Hardware Accelerator	GV & Associates, Inc.	AllianceCORE								
Virtex-II DSP Hardware Accelerator	GV & Associates, Inc.	AllianceCORE		V-II						
Virtex-II DSP Hardware Accelerator	GV & Associates, Inc.	AllianceCORE		V-II						
Math Functions										
1s and 2s Complement	Xilinx	LogiCORE								S
Accumulator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Adder Subtracter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Divider, Pipelined	Xilinx	LogiCORE		V-II	V-E	V		S-IIE	S-II	S
Floating Point Adder	Digital Core Design	AllianceCORE		V-II		V			S-II	
Floating Point Divider	Digital Core Design	AllianceCORE		V-II		V			S-II	
Floating Point Multiplier	Digital Core Design	AllianceCORE		V-II	V-E	V			S-II	
Floating Point Square Comparator	Digital Core Design	AllianceCORE		V-II		V			S-II	
Floating Point Square Root Operator	Digital Core Design	AllianceCORE		V-II	V-E				S-II	
Floating Point to Integer Converter	Digital Core Design	AllianceCORE		V-II		V			S-II	
Integer to Floating Point Converter	Digital Core Design	AllianceCORE		V-II		V			S-II	
Integrator	Xilinx	LogiCORE								S
Multiplier for Virtex, Variable Parallel	Xilinx	LogiCORE			V-E	V		S-IIE	S-II	
Multiplier, Constant Coefficient	Xilinx	LogiCORE								S
Multiplier, Constant Coefficient - Pipelined	Xilinx	LogiCORE								S
Multiplier, Dynamic Constant Coefficient	Xilinx	LogiCORE			V-E	V				
Multiplier, Parallel - Area Optimized	Xilinx	LogiCORE								S
Multiply Accumulator (MAC)	Xilinx	LogiCORE	V-IIP	V-II		V	S-3	S-IIE	S-II	
Multiply Generator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Registered Adder	Xilinx	LogiCORE								S

