ISSUE 41, FALL/WINTER 2001
XCELL JOURNAL
XILINX, INC.

Xcell journal

Issue 41
Fall/Winter 2001

THE AUTHORITATIVE JOURNAL FOR PROGRAMMABLE LOGIC USERS

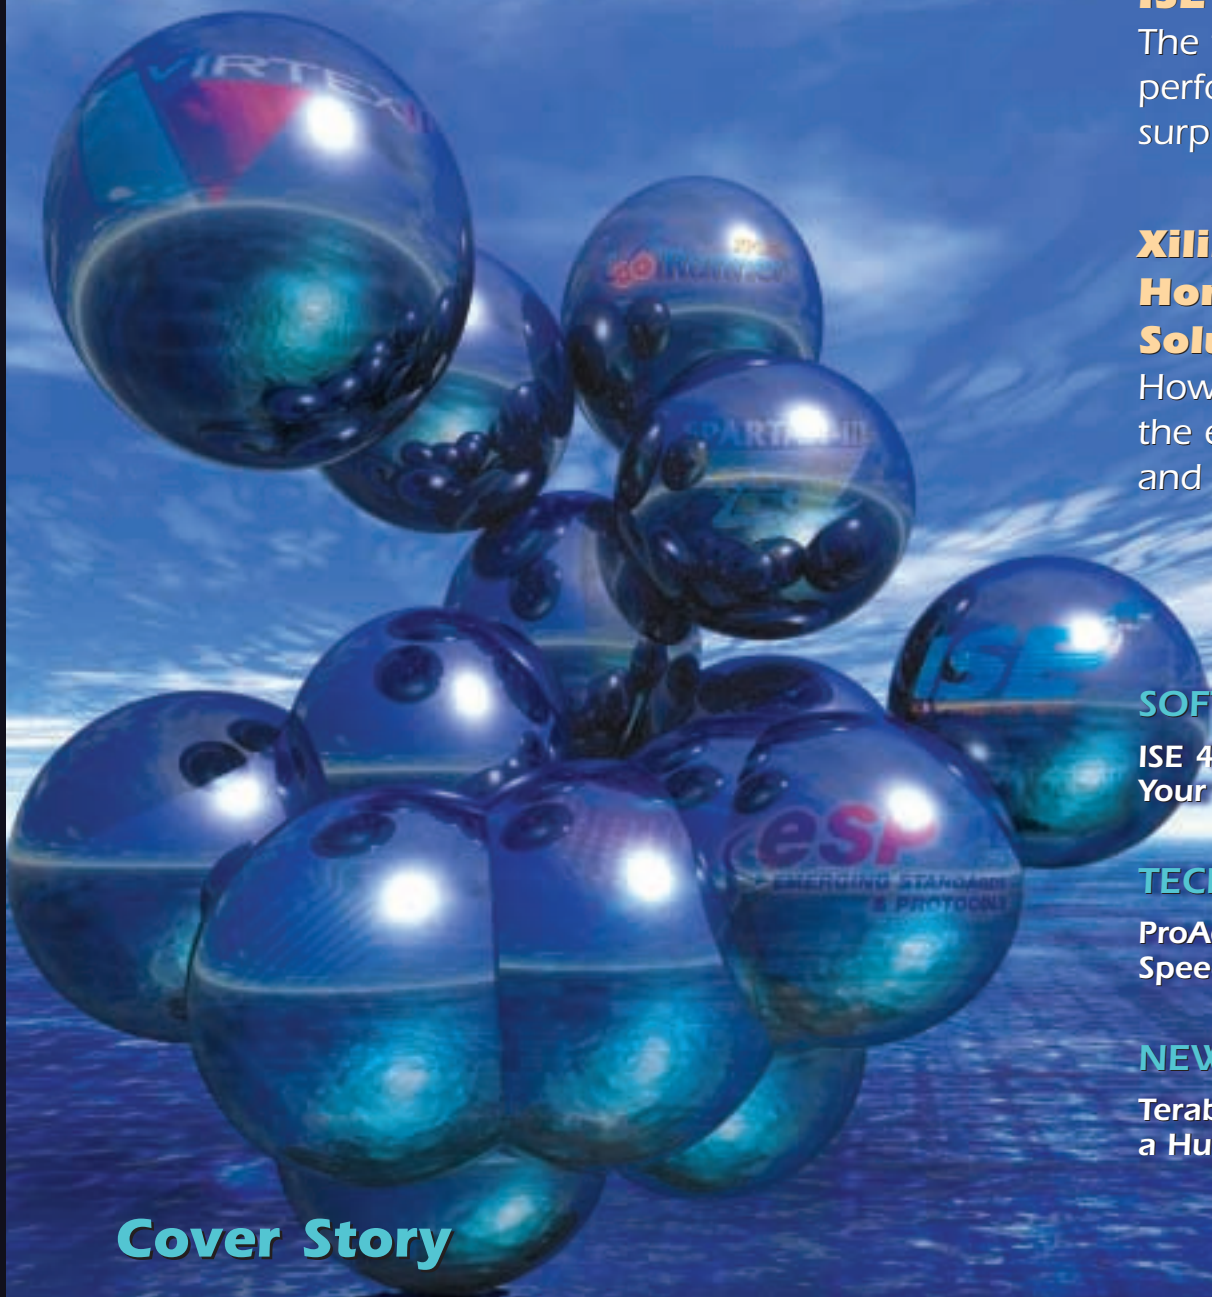

Xilinx Introduces Fast, Efficient ISE 4.1i Software

The features and performance will surprise you

Xilinx Creates Home Networking Solutions

How to keep up with the emerging standards and protocols

SOFTWARE

ISE 4.1i Accelerates Your Time to Market

TECHNOLOGY

ProActive Timing Closure Speeds Up Design Process

NEWS

Terabit Networking Forum a Huge Success

Cover Story

Synopsys Senior VP Sanjiv Kaul Leverages ASIC Expertise for FPGA Design

XILINX®

A Design Win for Xcell!

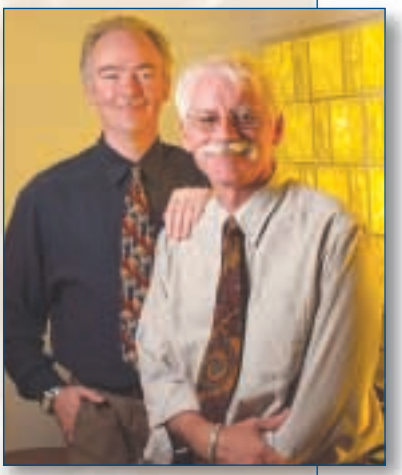

Xcell journal

Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124-3450
Phone: 408-559-7778
FAX: 408-879-4780

Those of you who have been reading *Xcell Journal* for a while have probably noticed a steady improvement in the quality and content of the magazine. This reflects the Xilinx cultural values of ever-escalating standards and the continuing pursuit of excellence.

We are especially proud of the graphic design of *Xcell Journal*, which has risen to the award-winning level. As of press time, *Xcell* had just won a prestigious Ozzie Award for Excellence in Magazine Design. *Xcell's* graphic designers – Scott Blair and Dan Teie of Teie, Gelwicks & Associates – won the Gold Award for Best Use of Digital Imagery. Scott and Dan's striking 3D imagery and imaginative layouts have proved that a diagram, graph, schematic, illustration, or picture really is worth a thousand words. Their blend of art and science help us all better understand the complex technology of programmable logic.

We are justifiably proud that Scott and Dan have brought *Xcell Journal* to the level of world-class quality that genuinely reflects the world-class leadership of Xilinx itself.

Xcell is more than just killer graphics, of course. We've also elevated the quality of the editorial content with each issue as well. This summer's Special Edition of *Xcell* was so well received by the industry that the Virtex™-II special section has gone into a second printing.

This issue of *Xcell* is reaching for new levels of excellence – in spite of the hard times our industry is enduring. As a concession to the economy, this issue is a combined fall/winter edition – but as no concession to quality, we are featuring, for the first time, two special sections:

1. **ISE 4.1i** – The new look and unified feel of the latest release of Xilinx Integrated Software Environment design tools not only present the easiest to use graphical interfaces to date, but the ISE 4.1i configurations are fully compatible with ASIC design and synthesis software from the leading vendors. ISE 4.1i makes it easier than ever to lure ASIC engineers into the superior design environment of programmable logic.
2. **Home Networking** – Xilinx continues to provide the most up-to-date information and analysis of the competing technologies vying for the lucrative home networking market. Log on to the emerging Standards and Protocols (eSP) Web portal to stay informed of who's who and what's what in home networking.

In addition to the two special sections, we have features on:

- SystemIO – the state-of-the-art solution for interconnecting the various standards and protocols of emerging terabit networking systems
- How Xilinx XPERTS Partner Dillon Engineering developed an ultra high-speed digital image-processing system using Virtex-II Platform FPGAs
- The debut of another Xilinx Web portal, Signal Integrity Central.

And that's not all. *Xcell* is now wired. Visit our new website at *www.xilinx.com/publications/xcellonline/*. Please, tell us what you think.

Tom Durkin
Managing Editor

View from the Top Page 4

Managing in Difficult Times

It takes more than technology to survive and flourish in this current economic slowdown.

Cover Story Page 8

Leveraging ASIC Expertise for Platform FPGAs

As Platform FPGAs encroach into ASIC territory, complex synthesis tools become essential.

New Products Page 10

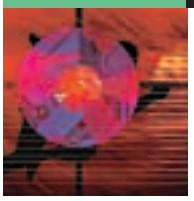

Lead Story – New ISE 4.1i Software

New ISE 4.1i software gives you the fastest time to market, by far.

Perspective Page 40

Lead Story – Networking Comes Home

A silent revolution is happening in home networking.

Perspective Page 56

Residential Gateways

A single device connects multiple broadband access and home networking technologies.

Technology Focus Page 62

Platform FPGA SystemIO Solution

Higher I/O performance to meet the growing demand for communications bandwidth.

Contents Fall/Winter 2001

ISE 4.1i

Home Networking

Xcell journal

For a Free Subscription to the Xcell Journal enter your request at: www.xilinx.com/forms/literature.htm

Visit Xcell Online at: www.xilinx.com/publications/xcellonline/

Managing in Difficult Times

It takes more than technology to survive and flourish in this current economic slowdown.

by Wim Roelandts
CEO, Xilinx

We are facing some of the most difficult challenges our industry has ever encountered. Yet, in spite of the current economic slowdown, Xilinx remains strong and we are very optimistic about the future. We continue to aggressively develop the next generation of programmable logic technology and to provide the industry's most comprehensive design solutions for both high-performance and low-cost applications. This year we will introduce more new products than ever before – our rate of innovation has not slowed.

In addition to our new products, we continue to develop new ways to help you succeed. For example, our eSP website provides a wealth of current information to help you create designs based on emerging Standards and Protocols – developing products for the home networking market has never been easier. We also hosted the highly successful Terabit Networking Forum, which provided much needed insight into the future of high-speed networking. These innovative marketing efforts, combined with the advantages of our technology, are the reason why the programmable logic sector of the semiconductor market consistently grows faster than other alternatives.

Even in these difficult times, when our revenues have been significantly reduced, Xilinx continues to be a profitable technology leader, as well as a great place to work.

The Xilinx Way

Xilinx is able to weather this current economic storm better than most companies, while continuing to bring unprecedented innovation to market, because of four primary factors:

- Our management philosophy
- Our company structure
- Our people
- Our technology.

The Xilinx Management Philosophy

I left Hewlett Packard in 1996, after 29 years of service, to become the CEO of Xilinx. I chose Xilinx, because one of my goals is to create a new, enduring style of management that other companies will want to copy. I believe that great companies treat people with respect – and also win in the market – by using creativity, innovation, and the common wisdom of all employees to make all of our jobs enjoyable and rewarding.

In Xilinx, I saw the potential to help create a company that would be a technology leader and a financial success, while also being an excellent place to work. I believed Xilinx could be a company that would one day become great – not just financially successful – and we are well on our way.

I follow four basic management principles that help to make Xilinx a great place to work while also creating an environment in which new technology can flourish:

- People want to be part of a supportive team. We all do our best in a supportive team environment. And, because we spend a large part of our time at work, it is very important that we enjoy our work. At Xilinx, everyone is respected and given plenty of opportunity to succeed. We provide both technical and management career paths so all employees can go as far as they choose.
- People want to do a good job and to contribute. We all want our contributions to be recognized, and we want to be proud of the work we do. Therefore, it's important to provide work that is interesting and exciting, and to let people make their own decisions whenever possible. This requires a lot of communication, so everyone understands our common goals and overall business conditions.
- People want to grow and improve. We all want to be better employees tomorrow than we are today. So, at Xilinx, we encourage all employees to continually learn, and we give them plenty of opportunities to do so.
- People want to have ownership. We all want to share in the responsibilities and rewards of success. As owners of the company, we will do a better job and have a more positive attitude. At Xilinx, everyone gets stock options when they join the company, and they also have an opportunity to participate in further ownership through the company's stock purchase plans.

The end result is that you get better products and services, and you get to work with highly skilled and motivated people who enjoy what they do.

The Xilinx Company Structure

We focus on the core business functions that are necessary for our success – design, marketing, and technical support. We leave the two most expensive operations – sales and manufacturing – to our business partners.

Because we outsource these functions, we don't have to worry about keeping a manufacturing facility busy, or worry about financing a large sales force during slow business cycles. Thus, we can focus all of our resources on bringing more innovation to market, while keeping our expenses at a minimum. This is one of the key reasons why we have been able to avoid layoffs at this time, even though our revenues are down by more than 50 percent.

By staying focused on our core functions, and maintaining a full workforce, we are able to keep producing the products and services you need to keep your business running smoothly.

The Xilinx People

Managing innovation is not difficult – just hire the best people and allow them to rock the boat and make some mistakes as they "push the envelope." What I've done at Xilinx is to create an environment and a management structure where people are encouraged to be creative and take risks; a place where decision making is fast and people feel they can contribute their best. As a result, our company has attracted the best people, we have consistently created the industry's most advanced products, and we have gained significant market share.

We describe ourselves as "competitors with heart" – that's part of the unique culture that we've created at Xilinx. Our employee attrition rate is a very low 5 percent, while other Silicon Valley companies average about 20 percent or more. In addition, we are among Fortune's "100 Best Companies To Work For" – an honor based largely on the reports from our employees. Our success demonstrates that you can build a company that is very aggressive in technology, and the first to market, while treating employees with respect.

When you choose Xilinx, you are not just buying the best devices, software, and services; you are also choosing a worthy partner. You'll find that Xilinx people are not only capable and friendly, we are also dedicated to your success – because we succeed together or not at all.

The Xilinx Technology

Most companies are lucky to be the best at one technology. However, due to the rigorous requirements of programmable logic, we must master both silicon and software – and in our industry, Xilinx is the recognized leader in both. We are producing some of the most advanced silicon in the world, and our software teams are advancing the state of the art with device programming tools and intellectual property (cores). We are achieving dramatic results on all fronts.

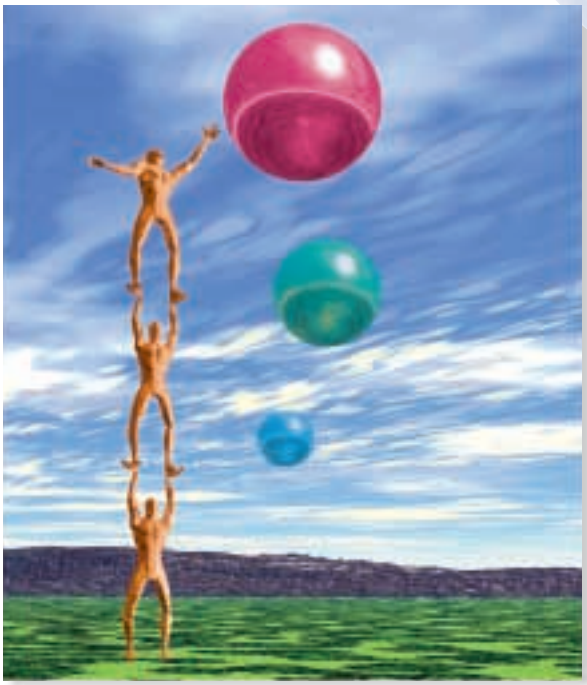

For example, five years ago our biggest device had 50,000 system gates, and was running at a clock rate of 20 MHz. Today, we are delivering devices with more than 6,000,000 system gates, running at a clock rate of over 200 MHz. Devices with 10,000,000 system gates, running faster than 300 MHz are on the horizon, and our technology plan calls for devices with over 50,000,000 system gates in the next few years. This is tremendous progress by anyone's standards.

While many people think of Xilinx as primarily a chip manufacturer, in fact more than half of our technical employees are software engineers. Each time we double the density of our devices, our software must also double in speed to keep your productivity high. A few years ago it took our tools about 60 minutes to compile a 10,000 system gate design. With the latest 4.1i release of our software, it takes only 30 minutes to compile a 1,000,000 system gate design – giving you a huge improvement in your productivity. Our goal is to continue to increase the efficiency of our software tools by at least 50 percent a year.

We can make this kind of progress because we have developed a tightly coupled product development system where different teams create silicon, software, and intellectual property in coordination. In each discipline alone, we could never have achieved these breakthroughs.

You can be assured that all of our devices, our software, and our intellectual property will give you the best possible performance, because they were all designed, from the beginning, to work together.

Conclusion

When I joined Xilinx, programmable logic technology was still in its infancy, but I realized it was going to have a very bright future. Today programmable logic is one of the fastest growing segments of the semiconductor business, and it's growing faster than the overall semiconductor market.

I expect Xilinx growth to continue at an average of 30 percent per year for the next several years. And in an odd way, we are actually benefiting from the current economic slowdown because programmable logic provides a cost-effective, flexible alternative to ASICs. As a result, we are gaining many new customers. We previously expected to be a $5 billion company by the year 2003, and now it will take a little longer. Yet, I'm still confident we will achieve that goal.

Xilinx is truly unique company. Because of our values, our trust, and our confidence in each other, we can get through these tough times and emerge stronger than before.

ONE PLATFORM. ONE SOFTWARE. TOTAL SOLUTIONS.

VIRTEX®-II

Over 300 MHz Design Capability
Densities up to 6,000,000 Gates Available Now
10 Gigabit Ethernet Plus Other High-Speed Interfaces

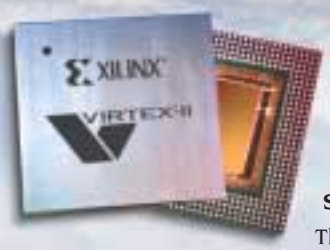

The Virtex®-II Platform FPGA family dominates the world of programmable logic system design. Once again Xilinx delivers. Right now we are shipping the fastest FPGA ever produced in a wide range of densities from 40K to 6,000,000 gates.

Software with All the Speed You Need

The superior fabric of Virtex-II Platform FPGAs is now complemented by the industry's fastest software: ISE 4.1i. Offering a single software solution for all Xilinx devices, ISE 4.1i delivers lightning fast clock speeds over 300 MHz. ASIC designers can also take advantage of the first and only formal verification tools for FPGAs. The performance speaks for itself: run times of 100K gates in one minute, or 3,000,000 gates in an hour.

The Platform for System Solutions

Our unique digitally controlled impedance technology—another world's first—has solved signal integrity issues. With on-chip multipliers, Xtreme DSP delivers over 600 billion MACs/s. SystemIO™ interfaces, including our 10 Gigabit Ethernet MAC core, bridge emerging standards and address all aspects of system connectivity. The Virtex-II IP Immersion™ architecture gives you the most flexible platform for embedding IP. And our new 70 D-MIPS MicroBlaze™ soft processor is the fastest on the market today.

The industry leader in speed and density, Virtex-II Platform FPGAs are available now, shipping now, and changing the world of system design . . . *now*.

XILINX®
The Programmable Logic Company℠

www.xilinx.com

FORTUNE® 2001
100 BEST COMPANIES TO WORK FOR

Leveraging ASIC Expertise for Platform FPGAs

As Platform FPGAs encroach into ASIC territory, the need for ever more complex synthesis tools and strategic alliances with industry-leading EDA vendors becomes essential.

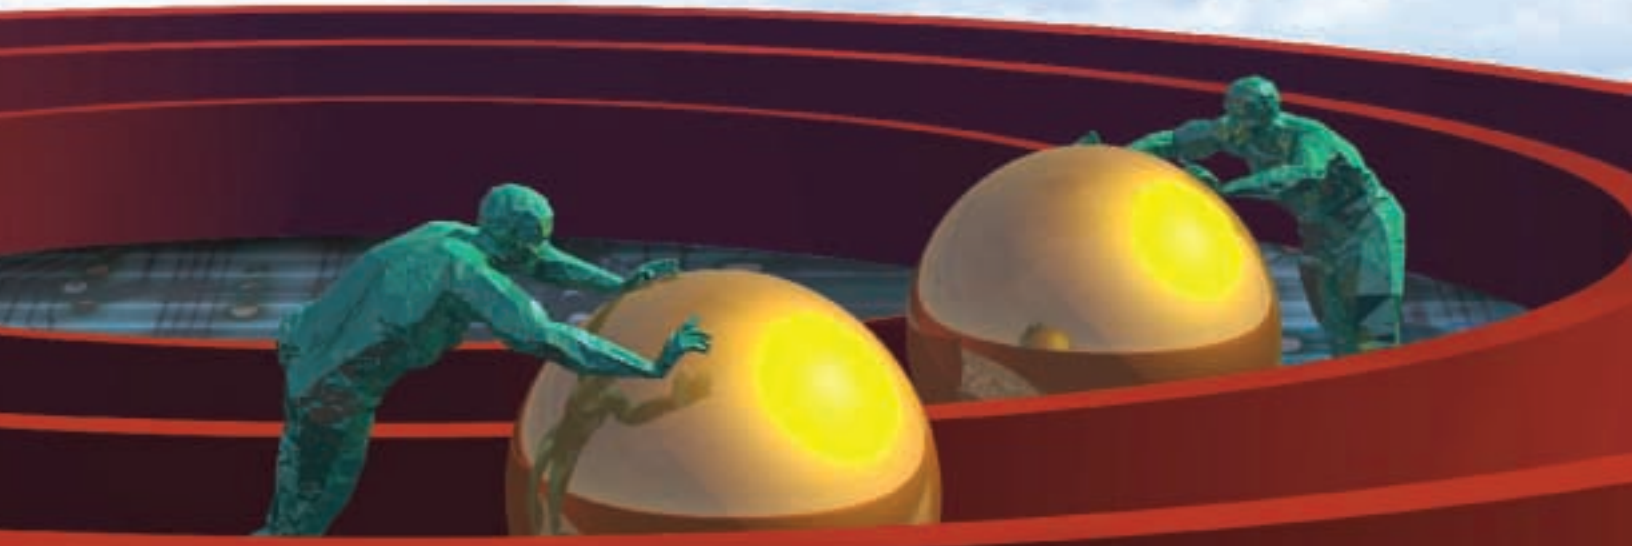

by Sanjiv Kaul
Senior Vice President and General Manager
Physical Synthesis Business Unit, Synopsys Inc.
www.synopsys.com

In a relatively short period of time, programmable logic has evolved from simple glue logic to system-on-a-chip (SoC) prototypes to Platform FPGAs for complex system applications. These Platform FPGAs, as exemplified by the Xilinx® Virtex™-II family, are being targeted for applications that were once solely the domain of ASICs. Now, designers from many different industry segments – communications, medical imaging, graphics processing, and consumer electronics – work with FPGAs.

As devices and applications grow in size and complexity, designers are increasingly applying ASIC-like design and verification techniques to take advantage of the new array of capabilities of Platform FPGAs. The reasons for this vary with the designer's perspective. In some cases, a team verifying an ASIC SoC wants the Platform FPGA prototype to go through the same flow. A designer who used an ASIC on the last project might want to use familiar and proven tools. Many designers who are used to simpler FPGA tools will want to know how ASIC designers have addressed the challenges now faced by Platform FPGA designers. For all these reasons, it's important to ask what has made ASIC and standard cell designers successful in the face of burgeoning complexity and how their EDA tools and flows can apply to Platform FPGA design.

More Sophisticated Synthesis

For smaller designs, push button synthesis tools have met the needs of FPGA designers. As FPGAs get into the million-gate-plus count and support complex clocking schemes, much more sophisticated synthesis technology is required. ASIC designers are accustomed to employing much more sophisticated synthesis methodologies than the current FPGA synthesis tools provide.

More often than not, ASIC designers realize that the fastest path to silicon is to assert a great deal of control over the synthesis process to meet performance goals. Designers of complex ASICs occasionally use a top-down synthesis methodology to handle time budgeting and to give them an overall idea if the design will meet timing goals. This top-down methodology can produce quick results if the design is not pushing the device's performance limits.

Designers at the leading edge of the performance curve, however, also want a bottom-up synthesis capability to focus on the blocks where timing and/or area are tight. During block synthesis, they need precise control of the design hierarchy and a wide

range of design constraints. As designers use FPGAs to implement more challenging designs, they find that only some of the features they need currently exist in today's FPGA synthesis tools. They are demanding more ASIC-like capabilities to get their jobs done.

Combining Synthesis and Placement

One of the hottest trends in standard cell design is physical synthesis, and this area is heating up for Platform FPGA design as well. True physical synthesis offers more productivity and better performance by integrating synthesis and placement into a single optimization step. By incorporating placement, the optimization algorithms get more accurate timing data for successful timing closure. Currently, FPGA designs are synthesized and placed in two distinct steps. However, even when data is passed back to synthesis after placement, it is difficult to achieve timing closure on the largest designs. A true physical synthesis solution is needed for FPGA designers to reach timing closure and realize fast time to market.

Successful FPGA and ASIC designers also strive to do two things with their HDL code: make their HDL code as technology-independent as possible, and verify functionality first – because it costs less to fix an error detected early in the design process. The latest in technology independence is power management in SoC designs. Many SoCs use clock-gating schemes in a standard cell implementation. To implement or prototype the same functionality in an FPGA requires clock enables. A complete ASIC tool flow allows the designer to write technology-independent code, which relies on power synthesis tools to insert the clock gating for the standard cell SoC. FPGA synthesis then uses that same code to target the clock enables on the FPGA without the need for any additional technology conversion steps.

Complexity Demands Robust Verification Flow

Platform FPGAs also need ASIC-like performance in the verification stages of the design. The team that verifies a design often doesn't care whether the final implementation is in standard cell or FPGA technology, they just want to prove it works. That means that the same speed and accuracy that they have come to expect from today's ASIC verification tools must be available, regardless of the final device's implementation. This verification process must also feature test bench generation to handle the soaring gate counts and complex vectors used to verify Platform FPGAs.

Many of the same verification bottlenecks brought on by deep submicron processes also apply to complex FPGAs. Verification engineers need and demand formal verification tools to avoid multiple iterations that can stall a design. Formal verification – in particular the equivalence checking method of formal verification – has recently entered the mainstream of high-capacity, complex designs. This mathematical proof of functional equivalence between two versions of a design can comprehensively verify in a matter of minutes or hours a design that would take weeks to verify using a gate-level simulator. However, to be truly effective as a replacement for gate-level simulation in today's design flows, a formal verification tool must offer more than capacity and speed. It must also provide fast and easy debugging of design errors, it must be usable in a broad array of design types and applications, and it must be easy to integrate into existing design flows.

All verification tools must also provide a feedback loop into implementation, or more valuable design time will be lost. Some of today's formal verification tools meet these requirements, but Platform FPGA designers must be sure that their entire verification environment is suited to verify complex designs.

Another way to reduce the need for gate level simulation is static timing analysis – a technique that both FPGA and ASIC designers have been using for quite some time. Static timing analysis checks the delay on all paths of a chip, providing an exhaustive check of timing to a given set of constraints. What's new for the Platform FPGA world is the need to model the timing of complex IP such as processors and ROMs. FPGA designers also need the application of advanced debugging aids to quickly locate the root cause of timing issues – and feed constraints back to earlier stages of the design process.

Conclusion

For designers to meet their goals with Platform FPGAs, EDA, and FPGA vendors must work together closely to ensure that the ASIC design success story can translate into Platform FPGA technology. Synopsys is working closely with Xilinx to ensure that designers who use Platform FPGAs and supporting EDA technology can get the same performance, reliability, and ease of use that the ASIC design community has enjoyed. As a premier EDA company and synthesis pioneer, Synopsys is well poised to apply its extensive knowledge of ASIC design tools and flows to the Platform FPGA arena. Working with Xilinx and its mutual customers, Synopsys will constantly strive to accommodate the myriad changes that will undoubtedly occur as Platform FPGAs take their place alongside ASICs for complex system design.

Xilinx ISE 4.1i Delivers the Speed You Need

Xilinx ISE 4.1i presents a new set of features and device support to give you the fastest time to market with the most advanced technologies available for FPGA design today.

by Lee Hansen
Software Product Marketing Manager
lee.hansen@xilinx.com

Steve Wenande
Senior Technical Marketing Engineer
steve.wenande@xilinx.com

David Blevins
Staff Technical Marketing Engineer
david.blevins@xilinx.com

Larry McKeogh
CPLD Software Marketing Manager
larry.mckeogh@xilinx.com

The new version 4.1i of the Xilinx Integrated Software Environment (ISE) design tool suite has hit the streets running. In this special section of Xcell, we will give you a detailed look at ISE 4.1i.

Among the improvements we have made are simplified product configurations and software runtimes up to two times faster than ISE 3.1i. We've added ProActive Timing Closure and expanded your choice of verification strategies to deliver the fastest performance possible and the best design flows available – despite the push towards ever-increasing design complexity and size. ISE 4.1i also includes upgrades of existing tools, such as ModelSim™ and ECS (Engineering Capture System) software.

ASIC designers looking to migrate to Platform FPGAs will be happy to find familiar ASIC tools from Synopsys®, such as PrimeTime™ static timing analysis and LEDA™ HDL analysis. Moreover, the ISE 4.1i release includes XPower – the industry's first power and thermal analysis tool for FPGAs and CPLDs.

With all this improvement and expanded functionality, ISE 4.1i will allow you to reliably get your product to market faster than your competition.

Simplified Configurations – A Single Look and Feel

ISE 4.1i has been simplified to four powerful configurations:

- ISE Alliance Series™ configuration – for EDA design flow integration
- ISE Foundation™ configuration – for the most robust single-vendor programmable logic design platform available
- ISE BaseX configuration – for a cost-effective, full-featured, PC-based programmable logic design environment for the standalone designer, and device size support up to 300K gates
- ISE WebPACK™ configuration – for a free, Web-delivered design environment for CPLD customers and entry-level FPGA designers, with device size support up to 300K gates.

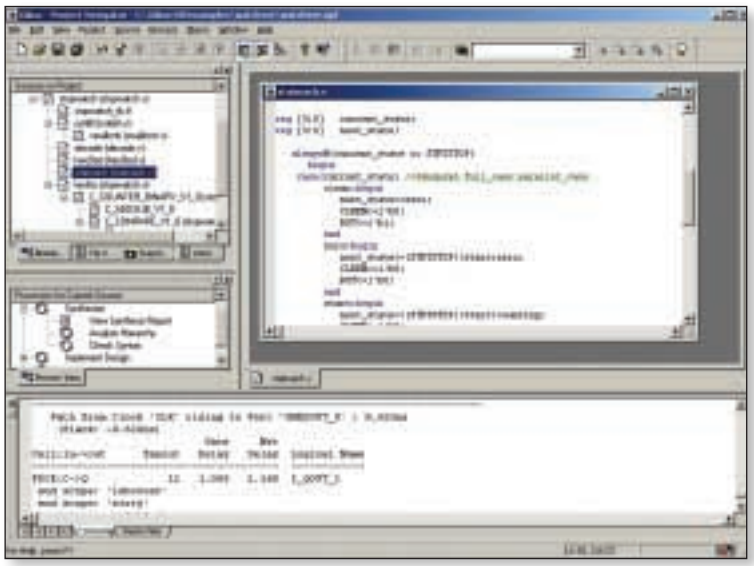

Figure 1 - ISE Project Navigator

Full Device Support in a Single Package

All Xilinx device families can be designed and programmed from the ISE family of products. This includes existing devices as well as those new FPGA and CPLD families scheduled to be released in the coming year.

ISE 4.1i also delivers a new, unified look and feel in the Alliance and Foundation configurations. By combining ISE Alliance and ISE Foundation configurations into a unified development effort, our customers realize the best performance and features in any of the ISE design configurations, and migration or upgrade between ISE packages is quick and easy.

Project Navigator Up Front

As shown in Figure 1, both ISE Alliance and ISE Foundation customers will now see Project Navigator for their design management tool. Project Navigator features several new powerful features, including integration with Exemplar's LeonardoSpectrum™ synthesis environment, Unix platform support, EDIF source support, and more. For Alliance customers who want to transition over the coming year, Design Manager will still be available in ISE 4.1i.

Installs Just Got Easier

ISE 4.1i contains significant improvements to the installation process. Immediately, you will notice ISE 4.1i searches for earlier versions of ISE and offers to remove the older versions of software, helping to free up disk space – or ISE 4.1i allows you to save your older ISE versions.

Adding device support is quick and easy. You can add support for new device families incrementally at a later date simply by re-running the 4.1i install program.

50% Faster Runtimes – Again

The Xilinx Intellectual Property, Services and Software Group operates under a mandate to deliver 2X better software execution with each major release, and ISE 4.1i is no exception. Benchmarks indicate you will see as much as 50% faster design compilation – that means as many as 100,000 gates compiled per minute. And that translates to millions of gates compiled per hour. This blazing design speed outperforms all other logic vendors.

Maximum Performance Through ProActive Timing Closure

If you take a design compiled in ISE 3.1i and run it on ISE 4.1i, you will see impressive gains in device performance (Fmax) through a new collection of timing technologies known as ProActive Timing Closure. Collectively, the advances in ProActive Timing Closure technology can deliver as much as 133% better Quality of Results (QoR) than our previous major software version:

- As the industry's most advanced timing-driven, place-and-route technology available, ProActive Timing Closure brings a new level of active control to design timing. Placement algorithms scan the design data paths and place critical paths first, helping to reduce timing delays. A new extra-effort routing mode remembers which paths in

an implementation pass successfully met timing requirements, and then transparently "re-places" and "re-routes" unsuccessful areas to attempt to meet overall timing goals. And our new Directed Routing brings predictable and repeatable place-and-route capabilities to IP cores.

- ProActive Timing Closure further expands physical synthesis for programmable logic by using physical timing information to ensure that optimization during synthesis is focused on the critical paths. ISE 4.1i works with Synplicity® Inc.'s Amplify™ software and new TOPS™ (Total Optimization Physical Synthesis) capability to further source this physical timing information. [See "ProActive Timing Closure Delivers up to 133% Better Device Performance" in this issue of *Xcell* for more information on ISE physical synthesis.]
- Timing and HDL interaction now occur at several levels of implementation, offering design-specific tips that suggest modifications to the design, constraints, or source code, to help meet the design's timing requirements.
- ISE 4.1i offers timing cross-probing from the timing analysis report to either the Xilinx Floorplanner or to third party synthesis "technology viewers." Cross-probing displays exactly where the problems are in the design in the physical floorplan of the FPGA, the logic gates which were created during synthesis, and even the specific HDL source code itself, significantly reducing debug time.

Expanded Verification Strategies

As designs grown in complexity and size, verification becomes a greater challenge. ISE 4.1i has been designed to expand the options you have for checkpoint verification. At each stage of the design cycle, you can check your design for accuracy.

First Formal Verification Capability for Programmable Logic

ISE 4.1i offers integration with Synopsys Inc.'s Formality™ and Verplex Systems Inc.'s Conformal™-LEC equivalency checkers, leveraging the same technology that was adopted in past years to check high-density ASICs. In the equivalence method of testing, design passes can be checked in "blocks of logic" against a previous known-good version. This can occur at any point in the design cycle, particularly in post-synthesis and place-and-route passes. This strategy offers a rapid checking method that is proving invaluable to high-density design work. For Virtex-II designers moving to 1-million gate designs and above, this new formal verification integration provides a valuable step in the Platform FPGA initiative. [See "High-Performance Platform FPGAs Now Need Formal Verification" in this issue of *Xcell*.]

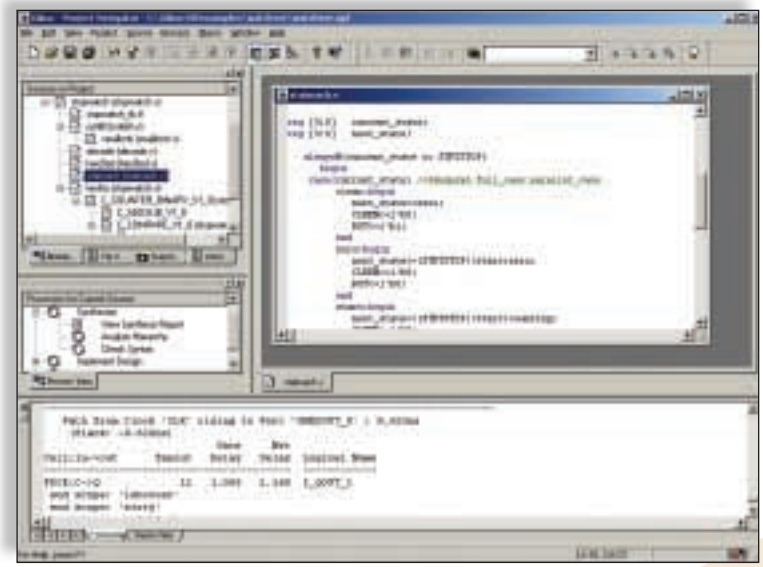

Figure 2 - ISE Engineering Capture System (ECS)

ISE 4.1i Supports ModelSim Version 5.5

ISE 4.1i also includes support for the 5.5 family of ModelSim HDL simulators from Model Technology, a subsidiary of Mentor Graphics Inc. The ModelSim Xilinx Edition (XE) simulator, has been upgraded to version 5.5 as well. These new versions of ModelSim software add valuable library updates, and they feature support to keep Xilinx ISE 4.1i at the forefront of HDL simulation use with:

- Library importer wizard
- New library view window
- An array of bug fixes.

For convenience, the HDL simulation libraries are now broken out as either CPLD- or FPGA-centric, based on either VHDL or Verilog hardware description languages. This separation of libraries allows the smallest possible download size for the desired design application.

ECS Schematic Capture

When upgrading to the ISE Foundation 4.1i software from Foundation Series, customers will notice a new change in their schematic capture and block diagram editor. Shown in Figure 2, ECS is now the schematic capture product for all software configurations. The Xilinx ECS application provides the designer with a powerful graphical input tool for creating schematics and block diagrams. Block diagrams can be created and used as the "top level" description of an HDL design. The designer can describe logic modules in the HDL editor, using VHDL or Verilog, and then use ECS to auto-generate graphical block symbols for each module. Designers can then instantiate and connect the symbols to create a block diagram representation of the design. Prior to synthesis and implementation, the diagram is converted into an HDL netlist that describes the blocks' interconnectivity, which (along with the HDL modules) will be used for subsequent synthesis and implementation in the target Xilinx device.

ECS also supports traditional schematic-based PLD design. Symbol libraries are provided for each Xilinx device architec-

ture, which allows the designer to define a design using flip-flops, counters, gates, and other mechanisms. Symbols come in two varieties – "primitive" components such as a two-input AND gate, and "macro" symbols for more complex functionality, such as counters. ECS now not only provides enhanced functionality and user friendliness, but it will also serve as a framework for future graphics-intensive features in ISE Foundation software.

Some of the new features in ECS include:

- **Multi-window interface** – Multiple schematic and/or symbol windows can be open simultaneously.
- **Enhanced object attribute handling system** – The designer can now arbitrarily and easily create custom attributes for objects (e.g. wires or symbols). Also, attributes for all currently selected objects can be viewed/edited at the same time.
- **Right mouse button menu support** – You can accelerate design tasks by making editing actions immediately available for currently selected object(s).
- **Enhanced design rule checking** – Objects that result in warnings or errors are highlighted on the schematic sheet so they can be easily found.
- **Preferences menu** – Rather than editing a text initialization file, you can now choose application settings and user preferences from a menu.
- **Customization** – Now you can create sheet sizes to your exact specifications.
- **Support for Windows® Clipboard** – For design documentation, you can cut-and-paste diagrams and schematics (or portions thereof) into other applications such as Microsoft Word.
- **Print Preview** – Now you can see – and adjust – how a printout will look before you print it out.

FPGA Tools Powerful Enough for ASIC Users

As the Platform FPGA initiative has reached full steam in the industry, we have seen an influx of ASIC designers converting to FPGAs – both for prototyping and for full-production logic. In order to facilitate the migration of ASIC designers to FPGAs, we have provided additional tools that ASIC designers have come to know and love.

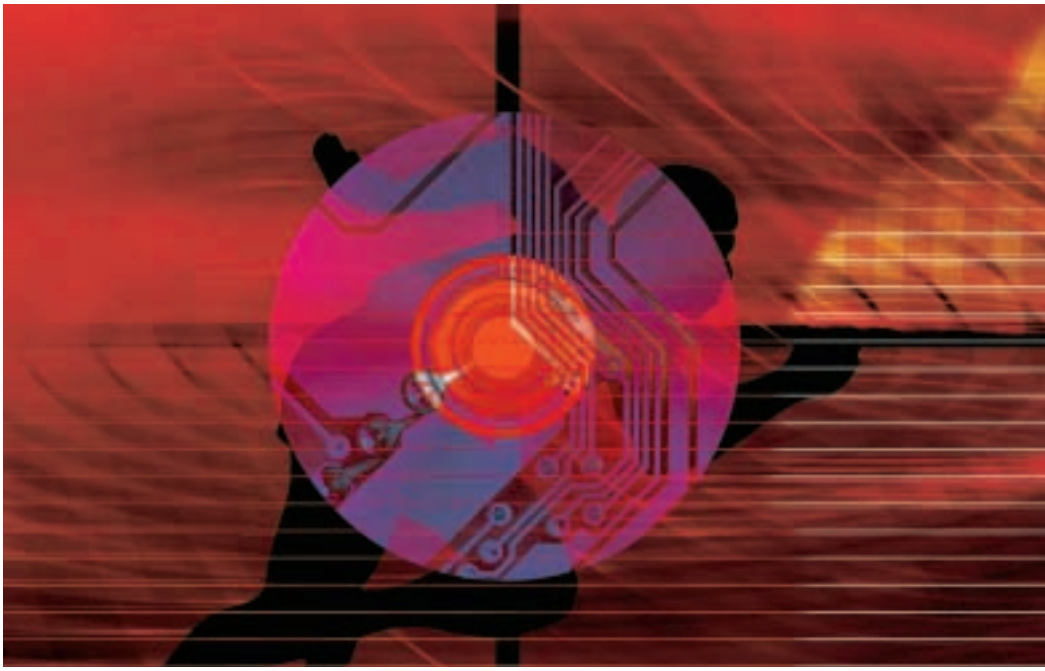

XPower

This release of ISE 4.1i includes the XPower power analysis product – the industry's first power analysis tool with thermal analysis and reporting capability. XPower lets you see early on in the design cycle how the logic activity will affect the thermal characteristics of Xilinx FPGAs and CPLDs. XPower software is proving invaluable in the high-density design world of Virtex-II devices where multiple FPGAs may be operating at 90% capacity or higher on a single board. Likewise, XPower excels in the low-power world of CoolRunner™ CPLDs where battery life and low power consumption are critical to product success.

Synopsys PrimeTime Support

ASIC designers who are more familiar with the Synopsys PrimeTime™ tool for static timing analysis can now use it for their FPGA designs in ISE 4.1i. The Xilinx Timing Analyzer, also delivered in ISE 4.1i, remains the golden signoff for static timing analysis of Xilinx FPGA designs.

Synopsys LEDA-HDL Analysis

ProActive Timing Closure includes HDL analysis capabilities from within XST (Xilinx Synthesis Technology) reports, and ASIC designers can use the Synopsys LEDA™ family of HDL products with ISE 4.1i if they prefer. The LEDA design analyzers not only provide HDL analysis and suggestions on how to improve HDL source code, but they also offer the flexibility of customization to support corporate design standards. Customers can use the ProActive Timing Closure HDL analysis capabilities and the Synopsys LEDA tools together to deliver programmable synthesis-ready source code earlier, reduce extraneous logic cells, and get to timing closure faster with better overall results.

Conclusion

If speed is what you need, Xilinx ISE 4.1i will put you on the fast track for time to market. Simply put, the ISE 4.1i design tool suite will enable you to realize the fastest device speeds and highest design performance available in the industry for both low-density and high-density designs – and that's not all. To learn even more about the enhanced capabilities of ISE 4.1i, or the software upgrade process, go to *www.xilinx.com/ise/xcell/*. Upgrade to ISE 4.1i before your competition does.

High-Performance Platform FPGAs Now Need Formal Verification

Verplex teams with Xilinx to include ASIC-level equivalency checking in ISE 4.1i for formal verification of large Platform FPGA designs.

by Bassilios C. Petrakis
Business Development Manager
Verplex Systems Inc.
petrak@verplex.com

Historically, FPGAs were small devices enabling ASIC designers to prove designs while being able to experiment with various programming options until they got it right. Design closure was relatively simple.

Today's Platform FPGAs, however, have reached ASIC proportions with as many as 6 million system gates running up to 400 MHz. High performance synthesis engines, floorplanning, and physical synthesis tools are now necessary to achieve timing closure.

These synthesis and physical synthesis tools perform a high number of optimizations that drastically change design structures. Although these drastic changes are necessary to meet the design requirements, it is practically impossible to exhaustively prove with traditional simulation alone that a chip will work under all conditions. It would just take too much time to write and simulate the vast number of vectors required to verify all the possible combinations of states, events, and inputs for a typical design.

The speed, volume, complexity, and size of Xilinx® Platform FPGAs therefore demand an ASIC-level solution to FPGA design validation. That solution is an independent verification engine to audit or validate the changes. The independence of the verification engine is critical. To obtain the highest degree of confidence in the design, the algorithms and methods used in verification must be from a different company with a different tool.

Formal Verification

Formal verification tools address the design intent validation and design implementation functional closure. In support of the growing needs of large FPGA design verification, Xilinx has established a pivotal working relationship with Verplex Systems Inc. for a formal verification solution. Xilinx and Verplex have collaborated to establish a design methodology that guarantees interoperability between the Xilinx Integrated Software Environment (ISE) 4.1.i and the Verplex Conformal™ Logic Equivalence Checker (LEC).

Now in the FPGA flow, instead of completely relying on simulation by applying stimuli to a design and comparing its responses with expected results, formal verification examines the design mathematically, and proves its functional properties without test vectors. Unlike simulation, which will most likely overlook bugs for the lack of appropriate vector patterns or only report a minimal number of bugs based on limited input stimuli, formal verification will exhaustively prove the design while isolating bugs. When errors are found, formal verification can generate counter examples to demonstrate error conditions, allowing designers to complete verification tasks faster.

For design implementation verification, Conformal LEC offers superior verification with speed, capacity, and most important, 100% coverage. It also offers a more productive debugging environment over gate-level simulators. Conformal LEC compares designs in a matter of minutes or hours instead of the days or weeks that is required using traditional gate-level-simulation techniques. Conformal LEC, being an independent verification tool, boosts the designer's confidence that the design has been synthesized and optimized correctly.

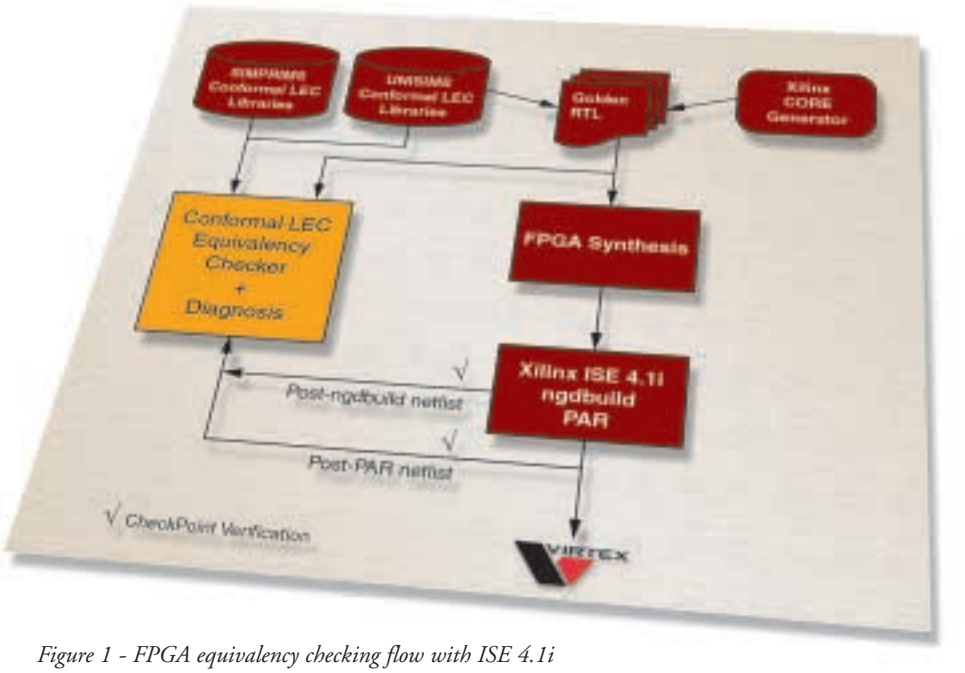

Figure 1 - FPGA equivalency checking flow with ISE 4.1i

Uses of Equivalence Checking in the FPGA Design Flow

Equivalency checking (EC) detects functional inconsistencies between two design representations. One of the two designs is considered to be the "Golden" design, typically described in register transfer-level (RTL) code. The other is the revised gate-level implementation. The goal of running equivalency checking on an FPGA design is to make sure the final design implementation does what the RTL code specifies it should do. EC can be used throughout the FPGA implementation process to:

- Compare a Golden RTL model against a partitioned design. This type of design partitioning occurs in ASIC prototyping using multiple FPGAs.
- Compare a Golden RTL model against a modified version of the model. It is often the case that an RTL model is altered to improve the performance of the design.
- Compare a Golden RTL model to the post-synthesis (post-ngdbuild) netlist. In the Xilinx flow, the netlist from the synthesis tool is first processed by ngdbuild in ISE 4.1i. The netlist generated, which is SIMPRIMS-based, is considered the post-synthesis netlist for the purpose of equivalence checking.
- Compare a Golden RTL model against the final post place-and-route (PAR) netlist (Figure 1). This is the final functional verification step before proceeding with programming the device.

Together with Xilinx ISE Foundation™ 4.1i, Conformal LEC can provide equivalence checking at the RTL and gate level of the design flow to functionally verify designs at every checkpoint.

Running Conformal LEC

It used to be that equivalency checkers were very difficult to set up and run. Typically, it would require a handful of expert verification engineers.

Verplex has set a precedent by creating a tool without the legacy limitations of earlier equivalence checkers. Conformal LEC was designed to tackle large designs (20M+ gates). It was created from the outset with ease of use and high performance in mind. The target user for Conformal LEC is the

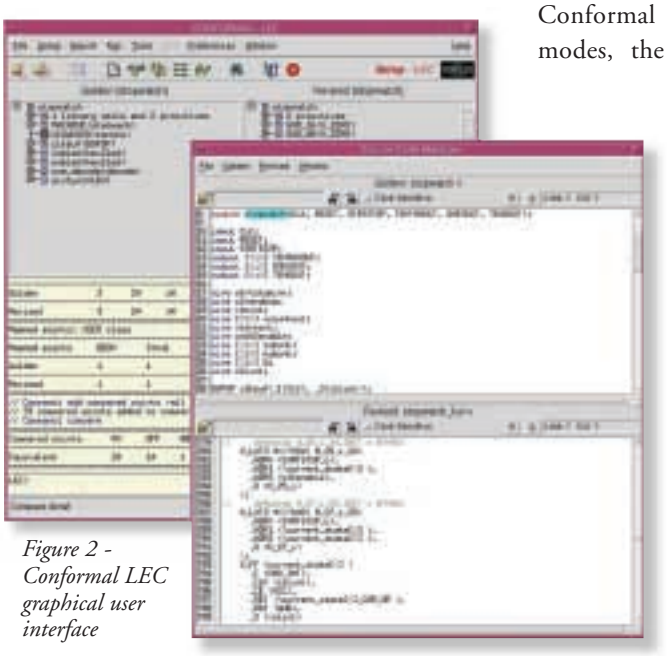

Figure 2 - Conformal LEC graphical user interface

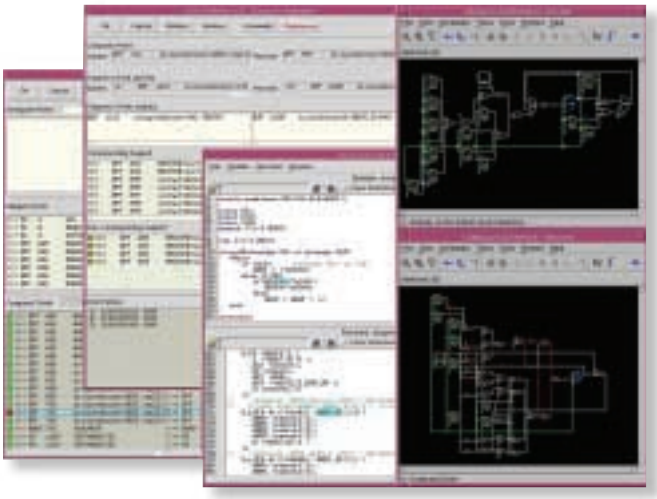

Figure 3 - Conformal LEC Diagnosis Manager

designer who is coding the entire design or a subset of a design. Conformal LEC requires minimal user setup. It runs in GUI interactive mode (Figure 2), command text mode, as well as in batch mode (using a command file).

Conformal LEC accepts Verilog, VHDL, and mixed-language RTL designs. Conformal LEC reads Verilog, VHDL, and EDIF gate netlists and supports Verilog simulation and Synopsys Liberty™ library models. Xilinx UNISIMS and SIMPRIMS libraries for Conformal LEC are now available in ISE 4.i software.

Conformal LEC operates in two modes, the "setup" mode and the "LEC" mode.

Setup Mode

In the setup mode, the Golden RTL and the revised (post-ngdbuild or post-PAR) designs are read into Conformal LEC along with the Xilinx UNISIMS and SIMPRIMS Conformal LEC libraries. The post-ngdbuild and post-PAR netlists can be generated by ndg2ver in ISE 4.1i

LEC Mode

In the LEC mode, the tool performs design flattening, remodeling, key-point mapping and compare.

Key-point mapping involves partitioning the two designs respectively by state element, input, output, tri-state, combinational feedback loop cuts, and black box instances, then automatically pairing up (mapping) the corresponding instances.

Once the design is fully mapped, Conformal LEC compares the corresponding logic cones of each map point. When all mapped points are found equivalent, then the whole design is equivalent. If one or more map or key-points is non-equivalent, the tool offers an intuitive debugging environment to help you understand the source of the mismatch (Figure 3).

The three RTL-to-Xilinx equivalent gate designs in Table 1 highlight the run time performance of Conformal LEC. The gates size for each design does not include memories.

Conclusion

RTL-to-gate equivalency checking has been used successfully in ASIC design verification for many years. It has been widely adopted as the mainstream verification step in the ASIC design flow. With the growth of FPGA devices for system applications, the need for verification for FPGA-to-ASIC and ASIC-to-FPGA conversions has become critical. The low cost of FPGA design will push the device capacities. "Reprogrammable" prototyping has reached it limits. Expanding into formal verification has now become an essential factor in the FPGA design flow.

Common uses for both ASIC and FPGA designs are growing. Commonality in both design flows is merging. One of the mutual meeting points in the design flows will be formal verification, and in particular, equivalence checking. Conformal LEC is available today to meet that challenge.

As Xilinx continues to meet the future demands in FPGA design, Verplex will provide independent verification tools to meet those same demands. For more information on formal verification and equivalency checking, go to *www.verplex.com.*

Test Case	Type	Design Size (gates)	CPU Time (sec.)	Memory (Mbytes)
1	RTL vs. Post-Route	50K	19.08	67.22
2	RTL vs. Post-Route (networking design)	720K	305.47	395.73
3	RTL vs. Post-Route	241K	297.89	180.16

Work Station: Ultra 80 Sun Solaris machine with 2G RAM

Table 1 - Test case examples and results

WebFITTER Gets a Face-Lift

The WebFITTER design tool has a new look and feel – and increased functionality.

by Larry McKeogh
CPLD Software Marketing Manager
larry.mckeogh@xilinx.com

As an online interface for Xilinx CPLD implementation tools, the latest upgrade of the WebFITTER™ design utility makes the WebFITTER interface more flexible and extensible. Accessible from anywhere on the World Wide Web, the free WebFITTER software supports all Xilinx CPLD devices, such as the high-speed 9500 family or the ultra low power CoolRunner™ devices. The WebFITTER tool accepts HDL design files, EDIF netlists, as well as Altera or Lattice netlist files.

A New GUI

As shown in Figure 1, the WebFITTER graphical user interface has a new tool layout and navigation scheme:

1. **Primary Tool Navigation** – *This navigation layout allows you to instantly access the main features of the WebFITTER design environment: displaying current projects, starting a new project, accessing tutorials, and context sensitive help.*

2. **Main work window** – *The WebFITTER work window facilitates the design submission, modification, and retrieval of pertinent design information.*

3. **Secondary Project Navigation** – *This navigation menu changes depending on whether you are submitting a design, browsing the projects available, or reviewing the results of a design submission.*

Figure 1- The new look of the WebFITTER interface

The Secondary Project Navigation menu provides the following options:

- Modifications to submitted designs and the addition of project notes are new for the WebFITTER tool. Other operations supported by the WebFITTER online utility are the ability to resave the project for another 14 days, delete a project that is no longer needed, view the output reports, and obtain online price quotes for CPLD devices.
- The project submission navigation menu leads you through the process of uploading the design files, selecting the desired CPLD device, modifying any design options applicable to the design type, and reviewing all options selected prior to submitting the design.
- Report viewing options allow you to quickly return to the project listing or access other operations while reviewing the design results. Such operations include navigation to the project listing, obtaining price quotes, viewing design notes, and downloading desired design files for storage on your local computer.

WebFITTER Features and Functions

Uploaded files are shown during the design submission process. The file listing can be dynamically adjusted, making large hierarchical types of designs easier to submit and modify.

The Web-based WebFITTER application presents tabular listing of current projects, the status of these projects, and the ability to modify these projects without having to resubmit the design files. This feature is useful when modifying the design fit options or adding files for a hierarchical design.

The improved WebFITTER functionality gives you the ability to add notes to the design either during the design process or after the design has been submitted. This feature also allows you to list the purpose of the design or keep track of any changes occurring between projects. Project design options can be saved and used again for future designs.

Conclusion

The enhancements to the free WebFITTER design environment make better use of the technology available on the Web today. These changes make it easier than ever to submit, modify, and review designs targeting Xilinx CPLDs. The new structure also allows Xilinx to implement new and innovative ideas over the Web in the future. We encourage you to visit the WebFITTER website at *www.xilinx.com/sxpresso/webfitter.htm* and test its ability to fit your CPLD needs.

ISE 4.1i

Xilinx Timing Analyzer Is Default Viewer for Static Timing Reports in ISE 4.1i

New, easy-to-use features have been added to the ISE 4.1i Timing Analyzer software to help you view and debug your static timing issues.

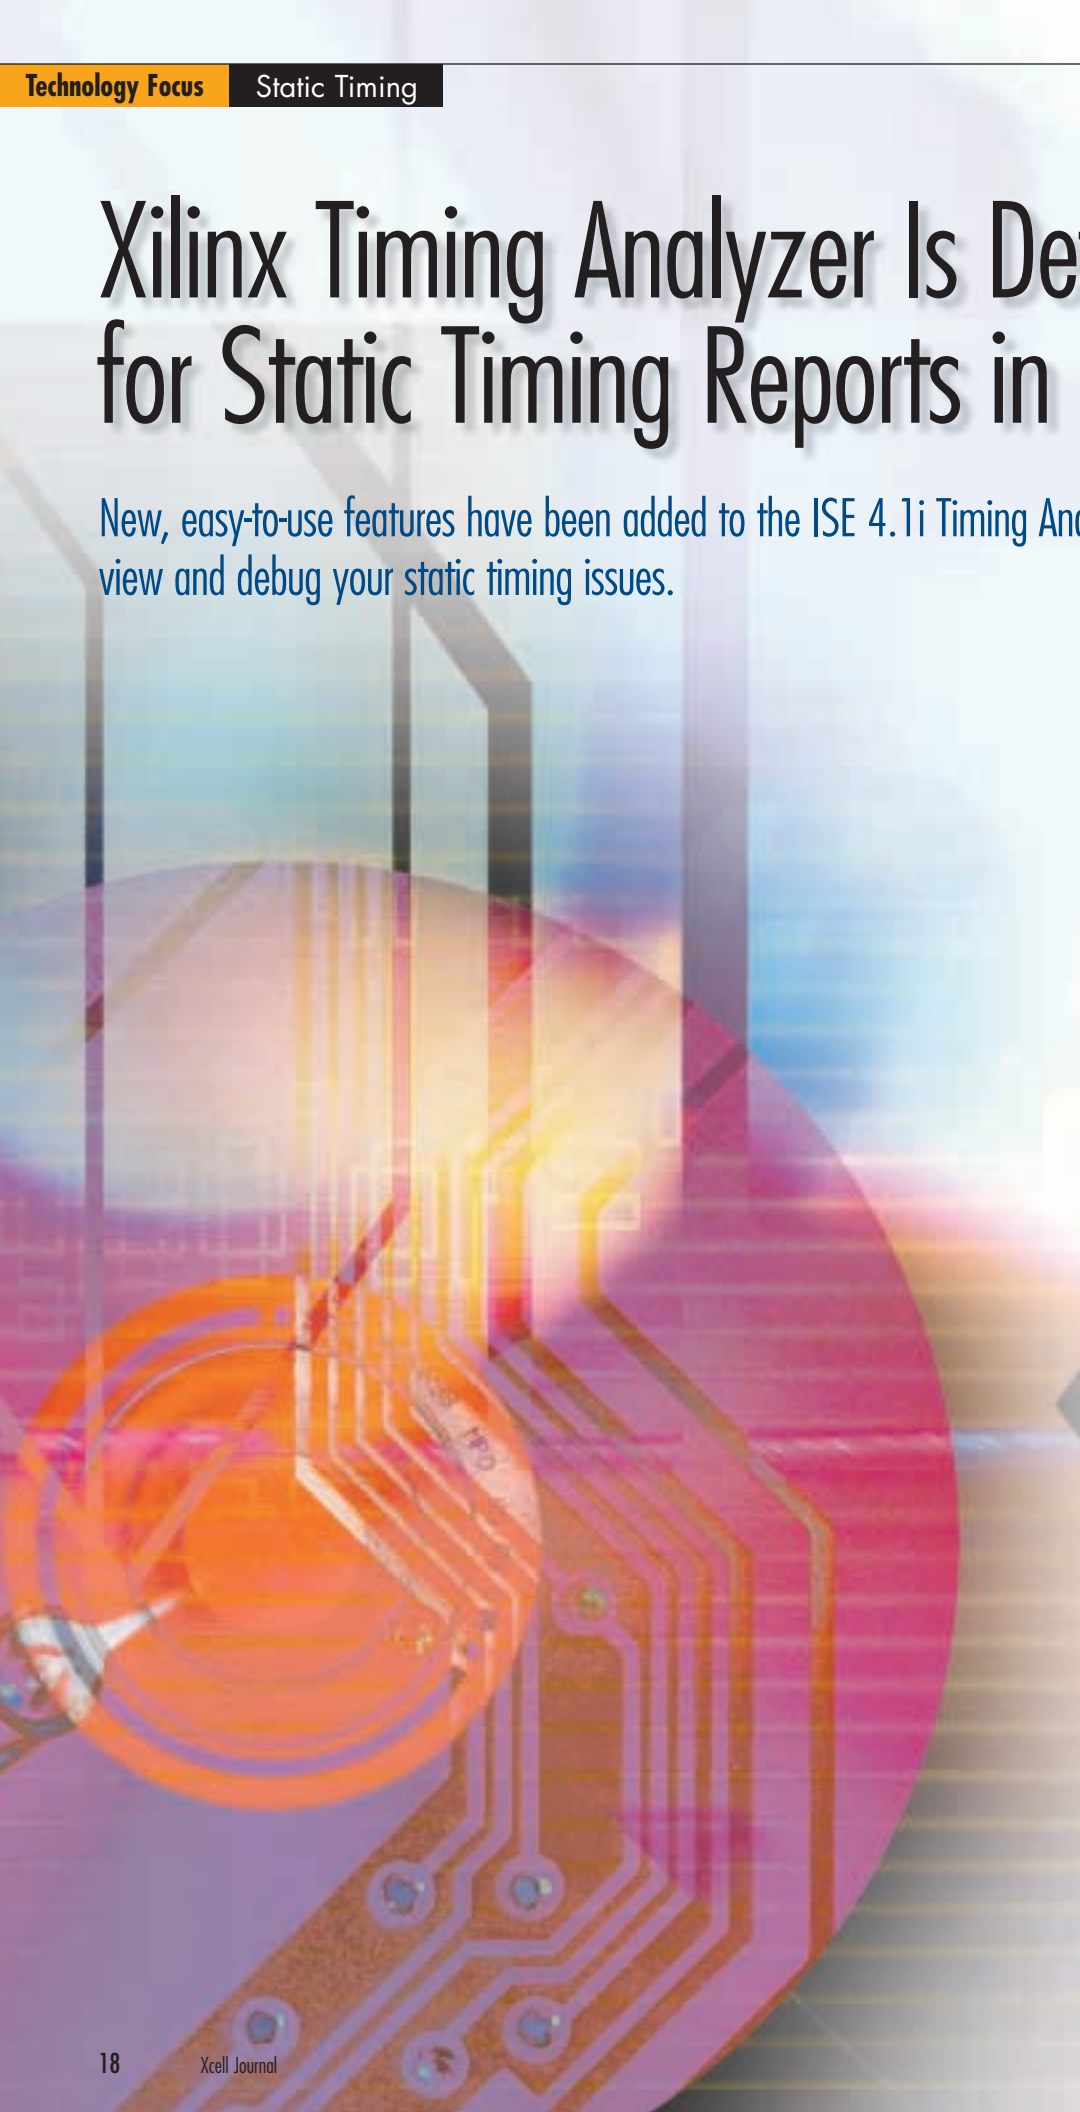

by Kate Meilicke
Technical Marketing Engineer
Kate.Meilicke@xilinx.com

Craig Cholvin
Technical Marketing Engineer
Craig.Cholvin@xilinx.com

The Xilinx Timing Analyzer software is a graphic user interface (GUI) for the static timing engine. In the past, Timing Analyzer was perceived as a tool for experts only. With the ISE 4.1i software release, the Timing Analyzer has become the default viewer for static timing reports – and a tool for every designer.

Timing Analyzer delivers reports that are easier to understand because important information is highlighted. Furthermore, Timing Analyzer has new ways to sort the data that help you find the critical paths and get suggestions on how to fix failing paths. Cross-probing timing paths to the Xilinx Floorplanner, Synplicity's Synplify™, and Exemplar Logic's LeonardoSpectrum™ design tools have also been enhanced. These and other new features greatly simplify the resolution of static timing issues.

Report: Just the Facts

In the past, too many customers complained that the ISE 3.1i static timing report had too much information and was too confusing. Responding to that feedback, we slimmed down the timing report to show, by default, only the most important information. Expert users who want to see all the details, however, can enable that information in the Timing Analyzer viewer.

Analyzing the Timing Report

Figure 1 shows the new timing report format with a period constraint that was created during the Translate (ngdbuild) step. TS_clk25 is the original TIMESPEC with a PERIOD constraint of 40 ns. In this case, clk25g has the same timing as the original clock (TS_clk25 / 1.000000), therefore, it also has a period of 40 ns.

The heading for each path has more information. The amount of slack and the equation used to calculate the slack is shown first. Next, you see the source and destination using the logical name (the name in the design). The source clock is clk50g, which is twice as fast as the destination clock, clk25g, in this design. Because clk50g rises at 20 ns and clk25g rises at 40 ns the new requirement is 20 ns (40 ns – 20 ns). Only negative clock skew is used in the equations when calculating setup times.

After the header, a hyperlink highlights the path in the Floorplanner. Following that is a simplified path. The physical CLB (configurable logic blocks) locations are not shown. Only logical names are used.

Looking for Details

The simplified static timing report omits certain details, but all the information is still there. If you're used to working with all the information in the ISE 3.1i report, you can find the details in the .twr file (ASCII report file), and you can turn on the details in Timing Analyzer under File > Preferences. The Timegroup information can be generated in the Timing Analyzer under Analyze -> Query Timegroup.

Improved Period Constraint

In ISE 4.1i, paths between "unrelated" clocks are not covered by a PERIOD constraint. Clocks can be identified as related in the timing constraints language (user constraints file) or through the Constraints Editor. Paths between clocks from a DLL/DCM are automatically related when the Translate (ngdbuild) creates PERIOD constraints for each output of the DCM when the input clock is constrained with a PERIOD constraint. There is a new constraint keyword, "PHASE", that defines the phase relationship between two clocks.

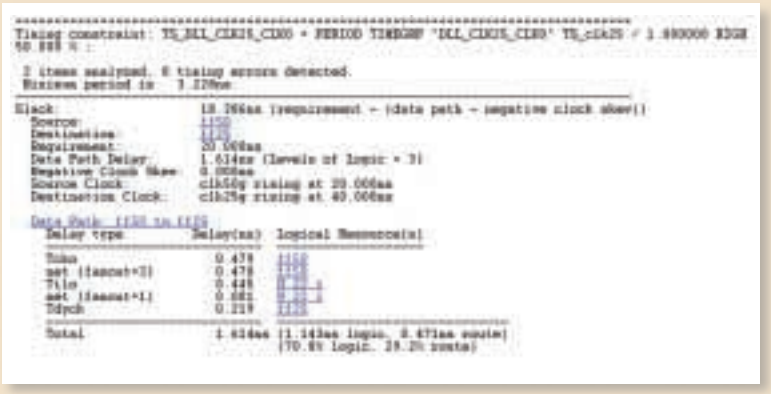

Figure 1 - New 4.1i Timing Report

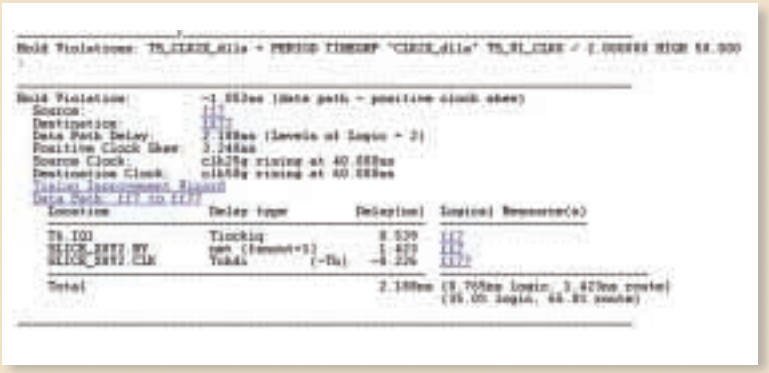

Figure 2 - Hold Violation Report

Related Clock Definition Examples

First clock definition is for the main clock:

```
TIMESPEC "TS_Clock1" = PERIOD
"Clock1" 20 ns HIGH 50 %;
```

The second clock is defined as 1/2 the period of Clock1:

```
TIMESPEC "TS_Clock2" = PERIOD
"Clock2" "TS_Clock1" / 2;
```

The third clock is defined as twice the period of Clock1 and offset by 5 ns:

```
TIMESPEC "TS_Clock3" = PERIOD "Clock3"
"TS_Clock1" * 2 PHASE + 5 ns;
```

Prevent Hold Violations

When the positive clock skew (delay to the destination clock – delay to the source clock) is greater than the data path delay, a "Hold Violation" is reported. The data from the source will no longer be valid at the destination register when the next clock edge arrives. See the example report in Figure 2.

Typically, a hold violation happens when a clock is not on a dedicated clock routing, such as the global clock networks or the low skew lines. The example in Figure 2 was created by routing the clock on local routing resource – which is poor design practice. All clocks must be on a clock resource to prevent hold violations.

Sorting Speeds Up Report Reading

"Slack" is the default sorting algorithm of timing paths in the updated Timing Analyzer. Sorting by slack allows the path with the worst timing to be presented at the top of the report. To accommodate requests for alternate sorting mechanisms, we have implemented an "index" window in the Timing Analyzer where you can select one or all of the timing constraints to sort (or reverse sort) on source name, destination name, or slack for paths. The index is useful when you want to see unique failing paths. A right mouse

ISE 4.1i

click over a time constraint will show the sorting options.

Debugging Static Timing Issues

When a failing path is displayed in the Timing Analyzer, a link to the Timing Improvement Wizard is presented. As shown in Figure 3, the wizard gives suggestions on how to improve the failing path based on information in the path.

Cross Probing

Cross probing is now available between the Timing Analyzer and the editor/viewers of Xilinx, Synplicity, and Exemplar Logic. Simply clicking on the path or element in the Timing Analyzer timing report (.twx) will "hilite" the path in the active editor, as seen in Figure 4 of the Floorplanner.

Tips and Tricks for ISE 4.1i Timing Analyzer

1. If you are looking at the timing report, but also want to find out more information about a specific path, the design must be loaded. This is done under File > Open Design and browse the design. Once the design is loaded, click on Analyze > Against User Specified Paths > by Defining Endpoints. This window will allow you to select specific sources and destinations with which to evaluate paths in a new timing report.
2. If there are no time constraints in your design, a report can be generated to show the longest path for each clock. Again, the design must be opened (File > Open Design), and then you must run Analyze > Against Auto Generated Design Constraints. This will automatically generate a timing report organized by clocks and their associated inputs and outputs.
3. A datasheet report appears at the end of the timing report. The first section of the datasheet shows the setup and hold times for each signal, and it is organized by clocks. The setup and hold times are reported with respect to the pins of the chip – not at the registers themselves. This means the clock and data delays have been incorporated into the numbers. A negative setup typically indicates a long delay on the clock net and usually a hold requirement. This means the data can show up after the clock at the pins. Because there is a shorter delay on the data path, the data will show up at the synchronous element before the clock.

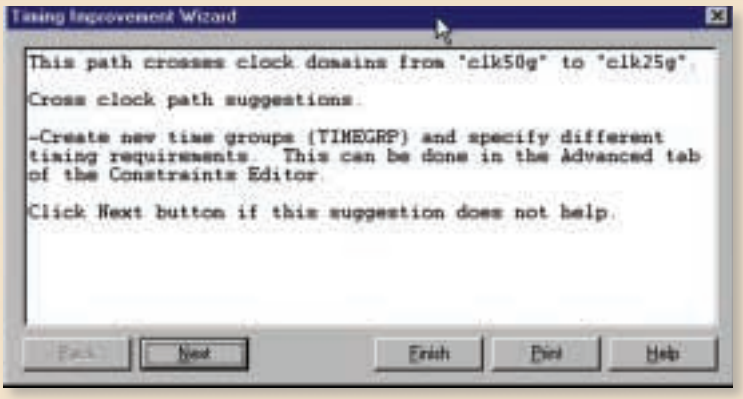

Figure 3 - Timing Improvement Wizard

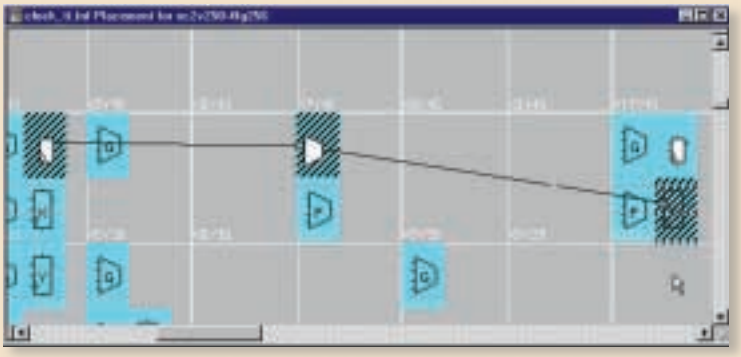

Figure 4 - Cross probing to the Floorplanner

The next section of the datasheet is the Clock to Pad . This shows the Clock to Out of all the output signals organized by clocks.

The last section of the datasheet is the Clock to Setup for each clock. All the source clocks that go to the destination clock appear in the first column. The setup time when both clocks are using the rising edge is shown in the second column. The next column shows when the source clock is falling, but the destination clock is rising and so on. This is a great place to see when data is crossing clock domains and/or half-cycle paths.

Paths are only reported in the datasheet if they are covered by constraints, or if unconstrained paths are included, or if a default or advanced analysis is run.

Conclusion

Whether you're a novice or an expert, the ISE 4.1i Timing Analyzer is very useful. Although several new functions, such as the Timing Improvement Wizard, are geared towards the new Xilinx user, we've kept all of the advanced features expert users like so much.

For more product training, please visit *www.xilinx.com/support/education-home.htm.* For more product information, please see the Help Topics in the Timing Analyzer tool.

Speed Up Your Design Verification

You can develop complete, timing-constrained HDL and Verilog test benches in minutes with the HDL Bencher component of Xilinx ISE 4.1i.

by Ricky Escoto
Manager, Product Line Marketing
ricky.escoto@xilinx.com

The HDL Bencher™ testbench generator delivers shorter time to market and increased engineering productivity. Using this automated testbench development tool, you get fast verification of HDL-based FPGA and CPLD designs. The new HDL Bencher generator is fully integrated with the Xilinx Integrated Software Environment (ISE) 4.1i family of design tools. HDL Bencher software is so easy to use that you don't have to know anything about hardware description languages to develop testbenches.

HDL Bencher Overview

With the HDL Bencher software, you can generate complete, self-checking VHDL or Verilog™ testbenches in just minutes. All you have to do is input a VHDL or Verilog design, along with the timing parameters that must be met by synthesis. You then describe the stimulus and expected behavior using the built-in pattern generator, the waveform editor, and the spreadsheet-based graphical user interface (GUI). The output is a simulation-ready testbench with library declarations, stimulus and check statements, and error reporting routines.

What's New in This Version

The HDL Bencher testbench generator has been enhanced to improve the overall testing process. This version features an improved user interface, faster testbench specification, and compressed testbench output.

Interactive Debug and Simulation

The HDL Bencher tool now includes automatic simulation through the Xilinx ISE framework. Simply draw the stimulus, then generate the expected results, and export a self-checking testbench automatically, as shown in Figure 1. Use the domain knowledge inherent in the HDL source, and in your head, to quickly produce regression level test cases. Mismatches between expected and actual values are highlighted automatically.

ISE Integration and Enhanced File Support

The HDL Bencher tool has been completely integrated with ISE 4.1i software. Integration includes automatic launching of the HDL Bencher software from ISE 4.1i, automatically adding waveform and testbench sources, and verification of large, multi-source file designs.

Faster and Better Performance

We've reduced File Open time by 20% and improved graphics display by 300%. The major effect of these speed enhancements is a faster, crisper application – ideal for large FPGA testbenches.

Faster Simulation

To improve regression run times, we've modified testbench output in a number of ways, resulting in testbenches typically 90% smaller than those generated by the previous version of HDL Bencher software:

- The clock process has been separated from assignments and assertions (clocked designs) to minimize code size.
- The delays between output assertions and input assignments have been packed to reduce code size.

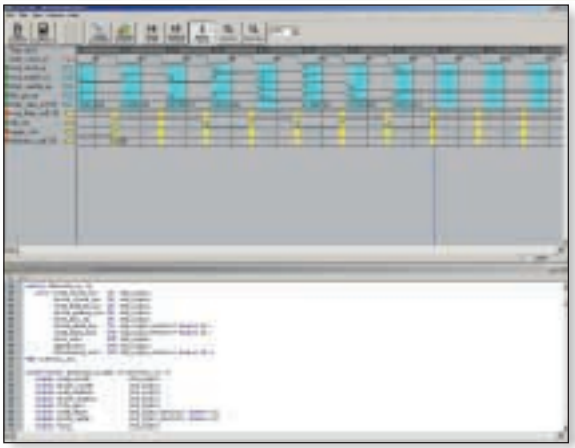

Figure 1 - The HDL Bencher 2.0 GUI

All-Level Verification

ISE 4.1i now features a waveform-based test environment. The same behavioral waveform file created in the HDL Bencher software is automatically updated to verify post-route and post-synthesis designs. Signal types are automatically remapped, and port definitions are resolved automatically.

Conclusion

The HDL Bencher testbench generator for ISE 4.1i supports integrated output simulation through ModelSim™ software, thereby closing the loop on its mission to provide seamless and fast testbench generation for Xilinx ISE applications

The HDL Bencher testbench generator is available at no charge to all Xilinx customers. It is fully integrated into Foundation™, BaseX, and WebPACK™ ISE 4.1i software. For more information on the ISE family of products, log on to *www.xilinx.com/ise/xcell/*.

ProActive Timing Closure Delivers up to 133% Better Device Performance

Take a look under the hood at one of the technology innovations embedded in ISE 4.1i and learn how it can meet your need for speed.

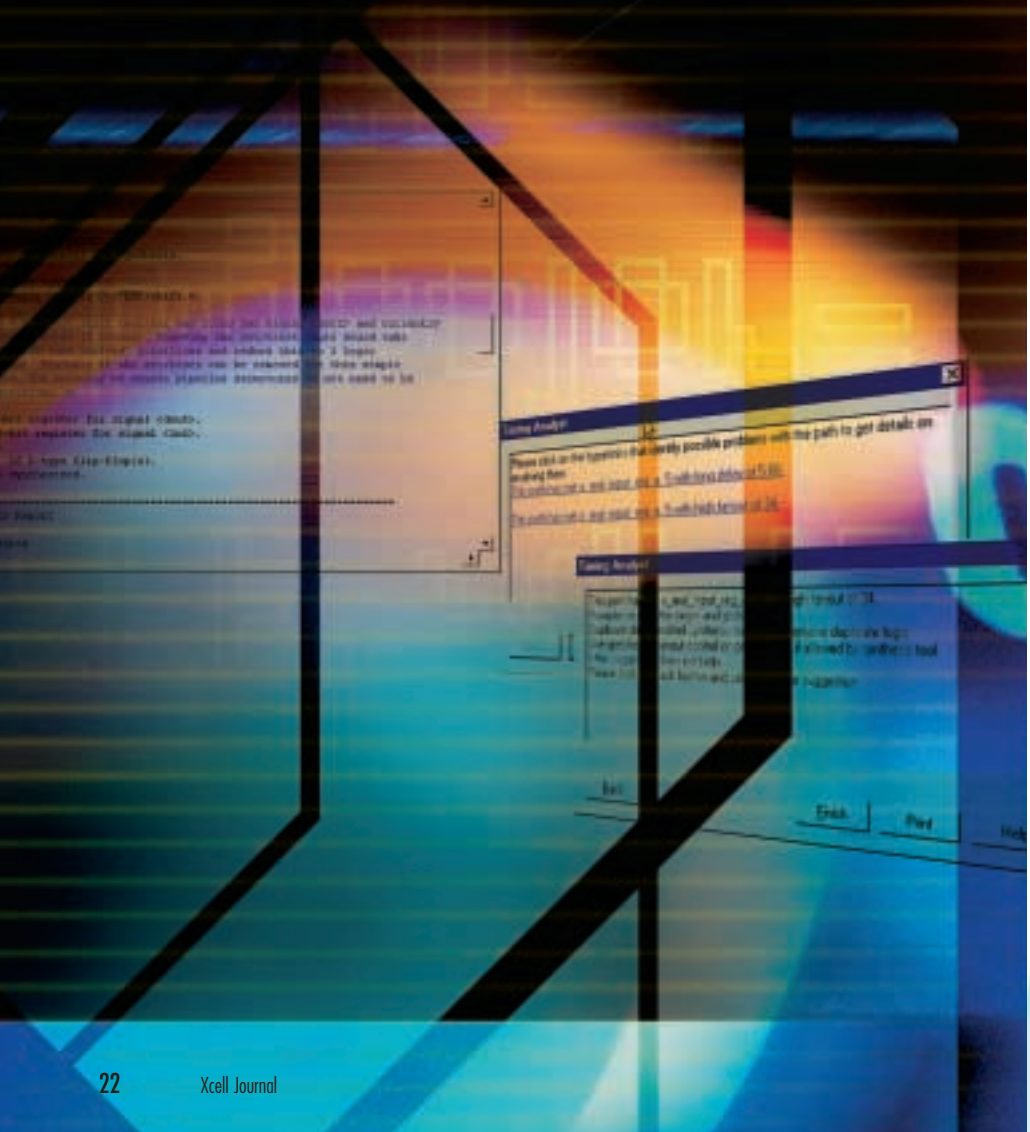

by Lee Hansen
Product Marketing Manager
lee.hansen@xilinx.com

As network and bandwidth capabilities continue to increase, the demands placed on design performance rise incrementally. These design pressures are leading to changes in programmable logic design tools – and Xilinx is leading the way with a new set of timing technologies, collectively known as ProActive Timing Closure.