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-3	Spartan-IIE	Spartan-II	Spartan
Registered Loadable Adder	Xilinx	LogiCORE								S
Registered Loadable Subtracter	Xilinx	LogiCORE								S
Registered Scaled Adder	Xilinx	LogiCORE								S
Registered Serial Adder	Xilinx	LogiCORE								S
Registered Subtracter	Xilinx	LogiCORE								S
Scaled-by-One-Half Accumulator	Xilinx	LogiCORE								S
Sine Cosine Look Up Table	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	S
Square Root	Xilinx	LogiCORE								S
Twos Complementer	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
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-3	S-IIE	S-II	
Block Memory, Single-Port	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
CAM for Internet Protocol	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II				S-IIE		
Content Addressable Memory (CAM)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Distributed Memory	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
FIFO, Asynchronous	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
FIFO, Synchronous	Xilinx	LogiCORE								S
FIFO, Synchronous	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Pipelined Delay Element	Xilinx	LogiCORE								S
Registered ROM	Xilinx	LogiCORE								S
Registered Single Port RAM	Xilinx	LogiCORE								S
SDRAM Controller	Eureka Technology	AllianceCORE				V				
SDRAM Controller, DDR	Memec Core	AllianceCORE		V-II	V-E				S-II	
SDRAM Controller, Pipelined	Eureka Technology	AllianceCORE	V-IIP	V-II			S-3	S-IIE		
Microprocessors, Controllers & Peripherals										
1 Gigabit Ethernet MAC w/PLB interface	Xilinx	LogiCORE	V-IIP							
10/100 Ethernet MAC Lite w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
10/100 Ethernet MAC w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
16450 UART	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
16450 UART	QualCore Logic, Inc.	AllianceCORE								S
16450 UART w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
16450 UART w/Synchronous Interface	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
16550 UART w/FIFOs	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
16550 UART w/FIFOs	QualCore Logic, Inc.	AllianceCORE								S
16550 UART w/FIFOs & synch interface	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
16550 UART w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
16-bit proprietary RISC Processor	Loarant Corporation	AllianceCORE	V-IIP	V-II			S-3	S-IIE		
2901 Microprocessor Slice	CAST, Inc.	AllianceCORE			V-E	V			S-II	
2910A Microprogram Controller	CAST, Inc.	AllianceCORE				V			S-II	S
2D Fabric Evaluation and Demo Board	Crossbow Technologies, Inc.	AllianceCORE							S-II	
2-dimensional multiprocessing interface fabric	Crossbow Technologies, Inc.	AllianceCORE	V-IIP	V-II				S-IIE	S-II	
68000 compatible microprocessor	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-3			
80186 compatible processor	eInfochips Pvt. Ltd.	AllianceCORE		V-II		V				
8051 compatible microcontroller	CAST, Inc.	AllianceCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
8051 compatible microcontroller	Digital Core Design	AllianceCORE		V-II		V			S-II	
80515 High-speed 8-bit RISC Microcontroller	CAST, Inc.	AllianceCORE				V				
8052 compatible microcontroller	Digital Core Design	AllianceCORE		V-II		V			S-II	
80C51 compatible RISC microcontroller	CAST, Inc.	AllianceCORE			V-E	V			S-II	
8237 Programmable DMA Controller	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
8250 UART	CAST, Inc.	AllianceCORE		V-II	V-E			S-IIE	S-II	
8254 programmable interval timer/counter core	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
8254 Programmable Timer	QualCore Logic, Inc.	AllianceCORE				V			S-II	S
8254 programmable timer/counter	eInfochips Pvt. Ltd.	AllianceCORE		V-II					S-II	
8255 programmable I/O controller	eInfochips Pvt. Ltd.	AllianceCORE		V-II					S-II	
8255 Programmable Peripheral Interface	QualCore Logic, Inc.	AllianceCORE				V			S-II	S
8255A Peripheral Interface	CAST, Inc.	AllianceCORE				V			S-II	
8259 Programmable Interrupt Controller	QualCore Logic, Inc.	AllianceCORE								S
8259A Programmable Interrupt controller	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
Arbiter and Bus Structure w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Arbiter and Bus Structure w/PLB interface	Xilinx	LogiCORE	V-IIP							
ATM Utopia Level 2 Master and Slave w/OPB Interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
ATM Utopia Level 2 Master and Slave w/PLB Interface	Xilinx	LogiCORE	V-IIP							
BRAM Controller w/LMB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
BRAM Controller w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
BRAM Controller w/PLB interface	Xilinx	LogiCORE	V-IIP							
BSP Generator (SW only)	Xilinx	LogiCORE	V-IIP							
CRT Controller	CAST, Inc.	AllianceCORE		V-II	V-E			S-IIE	S-II	
DCR Bus Structure	Xilinx	LogiCORE	V-IIP							
DR8051 RISC MicroController	Digital Core Design	AllianceCORE		V-II		V			S-II	
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-3	S-IIE	S-II	
External Memory Controller (EMC) w/PLB interface (Includes support for Flash, SRAM, ZBT, System ACE)	Xilinx	LogiCORE	V-IIP							
Generic compact UART	Memec Core	AllianceCORE	V-IIP	V-II			S-3			
Generic compact UART	Memec Core	Submitted	V-IIP	V-II			S-3			
GPIO w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
HDLC Controller (Single Channel) w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
I2C w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Interrupt Controller (IntC) w/DCR interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