Introduced in all configurations of the new release of Xilinx Integrated Software Environment (ISE) 4.1i, ProActive Timing Closure is setting new benchmark standards reaching design speeds up to 133% faster than ISE 3.1i. Furthermore, our internal benchmarks have shown that ProActive Timing Closure in ISE 4.1i can process on average as many as 100,000 gates per minute. That equates to millions of gates per hour for high-density designs.

Advanced Place-and-Route Algorithms

At the heart of ProActive Timing Closure are enhancements to the place-and-route algorithms. Placement now begins by scanning your design for critical data paths and attempting to place those data paths first – thus, improving timing for the critical portions of the design first, and improving overall timing.

Extra-Effort Routing Mode

The place-and-route tools now also include a new "extra-effort" mode. In this mode, the routing tool analyzes the design for nets that did and did not meet timing. Then the extra-effort mode automatically requests a new placement for areas that didn't meet timing. The router then attempts to route these newly re-placed areas – an operation that is transparent to the user.

The new place-and-route algorithm can result in fewer iterations, and more successful place-and-route passes earlier in the design cycle. Some customers are reporting they can use this advanced timing technology to generate several successful layouts of a particular design in the time it used to take to do one pass.

Directed Routing for IP

As IP has grown in speed and complexity, the need for more visibility into place-and-route directions for IP has grown accordingly. Today's POS-PHY and optical interface designs contain many critical data paths. And marrying your favorite synthesis tool to the increasingly complex and capable Xilinx device fabrics only increases the challenge.

The Directed Routing component of PTC is designed to help solve that challenge. Directed Routing brings a new level of visibility into the configurable logic blocks (CLBs) in the Xilinx device. Directed Routing describes in detail the placement and routing that should be used to achieve successful and repeatable IP implementation.

Physical Synthesis for FPGAs

Last year, Xilinx announced the first physical synthesis integration for FPGA design. ISE 4.1i continues the expansion of FPGA physical synthesis. With PTC, both Synplicity®'s Amplify™ and Exemplar Logic®'s TimeCloser™ physical synthesis tools now have more knowledge of the floorplan and early placement information. Thus, these tools make better place-and-route decisions that lead to better performance.

Also enhanced in ISE 4.1i, Synplicity customers can use the new Total Optimization Physical Synthesis (TOPS™) technology as part of the Amplify Physical Optimizer™ software to perfect their FPGA designs. Amplify physical optimizer software and the new incremental benefits from the TOPS technology are leading customers to better quality of results (QoR). [See "Understanding Physical Synthesis and Timing Closure" in this issue of *Xcell* for more information on ISE physical synthesis.]

HDL and Timing Interaction

Using traditional design techniques, much of a designer's time is spent trying to reach timing closure. It, therefore, stands to reason that improving timing-driven compilation – the timing results with more information and more interaction – will go a long way to help logic designers be more productive.

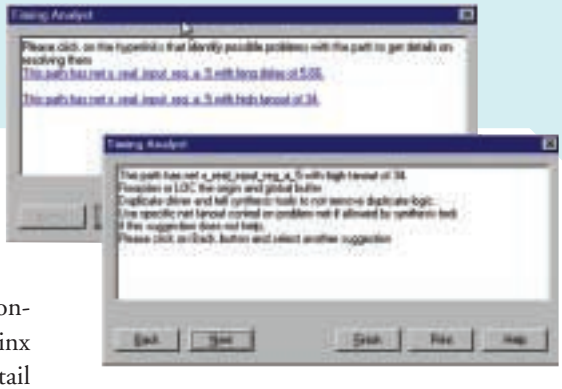

Figure 1 - ISE 4.1i Timing Wizard

Interactive Timing Wizard

ProActive Timing Closure includes a new interactive timing wizard, as shown in Figure 1. By selecting highlighted problem and warning areas in the timing report, the wizard suggests design changes and improvements that can help solve the timing problem and speed up the debug process.

HDL Analysis

Timing can be seriously affected by the "synthesizability" of the design code. In December last year, Xilinx announced the 1.0 coding style guide for Synopsys® Inc.'s LEDA™ HDL language checkers. In ISE 4.1i, customers can now use Synopsys' LEDA family of tools on their FPGA designs. The LEDA set of tools can verify your module against standard, good coding practices. This reduces the chance for problems cropping up during implementation due to bad coding (like introducing unnecessary latches into the finished module that cause timing analysis mistakes). The programmable LEDA tools are flexible to support corporate design techniques, to assure your source code meets your own specific standards.

As shown in Figure 2, Xilinx ProActive Timing Closure technology includes HDL analysis messages as part of the Xilinx Synthesis Technology (XST) report file. This new report feature suggests changes to the HDL code that would reduce design size and improve timing. This powerful capability is designed especially to help new HDL engineers and ASIC designers to write code that synthesizes well into an FPGA.

Timing Cross-Probing

Cross-probing for timing reports is yet another new enhancement delivered in the ProActive Timing Closure tools suite. You can select highlighted problem areas in the timing report and then see the problem net or area in either the Xilinx Floorplanner or the synthesis "technology viewer." Being able to find problem areas quickly and easily further reduces the time spent having to search through differing tools and GUIs for the same problem net, significantly speeding up design debug times.

Figure 2 - XST HDL analysis message

Conclusion

ProActive Timing Closure sets a new standard in the way programmable logic design tools can help you reach design closure faster and with better performance. Whether you're compiling six million gates or just a few thousand, you'll see the benefits in improved performance and reduced overall design times. ProActive Timing Closure will help you reach the fastest clock speeds available, and help you squeeze the most performance out of your devices.

Understanding Physical Synthesis and Timing Closure

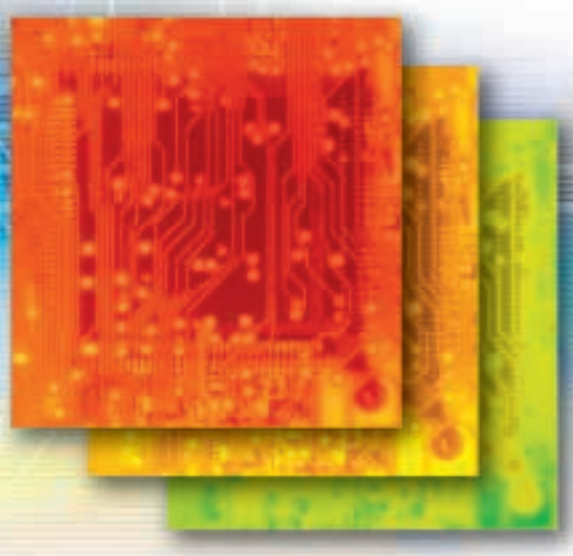

Using Xilinx Active Interconnect technology, you can achieve timing closure faster and with fewer iterations.

by Hamid Agah
Technical Marketing Manager
hamid.agah@xilinx.com

In the domain of deep submicron (DSM) and nanometer ASIC technologies (180 nm and below), the predictable timing relationship between logical (synthesis) and physical (place-and-route) design often breaks down. Designs cannot meet their realistic timing objectives, creating the well-known "timing closure problem."

Timing closure is currently the biggest area of difficulty for ASIC performance-oriented designs. The underlying reason for this problem is that circuit delays in the DSM realm are dominated by net delays, which are influenced by the placement of the cells. The traditional fan-out-based wireload models for estimating interconnect delay during synthesis become inaccurate at DSM levels, thereby causing the lack of timing predictability between post synthesis and post layout results. Clearly, logical synthesis and physical placement technologies must merge to create properly placed and routed designs that meet realistic performance goals.

Physical Synthesis Defined

Physical synthesis refers to the ability to create a properly placed and routable circuit from the register transfer level (RTL) code – that is, to create a circuit that meets the realistic performance goals of the design in one pass. In cases where there are very aggressive performance goals or tight physical constraints, a second pass may be needed to achieve the desired performance goals. A properly placed design meets the design rules of the target silicon technology and is routable by a detailed router.

Physical Synthesis in an ASIC Environment

Physical synthesis tools, such as Physical Compiler™ from Synopsys® Inc. have replaced the typical synthesis tools (such as the Design Compiler from Synopsys) that many ASIC designers are using today for performance-critical designs.

Designers typically start with a design-planning tool such as the Synopsys Chip Architect™ design planner or other design planning tools such as LDP from Cadence and Planet from Avant! to decide on:

- Physical area allocated to each synthesizable module in the design
- Physical location of each synthesizable module
- Physical locations of RAM, ROM, hard IP, and other non-synthesizable blocks in the design
- Pad (I/O) locations.

Once the design is planned, the next step is to perform a top-level routing and timing analysis based on the chip-level timing constraints. Using the analysis results, designers adjust the physical port location on synthesizable modules and modify the location and orientation of non-synthesizable blocks to derive a realistic timing budget for each synthesizable module in the design.

At this point, all the necessary information is available for every synthesizable module in the design, so the designer can proceed with the physical synthesis step. For each synthesizable module, the Synopsys Physical Compiler takes in:

- RTL code
- Timing constraints derived from the design planning step
- Physical constraints (area and port locations, for instance) derived from the design planning step
- Synthesis and physical libraries.

The compiler produces a netlist and physical information (such as the proper placement of all the cells) that meet the timing goal of the particular module.

The next step is to perform the detailed routing on the circuit, based on the timing analysis, to ensure that the fully placed and routed module meets its performance goal. Cadence and Avant! Corp. are the only electronic design automation (EDA) vendors that offer proven detailed routers trusted by major ASIC vendors (such as NEC, LSI, and Texas Instruments).

Physical synthesis requires thousands of strategies to be evaluated by the software while the circuit is being properly timed, and placed and routed. Detailed routing takes a very long time, and it is not suitable during the synthesis process. Therefore, following placement, the Physical Compiler performs an estimated routing and RC parasitic extraction to assess the net delay and the impact on the module's timing objectives. If the timing goal is not met, the compiler must then decide on the next synthesis strategy.

The single most important factor is the correlation between routability analysis (obtained from the congestion map) of the physical synthesis tool and the results of the detailed router. Without correlation, the placement created by the Physical Synthesis tool may be unusable by the detailed router, thus nullifying the result of the Physical Synthesis tool. The next important factor is to properly calibrate the Physical Compiler's R&C extraction with the final parasitic extraction tool, for example, StarRC from Avant!®. Without this calibration, Physical Compiler's assessment of net delays may be inaccurate, hence compromising the synthesis results.

Merging Synthesis and Placement

Obtaining a timing correlation between logical synthesis and physical placement is not cheap – or easy.

Cost

The total cost of setting up a physical synthesis environment is about $350K. First, you must purchase a design planner or floorplanner (such as the Synopsys Chip Architect) for about $150,000. Then you must invest another $200,000 for a single license to use a Synopsys Physical Compiler.

Interoperability

The success of a physical synthesis tool is highly dependent on the routability of the placed circuits that it produces and the proper calibration of its RC extraction. If the physical synthesis vendor also provides a proven detailed router, then there is high degree of certainty that a placed circuit can be routed by the vendor's detailed router.

The Synopsys Physical Compiler is the most successful physical compiler so far and has a proven track record to work with Cadence® and Avant! detailed routers as well as Synopsys' own detailed router, the Route Compiler, to complete the physical implementation of the design.

Layout Expertise

Synthesis designers require extensive training to become comfortable with physical design concepts and components. This involves either training current staff or hiring an expert.

Solving the Timing Closure Problem with Xilinx FPGAs

Xilinx addresses the timing closure issue for Virtex™-II Platform FPGAs and Spartan™-II devices by using a three-step process of "Predict, Control, and Improve." This process allows you to implement designs that can meet their realistic timing objectives with a minimum number of iterations.

Predict

In FPGA architectures, Xilinx Active Interconnect technology predicts routing delays. This characteristic makes it possible to create interconnect models that are not based on fan-out. These models can be used during the synthesis process to estimate the interconnect timing with a high degree of predictability with respect to the placed and routed design. Xilinx has partnered with leading FPGA EDA vendors to offer synthesis tools that are aware of the Xilinx FPGA architectural details. For example, Synplicity® uses our Active Interconnect technology to produce netlists with timing within 20% of placed and routed designs. The next two steps close the remaining performance gap.

Control

Guiding the Xilinx timing-driven implementation tools with realistic timing constraints, high quality netlists, and accurate physical constraints are the keys to closing the performance gap that may exist between the synthesized netlist and the placed and routed design. Xilinx worked closely with Synplicity to develop the Amplify™ design planner and physical optimization tool for FPGAs.

Amplify software can improve the netlist quality through physical optimization techniques, such as moving registers across physical boundaries to increase performance. The Amplify program can also provide accurate area constraints and physical grouping of critical paths to the Xilinx implementation tools. Nonetheless, it is still possible to have failing paths remaining after place and route. The next step addresses the remaining failing paths.

Improve

Xilinx, in partnership with leading FPGA EDA vendors, has developed a tight interface between the vendors' synthesis tools and the Xilinx implementation tools. This interface allows re-optimization of failing paths and creates engineering change orders for new circuits into the Xilinx implementation tools. In a majority of cases, this capability can enable timing closure in no more than two passes.

Conclusion

Unpredictable interconnect timing during logical synthesis is the main reason for inaccurate timing estimation for DSM (0.18µ and below) ASICs. Xilinx has successfully applied a three-step process of "Predict, Control, and Improve" to close the timing gap for Virtex and Spartan-II architectures (0.22µ to 0.18µ). The key success factors are:

- Accurate timing estimation during synthesis
- True timing-driven place and route
- Re-optimization of failing paths.

As process technology continues to shrink, with ever more complex and higher performance designs, Xilinx Platform FPGA architectures will continue to enable designers to close the timing gap between logical synthesis and physical implementation.

Xilinx Takes Power Analysis to New Levels with XPower

Analyze your power consumption earlier in the design cycle to better make trade-offs that will allow you to meet specifications and get your product to market faster.

by Steve Wenande
Senior Technical Marketing Engineer
steve.wenande@xilinx.com

Richard Chidester
Senior Technical Marketing Engineer
richard.chidester@xilinx.com

In recent years, programmable design and device sizes have grown to astonishing levels of complexity. Average design sizes were approximately twelve thousand gates a few years ago. Now, they have bloomed into hundreds of thousands and even multimillion gate arrays. And, as design sizes increase, so does power consumption.

Meanwhile, the demand has risen for battery-powered handheld devices that are increasingly smaller and sensitive to power usage. It's clear that power consumption can no longer be ignored in programmable logic design.

Anticipating designers' needs, Xilinx has added the XPower analysis tool to its suite of ISE 4.1i design tools. XPower graphical power-analysis software is the first of its kind for programmable logic design. Now, earlier than ever in the design flow, you can analyze total device power, power per net, and routed or partially routed designs. The XPower tool can be controlled from the graphical interface or through a command line-driven batch mode. XPower delivers data in either an easy-to-use graphical interface (Figure 1) or in ASCII reports. XPower offers unmatched device support and accuracy for achieving the lowest possible power consumption.

Inside the XPower Analysis Tool

XPower works on the principle of "activity rates." Activity rates are defined as the rate at which a net or logic element capacitance switches. For dynamic power calculation and display, activity rates are expressed as a frequency. An activity rate may be relative to a clock, in that the net or logic element switches at some percentage of the clock frequency. This is often referred to as a toggle rate. Activity rates are very useful because they enable you to recalculate the power by merely changing the system clock frequency. This allows you to use your original simulation data and save time. XPower supports any number of input clocks as well as DLL/DCM derived clocks. Expressed as a percentage, an activity rate of 100% means that a signal state change happens on average once every clock cycle with the resultant frequency being half the associated clock. For nets and logic that are not synchronized with a clock, the activity rate is just the switching rate.

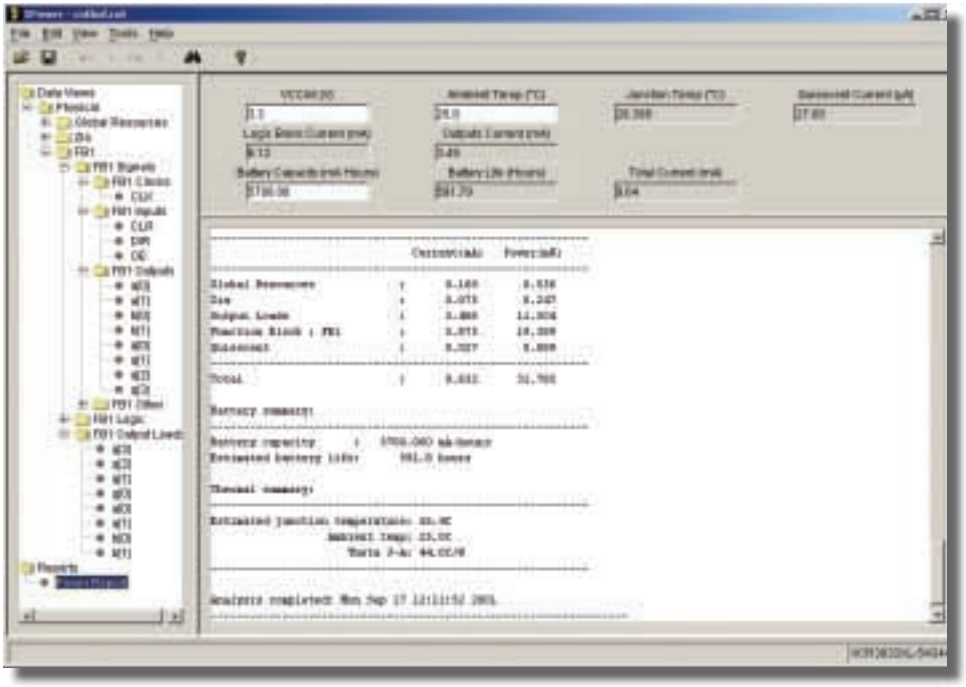

Figure 1 - XPower graphical user interface

XPower Analysis in FPGA Design

Power consumption in high-density design is becoming a serious issue as device sizes shrink and design sizes grow. For FPGA design, the XPower tool gathers design information from the following list of files:

- NCD – design topology and physical resource usage
- PCF – clock frequency and voltage
- VCD – detailed design activity rates for all nets
- XML – user setting file.

XPower accepts your post-route netlist (NCD) and physical constraints file (PCF) to determine the interconnect, I/O, and clock information of your design. With this information, XPower builds a hierarchical representation of your design – broken out by signals, clocks, logic, and outputs. In addition, XPower accepts value change dump (VCD) information from all ModelSim™ simulators. VCD information provides the XPower program with detailed design activity information that gives you a very accurate power estimate. The extensible markup language (XML) file saves the user settings data for subsequent tool usage.

The XPower program will estimate your power consumption for Virtex™, Virtex-E, Virtex-II, Spartan™, and the new Spartan-IIE FPGAs.

XPower Analysis in CPLD Design

Low-power design requires careful analysis of an application's power usage. To meet customer demands, battery-powered applications must have power budgets that can't be exceeded – deficits are unacceptable.

One of the pitfalls of power estimation lies in the maximum and minimum power corner cases. For instance, going with a maximum power estimation may require an oversized battery, creating a cost and packaging dilemma. Whereas, if you go with a minimum power estimation, an undersized battery could lead to short battery life and potential product return due to customer perceptions of poor quality. Therefore, using typical power estimation, combined with battery life tests, provides a good conservative approach.

XPower supports CoolRunner™ Fast Zero Power™ technology. Using FZP technology gives you an accurate, easy-to-use power estimation of your design. To enhance the ease of use, the application of activity rates to internal nodes is automated to minimize the data entry task. Refer to *xapp360.pdf* on the Xilinx website for additional guidance.

For CPLD design, the XPower program gathers design information from the following list of files:

- CXT – CPLD XML file that contains design topology and physical resource usage
- PCF – clock frequency and voltage
- VCD – detailed design activity rates for all nets
- XML – user setting file.

After completing and fitting your design in a CoolRunner CPLD, simulate your design using the ModelSim simulator in ISE 4.1i. To obtain the most accurate power estimation, your simulation should be as close as possible to the actual in-the-field usage of your product. The XPower analysis tool uses the CXT and VCD files to estimate the power used by your application. XPower generates a tabular report of the power used, and it lists the resources and signal names of your design. Performing end-of-life battery tests will ensure your success.

Conclusion

Increasing design sizes and the emerging handheld market are a reality. The combination of the two makes lowest possible power consumption a critical factor in programmable design today. The introduction of XPower into the ISE suite of tools gives you the new capability to address these issues and decrease your time to market.

The XPower component is now being delivered with the Alliance and Foundation Series ISE 4.1i design tool suites. It is also included in the free WebPACK™ configuration. For more information on ISE 4.1i with XPower, visit the Xilinx Design Tools Center at *www.xilinx.com/ise/xcell/*.

LeonardoSpectrum Now Supported in the Xilinx ISE Software

To take full advantage of the latest FPGAs and CPLDs, you need advanced, high performance development tools, and that's exactly what you get with the new Xilinx ISE software.

by Steve Wenande
Senior Technical Marketing Engineer
steve.wenande@xilinx.com

Xilinx ISE (Integrated Synthesis Environment) provides all the tools you need in a single, tightly integrated package for FPGA and CPLD design. With the release of our 4.1i software, ISE takes the next step in providing seamless integration by adding support for LeonardoSpectrum™ from Exemplar Logic Inc. The addition of LeonardoSpectrum to ISE gives you the ability to target all of the major synthesis technology providers, including:

- XST (Xilinx Synthesis Technology) from Xilinx
- FPGA Express™ from Synopsys
- Synplify™ and Synplify Pro™ from Synplicity

LeonardoSpectrum

LeonardoSpectrum is a market-leading field programmable gate array (FPGA) synthesis tool from Exemplar Logic Inc. It allows you to target Xilinx FPGAs and CPLDs using standard VHDL and Verilog flows. It includes schematic and critical path viewing, and support for CORE Generator™ files in the file list.

ISE integrates the LeonardoSpectrum Quick Setup (Level 1 and Level 2) flow when LeonardoSpectrum is properly installed on the same machine with ISE. LeonardoSpectrum combines push-button ease of use with the powerful control and optimization features associated with workstation-based ASIC tools. The advanced features for block-based and incremental design ensures the best results for your largest designs. The addition of this design flow gives you one more alternative when designing with ISE and brings the power of LeonardoSpectrum to your fingertips.

XST

For HDL optimization, you can use Xilinx XST (Xilinx Synthesis Technology) as an alternative to FPGA Express™, Synplify™ or Synplify Pro™, or LeonardoSpectrum. Xilinx is using XST as a proving ground for many of the innovative optimization ideas that Xilinx engineers are developing for improving HDL design flows. These improvements are then shared with Xilinx third-party synthesis partners to ensure that anyone targeting Xilinx FPGAs as their solution can benefit from the best optimization the industry has to offer.

FPGA Express

FPGA Express is a complete FPGA logic-synthesis and optimization tool. With it you can create optimized FPGA netlists from VHDL code and Verilog HDL code. FPGA Express core technology was developed specifically for FPGA and programmable logic device (PLD) architectures with the following features:

- Architecture-specific mapping and optimization for Xilinx programmable logic devices
- Easy-to-use design flows
- Integrated static timing analysis with TimeTracker
- Spreadsheet-style constraint entry for timing, placement, and synthesis constraints.

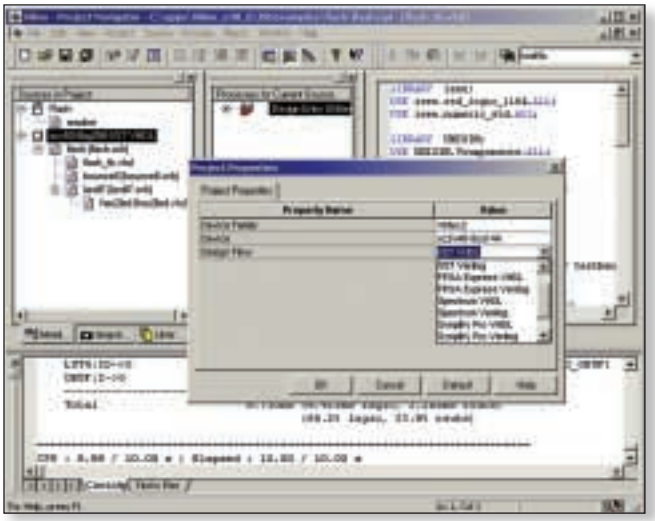

Figure 1 - Integrated synthesis options available in ISE

In addition, FPGA Express features BLIS (Block Level Incremental Synthesis), a first for FPGA synthesis, delivering the capability to:

- Modify and re-optimize individual blocks in a previously routed design
- Export these re-optimized blocks as distinct and separate netlists
- Execute a guided place-and-route only on the modified sections of the previously routed design.

This breakthrough technology automatically recalculates timing across the entire design, including taking into account the unmodified portions of the design. Because of this automatic, full-context timing, you are no longer required to re-synthesize the entire chip, repeat a full place and route, and conduct complex manual re-optimization of your design.

Synplify and Synplify Pro

The Synplify solution is a high-performance, sophisticated logic synthesis engine that uses proprietary Behavior Extracting Synthesis Technology (B.E.S.T.)™ to deliver fast, highly efficient FPGA and CPLD designs. The Synplify product takes Verilog and VHDL hardware description languages as input and outputs an optimized netlist specifically targeting Xilinx devices. The Synplify Pro™ software extends the capability of the Synplify solution to meet the needs of today's complex, high-density designs.

Synplicity and Xilinx have once again teamed together to offer an enhanced synthesis flow with ISE release 4.1i. ISE now runs with the latest offering from Synplicity, version 7.0. The combination of these two products offer advanced capability to users targeting Xilinx devices.

Availability

All synthesis integration will be delivered with Foundation ISE 4.1i products. In addition, XST, FPGA Express, and LeonardoSpectrum integration will be included in WebPACK, available online from the *xilinx.com* website.

XST and FPGA Express™ are delivered with the ISE 4.1i installation. To obtain the third-party synthesis products, you can refer to their websites: *www.exemplar.com*, and *www.synplicity.com*

Conclusion

Xilinx is committed to providing you with the most advanced design solutions available today. The addition of LeonardoSpectrum integration to the existing synthesis support already in ISE adds even more flexibility for targeting Xilinx devices. By combining the latest software technologies from Xilinx and our partners, you get the fastest and most productive development platform ever, and it just keeps getting better.

For more information on Xilinx ISE software go to:

www.xilinx.com/ise/xcell/.

ISE 4.1i

Synplicity Announces TOPS

A Second-Generation Physical Synthesis Technology for Xilinx FPGAs

Routing interconnect delays significantly affect your overall circuit performance, and therefore, your synthesis tools must account for these delays. TOPS (Total Optimization Physical Synthesis) makes physical synthesis even more efficient.

by Jeff Garrison
Director of Marketing, FPGA Products
Synplicity, Inc.
jeff@synplicity.com

As FPGA process technology gets faster and denser it poses new design challenges for your development tools – challenges that must be addressed to maintain the productivity you expect from programmable logic. Your synthesis tools, which were originally designed to optimize logic, must now also consider the effects of your interconnect delays, during the synthesis process.

Synplicity was the first company to address this problem with its Amplify™ Physical Optimizer™ software, which allows you to improve performance by using physical guidance in addition to normal timing constraints during the synthesis process. By assigning critical logic to physical regions of the device, you provide the Amplify tool with important information that it uses to not only create better logic placement, but also to aggressively optimize (change) the resulting netlist for better performance.

While this interactive method continues to be preferred for getting the very best results, Synplicity's new second-generation physical synthesis technology has demonstrated up to 15% performance improvement for Virtex™ devices in a totally automated flow. This new physical synthesis technology, TOPS, supports both a fully automated and an interactive physical synthesis methodology.

Total Optimization Physical Synthesis (TOPS)

After normal synthesis, place-and-route, and timing analysis, critical path information is used to interactively create physical regions on the device. The TOPS technology uses these physical regions to perform physical synthesis on each region. The improved netlist and detailed regional placement information is passed to the final place-and-route process.

Like Synplicity's first-generation physical synthesis technology, the TOPS technology performs simultaneous placement and optimization. It includes two methods, designed to meet your needs – Automated TOPS and Interactive TOPS.

- The Automated TOPS technology fully automates physical synthesis to help you achieve higher device performance with the push of a button. Using this methodology, you need not be familiar with the physical architecture of a specific device to improve performance.

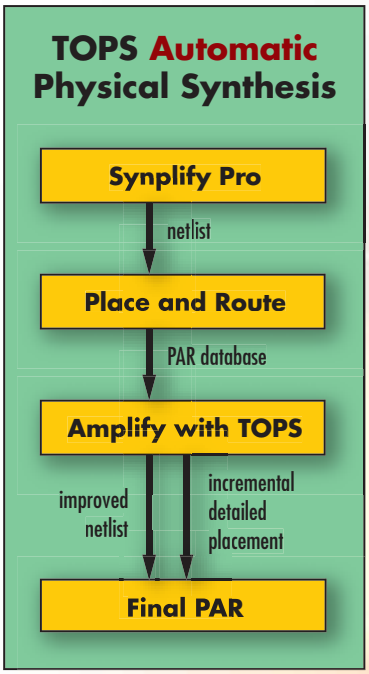

Figure 1 - Automatic physical synthesis using TOPS

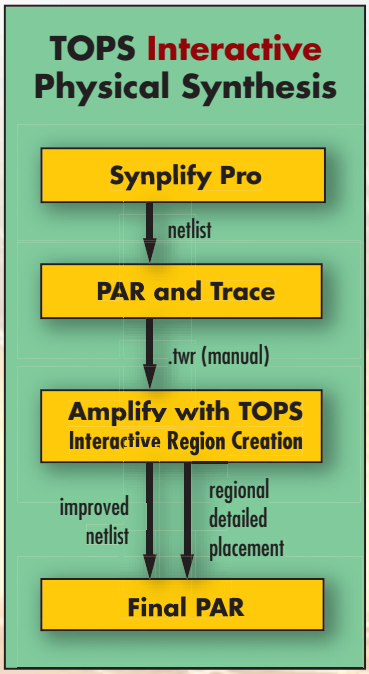

Figure 2 - Interactive physical synthesis methodology with TOPS

- The Interactive TOPS technology takes Synplicity's interactive physical synthesis approach to the next level by operating with exact placement information while re-optimizing and incrementally placing the design thereby significantly improving performance.

The automated and interactive TOPS phases may be used together to achieve maximum performance.

Automatic TOPS

After normal synthesis and place-and-route, detailed physical information from the Xilinx design database is used as input for the automated pass through the TOPS software. Amplify along with TOPS technology improves the netlist itself for performance and also generates incremental detailed placement which are passed on to the Xilinx place-and-rout software for final routing, as shown in Figure 1.

Interactive TOPS

The Interactive TOPS technology is designed to improve performance by focusing its physical synthesis algorithms on critical portions of the design as opposed to the entire chip – you create regions only for those timing-critical portions of your design as opposed to constraining the entire chip. If you are already familiar with the Amplify Physical Optimizer product's region-based interactive methodology, you know its impressive performance improvements. The new Interactive TOPS technology boosts those improvements by performing detailed incremental placement for critical regions. The interactive TOPS software then performs physical synthesis (including detailed placement) on those regions. The improved netlist and placement information for each region is then passed on to the Xilinx place-and-route software for the remainder of the placement and routing, as shown in Figure 2.

If required to meet timing performance, the Amplify software will automatically move (place) logic that you did not directly place (and report in a log file). All logic that you place in a region will remain in that region as long as it helps improve timing performance.

Conclusion

Physical synthesis is the answer to getting aggressive performance from your FPGAs – in recent designs, as much as 80% of the total timing delay is the result of routing, not logic. And, as devices get larger, interconnect delays have an increasing role in determining your chip's overall performance.

With Synplicity's second-generation physical synthesis technology (TOPS), the requirement for interactively creating physical regions is eliminated, and up to 15% performance improvements are possible with the push of a button. For even higher performance, interactive TOPS methodology uses incremental detailed placement to deliver up to 50% more speed than normal synthesis alone. For additional information visit: *www.synplicity.com.*

ISE 4.1i

The Amplify Physical Optimizer

The Amplify Physical Optimizer was introduced in March of 2000 and has helped more than one hundred companies quickly achieve aggressive timing performance in their programmable devices. Xilinx designers have seen performance improvements of as much as 45%, as a result of using the Amplify software. The average performance improvement (compared to logic synthesis alone), across the many designs tracked, is over 20%.

Until now, designers using the Amplify product have used an interactive methodology, where physical regions are created on a device, then logic from critical circuit paths is placed into the regions and then finally synthesized using Amplify's unique physical synthesis technology. Even though the Amplify software is very easy to use, it does require knowledge of the device architecture – it is important to know where to create the regions and how to assign logic to them, to get the best results.

Synopsys and Xilinx Unveil Next Generation Flow for Platform FPGAs

For Virtex Platform FPGAs, with gate counts comparable to ASICs, you need a design flow with code checkers and static verification technology.

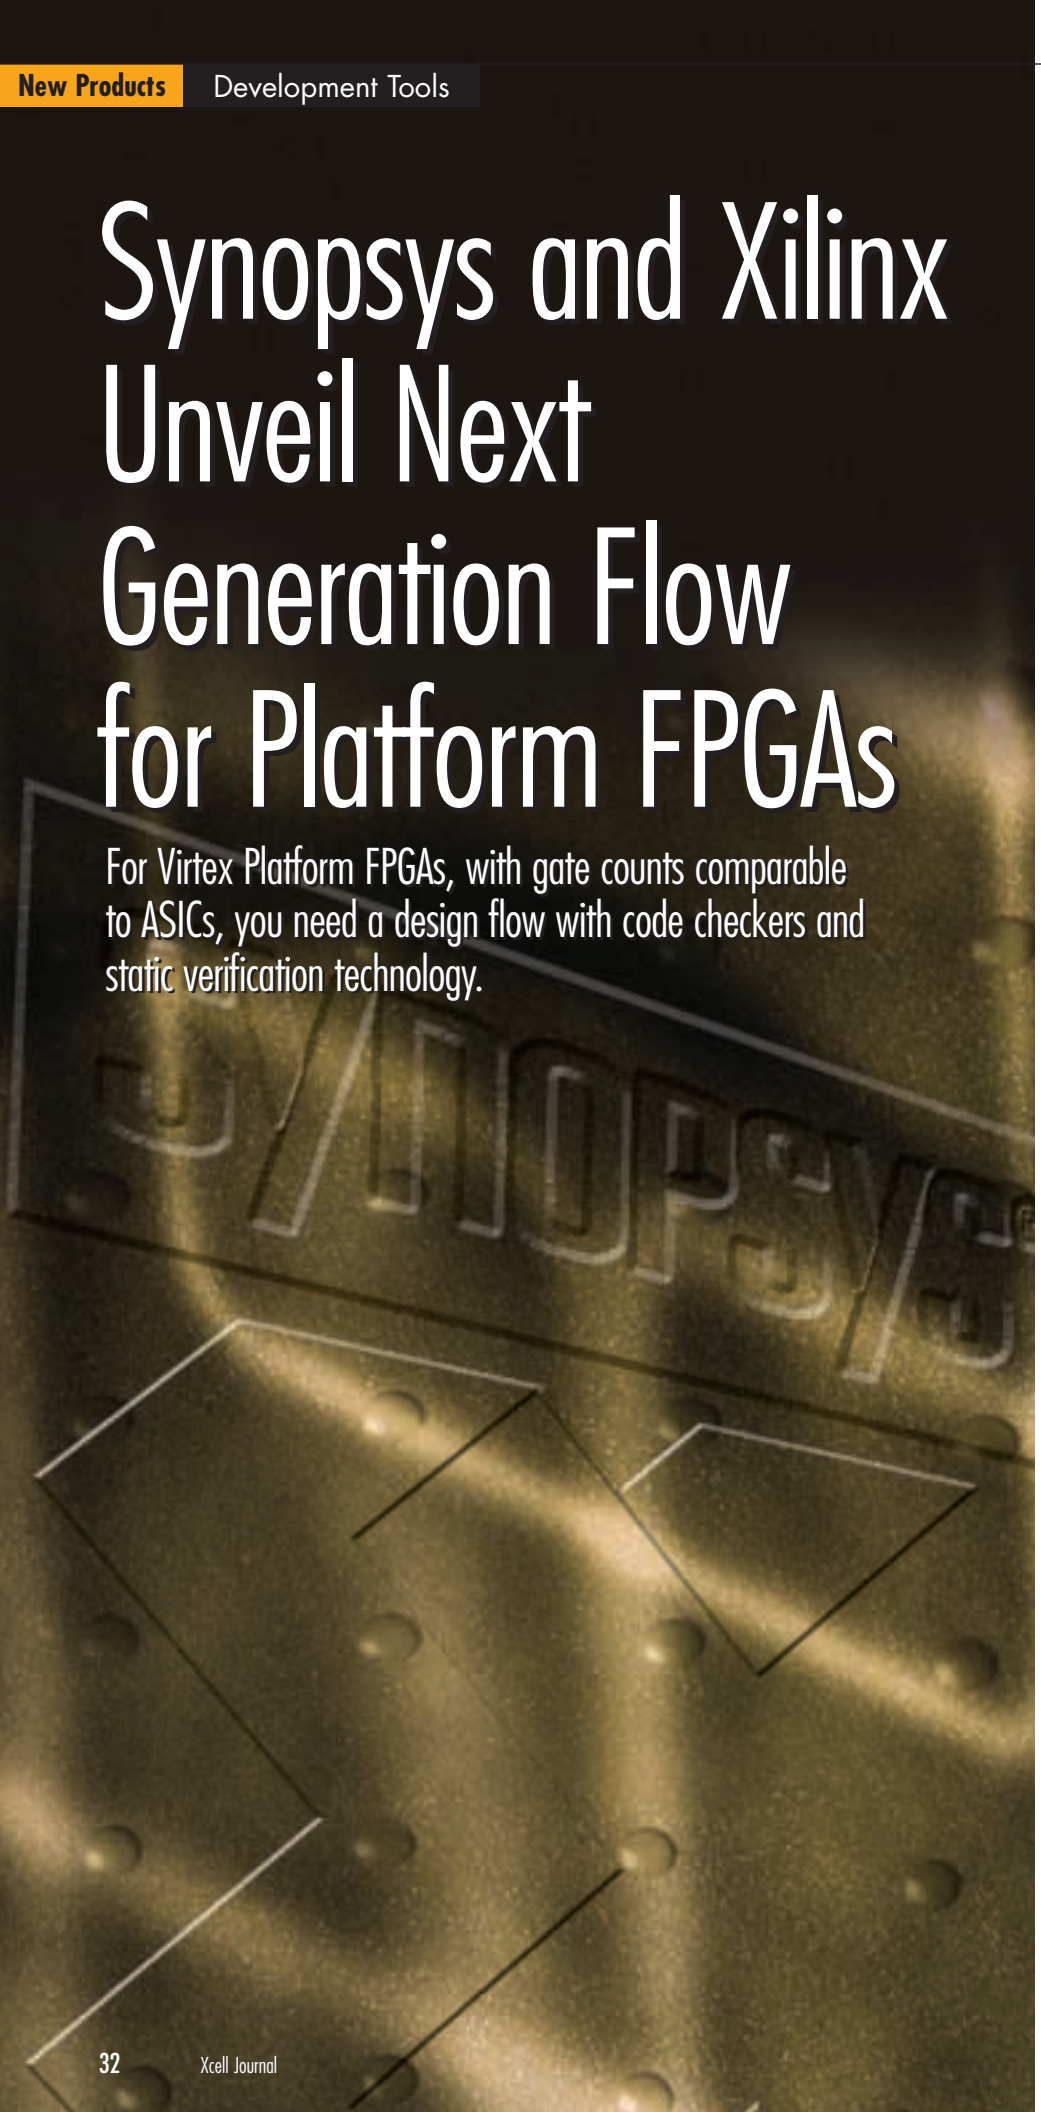

by Jackie Patterson
Director of Marketing Programs
Synopsys Inc.
jackiep@synopsys.com

Hamid R. Agah
Senior Technical Marketing Director
System Level and Alliance Marketing Group
hamid.agah@xilinx.com

As FPGAs grow in speed, size, and complexity, you need EDA tools and design flows that are similar to those used for ASICs. With 10-million system gates, 300+ MHz internal clock speeds, and gigabit serial I/O technology, the Virtex-II Platform FPGAs give you ASIC-like specifications and performance. And, because of far lower FPGA development costs, many designers are turning to complex FPGAs instead of ASICs for system-on-chip (SoC) designs. In fact, nearly one in two FPGA designs are SoCs.

While this FPGA complexity is good news for your SoC designs, you can face difficult design challenges unless you adopt some of the coding techniques and static verification methodologies used for developing ASICs. For high-density Xilinx® Virtex™-II Platform FPGAs, you can no longer rely on ad hoc methods and the simple design flows that are adequate for smaller FPGAs.

To meet this need, Synopsys and Xilinx recently announced an enhanced design flow that combines advanced Synopsys tools (LEDA, VCS, Scirocco, FPGA Compiler II, PrimeTime, and Formality) with the Xilinx ISE Alliance Series™ software, to reduce development time for Xilinx Virtex-II Platform FPGAs.

Enhanced Design Flow

The Synopsys tools give you an ASIC-like design flow that makes high density FPGA designs much easier and faster to develop. It begins with LEDA.

LEDA

The LEDA ProVHDL and ProVerilog programmable checkers enable you to create and automate the enforcement of design and coding-style guidelines, thereby improving performance and results. You

can create a specification in VHDL or Verilog, and LEDA checks your HDL code for correctness and performance, applying a Virtex-II Platform FPGA rule set that Synopsys and Xilinx jointly developed. This gives you the opportunity to fix problems up front where they are least expensive. Synopsys will also prepackage LEDA checkers with policies to improve performance and results of tools like VCS, Design Compiler, Formality, Scirocco, and so on.

VCS

The register transfer level (RTL) design is then simulated using a high-speed simulator such as VCS or Scirocco. VCS, the Synopsys Verilog simulator, uses special algorithms that boost performance to provide simulation run times that often outperform today's commercial cycle simulators. The simulator accurately analyzes negative timing constraints, which are critical in verifying skew between high-speed signals in deep submicron designs.

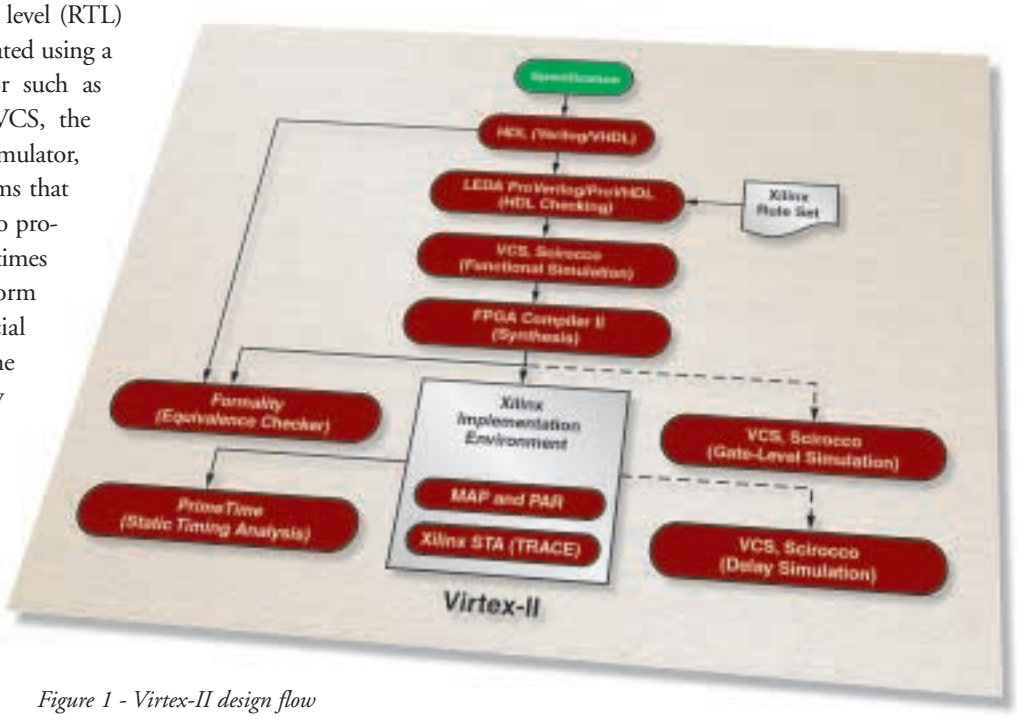

Figure 1 - Virtex-II design flow

Scirocco

Use the Synopsys Scirocco simulator for VHDL designs; it unites the performance potential of cycle-based optimization techniques with the flexibility of event-driven simulation in a single simulator. Its optimized VHDL language compiler generates memory-efficient simulation executables, enabling acceleration of complete system verification, providing capacity in excess of 10 million gates on a single workstation. This level of performance and capacity is essential if you are using the high density Xilinx FPGAs.

FPGA Compiler-II

Next, your design goes to synthesis with the Synopsys FPGA Compiler II, which optimizes the high-level logic description into the unique Xilinx Virtex-II Platform FPGA architecture. FPGA Compiler II delivers these architecture-specific synthesis capabilities with automatic register retiming, pipelining, and Block-Level Incremental Synthesis (BLIS). The fast Xilinx ISE Alliance software completes the implementation. The Xilinx ISE tools deliver the industry's fastest runtimes and highest clock rates for the Virtex-II Platform FPGAs.

PrimeTime and Formality

Once your design is implemented, you can save valuable development time by verifying and debugging your design before trying it in the lab. Xilinx supplies scripts, libraries, and application notes to enable the Synopsys PrimeTime and Formality verification tools. PrimeTime performs a comprehensive post-layout static timing analysis of the Virtex-II device, which enables a fast, accurate, and exhaustive timing verification for the design. You can also use PrimeTime's rich set of analysis features to easily analyze and debug any critical timing issues.

The Formality equivalence checker allows you to quickly confirm that the implemented design matches the RTL specification. Formality uniquely addresses the need for functional verification of large, SoC designs. Many design engineers report that verification tasks consume 60% to 80% of design resources. Formality offers a much faster (10-100X) and more thorough alternative. With this tool, you compare the functional equivalency of the RTL source to the post-synthesis (or the final post-place-and-route) gate-level netlist. You can employ Formality after each design step to maintain complete functional equivalence at the gate level during every stage of the flow. With Formality, the need to perform time-consuming gate-level simulations at intervals in the flow, or one massive simulation run at the end, is virtually eliminated. Thus, you save time and also verify your Virtex- II Platform FPGAs much more comprehensively than ever before.

Conclusion

With the addition of LEDA, PrimeTime, and Formality, Synopsys and Xilinx have created design flows that were previously only available to ASIC designers. As a result of these improvements, ASIC designers who want to use the Xilinx Platform FPGAs will have the tools they need and currently use. This means shorter learning curves and faster time to market. These improvements also mean that ASIC and SoC designers have an alternative to expensive silicon fabrication technology.

Speed up Verification of Long Transaction Sequences with MicroBlaze Soft Processors

Embedding soft processors and peripherals inside Xilinx Virtex-II Platform FPGAs has created a new set of challenges for designers. Here's one solution.

by Warren Miller
VP of Marketing, Avnet Design Services
warren.miller@avnet.com

Michael Scurry
Chief Engineer, Experience First
mscurry@expfirst.com

Ravi Pragasam
Marketing Manager, Processor Solutions
ravi.pragasam@xilinx.com

The Xilinx MicroBlaze™ soft processor core and the IBM CoreConnect™ bus architecture that the MicroBlaze processor uses can provide the hardware horsepower for you to implement an entire Internet appliance in a single Xilinx Platform FPGA.

The ability to run long sequences of processor code inside Virtex-II FPGAs, however, has made it more difficult to verify functionality and to test designs that connect to the processor within the FPGA.

Verification of a complex start-up initialization sequence or an Internet transaction – perhaps to download code to update the processor code or even to reconfigure the FPGA – would just take too long to verify in software. A software-only verification scheme could take hours or maybe even days to run, and it would involve hundreds of thousands of system cycles to accomplish verification.

In this article, we will show you an example of how you can use hardware techniques to augment software verification. This example reduces the time required to perform multiple test and modification cycles to verify a system reconfiguration from initialization and set-up to an actual transaction over the Internet.

MicroBlaze Applications

The ability to connect all sorts of electronics equipment to the Internet continues to create new applications and business models. With the availability of the MicroBlaze soft processor targeted to the Virtex-II family of Platform FPGAs, you can now include Xilinx Internet Reconfigurable Logic (IRL™) connectivity in equipment like vending machines, motors, heating and cooling equipment, industrial process controllers, and remote monitoring and communications appliances. This technology allows already deployed devices and appliances to be remotely reprogrammed, reconfigured, or upgraded without a visit from a technician.

One low-cost approach to implementing this type of equipment is to use a simple communications protocol to access the Internet. This complete device-networking solution requires only a processor with 4 Kbytes of memory and a UART to connect the embedded application to a gateway, and thus, to the Internet. In this application example, we will use a Virtex-II FPGA with an embedded MicroBlaze core, UART, and embedded memory (BRAM and distributed RAM) to provide all the resources required to connect to the Internet.

The rest of the Virtex-II FPGA, which has upwards of hundreds of thousands logic gates, is available for controlling the rest of the application. For instance, a remote control system for a modular factory floor could be upgraded and changed by reprogramming the FPGA over the Internet. An architectural level diagram of such an application is shown in Figure 1.

The MicroBlaze core can be configured in one of six combinations of busses, each with a configuration providing a different combination of performance and functionality. In particular, the amount of memory required by the application will determine which of the bus combinations is required. In our example application, we're assuming that all the instruction memory is located on-chip and that data memory may be required off-chip, perhaps as memory-mapped I/O.

The main on-chip peripherals needed for this application are a UART, a timer/counter, and an interrupt controller. The UART is a critical component in implementing the Internet protocol application. The UART is attached to the CoreConnect on-chip peripheral bus (OPB) and provides the serial communication link to the outside world. Once the UART is parameterized, you can select the base address for the internal registers, the number of data bits per character, and the type (if any) of parity supported. The registers accessible in the UART from the OPB are the read data register, write data register, read status register, and write status register. The UART must be initialized on startup to select the appropriate speed and data format. Once initial-

ized, it can be used to provide connectivity for the networking protocol.

MIcroBlaze Transaction Sequence

Now that we've set the stage with a description of the application and key system components, it's time to delve into an example problem that you might face when creating an embedded processor design. A good example would be to determine the length of time it would take for the system to reconfigure, initialize, and transact the first series of packets to establish communication over the Internet. If changes must be made in the code or in the Virtex-II FPGA design, it's necessary to establish timing for each of the main transaction sequences over the OPB.

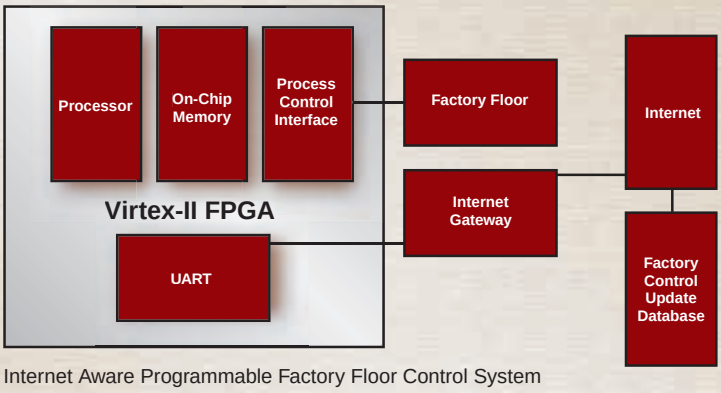

Figure 1 - Application block diagram

ChipScope Integrated Logic Analyzer

One method of timing the transactions would be to use the on-chip resources of the Xilinx ChipScope™ Integrated Logic Analyzer (ILA) to observe bus traffic in an FPGA on a prototype or development board. ChipScope ILA embeds logic analyzer cores into your design. These logic cores allow you to view all the internal signals and nodes within an FPGA. ChipScope ILA supports user selectable data channels from 1 to 256. The depth of the sample buffer ranges from 512 to 16,384 words in Virtex-II devices.

In our example, event triggers are changeable in real-time without affecting the logic or requiring recompilation of the design. ChipScope can capture OPB transactions and store them in on-chip memory. Bus transactions can be monitored, and with the trigger capability, specific transactions can be used to begin a capture event.

For instance, a specific address can be used to begin capture when a UART register is read or written. A status register read could indicate the beginning or end of a transaction sequence. You can use an ILA feature that counts the cycles between trigger events to determine the amount of time required for specific transactions – which is just what we need for our example application.

The ChipScope ILA provides the capability to observe signals on the OPB, but it does not have the ability to drive internal signals. Additionally, ChipScope ILA uses internal memory blocks to capture data on the OPB. If the application requires the use of internal memory, however, there may not be any memory left for use by ChipScope. In such a case, the ability to use off-chip memory would help solve this problem and provide an integrated software and firmware tool environment to accelerate the development, integration, and test of MicroBlaze-based Virtex-II FPGA designs.

The Raptor Solution

The Raptor Development Kit from Avnet automatically inserts logic around the MicroBlaze core to route inputs and outputs to real-time input signals, output signals, and buffer memory available on the development board. This allows the core to be exercised at "hardware speeds." Core output signals that are stored in the buffer memory can be read out over the USB port to a host computer and displayed on the waveform viewer, just as if the core output signals were software simulation results.

This "hardware speed" approach to verification reduces the design/debug cycle time dramatically. The Raptor Kit makes it easy to perform design or test set-up changes – and see the results immediately. Used in conjunction with a hardware-based input stimulus, verification of very complex and robust test suites are orders of magnitude faster than pure software simulation-based approaches.

The Raptor Kit provides the input signals from the UART to simulate the start-up transactions. The entire transaction sequence can be run at hardware speeds (25 MHz for our example design). The transactions on the OPB can be observed, and detailed measurements can be made, to pinpoint the key time delays for each portion of the start-up, initialization, and UART communications code phase. Once the code is optimized to remove unnecessary waits and loops, a total of only 236,000 clock cycles are required for initialization and first UART transaction with the gateway. The design now has enough time to reconfigure the FPGA without overflowing the FIFO at the application layer of the protocol.

Conclusion

Implementing embedded soft processors like the MicroBlaze core on Xilinx Virtex-II FPGAs presents designers with a new set of verification challenges. For long transaction cycles on the on-chip peripheral bus (which is part of the IBM CoreConnect architecture specification), software simulation alone may not be the best answer when time is of the essence. If observation of transactions is required, the ChipScope ILA capability of the Virtex-II FPGAs can provide a hardware-based assist that is much faster than software-only techniques. If observation and stimulus, on-chip or at the system level, is required, a MicroBlaze Development System like Raptor offers a hardware-based solution that you need to complete system verification on schedule.

For more information on the Raptor Development Environment, contact Bob Read at Experience First at (408) 985-9683 or *bread@expfirst.com*.

New Flexbus-4 Core

A Seamless Solution for 10 Gbps Networking Applications

by Paul Morrison
System Logic and Networking Engineer
paul.morrison@xilinx.com

Xilinx has developed a new LogiCORE™ core to give you a fully compliant solution for Flexbus-4, which is easily and quickly integrated into your networking system. Through user-configurable options, the Xilinx Flexbus-4 core provides maximum flexibility while seamlessly operating with the Applied Micro Circuits Corporation (AMCC) Application Specific Standard Products (ASSPs) to guarantee maximum bandwidth data transfers. The Flexbus-4 core is also fully compliant with the Optical Internetworking Forum's System Packet Interface-4 (OIF SPI-4) Phase I standard, a subset of Flexbus-4.

The Flexbus-4 core interfaces between Physical (PHY) and Data Link Layer devices within networking applications. The core communicates between devices using the Flexbus-4 interface standard, transferring data in excess of 10 Gbps, assuring compliance to the OC-192 data transfer standard.

AMCC developed Flexbus-4 as an interface to transfer data at 10 Gbps, the current leading-edge networking bandwidth. They have designed three different ASSPs (Ganges-I, Ganges-II, and Khatanga), also known as framer chips, that implement the PHY side of the interface. The Flexbus-4 core has been verified in hardware to operate with each of these PHY devices.

Core Functional Overview

Figure 1 shows the major blocks of the Flexbus-4 core. The source and sink Flexbus-4 blocks transfer data to and from the PHY device. The back-end interface enables transmitting and receiving data through a basic FIFO interface, and it's optimized for maximum flexibility and bandwidth. Finally, the scheduler gives you control over the sequence in which data is transferred out of the Flexbus-4 source interface.

Flexbus-4 Interface

The Flexbus-4 core communicates with the AMCC framer chips via the Flexbus-4 interface. To be compliant with the Flexbus-4 standard, the core supports the following:

- Data transfer on a 64-bit bus operating at 200 MHz, to transfer up to 12.8 Gbps.
- Symmetric interface on both the PHY and Link Layer devices, allowing the core to be used on either the PHY or the Link side of the bus.
- Point-to-point connection between a single PHY Layer device and a single Link Layer device.
- Source-synchronous clocking, where the source of the data provides a clock. This is used to simplify printed circuit board (PCB) design, by eliminating the need for complex clocking schemes.
- Out-of-band control signals, including Channel Address, Start and End of Packet, and Packet Error indications. The implementation of individual signals reduces overhead since they are not transmitted on the data bus.
- HSTL Class I I/O(EIA/JEDEC Standard EIA/JESD8-6), operating at either 1.5V or 1.8V.
- Multiple data transfer modes, including:

  Packet-Over-SONET (POS) – POS mode transmits data formatted in variable-length packets

  Asynchronous Transfer Mode (ATM) – ATM mode transmits data in 53-byte ATM cells.

  Direct-Data Mapped Mode (DDM) – DDM mode is a test mode specified by the Flexbus-4 specification, and is not supported by the OIF SPI-4 Phase I specification.

The Flexbus-4 interface provides you with configuration inputs to select between the three different data transfer modes independently for each channel. Transfer mode for each channel may be switched in-system, allowing real-time reconfiguration and the ability to operate in any networking system.

The Flexbus-4 interface also performs error checking across the data received from the AMCC framer chips. In ATM mode, if the cell received is not 53 bytes long then a cell length error is reported. Parity is checked in all modes, and an error is reported if parity is not correct.

FIFO Interface

The sink FIFO read interface is a standard FIFO interface that transmits data that has been received from the AMCC framer chips.

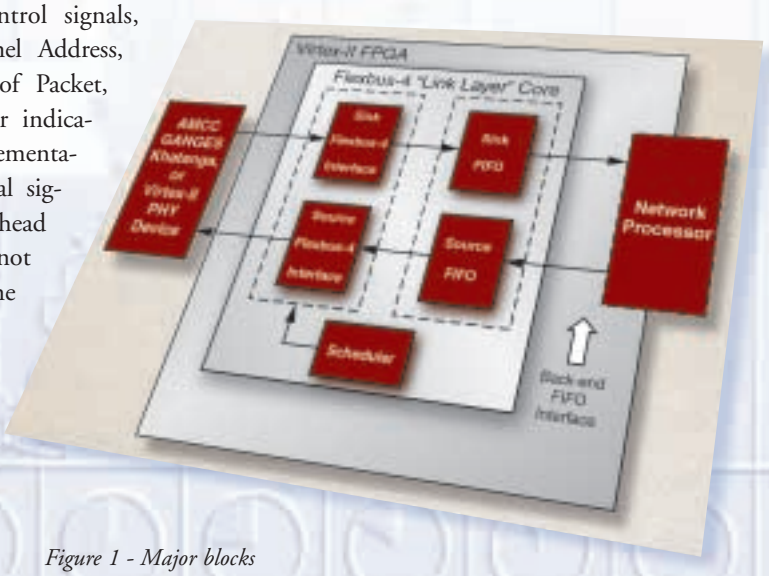

Figure 1 - Major blocks of the Flexbus-4 core

It has been optimized to support the following two options.

- Provide data for any given channel in any given clock cycle. This provides you with the flexibility to determine what channel gets serviced at what time, to ensure that the FIFO for a given channel doesn't overflow.
- Minimize overhead when switching between channels. To keep up with the Flexbus-4 sink interface, the FIFO interface needs to transfer data out as quickly as it is received. The sink FIFO read interface allows for switching between channels with no overhead, to meet this requirement.

The source FIFO write interface accepts data that is to be transferred to the AMCC framer chips. Data written into the source FIFO can target any channel on any clock cycle, providing flexibility in meeting bandwidth requirements for each channel within the system. Additionally, each channel has its own source FIFO, such that if the FIFO for one channel gets full, data transfers can still continue to all other channels.

Scheduler

The scheduler is a programmable sequence table, designed to specify the channel order that the Flexbus-4 source interface uses to send data out onto the Flexbus-4 bus. The scheduler is only used in the source direction of data transmission, and does not effect the data flow in the sink direction. The flexibility of the scheduler is extremely useful for meeting bandwidth requirements for different channels, as well as using available bandwidth efficiently.

Figure 2 shows an example of the programming of the scheduler, and how it translates to data transferred across the Flexbus-4 bus. The transfer rate per channel provided by the core is a ratio of the number of times a particular channel is written into the scheduler vs. the total number of locations used in the scheduler. If a particular channel requires four times more bandwidth than other channels (as is shown in Figure 2), it is written into the scheduler four times more often. Since the order and frequency of channels written into the scheduler is completely user-dependent, you are able to get any combination of transfer rates between channels that you require.

To take advantage of the maximum available bandwidth of 12.8 Gbps, the scheduler ignores channels that do not have data available to transfer, and splits up the available bandwidth between the channels that are ready to transfer data. If only a single channel within the system contains data, then it will get the full bandwidth with no overhead required when checking for data available on all other channels. If all of the channels within the system contain data, they will get bandwidth according to the programming of the scheduler.

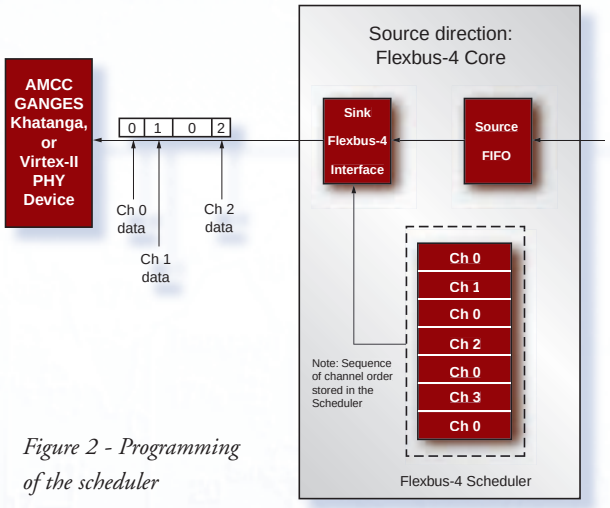

Figure 2 - Programming of the scheduler

What's Provided

To make implementation and integration quick and easy, Xilinx provides several additional support features in addition to the Flexbus-4 core. The following are included:

- **Loopback logic** – A design to loop data from the sink FIFO to the source FIFO is provided in both Verilog and VHDL. In addition to providing the loopback functionality required in many networking systems for error isolation, this design also provides an example interface that will be extremely useful when you design your interface to the core.
- **Instantiation Templates** – An instantiation template for the Flexbus-4 is provided in both VHDL and Verilog. This template provides an example of instantiating the core, as well as showing how to create source-synchronous clocking using Virtex-II Double-Data Rate (DDR) registers.
- **Build Script** – A push-button build script is provided that places and routes the Xilinx core and the loopback logic. This is useful for demonstrating that the core meets the 200 MHz timing requirement within the Virtex-II FPGA, as well as providing a sample build script for the core.
- **Demonstration Testbench** – A demonstration testbench to show toggling of the signals interfacing to the Flexbus-4 core is provided in both VHDL and Verilog. This can be easily modified to view operation of the Flexbus-4 core, and the different core configuration options available.

Conclusion

The Xilinx Flexbus-4 cores are currently available from the Xilinx IP Center at *www.xilinx.com/ipcenter/flexbus/flx4.htm*. There are three cores currently available, designed to interface to network systems that contain a single OC-192 channel, four OC-48 channels, or sixteen OC-12 channels. Each of these cores has been verified in hardware to operate with the AMCC Ganges-I, Ganges-II, and Khatanga framer chips at the OC-192 transfer rates.

Beyond being compliant with the Flexbus-4 specification, this core has been optimized to provide maximum flexibility, to ensure optimal bandwidth for your system's configuration. The Xilinx Flexbus-4 solution also enables you to quickly develop an OC-192 interface ensuring you the fastest time to market.

New Spartan-IIE FPGA Family for Digital Consumer Convergence Applications

Spartan-IIE FPGAs offer significant performance improvements for next-generation consumer products.

by Ashok Chotai
Senior Competitive Marketing Engineer
ashok.chotai@xilinx.com

The new Spartan™-IIE family is the latest generation of the highly successful Spartan series of programmable logic devices. It is targeted at digital consumer convergence applications where video, audio, communication, and data technologies converge to add more functionality and capability to the end product. The Spartan-IIE family supports system performance beyond 200 MHz and offers features that include distributed RAM, block RAM, 19 selectable I/O standards (including high-speed differential LVDS), and four DLLs, as shown in Table 1.

Distributed RAM

Distributed RAM is an ideal solution for designs that require multiple, small, fast, and flexible memories situated close to the logic. As with all other Xilinx FPGA families, the 4-input look-up table (LUT) in Spartan-IIE devices can be also used as memory, where it can be configured as ROM, and single-port or dual-port synchronous RAM. Each LUT is equal to 16 locations x1-bit wide, as shown in Figure 1. These memories can even be cascaded for various data widths or depths. The competing FPGA families do not have this feature.

Block RAM

In addition to the distributed RAM, Spartan-IIE devices offer dedicated blocks of synchronous RAM for designs requiring larger memory functions, as illustrated in Table 2. The memory block can be used as 4096x1, 2048x2, 1024x4, 512x8, or 256x16. These blocks can also be cascaded for a different configuration or to form larger memory functions.

True Dual-Port Operation

The memory block is dual-port with independent control signals for each port. Hence, these ports can be read from and written to simultaneously, independent of each other. All the control logic is implemented within the RAM block. The competing FPGA families have only 2-port RAM – one read port and one write port. To emulate Xilinx dual port capability, they would need twice the number of memory blocks and additional control logic – and our competitors would still operate with relatively lower performance.

Clock Management

As FPGA density grows, quality on-chip clock distribution becomes increasingly important. Clock skew and clock delay impact device performance. The Spartan-IIE family resolves this potential problem by providing four fully digital dedicated on-chip Delay-Locked Loops (DLL) in each device. These DLLs are used to remove on-chip and off-chip clock delays (de-skew), and to perform clock multiplication, division, and phase shifting.

As a fully digital implementation, the Spartan-IIE DLLs do not have the typical problems encountered with analog phase-locked loops (PLLs), which are extremely sensitive to noise and are difficult to migrate between process technologies. The Spartan-IIE DLLs are not sensitive to noise and process variations. Xilinx DLLS are a standard feature in all densities and all speed grades of the Spartan-IIE family.

Device	XC2S50E	XC2S100E	XC2S150E	XC2S200E	XC2S300E
System Gates	50K	100K	150K	200K	300K
Logic Cells	1,728	2,700	3,888	5,292	6,912
Distributed RAM Bits	24,576	38,400	55,296	75,264	98,304
Block RAM Bits	32,768	40,960	49,152	57,344	65,536
DLLs	4	4	4	4	4
Max I/Os	182	202	263	289	329

Table 1- Spartan-IIE Family Product Matrix

Device	No. of Blocks	Block RAM Bits
XC2S50E	8	32,768
XC2S100E	10	40,960
XC2S150E	12	49,152
XC2S200E	14	57,344
XC2S300E	16	65,536

Table 2 - Block RAM Amount

Flexible Logic Resources

The Spartan-IIE family provides flexible configurable logic blocks (CLBs), allowing implementation of simple or complex logic functions. Each CLB contains two slices. Each slice contains two 4-input Look-up Tables (LUTs), one carry logic chain, two

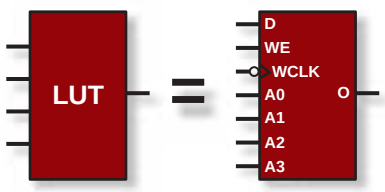

Figure 1- Distributed RAM

storage elements, and control logic (see Figure 2). In addition, each CLB has two dedicated 3-state buffers. These buffers can be used to implement on-chip bussing. The competing FPGA families do not have dedicated 3-state buffers; they emulate this functionality at the expense of logic resources.

Look-Up Table Flexibility

- **Wide input functions** – Multiple LUTs can be cascaded to implement functions of more than four inputs. Furthermore, each CLB has dedicated resources to implement 5-input and 6-input functions in just a single logic (LUT) level of delay. For example, you can implement an 8-to-1 multiplexer within a CLB with one logic level of delay, versus the competing FPGA solution with at least two logic levels of delay.
- **Shift Register** – The LUT can also be configured as a 16x1 single-port RAM, or as a 16-bit shift register. This functionality can be used to implement high-speed Linear Feedback Shift Register (LFSR) counters and other DSP applications.

Versatile I/O

The Spartan-IIE family provides 19 I/O standards that include differential standards such as LVDS, Bus LVDS, and LVPECL, allowing a single FPGA to interface with any other device without using external converters or translators (see Table 3). Differential signaling provides higher performance while reducing power, reducing noise, and lowering EMI emissions.

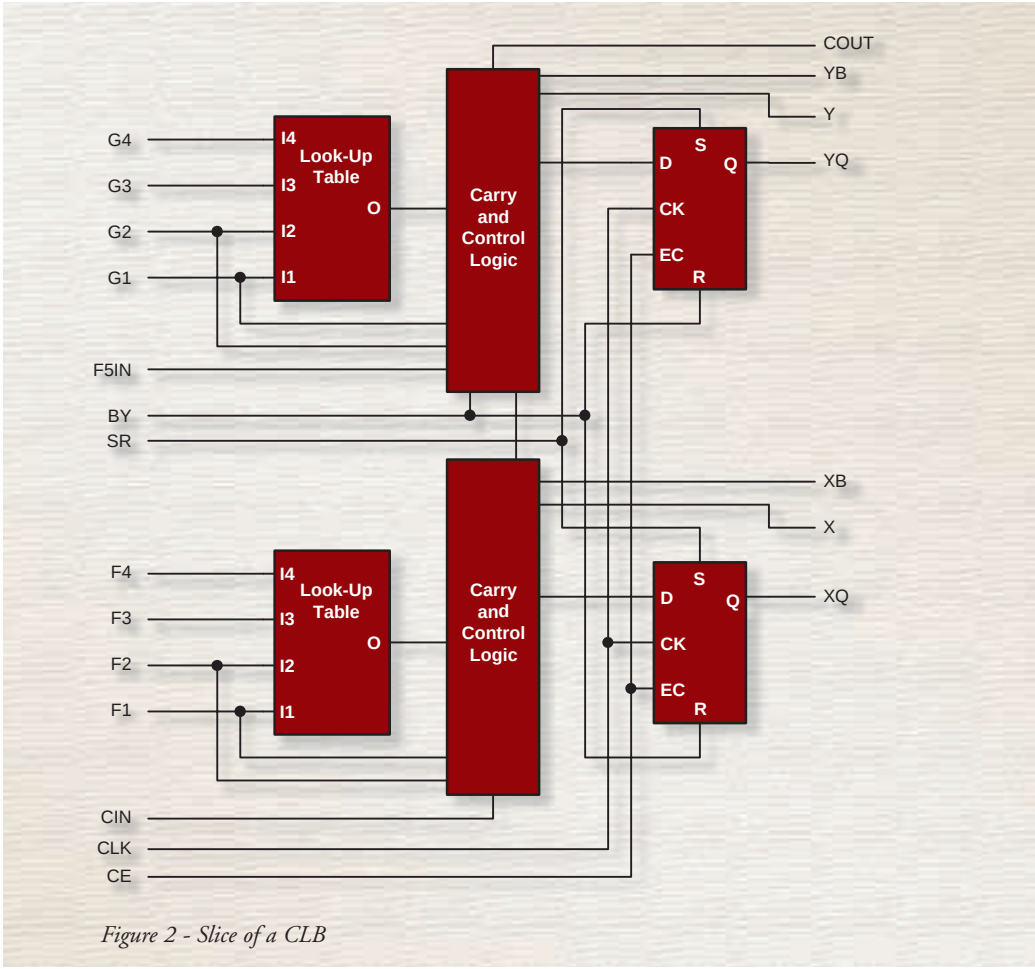

Figure 2 - Slice of a CLB

- **LVDS and Bus LVDS** – The Low-Voltage Differential Signaling (LVDS) standard exists in two implementations: LVDS and Bus LVDS. LVDS is used for unidirectional data transfer, and is optimized for high-speed point-to-point links. Bus LVDS is for bi-directional communication between two or more devices, and is optimized for multi-drop configurations. You can use up to 120 pairs of LVDS and Bus LVDS I/O pins on the largest Spartan-IIE device.
- **LVPECL** – The Low Voltage Positive Emitter Coupled Logic standard is used in 100+ MHz chip-to-chip interfaces. It is also used for transmission of clocks at frequencies over 100 MHz.

Chip to Chip Interface	Backplane Interface	Memory Interface
LVTTL	PCI 33/66 MHz, 3.3V	HSTL-I
LVCMOS2	GTL	HSTL-III & IV
LVCMOS18	GTL+	SSTL3-I & II
LVDS	AGP-2X	SSTL2-I & II
LVPECL	Bus LVDS	CTT

Table 3 - Spartan-IIE I/O standards and typical applications

Software and IP Cores

The Spartan-IIE family is supported in ISE WebPACK™ 4.1i and all other configurations of the ISE 4.1i software. ISE WebPACK is free software that can be downloaded from the Xilinx website. Xilinx also provides the most comprehensive set of IP Core solutions to assist you in reducing your design cycle time.

Conclusion

The new Spartan-IIE FPGA family helps you meet the requirements of digital consumer convergence systems. It gives significant performance and flexibility enhancements by providing on-chip features such as true dual-port memory for system integration, flexible logic resources for simple to complex logic implementation, DLLs for clock management, and single-ended and high speed differential I/Os for faster chip-to-chip speeds. For more information, please see the complete article on Xcell Online at: *www.xilinx.com/publications/xcellonline/.*

Networking Comes Home

A silent revolution is happening – the revolution of home networking. And while this revolution should give us no cause for concern, it does, however, hold the promise of forever changing and improving our lives.

by Robert Bielby
Senior Director, Strategic Solutions
robert.bielby@xilinx.com

Although much hype, confusion, and unfulfilled promises have surrounded networking in the home, it is clear that the technology is here to stay and will continue to grow and evolve. Consumers are now becoming sufficiently motivated to take on the numerous challenges of understanding and installing complex hardware and software to create home networks. Although this motivation has come primarily from the desire to share expensive resources such as Internet access, printers, and scanners with multiple PCs, it has set the stage for the revolution.

Living the Digital Life

Hand-in-hand with the home networking revolution is the invasion of the digital lifestyle. Consumers are becoming more aware of the higher quality, performance, and accuracy offered by anything that is "digital." The trend of digitizing all media, audio, video, and data has not only enabled a higher quality experience, but has also brought about the additional benefits of increased portability and the ability to easily share media.

The cost benefits from sharing resources may serve as a catalyst for consumers to build home networks, but the real benefits of home networking will be the enhanced entertainment experience and productivity that consumers will enjoy through the sharing of all digitized media.

Future Gazing

So, what is the networked home going to look like five years from now? Beyond the basics of sharing data, it's hard to say exactly which direction it will take. It is virtually impossible to fully understand the implications and benefits that will be derived from the networked home. Years ago, no one could have predicted that a collection of globally connected computers – the Internet/World Wide Web – would grow to become such an integral component of our everyday lives.

Can something as simple as the ability to share files or media really impact our lives? Ask the music industry (and soon the video industry) as they grapple with rampant sharing of media across the Internet. File-sharing programs like Napster and Gnutella are causing the entire music industry to rethink their business models.

So, why is Xilinx, the pioneer of programmable logic, dedicating such a large section of *Xcell* to home networking? Actually, nothing could be more natural or appropriate. Programmable logic has already played an essential role in catalyzing the growth and rapid evolution of the Internet. Used throughout the Internet switching infrastructure, reprogrammable FPGAs have allowed pioneering manufacturers to successfully navigate the ever-changing standards and protocols evolving in the virtual marketplace.

Because home networking is in its infancy, standards are still evolving. It's clear that programmable logic will play a central role in this "new" networking application. Programmable logic provides the flexibility and fast time-to-market required to enable this market to evolve and mature.

Home Networking Smorgasbord

In this issue of *Xcell*, we are covering many of the new and emerging home-networking technologies and products, including broadband access, residential gateways, information appliances, voice and data convergence, and wireless LANs. In each article, you will learn how Xilinx programmable logic solutions provide flexible, cost-effective solutions with all the traditional benefits of fast time to market.

In this issue, you will also find an article dedicated to Bluetooth™ networking – a new low-power wireless technology designed to "cut the cord" on the personal computer peripherals. Over recent time, Bluetooth wireless technology has become the leading wireless standard of a new networking paradigm – the personal area network. PANs allow all your information appliances to communicate with one another – not only in the house, but also on the road and in the workplace. Xilinx solutions are already playing a prominent role in the deployment of products based upon the Bluetooth standard, and Xilinx will continue to play an even greater role as this wireless protocol continues to evolve to support higher data rates and other feature sets in the future.

Because home networking technologies are evolving and changing almost daily, it is almost certain that once this issue of *Xcell* is published, many sections of the text will soon be either outdated or obsolete. As of this writing, new wireless standards, such as UWBD (Ultra Wideband) technology, the Bluetooth protocol, and IEEE 802.11b (also known as Wi-Fi), are challenging some of the "new" existing standards. Furthermore, the next version of USB standards, v2.0, has arrived on the scene providing data transfer rates that are 40 times faster than the previous version USB 1.1.

Conclusion

As an on-going supplement to articles published in *Xcell*, and as a means to remain current, Xilinx has established a Web portal called eSP (emerging Standards and Protocols) to keep design engineers up-to-date on this new and exciting technology. This Web portal is an industry first, and is dedicated to accelerating all phases of the product development schedule by providing education on these new standards in conjunction with intellectual property and key reference designs. Now, staying current on the most competitive home networking technologies is just a mouse click away at *www.xilinx.com/esp/*.

We hope you enjoy this special section in the *Xcell Journal*. In future issues, we plan to continue to keep you updated on the most significant and current events in home networking technologies – and all the other new and exciting applications where Xilinx solutions can be found.

Bluetooth Wireless Technology and Personal Area Networking – In Your Home and on the Road

Bluetooth wireless technology and personal area networking promise to radically change the way we live and work.