Interrupt Controller (IntC) w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
IPIF Address Decode w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
IPIF DMA w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-3	Spartan-IIE	Spartan-II	Spartan
IPIF Interrupt Controller w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
IPIF Master/Slave Attachment w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
IPIF Read/Write Packet FIFO w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
IPIF Scatter/Gather w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
IPIF Slave Attachment w/PLB interface	Xilinx	LogiCORE	V-IIP							
Java Processor, 32-bit	Digital Communications Technologies, Ltd.	AllianceCORE		V-II		V			S-II	
Java Processor, Configurable	Derivation Systems, Inc.	AllianceCORE		V-II		V				
JTAG UART w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Memory Test Utility (SW only)	Xilinx	LogiCORE	V-IIP							
MicroBlaze Soft RISC Processor	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
MicroBlaze Source Code	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
MIPS system controller	Eureka Technology	AllianceCORE		V-II	V-E					
ML300 VxWorks BSP (SW only)	Xilinx	LogiCORE	V-IIP							
Motor controller - 3 phase	MEET Ltd.	AllianceCORE						S-IIE	S-II	S
OPB2DCR Bridge	Xilinx	LogiCORE	V-IIP							
OPB2OPB Bridge (Lite)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
OPB2PCI Full Bridge (32/33)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
OPB2PLB Bridge	Xilinx	LogiCORE	V-IIP							
Operating System Accelerator	RealFast Operating Systems AB	AllianceCORE		V-II				S-IIE	S-II	
PIC125x fast RISC microcontroller	Digital Core Design	AllianceCORE		V-II		V			S-II	
PIC1655x fast RISC microcontroller	Digital Core Design	AllianceCORE		V-II		V			S-II	
PIC165X compatible microcontroller	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
PIC165x fast RISC microcontroller	Digital Core Design	AllianceCORE		V-II	V-E	V			S-II	
PIC16C55X compatible RISC processor	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
PicoBlaze (XAPP 213: PicoBlaze 8-bit Microcontroller for Virtex-E and Spartan-II/E Devices)	Xilinx	Reference Design			V-E	V		S-IIE	S-II	
PicoBlaze (XAPP 627: PicoBlaze 8-bit Microcontroller for Virtex-II and Virtex-II Pro Devices)	Xilinx	Reference Design	V-IIP	V-II						
PLB2OPB Bridge	Xilinx	LogiCORE	V-IIP							
PowerPC Bus Master	Eureka Technology	AllianceCORE	V-IIP	V-II			S-3	S-IIE		
PowerPC Bus Slave	Eureka Technology	AllianceCORE	V-IIP	V-II			S-3	S-IIE		
SDRAM Controller w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
SDRAM Controller w/PLB interface	Xilinx	LogiCORE	V-IIP							
SPI Master and Slave w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Systems Reset Module	Xilinx	LogiCORE	V-IIP							
Timebase/Watch Dog Timer (WDT) w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Timer/Counter w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
UART Lite w/OPB interface	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
VxWorks Board Support Package (BSP)	Xilinx	LogiCORE	V-IIP							
Xtensa-V Configurable 32-bit Processor	Tensilica, Inc.	AllianceCORE		V-II						
Z80 compatible programmable counter/timer core	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
Z80 peripheral I/O controller	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
Z80CPU Microprocessor	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-3		S-II	
Standard Bus Interfaces										
Arbiter	Telecom Italia Lab S.p.A.	AllianceCORE	V-IIP	V-II				S-IIE		
CAN 2.0 B Compatible Network Controller	Xylon d.o.o.	AllianceCORE	V-IIP	V-II			S-3	S-IIE		
CAN Bus Controller 2.0B	CAST, Inc.	AllianceCORE	V-IIP	V-II			S-3	S-IIE		
CAN with 32 mail boxes	Robert Bosch GmbH	AllianceCORE	V-IIP	V-II			S-3	S-IIE		
Controller Area Network	Memec Core	AllianceCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
HyperTransport Single-Ended Slave Core	Xilinx	LogiCORE	V-IIP	V-II						
I2C Bus Controller	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-3		S-II	
I2C Bus Controller Master	Digital Core Design	AllianceCORE		V-II	V-E				S-II	S
I2C Bus Controller Slave	Digital Core Design	AllianceCORE		V-II	V-E				S-II	S
I2C Bus Controller Slave Base	Digital Core Design	AllianceCORE		V-II	V-E				S-II	S
I2C Two-Wire Serial Interface Master-Only	Memec Core	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
I2C Two-Wire Serial Interface Master-Slave	Memec Core	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
PCI 64-bit/66-MHz master/target interface	Eureka Technology	AllianceCORE		V-II	V-E					
PCI Express x1 Core	Xilinx	LogiCORE	V-IIP							
PCI host bridge	Eureka Technology	AllianceCORE		V-II	V-E					
PCI, 64-bit Target Interface	CAST, Inc.	AllianceCORE		V-II	V-E			S-IIE		
PCI32 Interface Design Kit (DO-DI-PCI32-DKT)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
PCI32 Interface, IP Only (DO-DI-PCI32-IP)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
PCI32 Single-Use License for Spartan (DO-DI-PCI32-SP)	Xilinx	LogiCORE					S-3	S-IIE	S-II	
PCI64 & PCI32, IP Only (DO-DI-PCI-AL)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
PCI64 Interface Design Kit (DO-DI-PCI64-DKT)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
PCI64 Interface, IP Only (DO-DI-PCI64-IP)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V		S-IIE	S-II	
PCI-X 64/133 Interface for Virtex-II (DO-DI-PCIX64-VE). Includes PCI 64 bit interface at 33 MHz	Xilinx	LogiCORE	V-IIP	V-II						
PCI-X 64/66 Interface for Virtex-E (DO-DI-PCIX64-VE). Includes PCI 64 bit interface at 33 MHz	Xilinx	LogiCORE			VE					