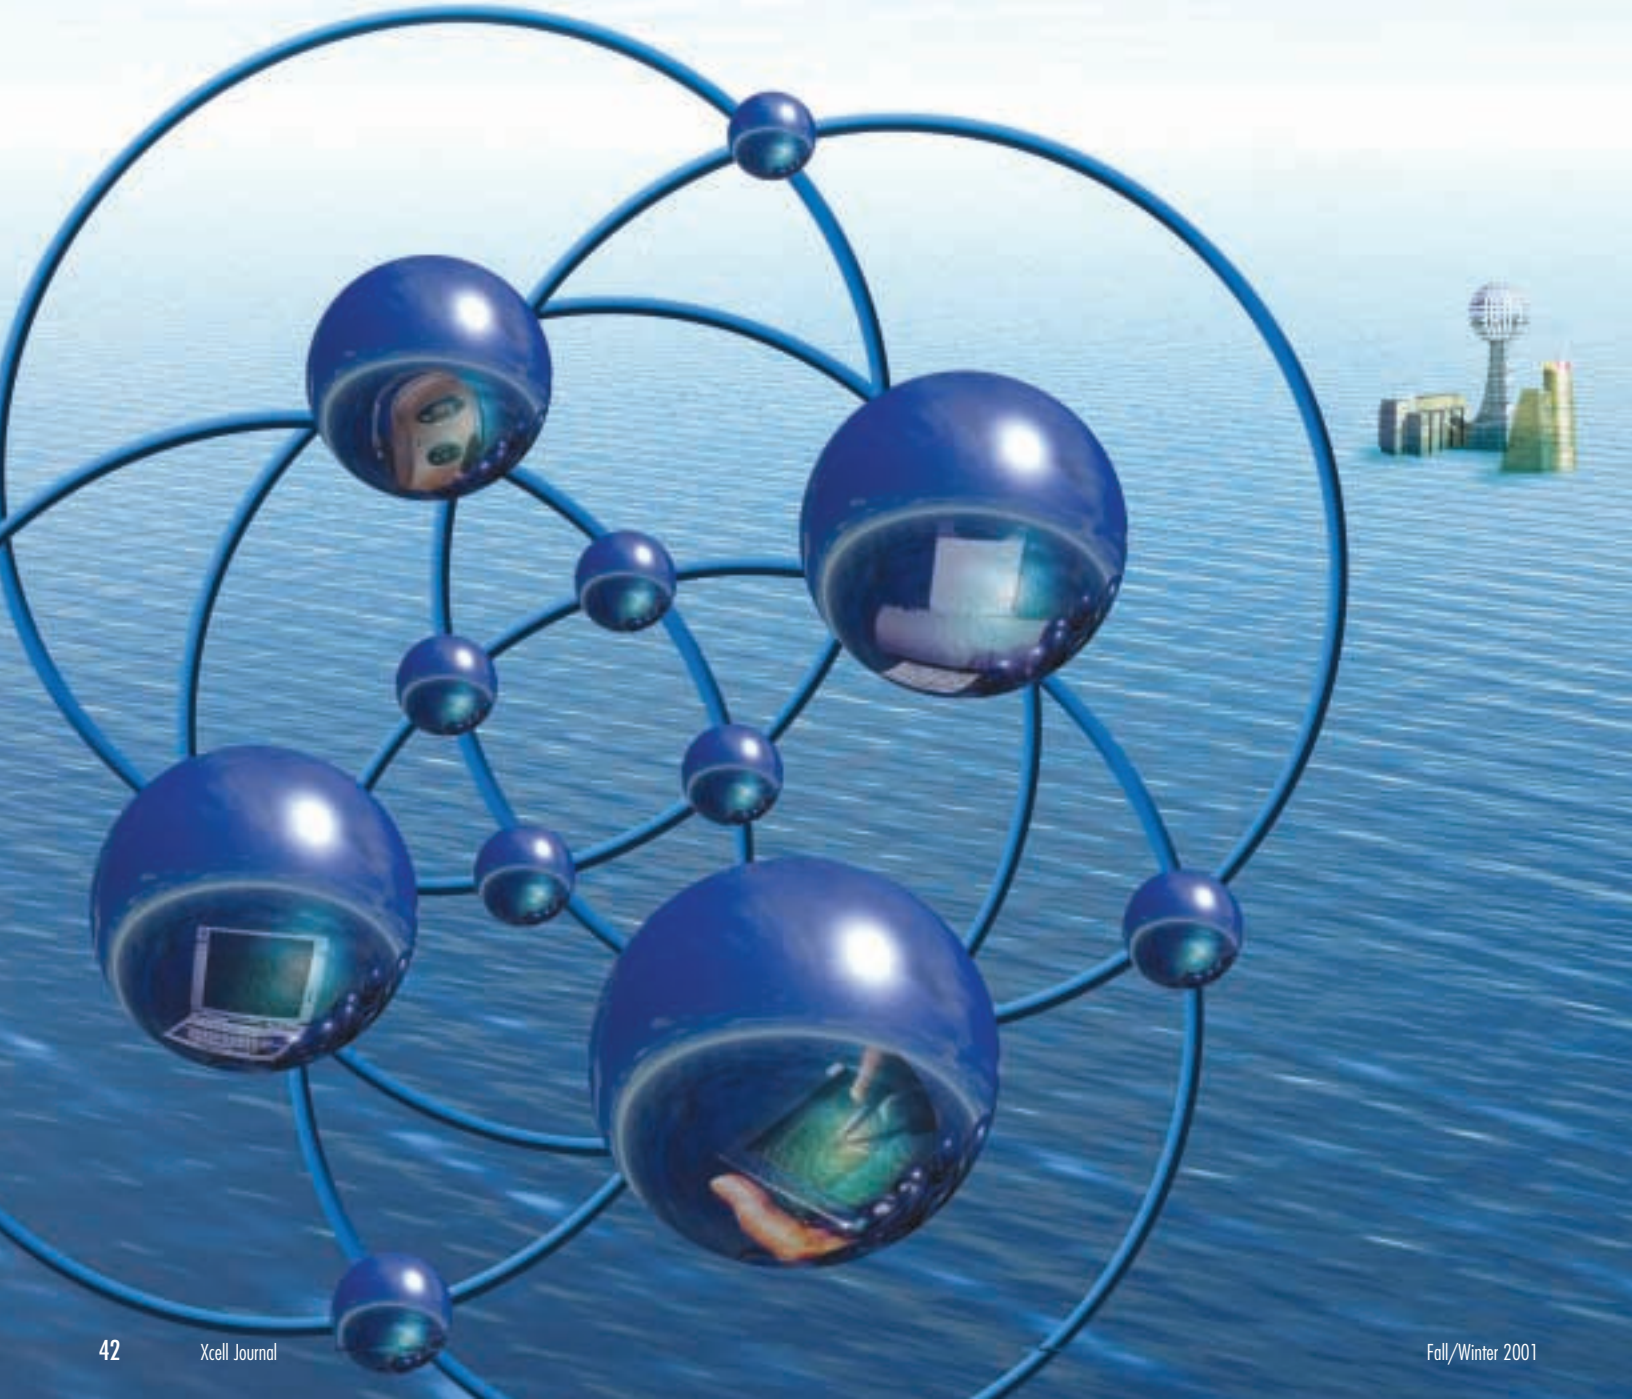

by Krishna Rangasayee
Senior Manager, Strategic Solutions
krishna.rangasayee@xilinx.com

Mike Nelson
Senior Manager, Strategic Solutions
mike.nelson@xilinx.com

Are you ready to cut the cords? Then wireless networking is the way you want to go – but with as many as a half dozen competing wireless technologies, which is right for you?

We don't know what is right for you, but we do know that Bluetooth™ wireless technology is the leading networking technology for enabling connectivity across a wide range of information appliances. Although the performance capabilities of Bluetooth wireless technology are somewhat modest (in comparison to the other technologies – see Table 1), they are acceptable for almost every networking application you need except high-quality video. At 720 Kbps over a range of 10 meters, the quality of voice and data transmission is equivalent to hard-wired networks. Furthermore, the unique functional capabilities, low cost, and support of Bluetooth wireless technology by more than 2,490 companies worldwide make this wireless networking system one of the highest volume home-networking applications available today.

The three main objectives of the Bluetooth wireless technology are:

- To eliminate the connecting wires associated with most consumer electronics and computer equipment
- To allow a collection of products to function as an intelligent whole
- To make personal area networking (PAN) seamless.

Personal area networking is a new connectivity paradigm supported by Bluetooth brand products. Bluetooth wireless technology and PANs will change the future of computing and consumer electronics. Using Bluetooth wireless technology, you can create dynamic, ad hoc PANs at home, at work, on the road, and almost anywhere else you choose.

How Bluetooth Wireless Technology Works

Instead of having devices connected by a plethora of cables and wires, each Bluetooth-enabled component has a small radio transceiver, similar to the one shown in Figure 1. These transceiver modules use protocols developed by the Bluetooth Special Interest Group *(www.bluetooth.com)*, of which Xilinx is a member.

Figure 1 - Ericsson Bluetooth transceiver module

Bluetooth Architecture Protocol Stack

The Bluetooth architecture protocol stack allows devices to discover, network, and exchange information with each other

	Technology					
	Bluetooth	HomeRF	802.11b	HiperLAN	802.11a	HiperLAN2
Frequency Band	2.4 GHz	2.4 GHz	2.4 GHz	2.4 GHz	5 GHz	5 GHz
Technology	Frequency Hopping Spread Spectrum	Frequency Hopping Spread Spectrum	Direct Sequence Spread Spectrum	Guassian Minimum Shift Keying	Orthogonal Freq. Division Multiplexing	Orthogonal Freq. Division Multiplexing
Performance	720 Kbps	1.6 Mbps	11 Mbps	23 Mbps	~50 Mbps	~50 Mbps
Range	<10 meters	50 meters	150 meters	150 meters	50 meters	50 meters
Power	Very Low	Medium	Medium	Medium	Medium High?	Medium High?
Relative Cost	Low/ Very Low	Medium/Low	Medium	Medium	High	High
Target Applications	Cable Replacement Wireless Data Wireless Voice Personal Networks	Wireless Data Wireless Voice	Wireless Data	Wireless Data	Wireless Data	Wireless Data
Fixed N/W Support	PPP, Ethernet	DECT, Ethernet	Ethernet	Ethernet	Ethernet PPP, 1394, UMTS	Ethernet, IP, ATM,
Key Features	Very Low Power Voice and Data Roaming Low Cost Good noise immunity	Voice and Data Moderate Cost	Good Performance	Good Performance	High Performance	High Performance
Promoters	2000+	<50	~100	<50	~100	<50
Regional Support	Worldwide	US	US/Asia	Europe	US	Europe
Shipping	Now	Now	Now	Now	2001	2001

Table 1 - Bluetooth specifications compared to competing wireless local area network technologies

seamlessly. It is a layering of functional modules as shown in Figure 2.

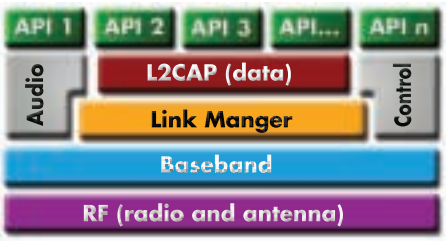

Figure 2 - Bluetooth architecture protocol stack

Application Program Interface Libraries

Software modules called API libraries connect the host application program to the Bluetooth communications system. They reside and execute on the same processing resource as the host system application.

Logical Link Control and Adaptation Protocol

The L2CAP provides the overall control of a Bluetooth system. It manages the high level aspects of each connection (such as determining who is connected to whom, whether to use encryption or not, what level of performance is required, and so on). In addition, L2CAP converts the format of data, as necessary, between the APIs and the lower level Bluetooth protocols. The L2CAP is implemented in software and can execute on either the host system processor or on a local processor in the Bluetooth system.

Link Manager

The Link Manager handles the physical details for Bluetooth connections. It creates the links, monitors their health, and terminates them gracefully upon command or failure. The Link Manager is implemented in a mix of hardware and software.

Baseband

The Baseband is the digital engine of a Bluetooth system. It constructs and decodes packets, encodes and manages error correction, encrypts and decrypts data for secure communications, calculates radio transmission frequency patterns, maintains synchronization, controls the radio, and takes care of all the other low-level details necessary to realize Bluetooth communications.

Radio

The Bluetooth radio converts the digital baseband data to and from a 2.4 GHz analog signal.

How Bluetooth Networking Works

Bluetooth devices connect into piconets – small networks comprised of a master device connected to anywhere from one to seven active slave devices.

When multiple piconets are interconnected, they create wireless networks called scatternets. Figure 3 illustrates a piconet made of nodes A, K, L, M, and G interacting with yet another piconet comprised of nodes H, E, C, K, and L. These two piconets share nodes K and L, and collectively, form a scatternet.

Bluetooth devices have four basic states. They can be any of the following:

- **Master** – In control of a piconet (nodes A and H in Figure 3)
- **Active slave** – Connected and actively monitoring/participating on a piconet
- **Passive slave** – Still logically part of a piconet but in a low power mode; occasionally monitoring, and still synchronized in an inactive state
- **Standby** – Not connected to a piconet, occasionally monitoring for inquiries from other devices, but not synchronized with any other devices (nodes B, J, I, and F).

Standby Mode

Initially, all Bluetooth devices are not synchronized or coordinated in any way. They are all listening at different times and on different frequencies. Bluetooth devices know only about themselves, and in this state, they are in standby mode. Standby is a passive mode where a Bluetooth device listens on an occasional basis, performing what are called inquiry and page scans (for 10 milliseconds out of every 1.28 seconds) to see if any other Bluetooth devices are looking to communicate. Passive behavior is inherent to half of Bluetooth states (standby and passive slave), and is a key mechanism for achieving very low power consumption. In standby mode, the Bluetooth device's occasional scans reduce power consumption by more than 98%.

Paging

Paging is how a Bluetooth device learns about other devices that are within its range. Node A in Figure 3 executes a page command and receives replies from devices within range. Through these replies, device A learns the explicit identity of these other devices (such as their unique Bluetooth device ID).

Piconets

In its general form, a page command establishes a formal device-to-device link between a master (the originator) and a slave. Master/slave connections in Bluetooth are referred to as a piconet. To create the piconet, Device A broadcasts the page command with the explicit device ID of the target slave (Device D in Figure 3). All Bluetooth devices except Device D will ignore this command, because it is not addressed to them. When the Device D replies, Device A will assign it an Active Member Address in the piconet. As an active slave, Device D will begin continuously monitoring for further commands from Device A, in synchronization with Device A's hopping pattern and clock offset. Furthermore, standard piconet activity continuously updates the clock-offset data, keeping the synchronization extremely accurate. Through successive page commands, a Bluetooth master can attach up to seven active slaves.

Scatternets

Each Bluetooth node is capable of maintaining multiple states simultaneously. This allows multiple piconets to combine into a structure called a scatternet. In Figure 3, two piconets combine into a scatternet through the common slaves K and L. Scatternets can evolve into extremely complex structures. Note that a node can potentially be a master, an active slave, and a "parked" slave on three different piconets, all at the same time.

The only hard limitation on scatternet configurations is that each Bluetooth node can only be the master of one piconet at any given time.

A critical feature making such configurations practicable is the support for Quality of Service (QoS) in the L2CAP. Through this mechanism in the L2CAP, Bluetooth devices are able to determine that the connections they are establishing are feasible and sustainable. Thus, a node would not agree to an additional connection if it would require bandwidth that it is unable to support. It may, however, negotiate to establish another connection, but at a more modest data rate that it can support.

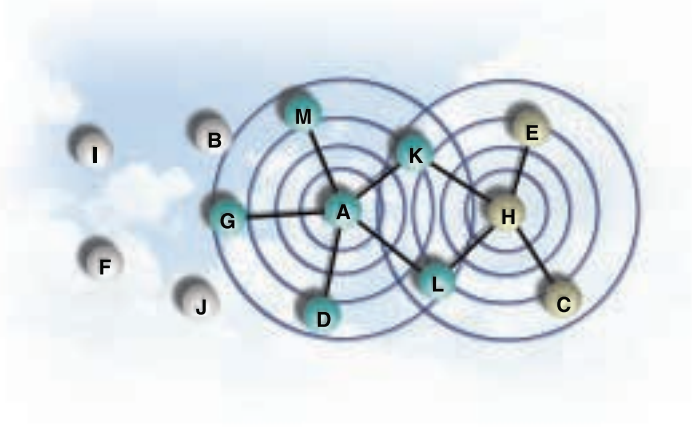

Figure 3 - Two Bluetooth piconets participating in a scatternet

Bluetooth PANs – Changing the Future of Connectivity

Bluetooth wireless technology will be widely accepted and deployed in the next few years – and Bluetooth personal area networks will revolutionize the way we interact with intelligent appliances in the future.

Bluetooth PANs will provide automatic visibility and access to Bluetooth-enabled products and services within a range of 10 meters (and up to 100 meters in high-performance systems). PANs will give you tremendous user-friendly control and flexibility. With PAN technology, you will be able to organize a collection of personal electronic products – your PDA, cell phone, laptop computer, desktop computer, MP3 player, and so on – to automatically work together. For example, the contact manager and calendar in your PDA, laptop computer, and desktop computer will all automatically synchronize whenever they are within range of each other.

Ultimately, Bluetooth technology will eliminate most cables related to consumer electronics in the home and on the road. Your PC, scanner, and printer will simply need to be within 10 meters of each other to work together. Your PDA, digital camera, and MP3 player will no longer need docking stations to transfer files or to get the latest information. And, your home stereo and other equipment will be wirelessly networked as well. What's more, your cell phone will synchronize its address book with your PC and function as a handset to your cordless phone in the house, enabling you to use the cheaper landline network.

One of the nice things about your PAN is that you can take it with you wherever you go. On the road, much of your Bluetooth PAN will go with you. Even when your laptop is in your briefcase and your cell phone is in your pocket, they will be able to collaborate to access e-mail. Moreover, next generation cell phones featuring Bluetooth and General Packet Radio Service technologies will function as wireless modems with Internet access speeds greater than 100 Kbps. With such performance, this may well be the "killer app" that ensures the widespread adoption and success of Bluetooth technology.

In your car, your cell phone will operate in a hands-free mode, allowing you to use the car audio system and an in-dash microphone while the phone itself rests comfortably in your pocket. Or you may use a wireless Bluetooth headset instead. For entertainment, your MP3 player will likely play music in eight-speaker surround sound, rip music right off an FM broadcast, or record your phone calls for later review – wirelessly.

When you're on foot, fixed landline access points, such as a pay phone in an airport or the desk phone in your hotel, will provide true broadband access. Also, look for pay phones to evolve and compete with your cell phone, because pay phone landlines – at the moment, anyway – are cheaper. In the world of deregulation and open competition, future smart phones may automatically put your calls out for bid and channel the traffic over the carrier that offers the lowest cost.

Xilinx eSP – Solutions at Your Fingertips

Earlier this year, Xilinx launched the eSP (emerging Standards and Protocols) Web portal *(www.xilinx.com/esp/)*. The first major project of the eSP portal has been to accelerate and facilitate the development of Bluetooth-based products. The site provides you with solutions ranging from the best Bluetooth-specific reference designs to the latest intellectual property resulting from our collaboration with the wide range of industry leaders participating in the Bluetooth Special Interest Group.

With the introduction of the eSP Bluetooth portal, Xilinx has taken a leading role in enabling you to achieve the successful development and deployment of products based on this fast-emerging wireless technology.

Conclusion

Although Bluetooth technology is wireless, its usage model and end-application space is quite different from wireless LANs. Bluetooth-enabled products provide a seamless interface among various information appliances and create mobile personal area networks to enrich the digital lifestyle. Bluetooth technology does not support the higher data rates of wireless LANs, but it does support both wired-quality voice and data transmissions. The voice/data capability, in combination with low prices, global acceptance, and go-anywhere PANs, bodes well for the success of Bluetooth-brand products in the marketplace.

Home Networking – Integrating Information Appliances with Personal Computers

While most "information appliances" perform a single function very well, the personal computer will remain the backbone and the gateway of the home network.

by Amit Dhir
Manager, Strategic Solutions
amit.dhir@xilinx.com

Krishna Rangasayee
Senior Manager, Strategic Solutions
krishna@xilinx.com

A plethora of new consumer devices – known as information appliances (IAs) – are invading the consumer market today, with a lot more to come. IAs such as MP3 players, PDAs, digital cameras, set-top boxes, gaming consoles, cellular phones, and digital VCRs are gaining popularity and acceptance among consumers. The market research firm Dataquest predicts that the worldwide unit production of IAs will explode from 1.8 million in 1999 to 391 million in 2003. Another research firm, IDC, predicts that in 2001, 18.5 million IA units will ship compared to 15.7 million personal computers (PCs). Despite the growth in the shipments of IAs, PCs will continue to penetrate more households for several years to come.

Dataquest reports PC penetration now exceeds 50% of US households – of the 102 million US households, 52 million own a PC. Dataquest also forecasts that multi-PC households will grow from 15 million in 1998 to more than 26 million in 2003.

In this article, we explore the underlying market dynamics and technological hurdles that face consumers and IA manufacturers in crossing the PC-centric home-networking threshold.

The Invasion of Information Appliances

The rapid growth in multi-PC households is creating the need for sharing broadband Internet access, files/data, peripherals (such as printers and scanners), and IAs among the multiple PCs in different rooms of a house. This need has given birth to the fast-emerging home-networking market.

The PC is the most important and widely used device for computing, Internet access, online gaming, chatting, e-mail, and data storage, and is hence, the residential gateway to network the home. With worldwide PC shipments totaling 134.7 million units in 2000 and predicted to exceed 200 million in year 2004, the PC remains in healthy demand. This is due to the ever-improving price per performance ratio and the position of the PC as a productivity tool. However, PCs have several weaknesses – such as being complex, buggy, and confusing.

On the other hand, IAs are a promising category of consumer digital electronics that provide users with low-cost, easy-to-use, instant-on devices that are lightweight, reliable, and offer special-purpose features and benefits. Some of these IAs include digital cameras, digital displays (PDP, LCD), digital TV, Internet audio (CD, MP3) players, DVD players and recorders, set-top boxes, gaming consoles, Internet screen phones, Web pads, security units, energy management units (automated meter reading), VoIP phones, smart handheld devices like PDAs and handheld PCs, Web terminals, e-mail terminals, mobile phones, automobile PCs, and "white goods" (dish washers, dryers, washing machines).

Smarter chips with increased functionality are being embedded into everyday consumer products and IAs. While these devices are quite useful as standalone units, they can provide consumers with even more value and convenience when they are networked with each other.

In the year 2001, unit shipments of IAs will exceed unit shipments for PCs in the US. Parks Associates predicts that in the US, there will be 22 million home IAs sold in 2001 (excluding Internet-enabled mobile phones and telematics systems) compared to 18 million home PCs.

The factors driving IAs are aggressive vendor marketing, consumer demands for Internet connectivity, advancing bandwidth capacity, and lower product costs. IAs are targeting three specific areas:

- Replacing PCs with robust Web browsing, e-mail, and interactive devices
- Supplementing PCs with coexisting peripherals like PDAs, printers, and scanners
- Sidestepping PCs with set-top boxes and cellular phones.

The PC Still Rules

Despite the IA invasion, the home PC is not going away soon, because PCs offer:

- **Compatibility** – Interoperability between documents for business, education, and government is essential. IAs often support only proprietary, non-PC supportive media standards.
- **Flexibility in the PC platform** – Video editing, music recording, Web authoring and hosting, gaming, e-mail, word-processing, financial analysis, and database management can all be done in one PC. Comparatively, most IAs are dedicated to only one or two functions.
- **Investment** – The massive corporate investment in the PC industry is simply too large to abandon – and there is no good reason to abandon it.
- **Momentum** – More than 50% of consumer homes in the U.S. have PCs – and the number is growing. Worldwide, PCs have a huge installed base that makes the World Wide Web not only possible, but necessary to modern society.
- **Pace of improvement** – Faster processors, bigger hard drives, better communication protocols, and other quantum leaps in the power of the PC, guarantee that it will remain the centerpiece of the Information Age for some time to come.

Also, as the price of PCs drop while their computing power skyrockets, the lines are becoming blurred between some PCs and IAs. In the under $500 price range, the handheld PC is eclipsing the PDA, for instance.

The higher intelligence and multifunctionality of the PC makes it the ideal platform for being the residential gateway for home networking. PCs are already being shipped with onboard devices designed for broadband Internet access and with ports to support home networking of multiple information appliances.

Figure 1 shows a generic PC configured for home networking. In addition to the standard processor, memory, hard-disk drive, and operating system, this PC provides a platform that is ideal for multiple functions such as wired Ethernet and wireless Bluetooth™ home networking. Future generation PCs will provide even greater home-networking capabilities to interconnect multiple PCs, PC peripherals, and a host of information appliances.

There are applications that information appliances perform well. A Web pad provides the capabilities of scheduling, ordering groceries, sending and receiving e-mail with portability. While very convenient within the home, a Web pad requires an access point and a gateway to a high-speed Internet connection. The PC can provide all the above functions and more, such as video editing, gaming, and word-processing. Hence, PCs will maintain a stronghold on consumer homes.

Conclusion

The growing ubiquity of the Internet, and the convergence of digital voice, data, and video are bringing interesting applications to the home. Lower PC prices have brought multiple PCs into the home. However, newer information appliances are arriving in the market that use the Internet to provide particular functions independent of a PC. The big question, then, is whether the information appliance will replace the PC in its entirety. We believe this is not likely. The information appliance provides a dedicated functionality, whereas the PC provides a platform to provide networking capabilities to a host of information appliances, PC peripherals, and other PCs. Hence, while information appliances will continue to penetrate homes – and some IAs will access the convenience of the Internet without the assistance of a PC – they will not replace PCs. Ultimately, to achieve their highest functionality, information appliances and the PC must be integrated.

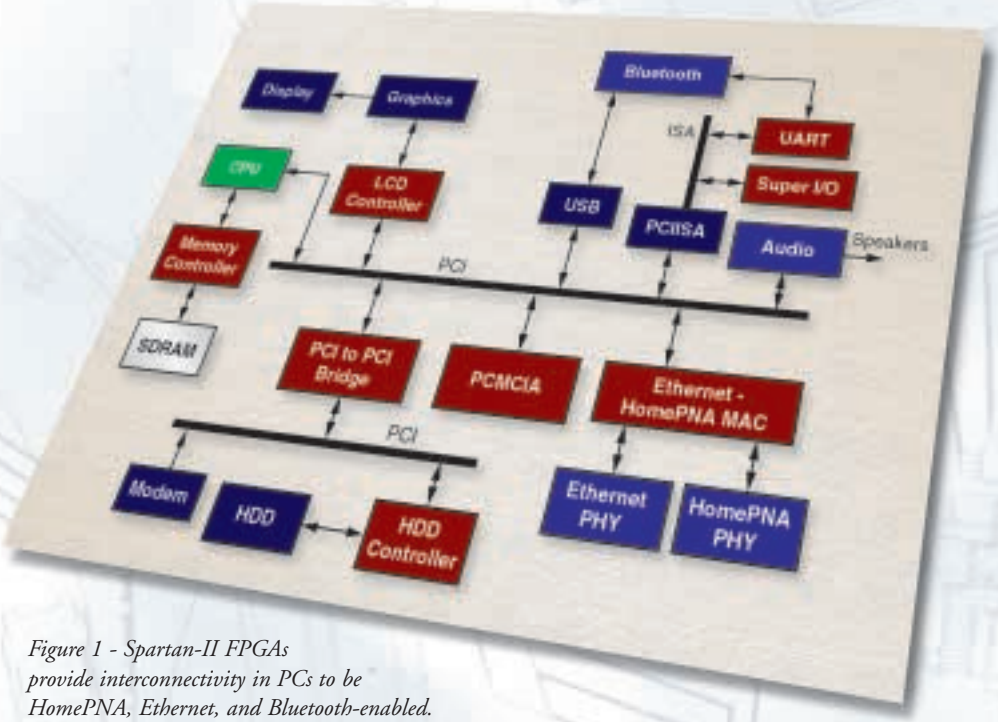

Figure 1 - Spartan-II FPGAs provide interconnectivity in PCs to be HomePNA, Ethernet, and Bluetooth-enabled.

Programmable Logic Will Drive Growth in Broadband Access and Home Networking

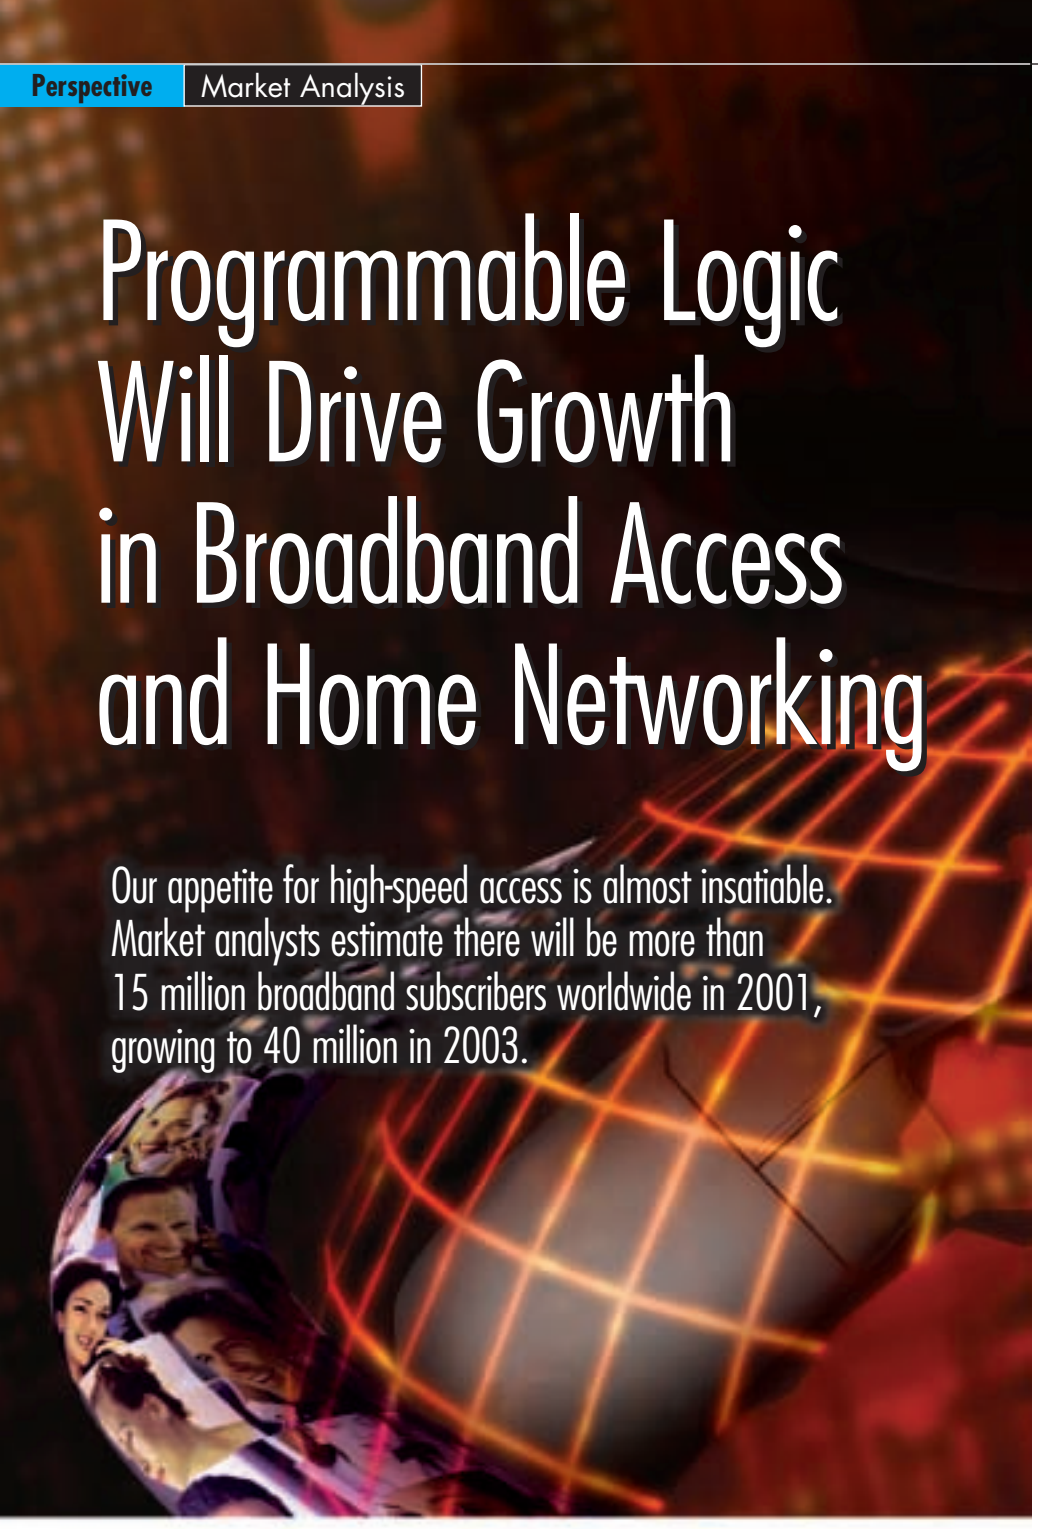

Our appetite for high-speed access is almost insatiable. Market analysts estimate there will be more than 15 million broadband subscribers worldwide in 2001, growing to 40 million in 2003.

by Robert Bielby
Senior Director, Strategic Solutions
robert.bielby@xilinx.com

The Internet has become an integral tool and resource for an ever-growing population of corporate users. As this population continues to expand to include remote corporate users – such as telecommuters, field engineers, and "road warriors" – their need for speeds equivalent to corporate LANs has increased, because access speed directly impacts worker productivity.

Likewise, reliance on the Internet and the need for greater bandwidth is driving consumers away from traditional analog-based modem technology in favor of higher speed broadband access technologies. Broadband not only delivers faster Internet access to the home, but it also serves as the catalyst for the expansion of the new home networking market.

We predict programmable logic will play an essential role in accelerating the growth and deployment of both broadband access and many of the new home networking applications and technologies.

Demand for Broadband

Broadband access describes a high-bandwidth digital communications link that provides access speeds typically greater than 128 Kbps. Broadband access not only provides significantly faster access times than older analog technologies (more than 100 times as fast), but it is also a conduit better suited for some of the newer, high-bandwidth Internet applications. These applications include streaming video, online shopping, MP3 music, VoIP (Voice over Internet Protocol), viewing of high-resolution images, and online gaming.

Broadband service is delivered to the home using a variety of different technologies, including cable, DSL, satellite, wireless, powerline, and ISDN. While each technology has its own set of benefits and drawbacks, in most cases, the existing home wiring infrastructure and the availability of the service itself usually plays the biggest role in determining the consumer's choice of broadband technology.

Traditionally, cost constraints have dictated the use of custom logic circuits, such as gate arrays, to implement additional logic circuits for products like digital modems, which have large annual unit volume shipments. While gate arrays may have proved to be cost effective in the past, the long development cycles can have a negative impact on a product manufacturer's time to revenue, and hence, success in the market.

In the past five years, advances in semiconductor technologies have dramatically affected the ability to create low-cost programmable logic families, such as Xilinx Spartan-II field programmable gate arrays (FPGAs). As costs are now beginning to reach parity with traditional gate arrays, low-cost FPGAs are rapidly becoming the technology of choice. They are used throughout a wide range of broadband access equipment because of their inherent flexibility and ability to accelerate the time to market delivery of new products. These benefits have not only accelerated the

growth and availability of broadband services, but the broadband services themselves are now spawning new markets, such as home networking.

Networking Comes Home

Although economics professors may cringe, supply can create its own demand. This phenomenon, called a "virtuous cycle," is perhaps best illustrated by the personal computer market, where faster microprocessors continue to drive more demanding software applications, which in turn drive the need for faster processors. In the case of home networking, more bandwidth, in the form of broadband access, is driving new applications such as MP3 audio players, online gaming, streaming video, and other amenities that in turn are driving the need for greater bandwidth that will eventually lead to more new applications.

This virtuous cycle, compounded by a voracious consumer desire to be entertained and to share information, leads market analysts to estimate the home networking market will swell to $6 billion dollars by the year 2004. First, however, several barriers must be overcome for this market to realize its full potential. While lack of consumer ease-of-use is one of the more problematic impediments to market growth, perhaps the most important issue today is the fragmentation of the home networking technologies themselves.

With over 22 disparate technologies – including phonelines, powerlines, IEEE 1394 (FireWire), Ethernet, USB 2.0, HomeRF, Bluetooth™, DECT, IEEE 802.11, and HiperLAN2 – it is unclear, which, if any technology, will be the winner. Each technology has its own pros and cons. Many of these technologies conflict with each other and are not interoperable. Furthermore, some of these technologies continue to be revised as problems are discovered after they are deployed.

Disparate Technologies, Programmable Solutions

With so much confusion and chaos surrounding home networking technologies, manufacturers are forced to overcome a difficult learning curve before developing products. Moreover, stiff competition and time-to-market pressures almost mandate a "ready, fire, aim" development model. Manufacturers attempt to design products for the longest life cycles, but they are up against the challenges of getting those products to market as fast as possible. Quite often, this dichotomy can lead to the development of products that become obsolete shortly after introduction.

Additionally, because of the sheer number of different home networking standards, leading manufacturers are driven to develop a large number of distinctly different products in order to garner an appreciable share of the market. Developing such a large number of products using traditional technologies, such as gate arrays, can prove to be both technically impractical and cost prohibitive. As we've seen in the digital modem markets, at this relatively early stage, highly integrated ASSP solutions are generally not available or feasible. The low unit volumes of these fragmented broadband technologies rarely justify the required engineering and financial investments.

And, as we have seen in the broadband access markets, the multiple standards and rapidly evolving features of home networking make programmable logic a natural fit for this new and changing marketplace. Home networking demands high performance, scalability, interoperability, upgradability, security, and a low cost of ownership. Low cost FPGAs, such as the Xilinx Spartan-II FPGA family, readily address these requirements and will serve to accelerate the growth and success of home networking.

Additionally, Xilinx is currently backing a new initiative that simplifies the task of designing products that can be remotely upgraded, modified, or fixed once they have been deployed in the field. Called Xilinx Online, this initiative uses Xilinx Internet Reconfigurable Logic (IRL™) to take advantage of the inherent reprogrammability of FPGAs to "future proof" products against premature obsolescence.

The ability to remotely upgrade a product will not only serve to accelerate the time-to-market of home networking applications, but it will also extend the useful life of existing systems – and significantly reduce production, maintenance, and support costs. Clearly, the benefits of the Xilinx Online initiative and IRL will prove to be a tremendous value to the new and emerging home networking applications.

Conclusion

It's clear the consumer desire for faster Internet access is the major catalyst fueling the demand for broadband access to the home. As broadband access technologies gain momentum and leaders emerge, they are already responsible for spawning a wide range of new home networking applications. We expect this trend will continue well into the future.

Where low cost programmable logic has played a key role in accelerating the growth and deployment of digital broadband access technologies, it will play an even greater role in bringing new home networking technologies to life.

For more information on the challenges and opportunities of home networking, visit *www.xilinx.com/esp/.*

FPGAs Enable Wireless LANs

Wireless local area networks (WLANs) provide mobility and portability with high-bandwidth data, voice, and video access – they are the ultimate solution for enterprise, SOHO, and home applications.

by Amit Dhir
Manager, Strategic Applications
amit.dhir@xilinx.com

Wireless local area networks (WLANs) are a rapidly emerging market. They combine data connectivity with user mobility and provide a connectivity alternative for a broad range of consumers and business customers. They have a strong popularity in vertical markets such as telecommuting, SOHOs (small offices, home offices), health care, retail, manufacturing, warehousing, and academia, where productivity gains are realized by using hand-held terminals and notebook PCs to transmit real-time information to centralized hosts for processing.

Spartan™-II FPGAs provide the flexibility needed in the WLAN market, allowing you to easily create products that can adapt to the evolving specifications and standards, and interface with other emerging technologies.

The WLAN Market

Business Research Group, a leading market researcher, predicts that revenues from WLAN products will exceed $2 billion in the year 2002, and show a steady growth of unit shipments and revenue. Demand for computing and telephony mobile devices will be one of the most influential market drivers, along with end users demanding higher data rates and ease of use to sustain growing Internet and data applications.

However, several issues remain unsolved for the industry. Although vendors have made great strides in achieving interoperability, a common wireless standard is far from reality (today, there are seven standards). Interference from competing 2.4 GHz technologies (like Bluetooth™ and HomeRF) threatens the already crowded band. In addition, uncertainties exist with several technologies migrating to the evolving 5 GHz frequency band.

WLAN Technology

WLANs focus on the PHY (physical) layer and the data-link layer – which includes the medium access control (MAC) and logical link control (LLC) sub-layers – of the seven-layer OSI network model. The physical layer defines the electrical, mechanical, and procedural specifications, which provide the transmission of bits over a communication medium or channel.

WLAN PHY layer technologies include narrowband radio, infrared, OFDM, and spread spectrum. The MAC sub-layer ensures error control and synchronization between the physically connected devices communicating over a channel. It is also responsible for determining priority and allocating access to the channel.

PHY Layer (Radio) Technologies

Infrared (IR) – Infrared technology uses very high frequencies, just below visible light in the electromagnetic spectrum to carry data. While popular with other wireless technologies, it is not gaining momentum with WLANs. It is good for inexpensive, directed, very limited range (up to 3 feet) systems, such as personal area networks.

Narrowband Radio – In narrowband radio, the user transmits and receives information on a specific frequency. This radio frequency (RF) bandwidth is kept as low as possible. Any undesirable crosstalk between communication channels is carefully avoided by coordinating the different users on different frequency channels. This technology requires the end user to obtain an FCC license for each site where this technology is to be deployed.

Spread Spectrum (SS) – SS was developed as a wideband RF technology for reliable, secure, and mission-critical military communications while trading off bandwidth. In this modulation technique, the radio transceiver spreads a signal's power over a wider band of frequencies. The spreading process makes the data signal much less susceptible to electrical noise than conventional radio modulation techniques.

The transmitted signal occupies a bandwidth considerably greater than the minimum necessary to send the information. Sacrificing processing gain (bandwidth) to gain signal-to-noise performance contradicts the desire to conserve frequency bandwidth. However, the tradeoff produces much louder and easier to detect signals, if the receiver knows the parameters of the broadcast signal.

Some of the SS modulation advantages include low power spectral density, interference limited operation, privacy, and random access possibilities. SS modulation techniques include:

- **Frequency Hopping Spread Spectrum (FHSS)** – A data signal is modulated with a narrowband carrier signal that hops from frequency to frequency as a function of time over a wide band of frequencies. It relies on frequency diversity to combat interference, which is accomplished by multiple frequencies, code selection, and FSK (Frequency Shift Keying). For example, a frequency hopping radio will hop the carrier frequency over the 2.4 GHz frequency band between 2.4 GHz and 2.483 GHz. If the radio encounters interference on one frequency, the radio will retransmit the signal on a subsequent hop on another frequency.
- **Direct Sequence Spread Spectrum** – DSSS is a more advanced, more recognized, and more used form of SS. The DSSS process is performed by effectively multiplying a RF carrier and a pseudo-noise (PN) digital signal. First, the PN code is modulated onto the information signal using one of several modulation techniques (such as BPSK and QPSK). Then, a doubly balanced mixer is used to multiply the RF carrier and PN modulated information signal. This process causes the RF signal to be replaced with a very wide bandwidth signal with the spectral equivalent of a noise signal.
- **Orthogonal Frequency Division Multiplexing (OFDM)** – OFDM technology is extremely effective in time-dispersive environments. The signals can take several paths to reach their destinations, thus resulting in variable time delays. At high data rates these time delays can reach a significant proportion of the transmitted symbol (a modulated waveform). This results in one symbol interfering with the next, thus making OFDM the only answer to "intersymbol interference" or ISI.

Direct Sequence (DS)	Frequency Hopping (FH)
Higher throughput	Interference immunity
Wider range	Echo resistant
Upgradeable to higher speeds at 2.4 GHz	Less expensive than DS systems
	Simpler installation
	More expensive product selection, more vendors

Figure 1 - Direct Sequence vs. Frequency Hopping

Data-Link Layer Technologies

LLC – The Logical Link Control sub-layer resides above the MAC sub-layer in the data link layer. It is responsible for the framing (or frame construction). The LLC inserts certain fields in the frame such as source and destination address at the head end of the frame and error handling bits at the end of the frame.

MAC – The MAC provides access control functions for shared medium PHYs in support of the LLC layer. The primary functions of the MAC layer include addressing (accessing the wireless medium), access coordination (joining the network), and frame check sequence generation and checking (providing authentication and privacy). The MAC layer uses access protocols such as carrier sense multiple access (CSMA) and time division multiple access (TDMA).