RapidIO 8-bit port LP-LVDS Phy Layer (DO-DI-RIO8-PHY)	Xilinx	LogiCORE	V-IIP	V-II						
RapidIO Logical (I/O) and Transport Layer (DO-DI-RIO8-LOG)	Xilinx	LogiCORE	V-IIP	V-II						
RapidIO Phy Layer to PLB Bridge reference design	Xilinx	Reference Design	V-IIP							
Serial Protocol Interface (SPI) Slave	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE		
USB 1.1 Function Controller	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
XAPP653: Virtex-II Pro/Spartan-3 3.3V PCI Reference Design	Xilinx	Reference Design	V-IIP				S-3			
Video & Image Processing										
BURST_PLL	Pinpoint Solutions, Inc.	AllianceCORE		V-II	V-E		S-3	S-IIE		
Color Space Converter, RGB2YCrCb	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
Color Space Converter, RGB2YCrCb	Perigee, LLC	AllianceCORE				V			S-II	S
Color Space Converter, YCrCb2RGB	Perigee, LLC	AllianceCORE				V			S-II	S
Compact Video Controller	Xylon d.o.o.	AllianceCORE		V-II		V			S-II	
Discrete Cosine Transform	eInfochips Pvt. Ltd.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-3	Spartan-IIE	Spartan-II	Spartan
Discrete Cosine Transform, 2D Inverse	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
Discrete Cosine Transform, Combined 2D Forward/Inverse	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
Discrete Cosine Transform, Forward 2D	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
Discrete Cosine Transform, Forward/Inverse	Telecom Italia Lab S.p.A.	AllianceCORE				V			S-II	
Discrete Cosine Transform, forward/inverse 2D	Barco-Silex	AllianceCORE		V-II	V-E	V			S-II	
Discrete Wavelet Transform, Combined 2D Forward/Inverse	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
Discrete Wavelet Transform, Line-based programmable forward	CAST, Inc.	AllianceCORE		V-II	V-E	V			S-II	
Huffman Decoder	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-IIE	S-II	
JPEG Fast Codec	CAST, Inc.	AllianceCORE	V-IIP	V-II	V-E	V	S-3			
JPEG Fast Codec	CAST, Inc.	Submitted	V-IIP	V-II	V-E	V	S-3			
JPEG Fast Decoder	CAST, Inc.	AllianceCORE	V-IIP	V-II	V-E	V		S-IIE		
JPEG Fast Encoder	CAST, Inc.	AllianceCORE	V-IIP	V-II	V-E	V				
JPEG, 2000 Encoder	CAST, Inc.	AllianceCORE	V-IIP	V-II	V-E					
JPEG, Fast color image decoder	Barco-Silex	AllianceCORE		V-II	V-E	V				
JPEG, Fast gray scale image decoder	Barco-Silex	AllianceCORE		V-II	V-E	V				
JPEG, Motion Codec core V1.0	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V				
JPEG, Motion Decoder core V1.0	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V				
JPEG, Motion Encoder core V2.0	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V				
Longitudinal Time Code Generator	Deltatec S.A.	AllianceCORE				V				
MPEG-2 HDTV I & P Encoder	Duma Video, Inc.	AllianceCORE		V-II						
MPEG-2 SDTV I & P Encoder	Duma Video, Inc.	AllianceCORE		V-II						
NTSC-COSEP	Pinpoint Solutions, Inc.	AllianceCORE		V-II	V-E		S-3	S-IIE		
Backplanes and Gigabit Serial I/O										
Aurora 401: single lane, 32-bit interface	Xilinx	Reference Design	V-IIP							
WP160: Emulating External SERDES Devices with Embedded RocketIO Transceivers	Xilinx	White Paper	V-IIP							
XAPP 649: SONET Rate Conversion in Virtex-II Pro Devices	Xilinx	Reference Design	V-IIP							
XAPP 651: SONET and OTN Scramblers/Descramblers	Xilinx	Reference Design	V-IIP							
XAPP 652: Word Alignment and SONET/SDH Framing	Xilinx	Reference Design	V-IIP							
XAPP660: Partial Reconfiguration of RocketIO Attributes using PPC405 core (DCR Bus)	Xilinx	Reference Design	V-IIP							
XAPP661: RocketIO Transceiver Bit-Error Rate Tester (BERT)	Xilinx	Reference Design	V-IIP							
XAPP662: Partial Reconfig. of RocketIO Attributes using PPC405 core (PLB or OPB bus) + RocketIO Transceiver Bit-Error Rate Tester (BERT)	Xilinx	Reference Design	V-IIP							
Basic Elements										
Binary Counter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Binary Decoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Bit Bus Gate	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Bit Gate	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Bit Multiplexer	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
BUFE-based Multiplexer Slice	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
BUFT-based Multiplexer Slice	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Bus Gate	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Bus Multiplexer	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Comparator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
FD-based Parallel Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
FD-based Shift Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Four-Input MUX	Xilinx	LogiCORE								S
LD-based Parallel Latch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Parallel-to-Serial Converter	Xilinx	LogiCORE								S
RAM-based Shift Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-3	S-IIE	S-II	
Register	Xilinx	LogiCORE								S
Three-Input MUX	Xilinx	LogiCORE								S
Two-Input MUX	Xilinx	LogiCORE								S
Prototype & Development Hardware Products										
CPU + FPGA (Virtex/Spartan-II) MicroEngine Cards	Intrinsyc, Inc.	AllianceCORE								
CPU + FPGA (Virtex-II) MicroEngine Cards	Intrinsyc, Inc.	AllianceCORE		V-II						
JockoBoard SOC Virtex-II Prototyping Platform	RealFast Operating Systems AB	AllianceCORE		V-II						
LogiCRAFT Evaluation System	Xylon d.o.o.	AllianceCORE			V-E				S-II	
PC/104-Plus Reconfigurable Module Board	Derivation Systems, Inc.	AllianceCORE								
uPCI Reference and Development Platform	Intrinsyc, Inc.	AllianceCORE								
XTENSA Microprocessor Emulation Kit	Tensilica, Inc.	AllianceCORE		V-II						