WLAN Types

The different types of wireless LAN standards are IEEE 802.11b, IEEE 802.11a, and HiperLAN2. IEEE 802.11b was among the first and currently most successful and popular wireless LAN standards to break onto the enterprise and home networking scenes. IEEE 802.11b is a 2.4 GHz standard using DSSS and is based on CSMA/CA. The future for wireless LANs is looking to migrate to 5 GHz OFDM-based PHY layers for both HiperLAN2 and IEEE 802.11a standards.

While having very similar PHYs and data rates (of 54 Mbps) the two will have dissimilar MAC layers, with HiperLAN2 supporting TDMA/TDD, and IEEE 802.11a supporting CSMA/CA. The growing popularity of wireless LANs is primarily because of the leap to high data rates and QoS (Quality of Service). Some industry pundits also believe that the wireless LAN market will pose a significant threat to the cellular (3G) space.

IEEE 802.11

IEEE 802.11 is the IEEE standard addressing the 2.4 GHz and 5 GHz WLAN market. The IEEE 802.11b extension employs a modulation scheme called complementary code keying (CCK) and operates in the 2.4 GHz ISM (industrial, scientific, medical) band. It is designed to enable data rates of 1 Mbps to 2 Mbps for FHSS networks and 1 Mbps to 11 Mbps for DSSS systems and provide interoperability between both DSSS and FHSS networks.

IEEE 802.11a, a recently formalized extension to 802.11, provides higher speeds and compatibility with existing standards. IEEE 802.11a employs the OFDM modulation scheme and uses the 5 GHz band. It provides a maximum optional speed of 40 Mbps and a range of 150 meters.

HiperLAN2

HiperLAN2 is an OFDM-based, variable bit rate PHY layer technology operating at 5 GHz. It has FEC error control, with dynamic sub-channel modulation allowing data transmission at higher rates with a strong SNR at lower throughputs in adverse conditions. HiperLAN2 provides high bandwidth up to 54 Mbps, with a range of over 150 meters. It has a generic architecture and supports Ethernet, IEEE 1394, ATM, PPP, and 3G.

The HiperLAN2 data-link layer/MAC provides QoS via dynamic fixed time slots. The time slotted structure allows simultaneous communication in both downlink and uplink in the same period. It is also a connection-oriented technology that allows negotiation of QoS parameters like bandwidth, bit error rate, latency, jitter, and delay requirements. This assures that other terminals will not interfere with subsequent transmissions. It provides ARQ (Automatic Repeat reQuest), dynamic frequency selection, power control and power save, cellular hand-over and, security (authentication and encryption).

WLAN Products Using Spartan-II FPGAs

WLAN products include network interface cards (or NICs/PC adapters), APs (end-user-to-LAN and LAN-to-LAN) and technology bridges for communications. NICs provide an interface between the end-user device (desktop PC, portable PC, or handheld computing device) and the airwaves via an antenna on the APs. APs act as transmitters/receivers between wired and wireless networks. They connect to the wired network via standard Ethernet cable (token ring is available, but less common) and use airwaves to transmit information to and from "connected" wireless end users.

Technology bridges exist at the periphery of each product and are the most susceptible to constant change and evolution. These products need a flexible, re-programmable and low-cost platform to accommodate for time-to-market pressures, specification changes, lack of clear direction and short product lifecycles.

The Spartan-II family is ideal for creating WLAN products. It provides increased densities (up to 200,000 system gates) and system-level features (such as DLLs, BlockRAM, and SelectI/O™) at a much lower cost. FPGA gates left over from programming the MAC may be used to customize the end products, and be used for additional functionality such as memory controllers, PCI controller, UARTs, and forward error correction (FEC).

NIC/PC Cards

Figure 3 shows the WLAN card, which consists of the antenna, radio/PHY, baseband controller and the MAC. The role of FPGAs is highlighted in red. The block diagram of the PC card shows that the Spartan-II FPGA provides functionality of a WLAN MAC (radio control, packet header generator, MAC protocol engine,

Characteristic	IEEE 802.11b	IEEE 802.11a	HiperLAN2
Spectrum	2.4 GHz	5 GHz	5 GHz
Maximum physical rate (approx.)	11 Mbps	40 Mbps	54 Mbps
Maximum data rate, layer 3 (approx.)	5 Mbps	28 Mbps	32 Mbps
Medium access control/Media sharing	CSMA/CA		Central resource control/TDMA/TDD
Connectivity	Connection-less	Connection-less	Connection-oriented
Multicast	Yes	Yes	Yes
QoS support	PCF	PCF	ATM/802.1p/RSVP/DiffServ (full control)
Frequency selection	DSSS	Single carrier	Single carrier with Dynamic Frequency Selection
Authentication	No	No	NAI/IEEE address/X.509
Encryption	40-bit RC4	40-bit RC4	DES, Triple-DES
Handover support	No	No	No
Fixed network support	Ethernet	Ethernet	Ethernet, IP, ATM, UMTS, FireWire, PPP
Management	802.11 MIB	802.11 MIB	HiperLAN/2 MIB
Radio link quality control	No	No	Link adaptation

Figure 2 - IEEE 802.11 vs. HiperLAN2

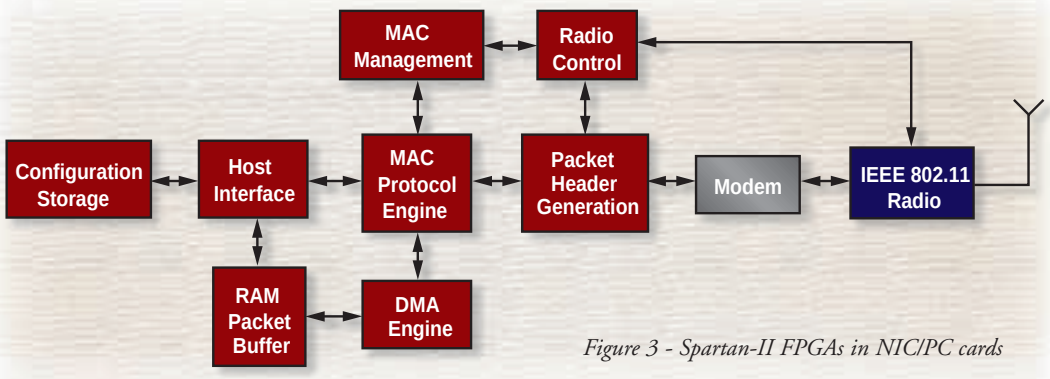

Figure 3 - Spartan-II FPGAs in NIC/PC cards

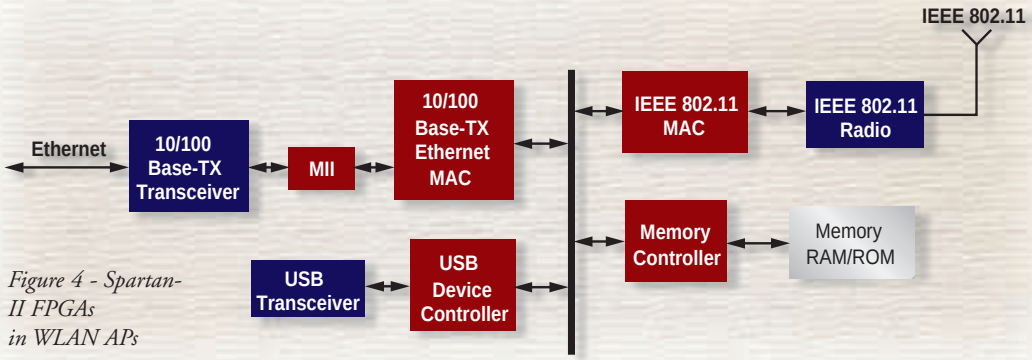

Figure 4 - Spartan-II FPGAs in WLAN APs

DMA engine, RAM packet buffer, host interface, configuration storage and MAC management), memory controller and as an interface to the PC.

Access Points

The APs are devices that provide a wireless hub or a gateway for non-wireless networks to wireless networks. They also act as the network police and perform network management. They receive, buffer, and transmit data between WLAN and the wired network infrastructure. APs function within a range of 100 to several hundred feet. They also connect WLANs to other technologies such as USB and Ethernet. In Figure 4, Spartan-II FPGAs provide solutions in APs such as memory controller, Ethernet MAC, USB device controller, and IEEE 802.11 MAC.

Technology Bridges

Conflicting specifications and lack of a clear direction create the need for FPGA-based technology bridges. It would be nearly impossible and cost-prohibitive for a supplier of home networking and WLAN products to meet all the various specifications and changing needs. At the same time betting on the success of one single product may preclude them from being successful in the marketplace.

Figure 5 shows an example of a WLAN to Ethernet technology bridge. In this example, Xilinx Spartan-II FPGAs are at the heart of the technology bridges, which usually connect unlike technologies – such as WLANs to Ethernet. While the Ethernet MAC has been around for a long time, the IEEE 802.11 specification that defines the MAC and PHY layers continue to evolve. With IEEE 802.11a and HiperLAN2 specifications still not defined, it seems ideal for the MAC and MII (media independent interface) to be programmed in an FPGA. Similarly, the HomePNA, HomeRF, FireWire™, USB, HiperLAN2, and Bluetooth are all technologies with evolving specifications.

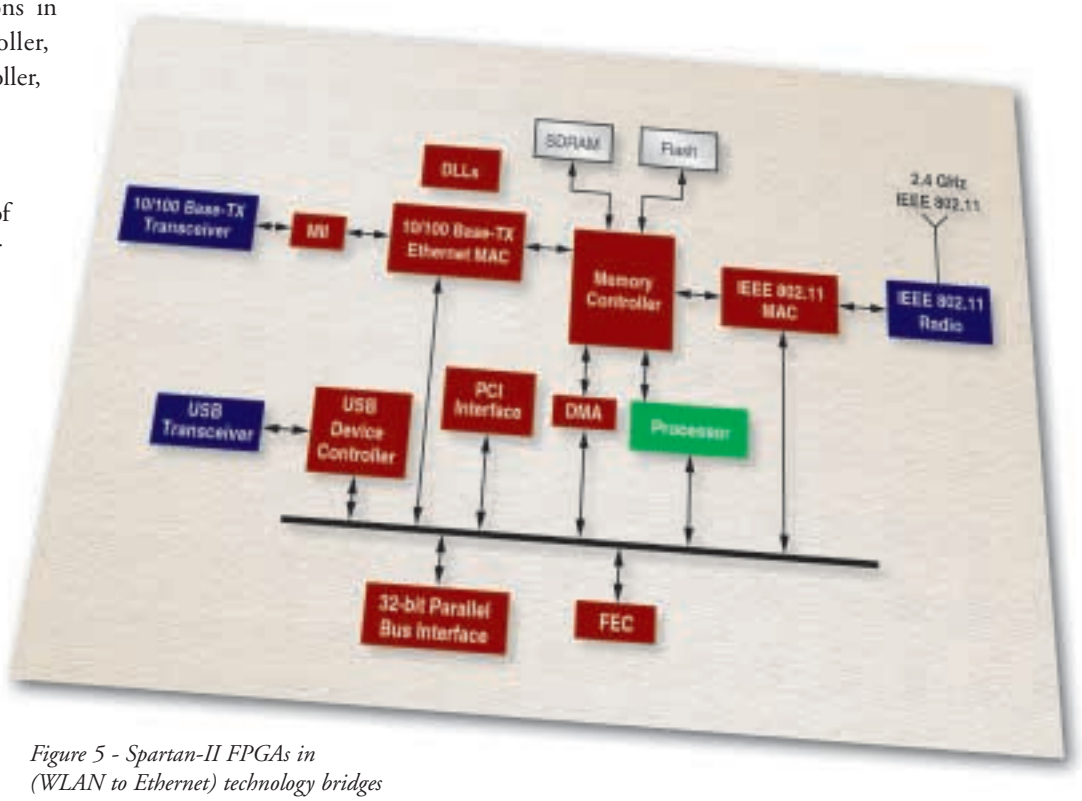

Figure 5 - Spartan-II FPGAs in (WLAN to Ethernet) technology bridges

WLAN products will extend beyond NICs, APs, and technology bridges, and will enable every device in the home, SOHO, and enterprise with WLAN capabilities. This includes such devices as digital TV, residential gateways, set-top boxes, digital modems, PC peripherals, gaming consoles, and other appliances.

Conclusion

The WLAN market is growing fast with a promise to penetrate homes, SOHOs, and enterprises in large volumes. Being a cost sensitive and evolving market, WLAN products require low-cost programmable logic solutions that allow customers to realize time-to-market and time-in-market advantages. Spartan-II FPGAs – with increased densities, system-level features, an extensive IP portfolio, and low costs – provide an ideal solution for WLAN products such as NICs, APs, technology bridges, and other products. Spartan-II FPGAs provide interoperability between different technologies, which is essential for success in this market.

Add Internet Connectivity with Spartan-II FPGAs and the UDP Stack Core

Insight's new VoIP Development Kit demonstrates the UDP stack core for Internet-based, point-to-point data applications.

by Jim Beneke
Director Technical Marketing
Insight Electronics
jim_beneke@ins.memec.com

In the past, if you needed to connect your application to a LAN or the Internet, you were forced to design-in additional circuits, which added complexity and increased cost. Often, these off-the-shelf networking solutions provided more functionality than you needed, and their cost was often a problem. The requirement for simple data transmission from one point to another via the Internet was further complicated by the need for processors and protocol stack software. Now however, the inherent advantages of programmable logic and the advanced architectural features of the Spartan™-II family have resulted in Internet connectivity inside a single chip.

The VoIP Development Kit

User Datagram Protocol (UDP)

Analysis of the OSI (Open Systems Interconnection) Seven Layer Model (application, presentation, session, transport, network, data link, and physical layers) resulted in the identification and implementation of an optimized set of protocols and processes. Nonessential functions were eliminated and exhaustive protocol options were streamlined. The resulting User Datagram Protocol (UDP) stack core was created, providing simple point-to-point communication and data streaming over a LAN or Internet connection. Figure 1 shows an example of the basic networking applications possible.

Building on the UDP stack core, Insight created a VoIP (Voice over Internet Protocol) Development Kit that demonstrates the UDP core in a simple voice over IP (Internet Protocol) application. Targeting the Spartan-II family, the VoIP application illustrates how basic networking capabilities can be cost effectively integrated into FPGA solutions. In addition to VoIP audio data, other data formats such as digital video or sensor data can also be implemented with minor modification.

VoIP Application Kit

The Insight VoIP Development Kit was created as an evaluation platform for the UDP stack core and the networking applications it targets. The kit provides two, identical Spartan-II based demonstration boards, so you can experiment with point-to-point communications over a LAN connection. And, the kit includes everything required to be up and running in minutes.

The VoIP demonstration board is based on the 150K gate Spartan-II FPGA. Figure 2 shows the block diagram of the board. On the application side, the FPGA interfaces to the Silicon Labs CODEC (COder/DECoder) device to receive and transmit digital voice data to and from the connected headset. On the network side, the FPGA connects to the LSI 10/100 Ethernet PHY. The other functions on the board include an XC18V01 ISP PROM to store FPGA configuration data, a JTAG port, a 50 MHz oscillator, an RS-232 port, two FPGA I/O expansion connectors, on-board power regulation for 2.5V and 3.3V, a serial EEPROM, a serial number ID chip, eight user-definable DIP switches, and a user LED.

The VoIP kit also includes a single project license for the UDP stack core. Although the UDP stack core comes as a netlist only, source code for the VoIP function is provided and allows for customization at the application level.

How It Works

The VoIP application uses the UDP/IP process to transmit voice data packets over the Internet. The CODEC device digitizes the analog voice signal from the headset, groups the data into packets, and then transmits it across the network. The packets are routed through the network as specified by the packet headers and undergo the reverse process at the receiving end. The UDP stack core plays a key role in this processing. Both the UDP core and the VoIP application are implemented in Xilinx Spartan-II FPGAs. The FPGAs provide the necessary functions in optimized hardware,

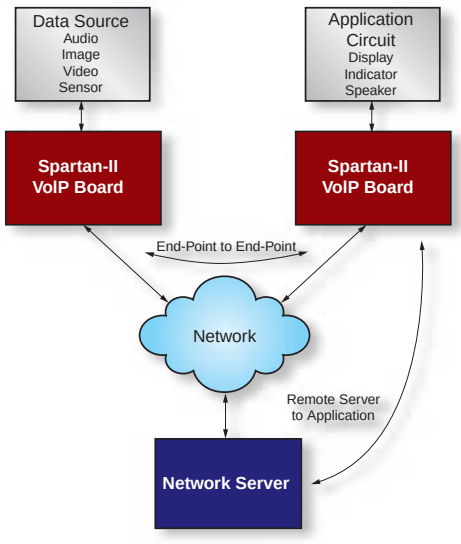

Figure 1 - UDP Stack Core Applications

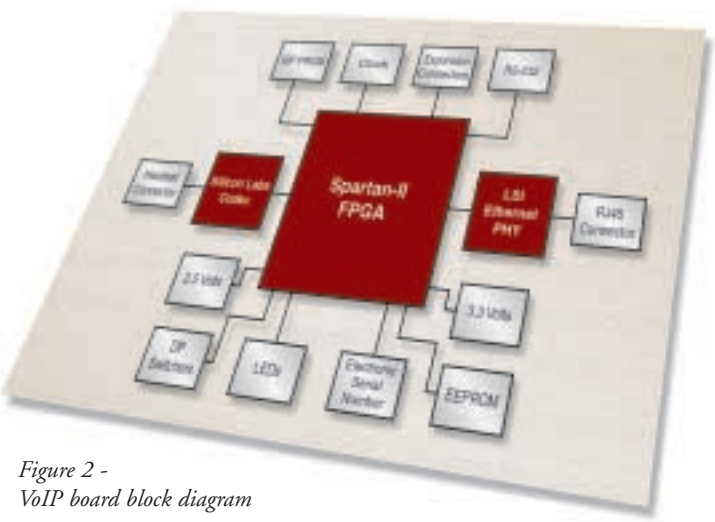

Figure 2 - VoIP board block diagram

making software unnecessary. The entire VoIP application and UDP stack consume less than 1,230 slices and six block RAMs.

Figure 3 shows the protocol layers implemented in the VoIP example. The application layer includes a serial interface to the CODEC chip, along with independent transmit and receive FIFOs for buffering data between the application layer and transport layer. The transport layer uses the Real-time Transport Protocol (RTP) to assemble/disassemble the audio payload and interface to the UDP.

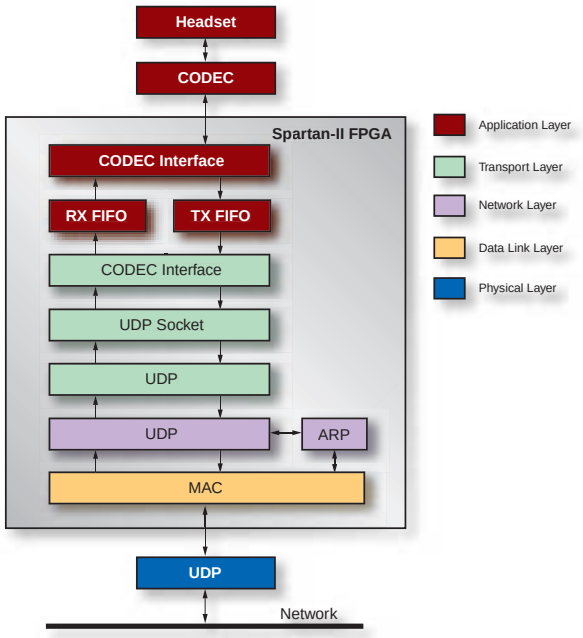

Figure 3 - VoIP protocol layers

The UDP Stack Core

The Xilinx XCoNet team developed the UDP stack core with MemecCore, the IP division of Insight Electronics. Unlike the more familiar TCP/IP stack, the UDP/IP stack is streamlined for simple point-to-point network connections and data streaming. As Figure 3 shows, the UDP stack spans the transport, network, and data link layers and includes the UDP, IP, and MAC (Media Access Control) functions.

The MAC implements the RFC894 Ethernet standard and manages the interface between the LSI physical interface chip and the UDP stack. Following the MAC is the IP function, which validates datagrams on the receive side and combines IP headers with UDP packets on the transmit side. If received data packets do not have valid source or destination IP addresses, the IP simply drops the packets. The included Address Resolution Protocol (ARP) translates the IP addresses to hardware addresses and responds to remote node ARP requests.

The UDP makes up the transport layer and has a receive section and a transmit section. On the receive side, the UDP checks for a valid destination port number and checksum, discarding packets that do not pass the criteria. For transmission, the UDP prepares the UDP header and combines it with the payload, writing the UDP packet to a buffer to wait for transmission.

The UDP stack core requires 1,000 slices, six block RAMs, and 166 I/Os. The majority of the I/Os reside on the design backend and are eliminated once the top-level application is integrated. In the case of the VoIP application, the total I/Os were reduced from 166 to just 28.

Conclusion

The VoIP Development Kit and the UDP stack core demonstrate cost effective networking solutions in a single chip FPGA. Both the UDP stack core and the VoIP reference kit are available from Insight Electronics. The VoIP reference kit is priced at $995 and includes a single project license for the UDP netlist core; the UDP stack core can be purchased separately for only $695. Both are available now.

Some applications that utilize the UDP stack may require UDP core customization. Insight Design Services can provide protocol enhancements or core modifications as needed. Call 888-488-4133 ext. 227 or go to *www.insight-electronics.com/spartan_IIVoIP* for more information or to order the VoIP kit.

Residential Gateways

A single device connects multiple broadband access and home networking technologies.

by Amit Dhir
System Architect, Strategic Applications
amit.dhir@xilinx.com

The Residential Gateway (RG) is a platform for the deployment of high-speed Internet access and integrated voice, data and video services over the same high-speed pipe to different nodes (appliances) throughout the home. RGs combine the functions of a digital modem, SOHO router, or hub for Internet access to multiple information appliances.

There are different types of RGs available based on the functionality required. Devices such as set-top boxes and digital modems provide broadband access to homes and SOHOs (small office-home office). These devices will incorporate home networking functionality and evolve into gateways. In addition to these devices, a number of dedicated RGs such as home servers and SOHO routers are emerging.

Market Trends

The home networking market is fragmented and includes the following four aspects:

- Broadband access
- Residential gateways
- Home networking technologies: no new wires (phone lines, powerlines), new wires (IEEE1394, Ethernet, USB 2.0), and wireless (HomeRF, Bluetooth, DECT, IEEE802.11, HiperLAN2)
- Information/Internet appliances.

Rapid implementation of integrated value-added services with home networking is creating an explosive market for RGs. Cahners In-Stat Group predicts the RG market will rise sharply from $100 million in 2000 to $5 billion in 2005.

RGs will evolve from providing basic broadband access to providing integrated services gateways such as remote management, home automation, home security, and video-on-demand. RGs will also network multiple home networking technologies and provide protocol translation between disparate technologies.

Consortiums like the OSGi (Open Services Gateway Initiative) are working to define and promote an open RG standard for connecting the coming generation of smart consumer and small business appliances with commercial Internet services. Hurdles inhibiting the mass deployment of RGs in households include the lack of:

- Clear business models (costs of ownership)
- Customer education and mass confusion
- Support mechanisms.

Widespread deployment of gateways into homes will come in three distinct phases. While the gateway is a new term, the first phase already exists in many homes. A good number of our homes have a set-top box for receiving television broadcasts and a cable modem to connect to the Internet. Phase two will include advanced features such as broadband connectivity, home networking interfaces, and IP telephony in the RG. The third deployment phase will be powerful RGs, capable of delivering video, voice, and data throughout the home. It will also provide other services such as home automation, energy management, security control, and so on.

The RG hardware architecture will be modular, which will allow support of multiple broadband and home networking technologies. Supporting multiple technologies makes the gateway less likely to become obsolete with technology advancements. Support for modularity will fuel the evolution of RGs into a type of application server that consumers will use to distribute broadband services throughout their homes. The gateway must have a reliable and robust hardware platform, and software that is not susceptible to errors. Unlike PC users, consumers will not tolerate having to reboot their gateways. Supporting multiple services such as voice, data, and video with complete security is essential. Functions such as e-commerce transactions, remote home control, and access to authorized service providers are critical. Providing quality of service to support multiple intelligent devices from different vendors is extremely important.

RG Components

The gateway provides a unified platform to satisfy the needs of most consumers, providing infotainment and communication. Because the underlying technologies are new and evolving, the RG will evolve in its

functionality. In its basic form it will combine digital modems, home networking chipsets, and software. Additionally, all gateways will contain computing resources to support the software required to operate the device.

Various types of devices are competing for a share of the RG marketplace.

PC-Based RG (Home Servers)

The PC-based RG could be the easiest path to the RG because of PC penetration in today's affluent homes. Also, it has the processing power, includes a digital modem, and is an ideal platform for the deployment of home networking services. However, the PC has certain disadvantages. The operating systems are traditionally not robust, and PC rebooting is not acceptable in a gateway, which provides critical features such as security. Some companies have pursued PC-based architectures even within a separate RG.

Modem-Centric RGs

Traditional methods to connect to the Web have hit a technological upside. Analog modems using phone lines cannot provide bandwidth beyond 56 Kbps. Digital modems, such as cable, DSL, and satellite, provide high-speed Internet access. Satellite modems offer speeds of 400 Kbps to 38 Mbps, DSL modems offer 1.5 Mbps to 52 Mbps and cable modems process multimedia content up to 30 Mbps. Usually, the RG is sold in conjunction with value-added services. Service providers subsidize the cost of the gateway by charging for services over one to two years. Typically, the gateway is easy to install and the service provider saves money by minimizing truck rolls. The following types of broadband modems will evolve into gateways:

Cable Modems

Cable modems are client devices that provide high-speed, two-way communication (data, voice, and video) over cable TV networks to access this mix of applications.

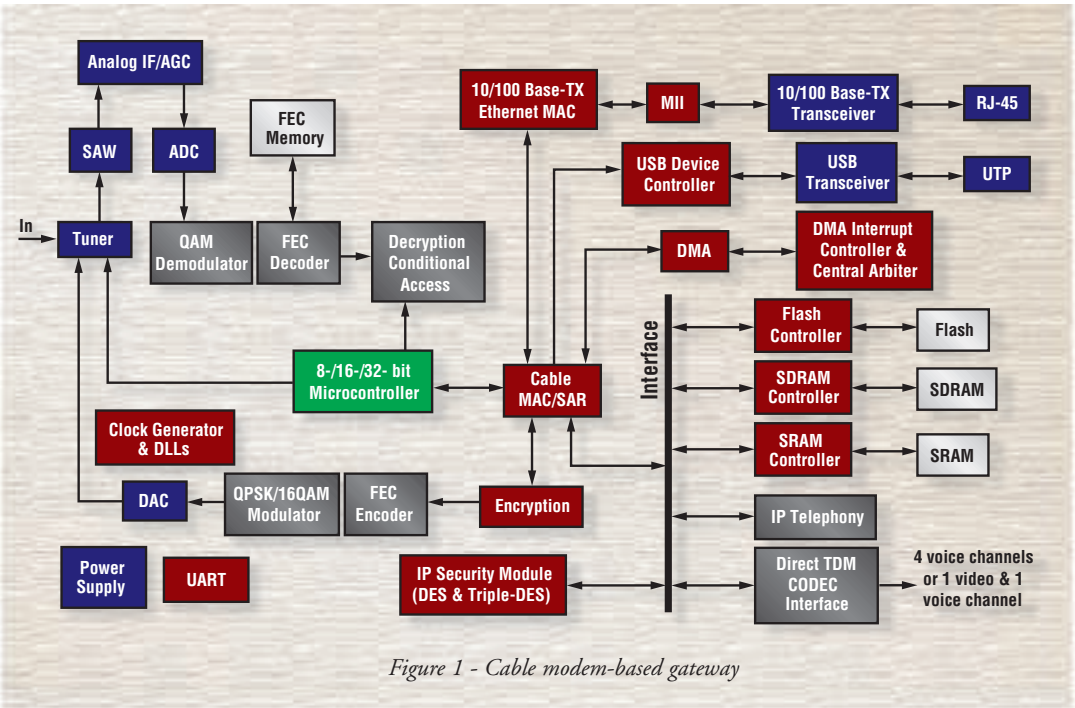

Figure 1 - Cable modem-based gateway

Cable modems provide an always-on connection, thus eliminating the slow and frustrating procedure of establishing a dial-up connection.

A standard cable modem has two connections; one port connected to the TV outlet on the wall and the other to the home network. The cable modem communicates over the cable network to a device called a CMTS (Cable Modem Termination System), which is a central device for connecting the cable TV network to a data network (Internet).

Theoretically, cable modems are capable of receiving and processing multimedia content at 30 Mbps. However, subscribers can expect to download information at 1.5 Mbps, because multiple users in the neighborhood share the bandwidth. Operators define a portion of the frequency spectrum to carry the data. The downstream path (head-end to home network) lies between 50 MHz and 750 MHz. The 5 MHz to 42 MHz frequency range transmits information from the home network to the head-end. Cable modems receive and transmit TCP/IP traffic, including multimedia content and Internet access.

The tuner in the cable modem receives a digital signal from the network and isolates a particular channel that contains the Internet data. It then converts the signal from RF levels back to baseband. The baseband output signal from the tuner is then forwarded to a demodulator. The demodulator samples the signal and converts it to a digital bit-stream. The bit-stream contains video, audio, and IP data. Once the bit-stream is recovered it is forwarded to a forward error correction (FEC) unit and checked for problems. The signal is then passed into the control unit and further to the home network using one of the high-speed data port interfaces.

In a cable modem-based gateway, as shown in Figure 1, programmable logic solutions provide system interfacing (PCI, UART, DMA), encryption, interfacing to several home networking technologies (Ethernet, USB 2.0, HomePNA), ASSP functionality (such as cable MAC), and glue logic.

Satellite Modems

Similar to the cable operators, satellite service providers deliver broadband connections to their subscribers. Consumers who connect their home networks to a satellite broadband platform are accessing a shared medium that is broadcasting at rates up to 38 Mbps. The transmission speeds from the service provider to the home is high, but the upload speeds are limited to a dial-up, telephone modem transmission rates.

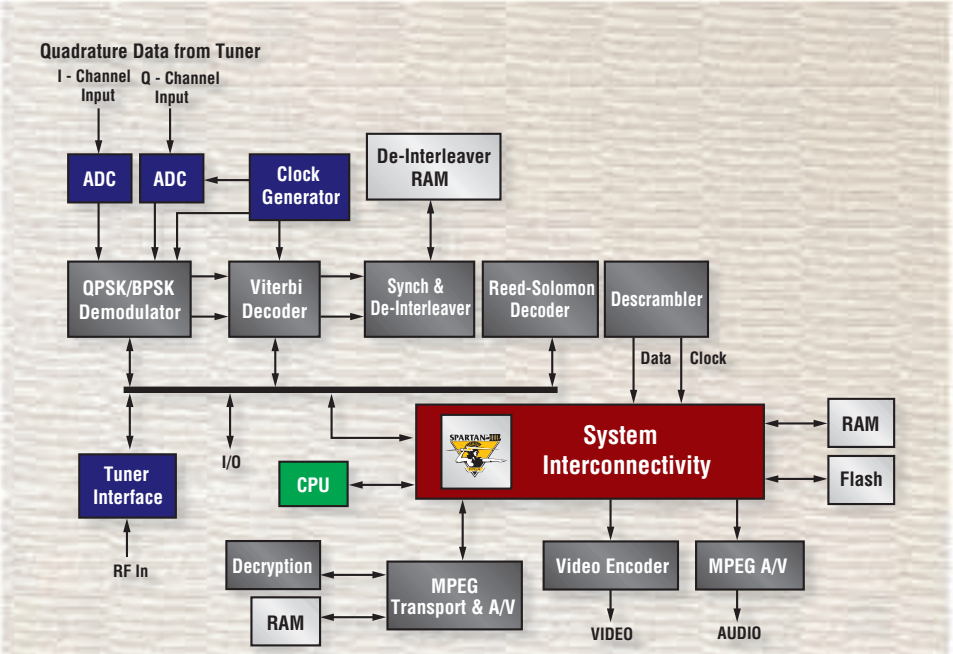

Figure 2 - Satellite modem-based gateway

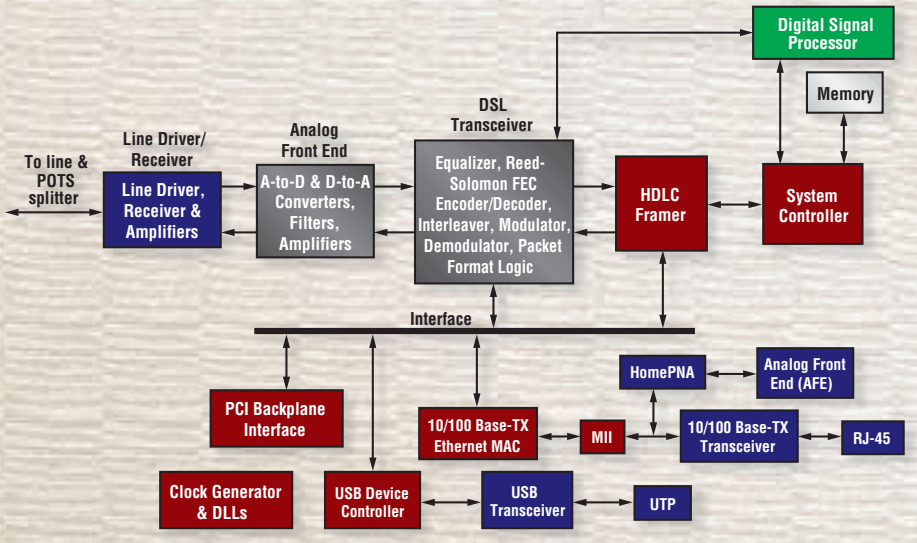

Figure 3 - DSL modem-based home gateway

This high-speed connection is facilitated through a dish and satellite modem.

In a satellite modem, the interface receives digital signals from the satellite network and isolates channels containing Internet data. The host interface provides an interface between the satellite modem and home network, using USB and PCI host interfaces. System glue interconnects the CPU and memory components with the host and satellite interfaces. Satellite modem-based gateways, shown in Figure 2, include satellite modem ASSPs and home networking ASSPs. These ASSPs provide high-speed (data, voice, and video) access (using the satellite broadband network) and network these information appliances. Programmable logic provides system interconnectivity solutions to network appliances using disparate technologies such as phone lines or Ethernet connections.

DSL Modems

Broadband modems based on DSL technology support data transmission over standard telephone lines up to 50 times faster than the analog modems. An Asymmetric Digital Subscriber Line (ADSL) circuit consists of two ADSL modems connected by a copper twisted-pair telephone line.

To maintain backward compatibility with the standard telephone system and to avoid disruption of service due to equipment failure, the voice part of the frequency spectrum is separated from the digital modem circuitry by means of a passive filter called a "POTS splitter." Hence, if the ADSL service fails, the POTS service is still available. Under this configuration, voice calls and Internet data are transmitted simultaneously over the same broadband DSL pipe. When an ADSL transmission is received at the central office, a more advanced POTS splitter is used to send the voice traffic to the public telephone network and data to the Internet.

DSL gateways provide DSL functionality and network information appliances. Programmable solutions are ideal for interfacing multiple home networking technologies (such as HomePNA, Ethernet, and USB 2.0) and system interfaces such as PCIs.

Digital Set-Top Boxes (STBs)

In the early 1970s, the only piece of equipment that people needed to watch TV was a standard television that they were able to purchase at a local store. In the 1980s, this simple model began to change. Cable and satellite providers required the consumer to connect their TVs to their networks. Also, operators decided to scramble TV signals, requiring a special box to de-scramble the signals at the consumers home.

Today digital television requires an STB to receive and decode digital transmissions into a form suitable for display on analog television sets. They are installed and configured by the local cable, terrestrial, or satellite service provider.

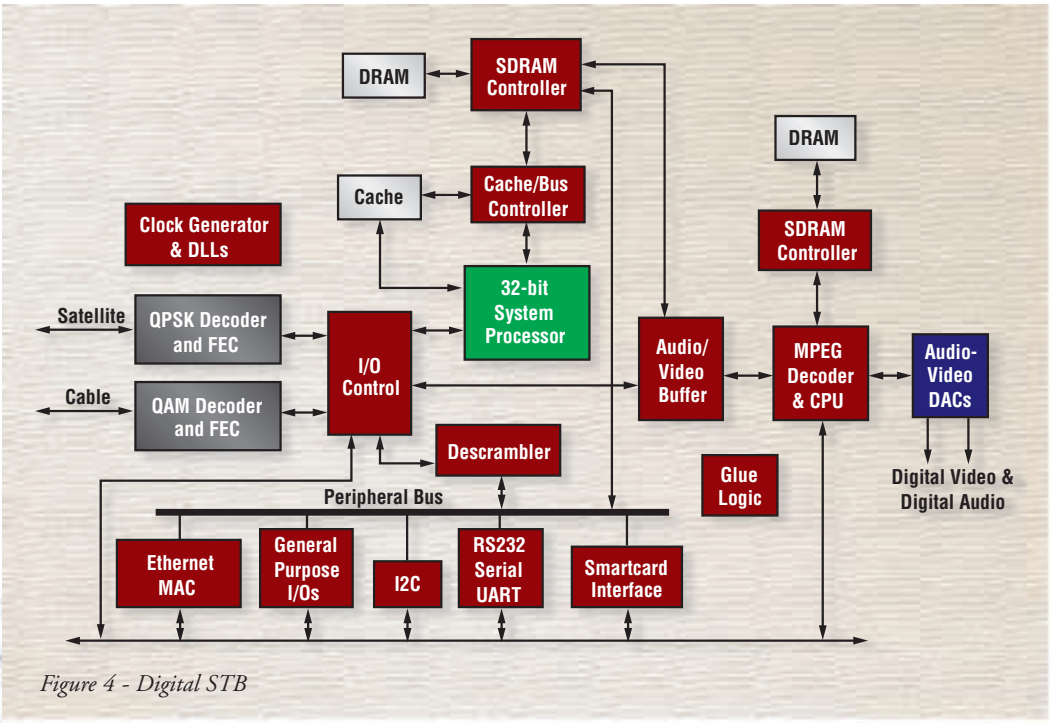

Figure 4 - Digital STB

STBs usually have cable or satellite modem chips to bring broadband and TV signals to the home. Other ASSPs/ASICs handle and process digital video and audio services. These application-specific devices communicate with each other via buses on the system board. The CPU is responsible for coordinating the different components. With additional features, STBs will require higher performance CPUs to keep pace with increased data throughput.

Programmable logic devices address the fundamental disconnect between application-specific devices. Time to market and the ability to upgrade quickly after the sale is imperative for a successful product. Programmable devices provide the cutting edge advantage over ASSPs by bringing products first to market – and having the ability to add features remotely to STBs already deployed in homes.

Need for Programmable Solutions

For RGs to be successful, programmable solutions will have to be at the heart of the system. While programmable logic solutions can perform the functions of cable MAC SAR, DSL receiver chipset, and satellite modem chip, their advantage is in interfacing the different broadband and disparate home networking technologies. They provide an ideal interface to access home networking technologies for protocol translation.

Programmable solutions provide the ability to interface multiple hard-disk drives and proprietary interfaces, and provide encryption capabilities using DES, triple DES, and even proprietary encryption schemes. They also provide system interface functions such as PCI, USB, and so on.

Conclusion

The primary function of the RG is to provide broadband connectivity to the home through cable, xDSL, satellite, and wireless technologies. Secondly, RGs provide home networking capabilities and distribute broadband information through appliances using technologies such as phone lines or wireless LANs. They also provide a unique platform for the deployment of value-added services and aftermarket upgrades. Programmable logic is necessary for the success of RGs, providing time-to-market and time-in-market advantages of interfacing disparate technologies and system interfaces.

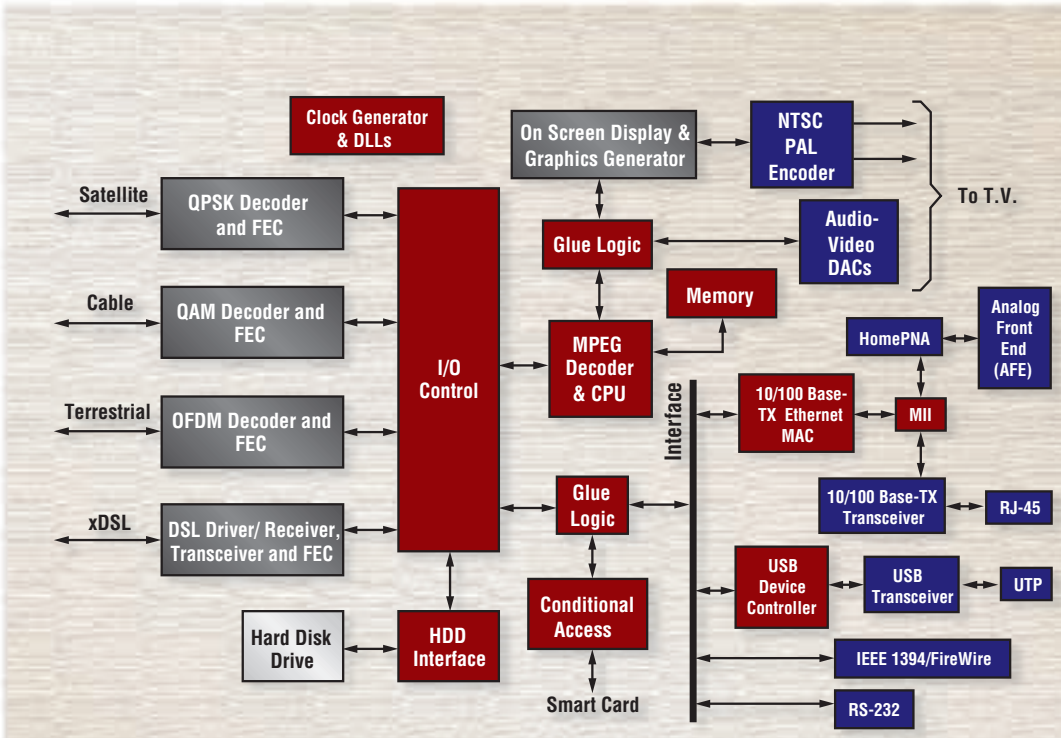

Figure 5 - Residential Gateway

Xilinx Publishes Free Comprehensive Guide to Home Networking

If you're designing and developing products for the home networking market, this book is a must-read.