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

XILINX®
The Programmable Logic Company℠

XILINX GLOBAL SERVICES

http://support.xilinx.com/support/gsd/index.htm

Part Number	Product Description	Duration
Education Service		hours
FPGA13000-5-ILT	Fundamentals of FPGA Design	8
FPGA23000-5-ILT	Designing for Performance	16
FPGA33000-5-ILT	Advanced FPGA Implementation	16
LANG11000-5-ILT	Introduction to VHDL	24
LANG21000-5-ILT	Advanced VHDL	16
LANG12000-5-ILT	Introduction to Verilog	24
PCI8000-4-ILT	PCI CORE Basics	8
PCI28000-4-ILT	Designing a PCI System	16
ASIC25000-5-ILT	Designing for Performance for the ASIC User	24
DSP2000-3-ILT	DSP Implementation Techniques for Xilinx FPGAs	24
DSP-10000-5-ILT	DSP Design Flow	24
RIO22000-5-ILT	Designing with RocketIO MGT	16
PROMO-5003-5-LEL	Designing for Performance, Live Online	9
EMBD-21000-5-ILT	Embedded Systems Development	16
Platinum Technical Service		hours
SC-PLAT-SVC-10	Seat Platinum Technical Service w/10 education credits	N/A
SC-PLAT-SITE-50	Platinum Technical Service site license up to 50 customers	N/A
SC-PLAT-SITE-100	Platinum Technical Service site license for 51-100 customers	N/A
SC-PLAT-SITE-150	Platinum Technical Service site license for 101-150 customers	N/A
Titanium Technical Service		hours
PS-TEC-SERV	Titanium Technical Service (minimum 40 hours)	N/A
DC-DES-SERV	Design Services Contract	N/A
Xilinx Productivity Advantage		hours
DS-XPA	Custom XPA Packaged Solution	N/A
DS-XPA-10K	Custom XPA for $0 – $10,000	N/A
DS-XPA-25K	Custom XPA for $10,001 – $25,000	N/A
DS-XPA	Custom XPA for $25,001 – $100,000	N/A
DS-XPA-10K-INT	Custom XPA (International) for $0 – $10,000	N/A
DS-XPA-25K-INT	Custom XPA (International) for $10,001 – $25,000	N/A
DS-XPA-INT	Custom XPA (International) for $25,001 – $100,000	N/A
DS-ISE-ALI-XPA	XPA Seat, ISE Alliance	N/A
DS-ISE-FND-XPA	XPA Seat, ISE Foundation	N/A
Promotion Packages		hours
PROMO-5004-5-ILT	FPGA Essentials: Includes both Fundamentals of FPGA Design and Designing for Performance classes	24
DS-SYSGEN-4SL-T	DSP Package: Includes SysGen for DSP and 15 Training Credits	24
DO-EDK-T	EDK Package: Includes EDK Kit and Embedded Systems Development class	16