WITHOUT A DOUBT, THIS IS THE MOST THOROUGH PIECE WRITTEN ON THE BURGEONING HOME NETWORKING INDUSTRY. THE CONTENT AND FORMAT IS SUITABLE FOR A DIVERSE AUDIENCE INCLUDING ENGINEERS, INVESTORS, AND HIGH-TECH AFICIONADOS. IT IS A MUST-READ FOR ANY CONSTITUENT OF THE HOME NETWORKING INDUSTRY.

– TIM MAHON,
DIRECTOR, CREDIT SUISSE FIRST BOSTON

by Amy Hills
Manager, Vertical Marketing
amy.hills@xilinx.com

In our continuing commitment to help our customers accelerate the design and development of home networking products, Xilinx has published a comprehensive guide called *The Home Networking Revolution, A Designer's Guide*. The 259-page book explores digital convergence and home networking dynamics, and it provides a thorough overview of the key technologies and components.

Written by Amit Dhir, manager of Strategic Solutions Marketing, and edited by Tom Pyles, editor of Corporate Technical Publications, the guide covers topics ranging from broadband access technologies to residential gateways to home networking options. Through tutorials on current and emerging standards and protocols (eSP), system block diagrams, and reference designs, the guide will help you perfect the design and development of consumer products. You will also find information on how programmable logic solutions can accelerate your time to market of home networking products.

You can order your free copy of *The Home Networking Revolution, A Designer's Guide* by registering on the Xilinx eSP Web portal at *www.xilinx.com/esp*.

Through *The Home Networking Revolution, A Designer's Guide*, in conjunction with the eSP Web portal, you will find that Xilinx provides you with the resources you need to:

- Build best-in-class products
- Get your product to market far ahead of the competition
- Extend the life of your product with Internet Reconfigurable Logic.

iMPACT

New Configuration and Programming Software

The new iMPACT software helps you configure and program all Xilinx FPGAs, CPLDs, and ISP PROMs.

by Frank Toth
Marketing Manager, Configuration Solutions
frank.toth@xilinx.com

The new iMPACT software features a series of "wizard" dialogs that guide even the most novice user through the every step of the configuration process. This software is fully integrated and easy to use, combining three previously separate tools: JTAG Programmer, Hardware Debugger, and the CoolRunner™ PC-ISP programmer. With iMPACT:

- You can quickly shift among the various programming modes.
- You receive instant visual feedback on all operations such as Erase, Program, Verify Blank Check, and various USER and ID codes.
- You can watch the chains being built graphically, so you can see how the cables should be connected for the supported modes.
- You can use a wide range of input and output file types including JED, BIT, BSDL, MCS, and SVF.

Programming Modes

The iMPACT configuration software enables you to easily configure all Xilinx FPGAs, using three different modes:

- Slave Serial (Single Bit Serial) – A two-wire implementation with performance up to 66 Mbits.
- SelectMAP (8-bit parallel) – The fastest mode (500 Mbps in system design performance), requiring eight data lines and several control lines.
- Boundary Scan Mode (commonly known as the JTAG mode) -The traditional four wire interface used for point-to-point testing and diagnostics by various standard third party tools. A special function in the JTAG mode allows you to test both the operation of the cable and the robustness of the JTAG chain. You can test chain operation by instructing iMPACT to write to and read back from the IDCODE multiple times. It then counts the number of errors that occur in this operation. Thus you can evaluate the relative robustness of the JTAG chain and the susceptibility to noise and other influences such as board layout.

Download Cables

The iMPACT software also supports two different cables:

- Parallel Cable III (Parallel Port) – Primarily used as a low cost solution for interface to a standard parallel printer cable on any PC.
- MultiLINX (USB and RS-232 ports) – Can be used with a USB port for the ultimate in download performance when used in the Slave Serial or SelectMAP mode. It also features a standard RS-232 port for users requiring this traditional interface for workstation support.

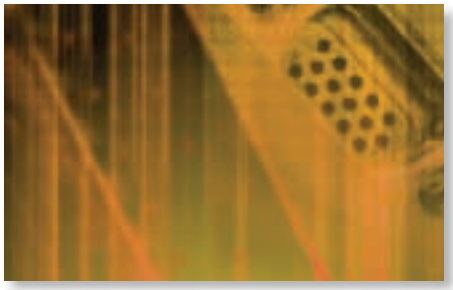

Conclusion

The new iMPACT software greatly simplifies the configuration and programming process. The sophisticated graphical interface and wizard dialogs guide you through the configuration process every step of the way, while providing instant feedback on various configuration modes and help with exact cable hookups.

For more information on iMPACT, visit *www.xilinx.com*

Platform FPGA SystemIO Solution

Enabling the Emerging High-speed Connectivity Standards

Higher I/O performance to meet the growing demand for communications bandwidth.

by Peggy Abusaidi
Product Marketing Manager
peggy.abusaidi@xilinx.com

Traditional system interfaces, such as the existing PCI and VME parallel bus schemes, cannot keep up with today's bandwidth requirements. Therefore many new I/O and communications standards, including RapidIO, HyperTransport™, InfiniBand™, 3GIO and others, have been developed to solve the I/O bottleneck challenges. Figure 1 illustrates the variety of applications that drive the need for more bandwidth.

How do you choose the right I/O interface standards for your systems, and how do you keep current with the evolving I/O standards? In addition, how do you meet your time-to-market and cost goals with minimal risk, while the interface standards rapidly evolve? The Xilinx® Platform FPGA SystemIO solution solves the I/O bottleneck problem by giving you the ability to implement designs using the latest I/O standards. You can accelerate your time to market and be assured that your designs will remain current as the I/O standards evolve. The Xilinx Platform FPGA SystemIO solution fully addresses all aspects of system connectivity in high-performance designs.

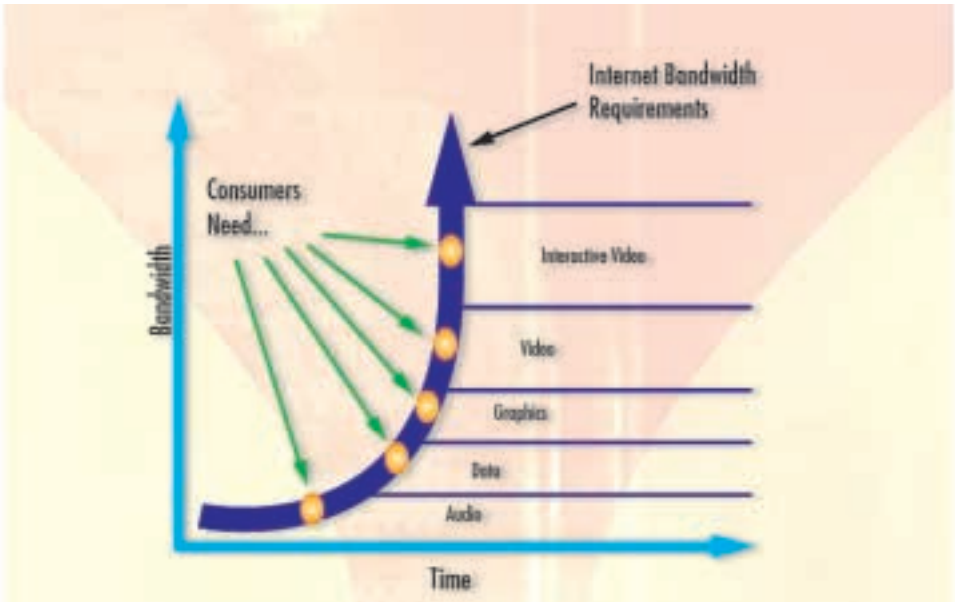

Figure 1- Internet bandwidth trend

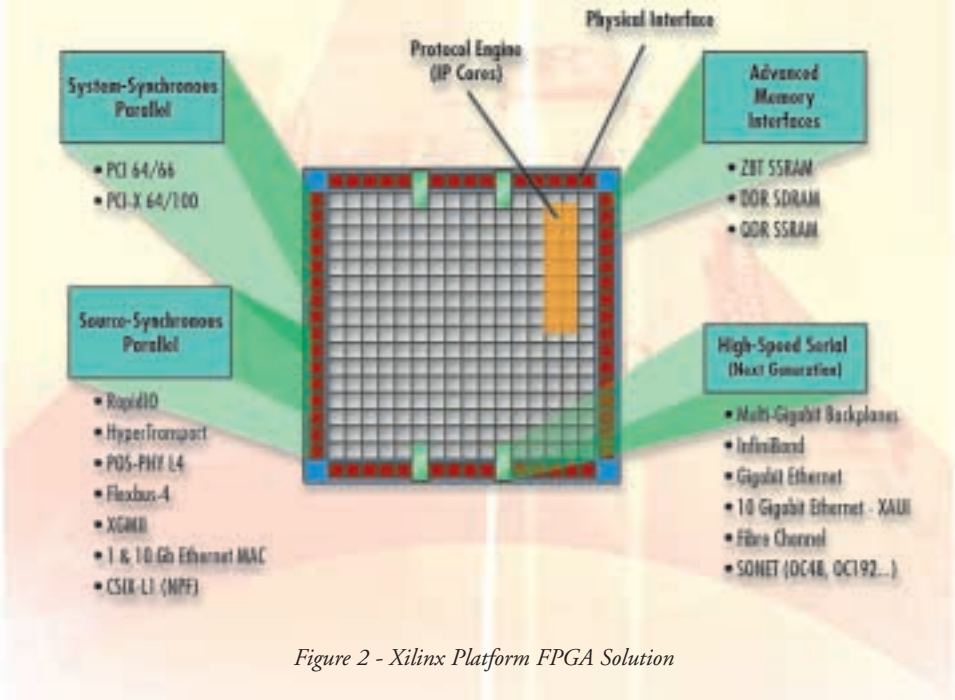

Figure 2 - Xilinx Platform FPGA Solution

The Bandwidth Problem

Today, most computer, embedded processing, and telecommunications equipment is parallel-bus oriented. However, as device performance increases, and demand for bandwidth rises, these multi-drop bus structures are reaching their performance limitations. Plus, in most cases, only a single module can use these buses at a time, with the obvious result that other subsystems on the bus are temporarily idle, waiting for bus access.

The industry's response has been the development of a host of new interconnection schemes; some with data transfer rates exceeding 10 Gbps. The new proposed standards have backers like Intel, AMD, Microsoft, Compaq, Dell, HP, Xilinx, and Sun Microsystems. They have names like RapidIO, XAUI, HyperTransport, and InfiniBand. Other standards already in use, or evolving toward higher levels of performance, include POS-PHY4, Ethernet, and various optical ATM formats like OC-12, OC-48, and OC-192. These new or evolving standards cover all architectural environments such as motherboards in chip to chip communication, backplanes for communications between subsystems, and Storage-, Local-, and Wide-Area Networks (SANs, LANs, and WANs).

The parallel standards are divided into two general categories:

- System-Synchronous Parallel – the venerable PCI family of buses, including PCI-X and Compact PCI.
- Source-Synchronous Parallel – RapidIO, HyperTransport, SPI-4 / Flexbus-4, POS-PHY 3, and others.

Because these standards are new and subject to change, FPGAs are the ideal way to ship a product without fear of obsolescence. Plus, the ability to use IP cores that implement these complex and timing-critical busses gives you a fast, risk-free method to address a rapidly changing market.

Clearly, the challenges you face are greater than ever. Not only are product life cycles shorter, simple evolutionary product steps are no longer sufficient with the multitude of standards hitting the market. When you consider that most of these standards are not final, the challenge to get your product right on the first iteration is daunting. That's why the ability to rapidly change your design to meet changes in the standards and markets is incredibly valuable. Further, the ability to leverage your design efforts with pre-engineered, supported IP Cores allows you to create more functionality in less time that ever before.

The SystemIO Solution

The Xilinx Platform FPGA SystemIO solution uses the unique Virtex-II SelectI/O™ Ultra to provide the fastest and most flexible electrical interfaces available. Each user I/O pin is individually programmable for 19

single-ended I/O standards or six differential I/O standards, including LVDS, SSTL, HSTL, and GTL+. The SystemIO solution is capable of delivering 840 Mbps LVDS performance using dedicated Double Data Rate (DDR) registers. Furthermore, any two I/O pins can be used as a differential pair, providing maximum board layout flexibility.

Using Virtex-II 840 Mbps LVDS performance and an abundance of memory and logic resources, you can easily create designs using 1GE and 10GE MAC, PCI and PCI-X, POS-PHY Level 3 and 4, RapidIO, HyperTransport, and Flexbus-4. Free reference designs for implementing interfaces such as LVDS and CSIX are available from the Xilinx website.

Table 1 shows the SystemI/O summary.

With the variety of interface reference designs and cores available in the SystemIO solution, you can easily customize your application. Figure 3 shows a 10 Gigabit Ethernet LAN/WAN line card example in which several Platform FPGA SystemIO solutions are used to provide seamless interfaces to external PHYs and Network processors. All of these interfaces are pre-engineered by Xilinx for easy drop-in functionality, enabling you to reduce your design cycle time.

Core	Standard Compliance	Aggregate Bandwidth	I/O Bus	Availability
POS-PHY L3	OIF-SPI3-01.0 Saturn POS-PHY L3	2.48 Mbps 104 MHz	32b	Now
POS-PHY L4	OIF-SPI4-02.0 Saturn POS-PHY L4	11.2 Gbps	16b LVDS 700 Mbps per LVDS pair 350 MHz DDR	Now
Flexbus-4	OIF-SPI4-01.0 AMCC Flexbus-4	12.8 Gbps	64b HSTL 200MHz	Now
10 Gb Ethernet MAC w/ XGMII	IEEE P802.3ae draft D3.1	10 Gbps	32b XGMII HSTL 312.5 MHz DDR	Now
XGMII Reference Design	IEEE P802.3ae draft D3.1	10 Gbps	32b XGMII HSTL 312.5 MHz DDR	Now
RapidIO PHY	RapidIO Interconnect Specification v1.1	8 Gbps	8b LVDS 500 Mbps per LVDS pair 250 MHz DDR	Now
PCI 64/66	PCI Spec V2.2	528 Mbps	64b 3.3V PCI	Now
PCI-X 64/100	PCI-X Spec V1.1	800 Mbps	64b 3.3V PCI-X	Now
CSIX Reference Design	CSIX-L1	6.4 Gbps	32b HSTL 200 MHz	Now
HyperTransport Single-Ended Slave	HyperTransport V1.01a	3.2 Gbps	8b LVDS 400 Mbps per LVDS pair 200 MHz DDR	Q1 '02
1 Gb Ethernet MAC	IEEE 802.3z	1 Gbps	GMII	Now

Table 1 - Platform FPGA SystemIO solution summary

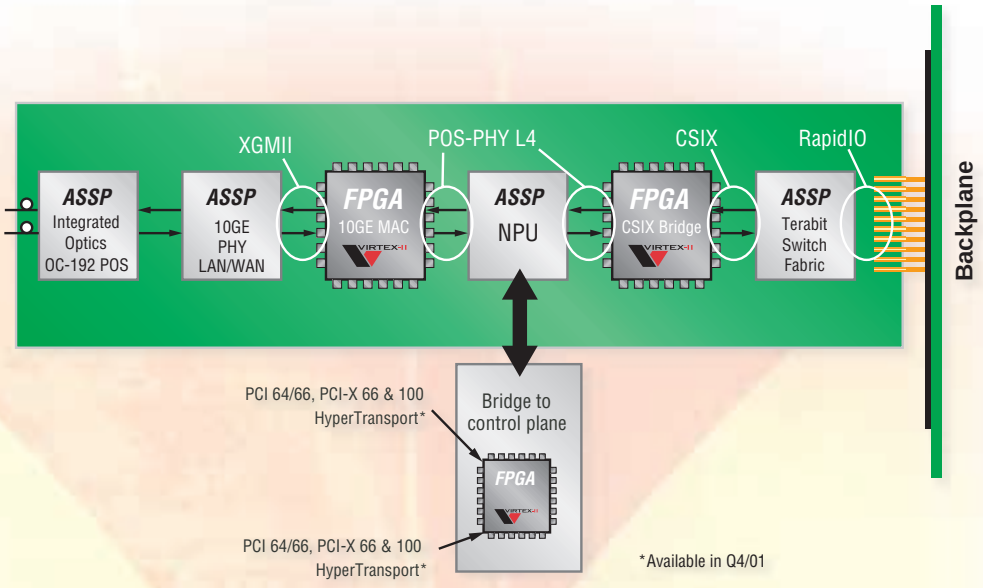

Figure 3 - 10-Gigabit Ethernet LAN/WAN card application example

In the future, we will incorporate Mindspeed's Skyrail™ 3.125 Gbps IO technology into our next generation Virtex-II family to provide support for Gbps interfaces such as 10GE, 3rd Generation IO (3GIO), Infiniband, XAUI and Fibre Channel.

Conclusion

The flexibility and high performance of the Virtex-II SelectI/O Ultra technology combined with the proven pre-engineered LogiCORE™ cores provide the complete connectivity solution to address the high speed challenges of tomorrow's networking and telecom systems. For more information go to: *www.xilinx.com*

Xilinx Launches Signal Integrity Website

Xilinx recently introduced the first website dedicated to solving signal integrity problems.

by Rob Schreck
Senior Marketing Manager
IC Solutions Marketing
rschreck@Xilinx.com

You face far more difficult engineering challenges today than you did just a few years ago. Feature size reduction and the need for reduced power consumption have driven device core voltages down from 5V to 1.5V and below. This change in voltage and signal frequency content requires us to use new design practices and take into account electrical effects that could previously be ignored. Xilinx is addressing these technical issues to help you get over the technical hurdles and complete your designs quickly, by providing an online reference for information on the fundamentals of signal integrity, PC board design and power supply considerations, simulation tools, and thermal characteristics.

Signal integrity issues in high performance designs have become a major concern in many companies. With frequencies above 300 MHz and rise times of a few hundred picoseconds, you face noise and signal reflections that severely complicate PC board design. These noisy signals affect performance, system development, and product introduction schedules. Without proper signal termination, you can experience a lot of ringing and reflection such as that shown in Figure 1.

To help solve these signal integrity problems, Xilinx offers the unique Digitally Controlled Impedance Technology (DCI) in the Virtex-II FPGA family. DCI allows you to attain far higher signal integrity with far fewer external resistors. With DCI implemented in Virtex™-II FPGAs, the signal shown in Figure 1 can look much cleaner, like the one shown in Figure 2. DCI eliminates most of the cumbersome termination resistors in your design and automatically adjusts to impedance mismatches caused by temperature and voltage fluctuations.

Figure 1 - Noisy signal without Xilinx DCI (Digitally Controlled Impedance)

Figure 2 - Clean signal resulting from the use of DCI

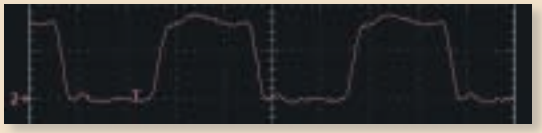

On the Xilinx Signal Integrity Central website, you'll find a wealth of information, including:

- A tutorial on signal integrity
- A glossary about signal integrity
- Information on IBIS models
- Information on simulation tool vendors
- Information on high-density fine-pitch BGA packages
- Information on how to route the I/O
- Information on power supplies, bypassing, heat sinks
- Power estimators
- Information on multi-gigabit signaling
- And much much more.

Conclusion

Xilinx DCI technology will make your PC boards easier to develop, less expensive, and more reliable. Visit the Xilinx Signal Integrity Central website at: *www.xilinx.com/xlnx/xil_prodcat_landingpage.jsp?title=Signal+Integrity*

Co-Processor Acceleration and Design Reuse: Extending the Market for Platform FPGAs with DK1 Design Suite

Celoxica's DK1 design suite accelerates co-processing between a CPU and a Platform FPGA.

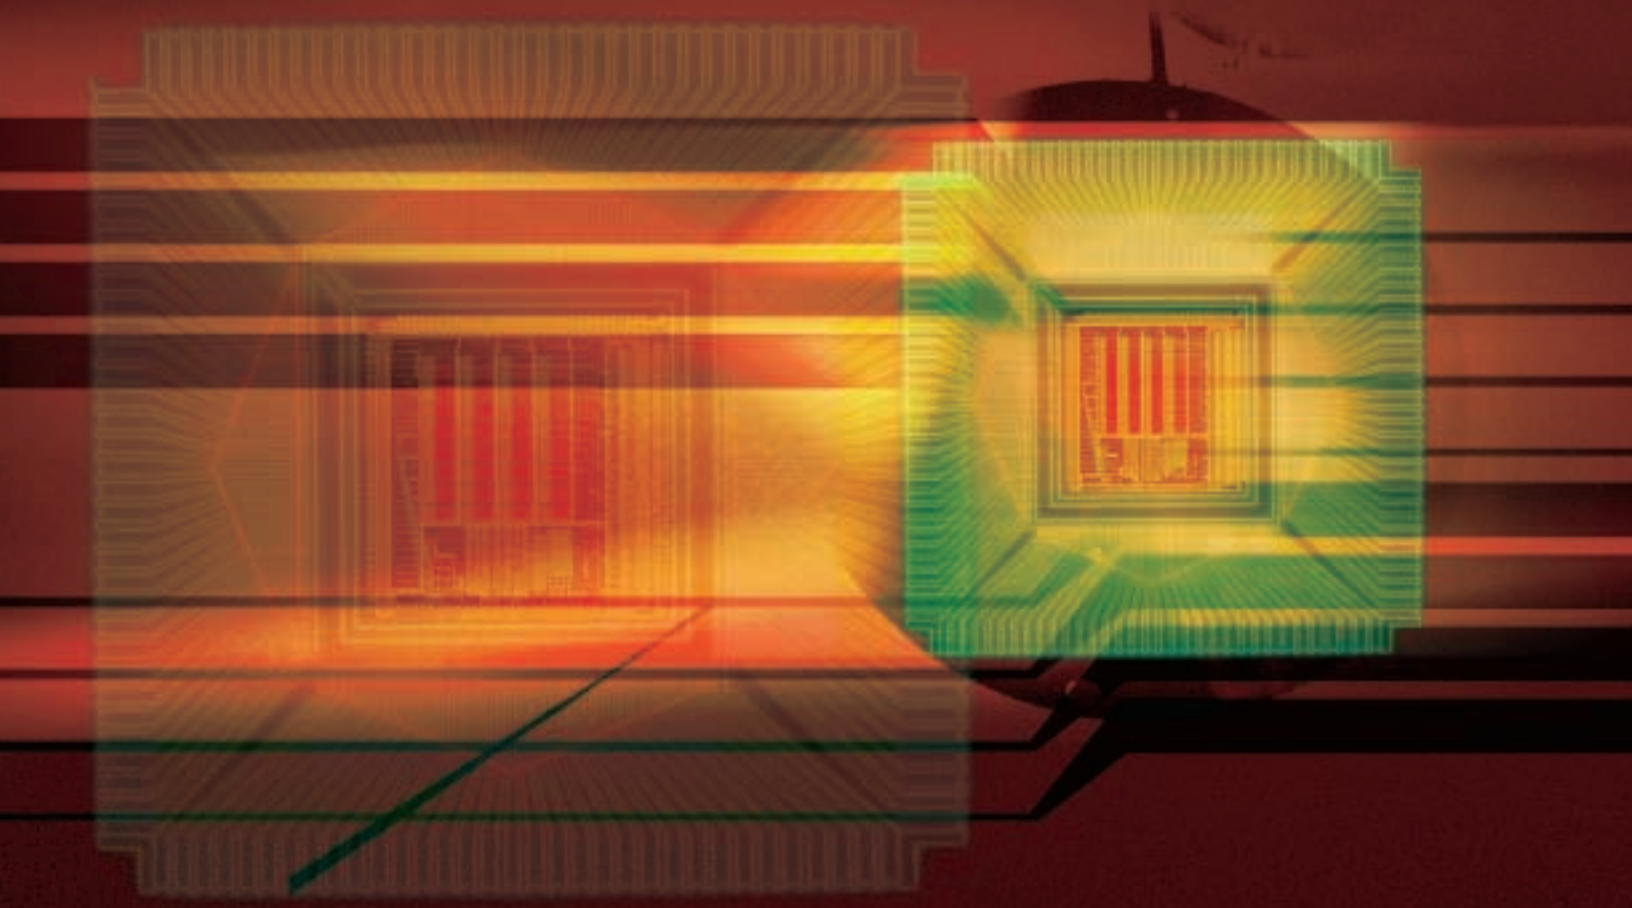

by Dennis Nye
Senior Vice President, Worldwide Sales and Marketing
Celoxica Limited
dennis.nye@celoxica.com

Few embedded applications will ever be fully created in programmable logic. Sequential algorithms are frequently best performed by software implementations using conventional CPU architectures. Optimal hardware acceleration comes most often from functions performed in parallel, normally in the physical form of a hardware co-processor closely coupled with the CPU.

In combination, the Xilinx Platform FPGA and the Celoxica™ DK1 system-level design suite represent a significant new approach to creating embedded solutions for a broad range of leading-edge applications, such as base stations, network compression, encryption, network traffic control, digital television, and more.

To exploit the multimillion-gate potential of advanced Platform FPGA architectures, companies need complementary methods that accelerate front-end logic design. C-based languages excel at implementing the algorithmic functions needed for system-level designs.

The DK1 design suite is a bridge across which Platform FPGAs can move into new markets. By addressing the needs of software engineers who want to benefit from the advantages of hardware acceleration, the DK1 suite speaks their language – Handel-C, a syntax based on standard ANSI-C. Thus, it is relatively simple for software engineers to learn and begin the process of converting C algorithms into Handel-C code to accelerate compute-intensive functions in parallel hardware.

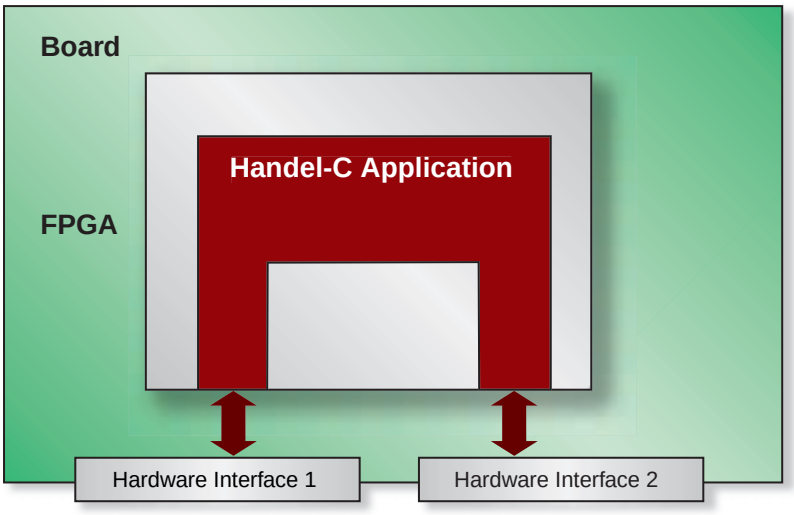

Figure 1 - Migrating software into an FPGA is just the first step.

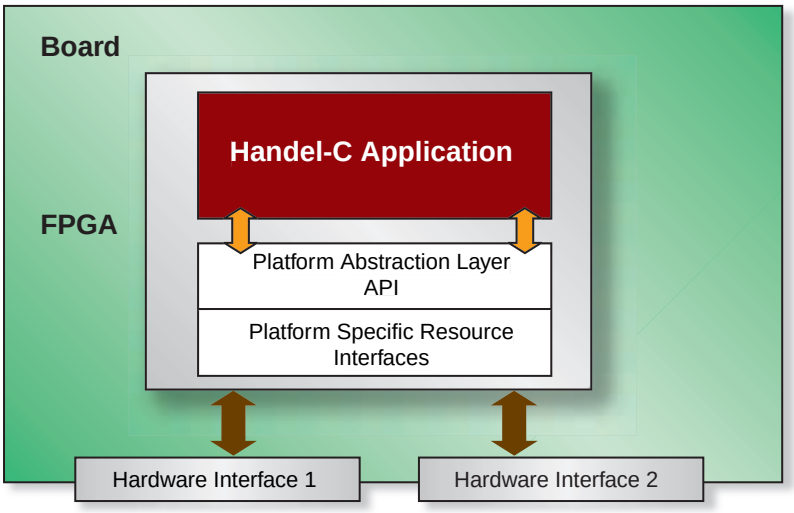

Figure 2 - Platform Specific Resource Interfaces reduce porting times.

PAL for FPGAs

Facilitating the migration of software into an FPGA is a useful first step (Figure 1), but engineers developing acceleration solutions are looking for further support. They expect the hardware independence offered by modern operating systems. Introducing this hardware independence into an FPGA environment requires the implementation of a layered and consistent approach to hardware interfacing – not dissimilar to the concept of device drivers within an operating system. Celoxica's recently announced Platform Abstraction Layer (PAL) strategy for FPGA implementations provides exactly this type of hardware independence.

PAL shields development engineers from low-level hardware interfaces in order to ease the integration of FPGAs with physical resources (Figure 2). It does this by providing a library of low-level interfaces to specific platform resources, such as I/O or memory. This library, called the Platform Resource Support Package (PRSP), is then accessed from the applications on the FPGA using a simple and consistent Application Programming Interface – the PAL API.

PRSPs are written once for each embedded system. Applications using the API to access platform resources are then portable between systems utilizing the PRSP concept. This approach can reduce porting times significantly. Development engineers can rapidly port designs from their prototype platform to their final system, or perform design re-use from one gener-

ation of a product to the next. So, at the same time as accelerating the development cycle for products utilizing FPGAs, the PAL approach allows companies to make more effective use of their design resources.

MPA Integrates CPU and FPGA

While PAL simplifies the implementation of solutions requiring access to the physical resources of the embedded solution, co-processing solutions will require sophisticated integration between those parts of the application running on the microprocessor and the accelerating routines provided on the FPGA. Developing this integration on a case-by-case basis would be unwieldy and complex (Figure 3).

Celoxica's Multi-Port Adaptor (MPA) technology provides the missing integration between the microprocessor and FPGA elements of a solution. By itself, MPA – using the PAL API – is also independent of the specific hardware solution.

The MPA provides a simple mechanism for function calls to be made from software within the microprocessor to functions in the FPGA or vice-versa (Figure 4). The developer may safely make multiple calls to the same or multiple functions as the MPA architecture ensures the management of the sequence of calls, delivering and returning data in the correct sequence, and to the correct function.

Conclusion

The benefits of a simple, layered approach for co-processor acceleration are:

- **Reduced solution complexity**
- **Reduced time to market**
- **Reduced cost**
- **Improved design re-use**

To learn more about Celoxica DK1 design suite, go to ***www.celoxica.com/products/design_suite/.***

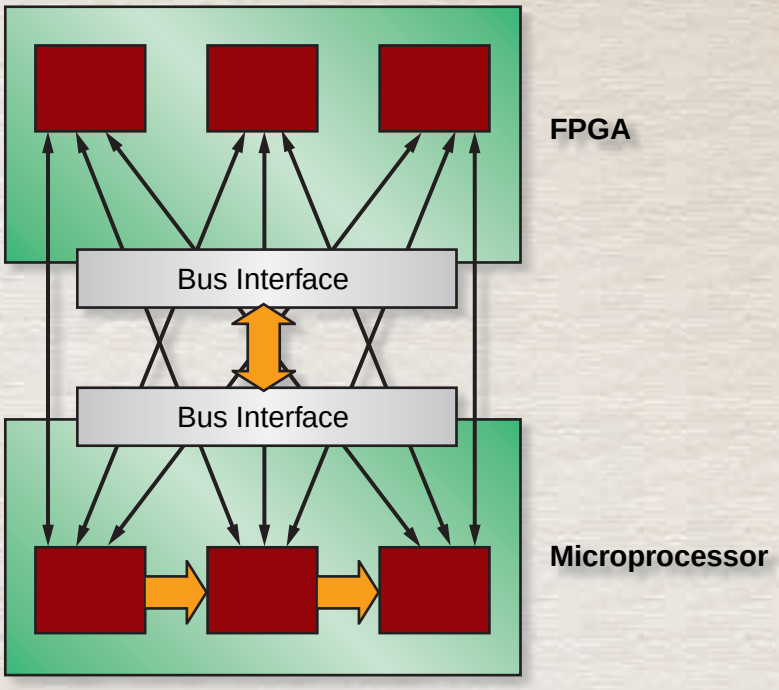

Figure 3 - FPGA-to-microprocessor without MPA support is too complicated.

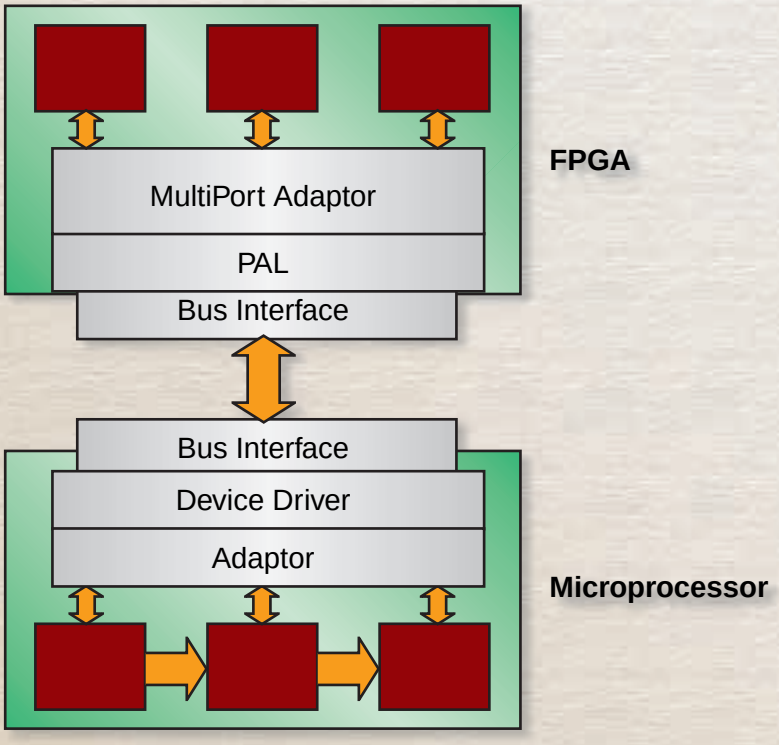

Figure 4 - DK1 design suite enhances and accelerates FPGA co-processing with a microprocessor.

And the Winner is... Cool Module Design Contest.

CoolRunner CPLDs and Insight make it easy to create low power designs for Visor PDAs.

by Steve Prokosch
Product Marketing
sprokosch@xilinx.com

We initiated the Cool Module Design Contest at the January 2001 Portable Design show in Santa Clara, California, to generate interest in the Insight Springboard™ development platform for Visor™ handheld PDAs (Personal Digital Assistants). The platform is based on the Xilinx CoolRunner™ XCR3256XL CPLD which enables ultra low power designs. With the Visor Springboard slot, you can develop a multitude of applications that run on Handspring PDAs, including communications and wireless connections, multimedia, ebooks, games and sports, and productivity tools. With the Insight Springboard development kit, you can quickly develop these Visor add-on designs.

The Contest

We reviewed over 100 entries from which ten finalists were chosen – it was a difficult task to select the most ingenious, complete, and comprehensive idea. The designs were divided into the following categories:

- Biometrics
- Medical Electronics
- Sports
- Communications and Wireless
- Audio
- Video and Display Technology
- Electronic Test
- Manufacture Test
- Physical Measurement
- Interface
- Other Unique Implementations

The ten finalists were shipped an Insight Springboard development kit and a Handspring Visor Platinum PDA. These finalists then implemented their conceptual ideas and competed for the grand prize of $10,000. The winner was chosen by a team of three industry technologists from Xilinx, Handspring, and Portable Design Magazine.

The Finalists

The ten finalists included these selections:

- A device for the recording, display, and distribution of a fingerprint. This focuses on child safety, which is now a national campaign for the recovery of lost children.
- A multipurpose device for cycling enthusiasts.
- A Springboard idea to connect to the Family Radio Service (FRS). FRS devices are readily available from manufacturers such as Motorola (T6300 radio).
- A musical instrument tuner.
- A musical Instrument Digital Interface (MIDI) recorder and sequencer.
- A spectrum analyzer.
- A logic analyzer.
- An engine analyzer for an onboard diagnostic (OBD) to monitor car engine information.
- A portable MIL-STD-1553 bus analyzer.
- An ultrasonic distance measuring device.

The Grand Prize Winner

The winner was "Cool Trak," created by frog design; it's an add-on module that collects and displays many crucial statistics for the bicycle enthusiast. The module measures speed, distance traveled, pedaling rate, air temperature, humidity, barometric pressure, altitude, and the rider's heart rate.

"The Cool Trak module is designed for recreational cyclists with an eye on fitness, as well as competitive or professional riders looking to achieve maximum performance," said Preston Brown, director of electrical engineering at frog design (*www.frogdesign.com*). "CoolRunner CPLDs are ideal for low-power and low-cost applications such as the Cool Trak module, and the free Xilinx WebPACK™ VHDL design environment allowed us to rapidly complete the project."

From left to right: Jeff Hawkins: founder, chairman, and chief product officer at Handspring, Inc. Bill Carter: former Xilinx chief technology officer. Richard Nass: editor-in-chief of Portable Design Magazine

"We were impressed by the innovation put forth by the teams submitting entries into the Cool Module design contest," said Jeff Hawkins, founder, chairman and chief product officer at Handspring, Inc. "Our hats go off to Xilinx for sponsoring the event, Portable Design for promoting the event and playing a crucial role in the judging process, and most importantly to the numerous teams that demonstrated incredible ingenuity in designing new Springboard module concepts for the Handspring Visor."

Conclusion

Creating low power Springboard designs for the Visor PDA is easy with the Insight Springboard development kit and Xilinx CoolRUnner CPLDs. Visit the Cool Module Design Contest home page and read all about it at: *www.xilinx.com/contest/index.htm.*

Two Virtex-II FPGAs Deliver Fastest, Cheapest, Best High-Performance Image Processing System

Dillon Engineering not only exceeded their client's performance specifications, but they also delivered the solution under budget by an order of magnitude.

by Tom Dillon
President, Dillon Engineering Inc.,
tdillon@dilloneng.com

In the early part of 2001, one of our clients asked us to consult on a sophisticated digital image processing system requiring a combination of high resolution and high frame rates. A key element in this application was to perform FFTs (Fast Fourier Transforms) on a huge amount of data at very high speed.

The client's existing solution for the FFT portion of the system was based on more than 40 high-end, fourth-generation (G4) IBM® PowerPC™ CPUs. This system had been intended to achieve at least 30 fps (frames per second), but in fact, it never delivered more than 15 fps.

After evaluating a number of design alternatives, we determined an implementation based on two Xilinx® Virtex™-II FPGAs would satisfy our client's immediate requirement for 120 fps – and would be scalable to 240 fps in the future. In addition to being one of the fastest – if not the fastest – FFT processor in the world, our solution would also cost only a fraction of our client's existing implementation.

In this article, we will first describe the scale of the challenge, then look at the alternatives we considered, and finally, reveal the solution we implemented, using our proprietary ParaCore Architect™ core generation utility, Mentor Graphic's LeonardoSpectrum™ 2001 synthesis technology, and two Xilinx Virtex-II XC2V6000 Platform FPGAs.

The Challenge

Unfortunately, our nondisclosure agreement with our client bars us from revealing their identity or even the nature of their extremely competitive business. Thus, we cannot describe the actual application in great detail or show images of the final hardware implementation. We must, therefore, leave any possible applications to your imagination.

What we can say is the most challenging part of this image processing system was to accept high-resolution digital images at 120 fps and generate a corresponding two-dimensional (2D) FFT for each frame. By generating these 2D FFTs, the images were transformed into frequency domain representations that can be used to analyze various features of the object being viewed.

The processing requirements associated with such a system are tremendous. The combination of 16-bit pixel data, a resolution of 2K x 2K pixels (2,048 x 2,048 = four megapixels), and a required frame rate of 120 fps results in 480 megasamples of 16-bit data per second. This huge amount of raw data must undergo extensive processing to convert it into the final 2D FFT, as depicted in Figure 1.

The system comprises two FFT processors. The first processor is used to generate a one-dimensional (1D) FFT of the frame. The second processor then converts this 1D FFT into a 2D FFT. The first FFT processor works on a row-by-row basis. The FFT for each row also contains 2,048 pixels, but in this case each pixel now represents a component in the frequency domain. Each of the 2,048 rows forming the frame requires an associated FFT to be generated. The result is a 1D FFT of the whole image.