support.xilinx.com

- Sign up for personalized email alerts at mysupport.xilinx.com
- Search our knowledge database
- Troubleshoot with Problem Solvers
- Consult with engineers in Forums

Xilinx Productivity Advantage (XPA) Program

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 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."

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.

Titanium Technical Service

- Dedicated Applications 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 a dedicated team of 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 is 2x the 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: +81-3-5321-7750
http://support.xilinx.co.jp/support/education-home.htm

Asia Pacific: +852-2424-5200
education_ap@xilinx.com
http://support.xilinx.com/support/education-home.htm

Design Services Contacts

North America:
Richard Fodor: 408-626-4256
Mike Barone: 512-238-1473

Europe:
Alex Hillier: +44-870-7350-516
Martina Finnerty: +353-1-403-2469

Asia Pacific:
David Keefe: +852-2401-5171

designservices@xilinx.com

XPA Contacts

North America: 800-888-FPGA (3742)
fpga.xilinx.com

Europe:
Stuart Elston: +44-870-7350-532

Asia Pacific:
David Keefe: +852-2401-5171

Titanium Technical Service Contacts

North America:
Telesales: 800-888-3742

Europe:
Stuart Elston: +44-870-7350-632

Japan:
+81-3-5321-7730
or japantitanium@xilinx.com

Asia Pacific:
David Keefe: +852-240-5171

Xilinx Global Services
Finish Faster

Avoid messy timing mistakes. Use Mentor Graphics® FPGA design tools.

FPGA Design

PRECISION™ SYNTHESIS | As any FPGA designer can tell you, achieving precise timing is their single most critical challenge. That's why Mentor Graphics® created Precision™ Synthesis, a powerful new tool suite that lets designers close on timing faster than ever. Precision™ Synthesis provides the most comprehensive analysis for complex FPGAs with the only complete, built-in incremental timing analysis. Quickly find and optimize the most critical paths and avoid frustrating trial-and-error design iterations. Visit **www.mentor.com/fpga** today or call **877.387.5873** for more information on how Precision™ Synthesis can help you find the fastest path to a completed design.

Mentor Graphics®

©2003 Mentor Graphics Corporation. All Rights Reserved. Mentor Graphics is a registered trademark of Mentor Graphics Corporation.

To maximize savings on your next FPGA design, just keep two things in mind.

In today's challenging business environment, cost reduction is critical. Together, Xilinx and Synplicity deliver everything you need to optimize your FPGA designs — and your budget.

By combining the performance of the best synthesis tool available, Synplicity's Synplify Pro,® with industry-leading Xilinx FPGAs, we give you the power to quickly meet critical design timing demands while minimizing area, resulting in a smaller device or lower speed-grade requirement. A slower speed grade alone could result in device savings of up to 25%.

This unique partnership is the latest advancement in our efforts to minimize system design costs through technology innovations such as 90nm design rules, 300mm wafers for manufacturing, high-level function integration and ISE-tool productivity enhancements. With Xilinx and Synplicity, you get the best of both worlds — increased FPGA performance and decreased design costs. Visit ***www.xilinx.com/savings*** to learn more.

XILINX®
The Programmable Logic Company℠

Synplicity®

© 2003 Xilinx, Inc., 2100 Logic Drive, San Jose, CA 95124. Europe +44-870-7350-600; Japan +81-3-5321-7711; Asia Pacific +852-2-424-5200; Xilinx is a registered trademark and The Programmable Logic Company is a service mark of Xilinx, Inc. The Synplicity "S" logo and Synplify Pro are registered trademarks of Synplicity Inc. All other trademarks and registered trademarks are the property of their respective owners.

PN 0010746