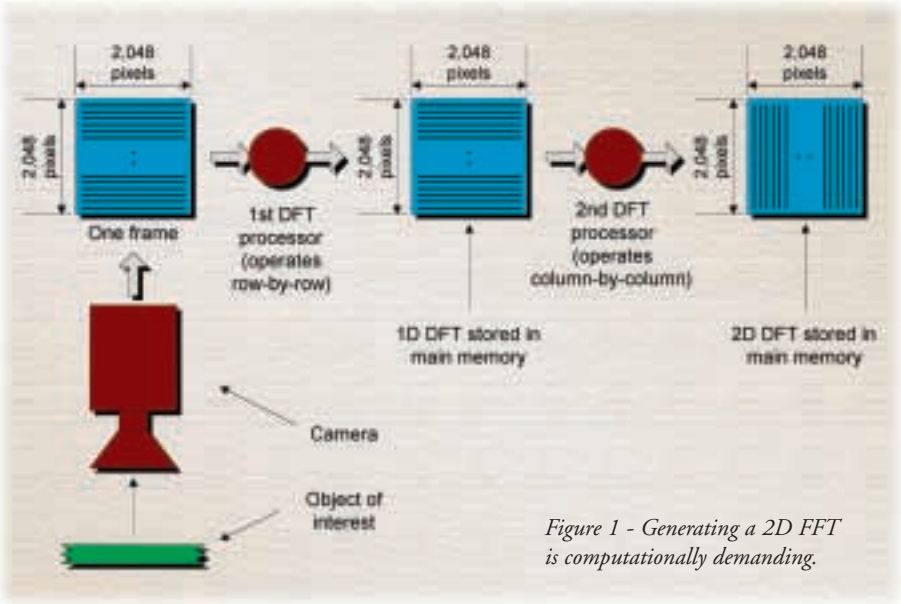

Figure 1 - Generating a 2D FFT is computationally demanding.

Once the 1D FFT associated with a frame has been produced, it is stored in main memory. If we visualize the first FFT processor as working "horizontally" across all of the pixels forming a row, we can consider the second FFT processor to work "vertically" and process the columns formed by taking the equivalent pixels in each row. That is, the second FFT processor will start working on the first column formed by pixel 0 on row 0, pixel 0 on row 1, pixel 0 on row 2, and so on. Once the second FFT processor has completed this column, it will commence working on the column formed by pixel 1 on row 0, pixel 1 on row 1, pixel 1 on row 2, and so on.

Note that as soon as the 1D FFT has been generated and stored in main memory, the camera is free to take another image. The first FFT processor then starts working on this new frame at the same time as the second FFT processor is converting the 1D FFT into its 2D equivalent.

The smallest computational element used to generate an FFT is called a "butterfly," which consists of a complex multiplication, a complex addition, and a complex subtraction as shown in Figure 2. In turn, the complex multiplication requires four simple multiplications and two simple additions, while the complex addition and complex subtraction each require two simple additions. This means that each butterfly requires a total of four simple multiplications and six simple additions.

Processing a single 2,048 pixel row (or column) requires a total of 11,256 butterflies organized in eleven "ranks," where the outputs from the butterflies forming the first rank are used to drive the butterflies forming the second rank, and so forth. Thus, processing a single row requires 45,025 simple multiplications and 67,536 simple additions. In order to generate the FFT for an entire frame, this process has to be repeated for each of the 2,048 rows (or columns) forming the frame. This means that in order to achieve a frame rate of 120 fps, the processing associated with each row (or column) must be completed within 4 µs (microseconds). This leads to a time budget of 90 ps (picoseconds) per simple multiplication and 60 ps per simple addition.

Alternative Solutions

The first stage of this project involved doing our "homework" into the current state-of-the-art for FFT processing. We had to become intimately familiar with the myriad reduction techniques and computational "tricks" available. We then started to evaluate a range of alternative implementation strategies.

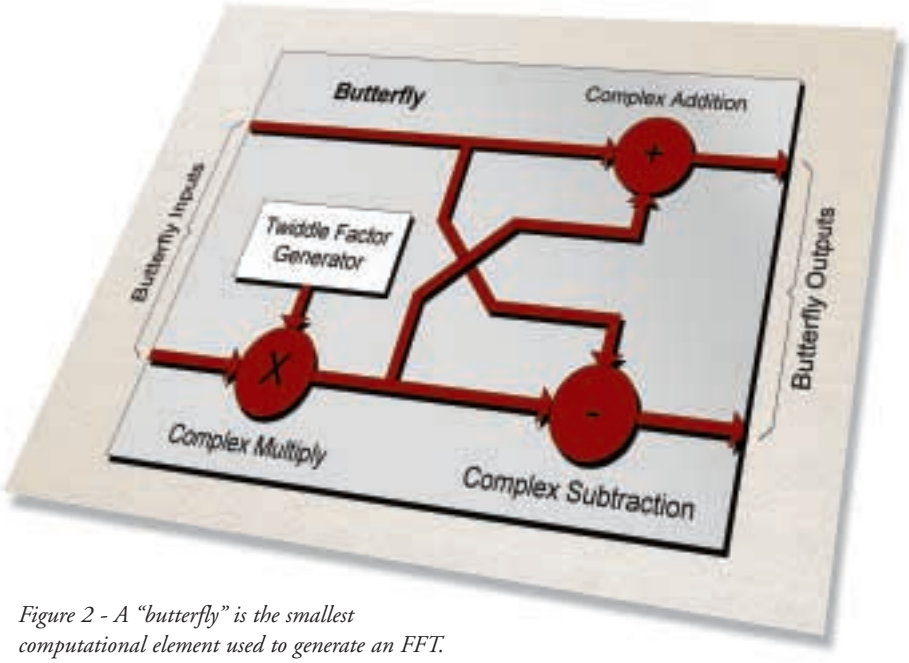

Figure 2 - A "butterfly" is the smallest computational element used to generate an FFT.

The first of these was to extend the customer's existing implementation, but we calculated that it would require more than 100 of the most powerful PowerPCs currently available to achieve 120 fps. In addition to being cost-prohibitive, this solution would not have been easily scaleable up to 240 fps in the future. Having ruled out a PowerPC solution, we moved on to consider off-the-shelf approaches – FFT processing boards, DSP-based solutions, and ASIC-based solutions – versus a custom FPGA-based design.

Off-the-Shelf Approaches

In the case of an off-the-shelf (commercially available) FFT processing board, the most appropriate option required 16 VMEbus boards, each costing $30K.

With regard to a DSP-based solution, high-end alternatives like the C67™ device from Texas Instruments or the Shark™ chip from Analog Devices took 300 µs and 900 µs to perform a 2k-point FFT, respectively (remember that our requirement was for 4 µs). Even using multiple devices, the end result would be slower than a PowerPC-based solution.

In the case of an off-the-shelf ASIC-based approach, the best options to implement a 2k-point FFT were the Pathfinder-2™ device from Catalina Research (31 µs), the PowerFFT™ chip from doubleBW (41 µs), and the DSP24™ component from DSP Architects (60 µs). A solution based on the fastest of these – the Pathfinder-2 device – would have required 32 of these components at $10K each, plus the cost of any supporting ICs.

Custom FPGA-based Solution

And so we came to consider a custom FPGA-based solution. We've been a member of the Xilinx XPERTS program since day one. Thus, it was no surprise to us that the Virtex-II family of Platform FPGAs leads the field for this class of computationally demanding, data-intensive application. In fact, using a Virtex-II XC2V6000, we managed to execute a 2k-point FFT in only 2.8 µs (well within the 4 µs budget required to satisfy our 120 fps design criteria). What's more, we achieved this with a clock frequency of only 125 MHz. The final product was a single VME board with two Virtex-II XC2V6000 devices (one each for the 1D and 2D FFTs) – at only one-twentieth the cost of the best off-the-shelf alternative as illustrated in Table 1.

ParaCore Architect

A critical factor in the design of the 2k-point FFT was our internally developed ParaCore Architect technology. Developed over a number of years, this technology facilitates the design of parameterized cores. The process begins by creating a source file containing a highly parameterized description of the design at an extremely high level of abstraction. The ParaCore Architect utility takes this description, combines it with parameter values specified by the user, and then generates an equivalent HDL representation. The resulting HDL is guaranteed suitable for use in any simulation and synthesis environment, so it isn't necessary to run any form of HDL rule-checking program.

The beauty of this type of highly parameterized representation is that it's extremely easy to target it toward a new application or an alternative device. For example, if we decide to change the length of the FFT from 2k to 1k points, setting a single parameter takes care of all of the details, including re-sizing the RAMS used to store

Solution	Comments	Total Cost
Off-the-shelf FFT processing board	Required 16 Cheetah™ boards from Catalina Research	$480,000
DSP-based solution	Required 5 PowerPC G4 VME boards from Mercury Computing (75 CPUs with DSP functionality)	$750,000
Off-the-shelf ASIC-based solution	Required 32 Pathfinder-2 ASICs from Catalina Research (plus supporting memory and logic)	$320,000+
Custom FPGA-based solution	Required 2 Xilinx Virtex-II XC2V6000 FPGAs plus SRAM on a single VME board	$20,000

Table 1 - Cost comparison of alternative implementation options

any internal results, and so forth. Similarly, another parameter can be used to select between fixed- and floating-point math formats (in the latter case, two further parameters are used to specify the size of the exponent and the mantissa).

Of particular interest is the way in which our FFT algorithm is implemented. Consider the 11,256 butterfly operations required to implement a 2k-point FFT. If execution time were not a major factor, it would only be necessary to use a Virtex-II XC2V40 device with its 4x multiplier blocks, create a single butterfly structure (four simple multipliers and six simple adders) and to cycle all of the butterfly operations through this structure. The resulting structure would take 90 µs to generate each 2k-point FFT. Although this is extremely respectable, it falls well short of the 4 µs time budget required by our image processing application.

The easiest way to increase the speed of the algorithm was to increase the number of butterfly structures instantiated in hardware and to perform more of the processing in parallel. In the case of XC2V6000 devices with 6 million system gates, 144 x 18-bit multipliers, and 144 x 18-kilobit RAM blocks, it was possible to perform an entire 2k x 2k-point FFT fast enough to achieve our 120 fps system requirement. And using XC2V10000 components, we will be able to scale the system to achieve 240 fps. To target these different devices only requires setting a single ParaCore Architect parameter to specify the number of butterfly structures we require to be instantiated in hardware.

LeonardoSpectrum 2001

When you're designing an FPGA with 6 million system gates, it's very difficult to achieve optimum performance without (a) spending decades working on the problem by hand or (b) having great tools. Our client would not have accepted the first approach, so we decided to go with option (b).

Thus, another vital factor in our design was to use the LeonardoSpectrum 2001 synthesis tool from Mentor Graphics. We've been using Mentor's tools for eight years and have always been satisfied with their power and functionality, but we still took the time to consider other options that were available. In the case of this project, there were a number of factors that made Mentor the vendor – and LeonardoSpectrum the tool – of choice:

- The excellent relationship between Mentor and Xilinx means that LeonardoSpectrum synthesis technology is at the cutting edge of the Xilinx product offerings. In this case, it was ready for Virtex-II devices pre-silicon, which allowed us to start designing well ahead of the physical components becoming available.
- We didn't have to create any special constructs in our HDL, because the LeonardoSpectrum tool inherently understands functional elements like multipliers, adders, block RAM, and the like. The tool automatically instantiates the functional elements' equivalent hardware counterparts in Virtex-II devices. (Other tools we evaluated would have required us to go to extremes to code our HDL in strange and amazing ways to achieve the same effect).
- For initial design evaluations, the LeonardoSpectrum interface has a "single pushbutton" mode that we used to good effect. Later in the design process, we moved to using the more advanced modes that allow every synthesis attribute to be controlled step-by-step.
- Furthermore, the LeonardoSpectrum interface can be used to drive the Xilinx Foundation™ place-and-route software. Having a single user interface to drive both tools made our lives much easier.
- Last, but not least, LeonardoSpectrum's TimeCloser™ technology allowed us to "close the loop" by seamlessly importing timing data from Foundation place-and-route back into synthesis to re-synthesize critical paths.

Conclusion

As you can imagine, everyone at Dillon Engineering is extremely happy with the outcome of this project. In addition to a significant design win, we currently have what we believe to be the fastest 2k-point FFT processor in the world. Besides leveraging our core competency in FPGA-based, high-bandwidth, real-time digital signal and image processing, this project also allowed us to take full advantage of our ParaCore Architect utility.

Furthermore, in addition to solving our client's immediate problem by providing them with a 120 fps solution, this solution is an order of magnitude less expensive than their existing implementation (which can only achieve 15 fps on a good day).

Finally, our solution is scalable and can be enhanced to provide 240 fps using Virtex-II XC2V10000 devices – just as soon as they hit the market.

Reconfigurable Vehicles Are Just Around the Corner

Automotive in-car applications are a fertile field for programmable logic devices. FPGAs and CPLDs offer designers the ability to evolve with new and emerging automotive standards and protocols, giving manufacturers faster time-to-market and longer time-in-market.

by Karen Parnell
European Marketing Manager
High Volume Products
karen.parnell@xilinx.com

When we envisage automotive electronics, we automatically consider electric windows, central locking systems, climate control, and electronic ignition systems, all of which require stringent qualification, temperature cycling, and certification. However, the emerging automotive electronics market has shifted from under the hood (or bonnet) to in-cabin multimedia applications. The inexorable trend toward mobile offices, onboard entertainment, and real time information on the move means the market is wide open for complex programmable logic devices (CPLDs) and field programmable gate arrays (FPGAs).

Historically, the lead time from a new vehicle conception to production has been about five to six years. Now, the shift is toward vehicle turnaround times as short as two years. This shift is forcing a phenomenal change in automotive design practices and techniques. Now, more than ever before, time to market is critical.

It's predicted that intelligent cars of the next decade will be able to perform the following functions:

- Provide a safe, informative, productive, and entertaining environment.
- Cruise intelligently by automatically keeping pace with the vehicle in front.
- Display their exact locations using satellite-based navigation.
- Summon emergency services to the precise location of a breakdown or accident with, or without, driver action.
- Illuminate long distances in the dark with infrared and ultraviolet headlights.

Figure 1 shows some of the emerging "semiconductor rich" products that will become standard options available in vehicles over the next few years.

By utilizing the flexibility, time to market advantage, and after-sale reconfigurability of Xilinx devices and software, manufacturers can not only be first to market but remain best in class long after the sale.

The Automotive Multimedia Platform Concept

As information and entertainment systems are added to automobiles, we're seeing the inevitable conflict of digital standards and protocols. A bewildering array of emerging standards and protocols are being tried and tested for use in the latest in-car systems, including: Bluetooth™, BlueCAN, MP3, Java™, AutoPC, AMIC (Automotive Multimedia Interface Consortium with JINI interface), WAP (Wireless Application Protocol), HTML, XML, MOST (Media Oriented System Transport), FireWire, CAN (Controller Area Network), TCP/IP, and more – but which one(s) will prevail?

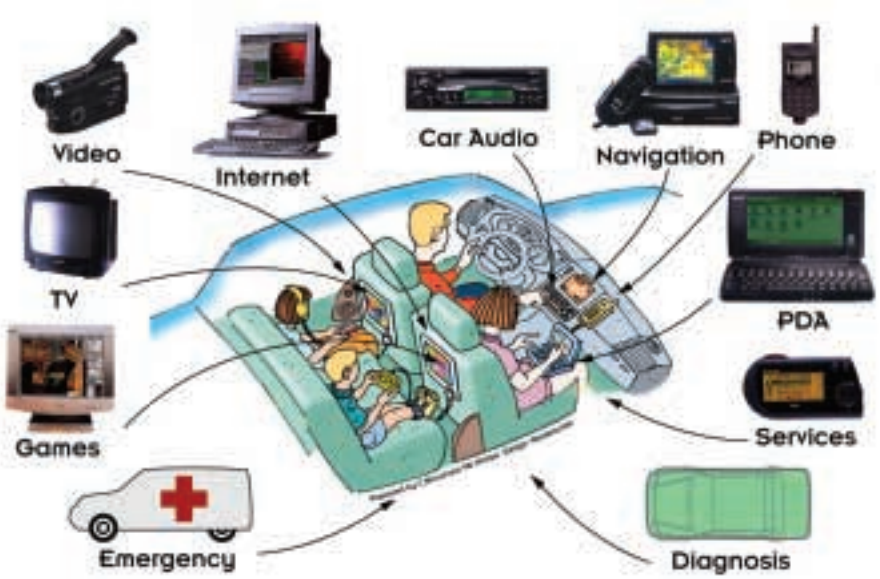

Figure 1 - Emerging in-car systems (source: Philips)

Overcoming the Design Challenges

Designers of the new wave in-car multimedia systems must include traffic information systems, Internet/Web access, electronic game consoles, MPEG music download capability, digital radio reception, and mobile commerce services. These designers must also have the flexibility to provide all or just a few of these functions.

Other challenges faced by the designers are:

- Small physical space available to place devices in the dashboard or seat backs
- Low cost units that can fit into more than one model of car
- Aesthetically pleasing look and feel
- Ease of driver use
- Ability to upgrade the unit when automotive standards and protocols change.

Designers must also ensure that the in-car multimedia system can "talk" to other devices introduced into the cabin space. For example, if a mobile phone is brought into the car environment, it should be automatically detected and able to communicate with the car's communications network. This automatic connectivity could also include connection of PDAs, portable PCs, MP3 players, and other personal portable electronic equipment to provide a truly connected and functional in-car environment.

Some after-market, third-party manufacturers have provided short-term solutions to some of these problems by producing audio systems based around MP3 technology. When in the cabin, the MP3 player is seated in the car audio system, but when you leave the vehicle, you can take it with you as a portable MP3 player

The New Way of Designing

As in-car and consumer functions converge – and new automotive and consumer standards and protocols emerge – manufacturers are starting to prototype multimedia platforms that can provide as much, or as little functionality, as required. Such multimedia platforms can be best realized by designing with reconfigurable hardware. Ideally, the multimedia platform should be based around one human-machine interface that allows all functions to be accessed via a menu-driven touch screen.

Reconfigurable hardware can be programmed late in the production flow to provide custom functionality on a standard

hardware platform – and can also be configured to accommodate the newest emerging in-car standards and protocols.

Figure 2 shows the automotive multimedia platform design approach. This concept allows hardware upgrades throughout the lifetime of the car. These upgrades can be implemented remotely by utilizing the wireless communications/Internet connectivity provided by the unit. Xilinx terms this remote hardware upgrading process Internet Reconfigurable Logic (IRL™) via the Xilinx Online program.

The Xilinx Online program is designed to enable, identify, and promote network upgradable systems, hence "future proofing" a system and avoiding product obsolescence. These systems can be upgraded, modified, or fixed long after they have been deployed in the field.

While many designers have been building upgradable devices based on Xilinx technology for years, the explosion of networked devices has dramatically increased consumer demand for these user-configurable, adaptable products. For example, an engineer can use the existing wireless communications/Internet infrastructure to reprogram an in-car multimedia system to include extra functionality, such as adding an MP3 player or upgrading the system to take advantage of the latest protocol or standard.

Multimedia System Design Flow Using FPGAs

The new way of developing in-car systems is to prototype using FPGAs in a generic development environment. The elements can then be developed quickly and easily without the need to fix specifications. This initial prototyping phase can be realized using Virtex-II Platform FPGAs. At this early stage, the different standards, protocols, and functions can be tried, tested, and debugged utilizing the headroom that a large FPGA gives.

As the design firms up, the specifications are "chilled" or "frozen," and in engineering terms, the boards are "productionized." In other words, the desired standard, protocol, or function is chosen from the many tested. This move from prototype board to productionized printed circuit board (PCB) enables the design to be optimized and fitted into a smaller, low cost FPGA, such as a device from the Xilinx Spartan-IIE family, while still leaving room for future system upgradability. Once in production, the FPGA can be used to aid total PCB testing using JTAG techniques. If necessary, designs can be 'tweaked' or enhanced at this stage.

The final stage is the look and feel or aesthetics of the product, which can be designed for each car manufacturer to fit into their specific dashboard (or fascia). All of the production multimedia units are built up around the standard FPGA-based platform. This standard platform can be programmed with its "personality" late in the production flow to accommodate last-minute design changes or end-user preferences.

Having a common platform cuts down on inventory and takes advantage of the cost savings associated with producing one platform for all units. The benefits of using this reconfigurable hardware-based multimedia platform approach are:

- Easier control of software and hardware development
- Re-use of components, be they software or hardware (design re-use)
- Increase in the time developers can spend on creating value as opposed to creating system interconnect structures
- Reduction of risk through clearer understanding of the basic components
- Choice of standard or protocol can be deferred and simply reconfigured later
- Increased productivity through the application of modern development tools
- Increased openness of the system, supporting the use of standards and integration of third-party components.

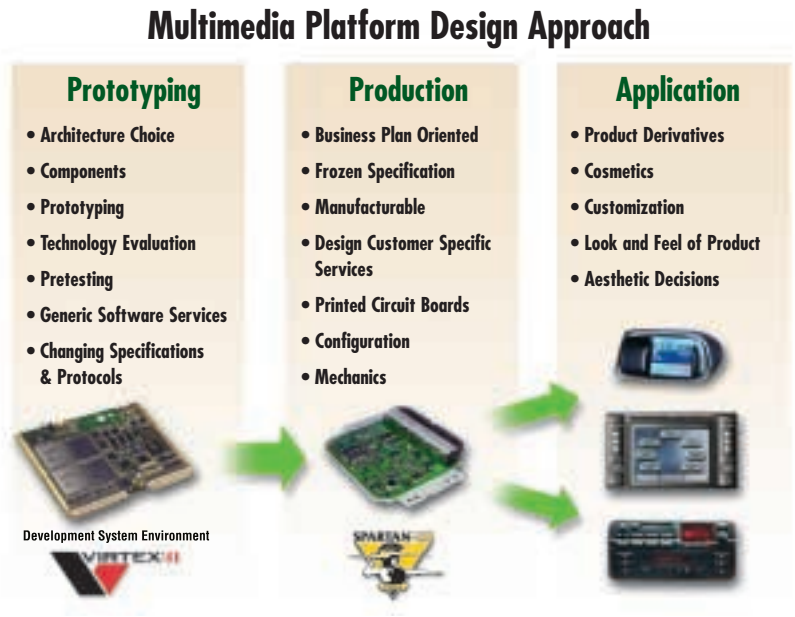

Figure 2 - Automotive multimedia platform design approach

The Xilinx In-Car Multimedia Solution

We predict the heart of the reconfigurable multimedia platform will be a Xilinx programmable logic device. The common platform approach enables one PCB to be produced for all customers – with the only change being to the style, shape, and color of the unit front panel to satisfy the need for product differentiation.

Xilinx is the leading provider of complete and innovative programmable logic solutions. Our products help minimize risks for manufacturers of in-car electronic equipment by shortening the time required to

develop products and take them to market – and keep them there long after the sale. You can design and verify your proprietary circuits in Xilinx CPLDs much faster than you possibly could by using traditional methods, such as manufacturing ASICs (application specific integrated circuits) and ASSPs (application specific standard parts).

Xilinx builds programmable integrated circuits, develops software, IP (intellectual property) cores, and other tools to provide complete solutions to our customers. We also provide world-class application support and design services to customers developing their own proprietary designs.

To "drive" the in-car digital convergence in the face of emerging standards and protocols, Xilinx recommends the Spartan-IIE FPGA family and the CoolRunner and 9500XV CPLD families. The Xilinx CPLD families are qualified to meet automotive temperature range parameters. The Spartan-IIE FPGA family is based on the very popular Virtex family. Spartan-IIE extends the legacy of the Spartan series, with more gates, better performance, and enhanced features. The Spartan-IIE family offers digital delay locked loops, programmable I/Os, on-chip block memory, and densities up to 300,000 system gates. These features and densities, coupled with enhancements to software and an increased number of available IP cores, provide a reduced time-to-market and increased time-in-market at a much lower cost. Xilinx also has a range of IP cores such as memory controllers, system interfaces, DSP, communications, networking, and microprocessors.

Looking Down the Road

Consumers are demanding the comforts of home, the facilities of the office, and the state-of-the-art in information and safety systems in their cars.

This digital product convergence scenario requires that the latest interface standards and protocols interconnect and interoperate. These standards and protocols are still emerging – and may get more complicated before they are standardized.

So, should manufacturers make an educated guess as to what standard will prevail and produce at risk, or should they wait for the standards to be fixed and get left behind?

With Xilinx devices you can win the time to market race at minimum risk by taking advantage of the capability to reconfigure your products to accommodate almost any standard or protocol – now and in the future. By using reconfigurable logic in production, the units can be reconfigured in-car to produce new hardware-based features, thus extending the life of the product.

Xilinx high-volume FPGA and CPLD devices provide you with cost-effective solutions that retain the traditional CPLD time to market advantage. Today, Xilinx programmable logic devices are employed by a large number of telematic and infotainment product manufacturers who recognize the added flexibility and time to market benefits achievable through the use of programmable logic solutions.

Conclusion

In the future, we expect programmable logic technology will play an even bigger role in automotive applications. This will come as a result of our ability to provide advanced features, such as Internet Reconfigurable Logic. By using IRL, you can build in upgradability and thus, reduce the risk of obsolescence.

For example, you could upgrade your car for the weekend by paying for the engine management system to be "race-tuned" to enhance the performance of your car. This would give you the ability to purchase a "racing car" for the weekend in your standard family vehicle. The engine management system would be reprogrammed via a wireless Internet connection to change the car's personality. At the end of the pre-paid time period, your vehicle would be reprogrammed back to its original settings. This is not just a dream about the future of the in-car environment but a realizable vision enabled by reconfigurable logic.

Car Driver Assistance and Active Safety

Car safety is one of the primary concerns in the automotive industry. Over the years, we have seen the addition of airbags, seat belt pre-tensioners, side impact beams, and other structural enhancements. We are now seeing the emergence of driver assistance systems to augment the more "physical" safety enhancements such as forward collision warning and night vision.

The latest driver assistance enhancement is based around a video-processing system located in the rear view mirror and dashboard of the vehicle. This system processes data from both inside and outside the car. It analyzes any perceived erratic vehicle movements, such as the car leaving the lane if the driver has fallen asleep, and automatically corrects and brings the car to a gentle stop. The system also analyzes outside light levels compared to road conditions, such as curves, and adjusts the headlamp luminosity and direction accordingly. This driver assistance system also measures distance to the car in front, and if an accident is imminent, it will ensure activation of the airbag and braking systems. The distance sensors can also aid when parking your vehicle. Xilinx FPGAs are ideal for implementing complex imaging and fast digital signal processing functions and are already being used in driver assistance systems.

Giga-Sample DSP Board Using Virtex-E FPGAs

Central Research Laboratories Limited (CRL) uses off-the-shelf FPGAs to create a wideband RF pulse capture and analysis board.

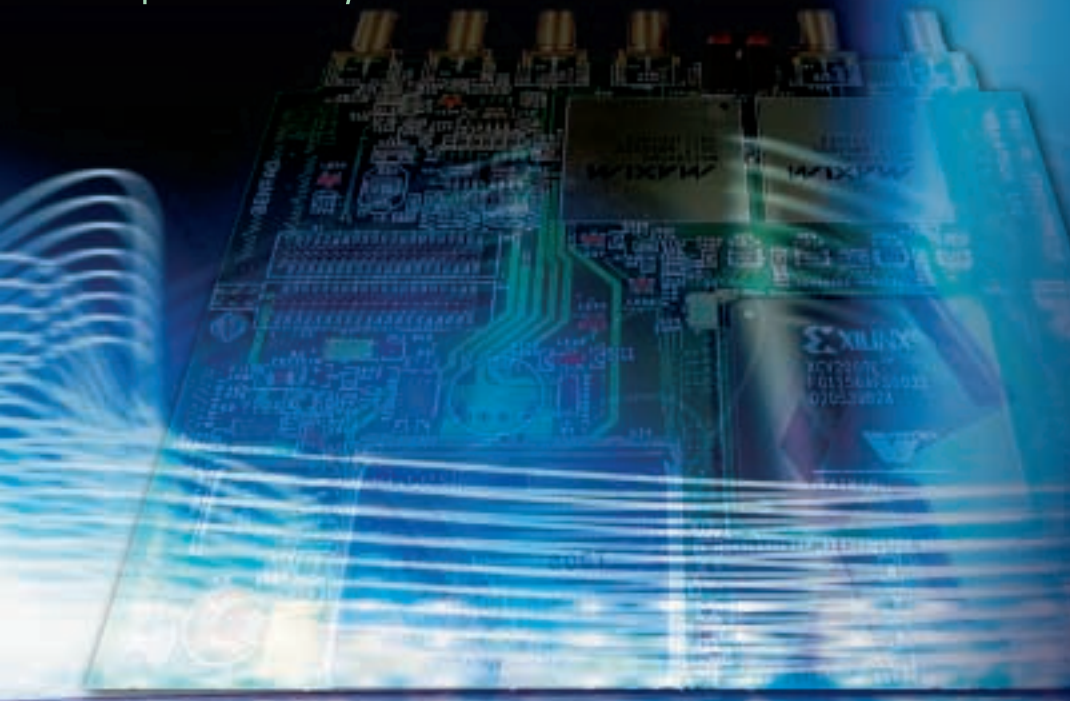

by Dr. Stephen King
Business Development Manager,
Central Research Laboratories Limited
sking@crl.co.uk

CRL considers Xilinx FPGA technology a key element for creating very wideband digital signal processing systems, bringing new capabilities to applications such as radar, spectral analysis, very broadband wireless communications, vibration analysis, and ultrasonics. To facilitate the development of these systems, CRL developed a high speed digitizer module in collaboration with Xilinx Xpert partner, Nallatech.

Using Virtex™-E FPGAs in conjunction with state of the art analog to digital (ADC) technology, CRL has combined the processing power conventionally associated with full custom ASICs, with data sampling rates normally associated only with fast sampling digital storage oscilloscopes. Key to the success of this program were:

- The advanced I/O capabilities of the Virtex-E FPGAs
- The System Generator tools for rapid DSP algorithm development
- The use of DIME-II™ FPGA platform from Nallatech.

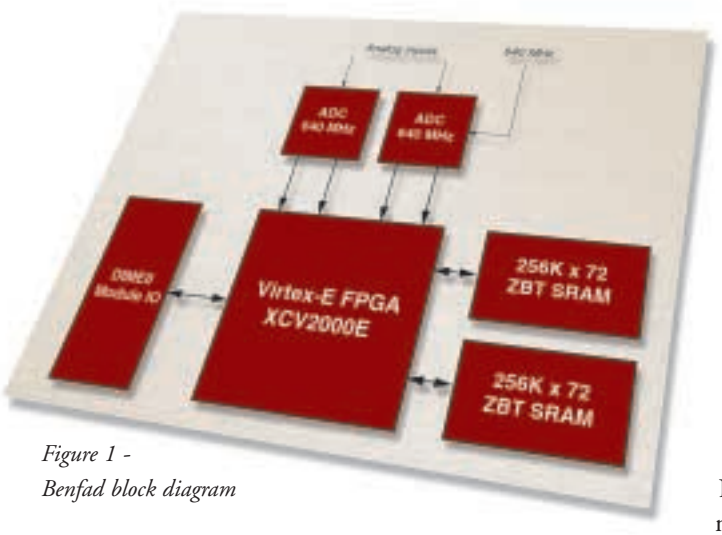

Figure 1 - Benfad block diagram

The "Benfad" DSP Module

The architecture of the CRL Benfad™ module, shown in Figure 1, relies heavily upon a number of advanced Virtex-E FPGA features. The primary advantage of these FPGAs is their support for LVPECL I/O capability. This enables the two 640 MHz ADC chips to be connected directly to the FPGA. Other key FPGA features include the 320 MHz DLLs, the high-speed on-chip RAM, and the high performance DSP capability of the extensive logic resources. These FPGA features, and the power of Nallatech's DIME-II module architecture, have resulted in the highly compact design shown in Figure 2.

Figure 2 - Benfad module

Giga-Sample DSP Algorithm Development

To get 640 MHz performance, you must broaden and slow down the data, and process it in parallel within the FPGA. However, care must be taken to use optimal DSP implementations to prevent the design from expanding and using excessive logic capacity.

Figure 3 illustrates how the data is broadened and processed using multiple parallel data paths at a practical clock rate. This design shows a broadband down converter and demodulator suitable for receiving broadband RF pulses. Two instances of this design readily fit into the 2 million-gate Virtex-E FPGA on the Benfad module, processing a total of 1.2 giga-samples per second.

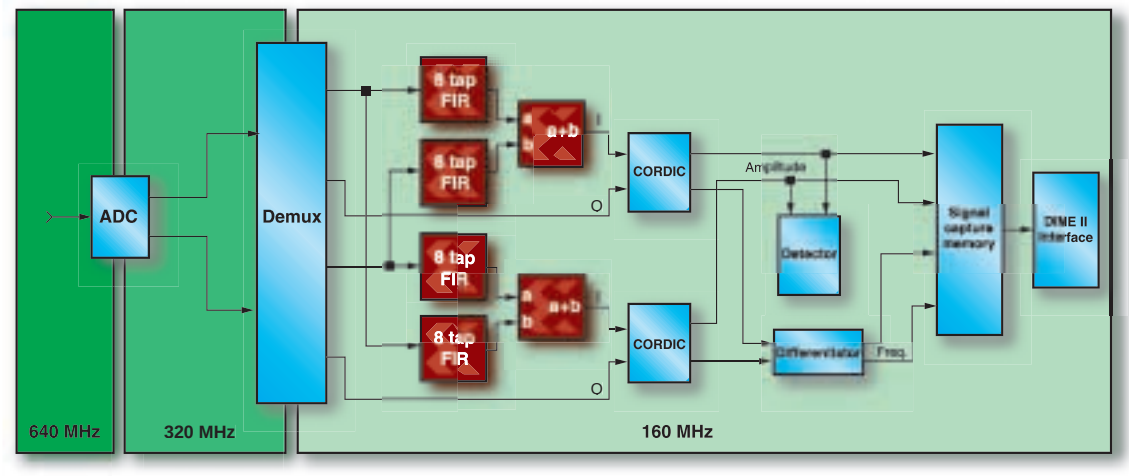

Figure 3 - Wideband down converter

Figure 4 shows how a simple chirped RF burst of 125 ns duration is processed by the system to yield its amplitude and frequency characteristics. Circuits may then be added to measure statistics of the pulse, such as time of arrival, duration, mean amplitude, and frequency. These statistics may then either be passed to the user or used as trigger events to pick out specific signals of interest in complex signal environments for further analysis.

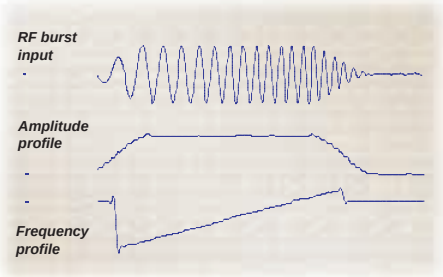

Figure 4 - RF pulse analysis

Tools such as the Mathworks' Simulink™, the Xilinx System Generator, and Nallatech's DIME environment, significantly reduce the time to design and implement such powerful DSP solutions. These tools unlock the potential of platforms such as Benfad for the real-time capture and analysis of a wide range of broadband signals.

Conclusion

Virtex-E FPGAs are a key technology for implementing sophisticated giga-sample data capture and analysis systems such as Benfad. In the future, migration to Virtex-II FPGAs will permit the ADC sample rate to be raised to 1.3 GHz, doubling the instantaneous bandwidth that can be processed. Pushing digital signal processing technology to these limits, using off-the-shelf boards, offers exciting prospects for the rapid development of new application areas.

For more information on the Benfad DSP module, e-mail CRL at: *sking@crl.co.uk,* or call +44 (0)20 8848 6452.

Internet Reconfigurable Logic with Alpha Data's ADM-XRC

Here's how to create a remotely upgradeable system on a Motorola MCP750 single board computer fitted with an Alpha Data ADM-XRC card.

by Tomas Whitlock
Design Engineer
Alpha Data Parallel Systems Ltd.
tw@alphadata.co.uk

It would be very advantageous if you could add new features, fix bugs, or make modifications to your existing designs, in the field, without replacing any hardware or sending a technician. That's the motivation behind the Xilinx® IRL™ (Internet Reconfigurable Logic) effort, which encompasses their PAVE (PLD API VxWorks™ Embedded) framework.

IRL and PAVE

Internet Reconfigurable Logic allows you to modify your FPGA designs after they have been deployed in the field, over any network, using well-known and standard methods of connectivity such as TCP/IP. And, unlike the flash BIOS on most PCs, IRL-based systems can automatically revert to a default known-good configuration.

For example, if you have a single board computer (SBC) with an FPGA embedded on it – it doesn't matter whether the FPGA is on a PCI Mezzanine card or coupled tightly to the CPU – you can incorporate the IRL framework into your VxWorks application. Then, when you deploy your SBC on a network, it starts listening for upgrade requests from a remote host or it sends an upgrade request when necessary. In the event of a failed upgrade, this type of IRL system can remain fully functional, because a read-only fallback configuration can be programmed into nonvolatile memory, at the factory, and used to reload a working configuration.

PAVE is a well-defined application-programming interface for loading bitstreams into FPGAs. It presents a common interface across all FPGA cards that it supports so you don't need to understand the process or figure out how to access the memory-mapped registers – the board support packages for the cards that you're using know how to do this, much like a device driver.

A Practical IRL System

A Motorola SBC750 SBC and an Alpha Data ADM-XRC can be combined to form an IRL system.

The MCP750 SBC

The MCP750 is a single board computer from Motorola. The important features, as far as IRL is concerned, are:

- A small NVRAM (nonvolatile RAM)
- On-board flash memory fitted as standard
- The ability to connect EIDE CompactFlash cards.

The ADM-XRC

The ADM-XRC is a PCI Mezzanine card with a Virtex™, Virtex-E, or Virtex-II FPGA and some local memory. A fast PCI interface is provided, capable of a real throughput of about 100 Mbps along with a flash memory for storing FPGA configurations.

The IRL System

As shown in Figure 1, we use the on-board flash memory (1) of the MCP750 to hold default known good images for the VxWorks kernel and VxWorks application. These images are written at the factory and never overwritten during the lifetime of the system in the field. Also contained in this flash memory is a bootloader, described below.

We use the flash memory (2) on the ADM-XRC to hold the default known good bitstream for the Virtex-II FPGA. This bitstream is written at the factory and is never overwritten during the lifetime of the system. On power-up, the default bitstream is automatically loaded into the FPGA.

- We put the upgradeable images of the VxWorks kernel, the VxWorks application, and the FPGA bitstream into the flash card (3). This has the benefit of making manual upgrades easy.

In summary, we use the flash memories as follows:

MCP750 on-board flash (1)	Flash on ADM-XRC (2)	EIDE flash card (3)
• Boot loader • Default VxWorks kernel • Default VxWorks application	• Default FPGA bitstream	• Upgradeable VxWorks kernel • Upgradeable VxWorks application • Upgradeable FPGA bitstream

Table 1. Flash memory usage

The Boot Loader

An important firmware component is the boot loader that resides in the on-board flash memory of the MCP750 (see Table 1, column 1). The boot loader's job is to decide whether to use the default or the upgradeable configuration. It makes several checks to decide this, including:

- Are the checksums of the upgradeable components valid/correct?
- Is the "clean shutdown" flag in the NVRAM of the MCP750 set to TRUE?

If these checks fail, the boot loader selects the default configuration.

The Upgrade Process

When we want to upgrade the system, we typically establish a TCP/IP session (for example, TFTP) with the target system to download new images and bitstreams for the components listed in column 3 of Table 1.

At the end of a successful download, the target system marks the upgraded components as valid. We then issue a reset command to reboot the system with the new software.

Conclusion

Using IRL, PAVE, and Xilinx FPGAs, it's easy to create a system with upgradeable hardware that is guaranteed to work reliably. For more information on this application, go to: *www.alphadata.co.uk*

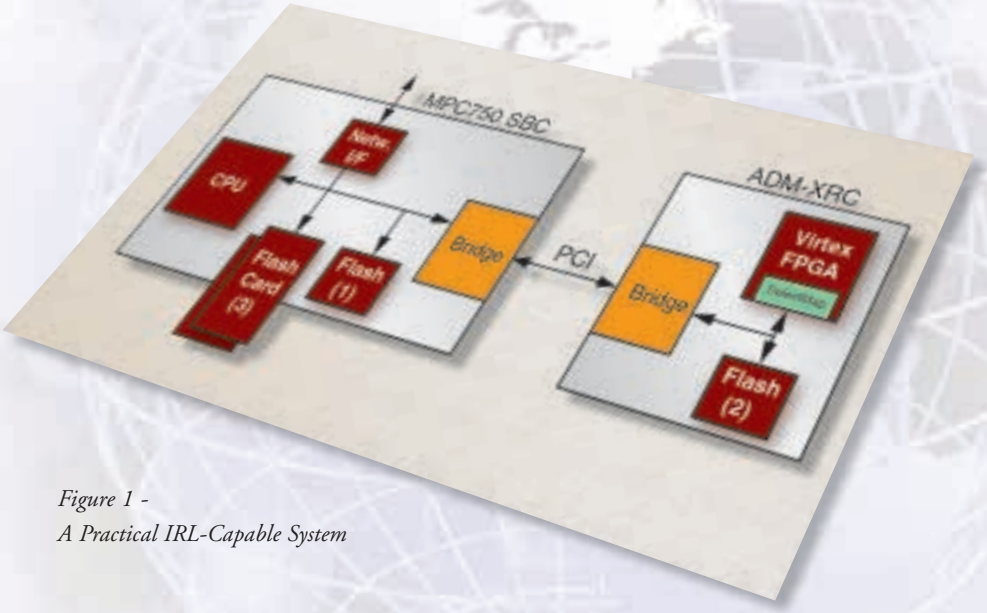

Figure 1 - A Practical IRL-Capable System

Virtex-II DSP Engines Enable Software Defined Radio

Use Virtex-II FPGAs to create high-performance, flexible SDR systems.

by Katie DaCosta
DSP Solutions Marketing
katie.dacosta@xilinx.com

Migrating an existing communication infrastructure from one generation to another is enormously expensive, requiring service providers and operators to make significant investments in new hardware equipment. To address this problem, the industry has developed a software upgradeable communication system known as Software Defined Radio (SDR). This concept is enabled by new and extremely fast FPGAs, such as the Virtex™-II family, that are designed for programmable, high performance, digital signal processing.

SDR solutions require high data sample rates and channel integration, creating the need for very high-performance, yet fully programmable, digital signal processing designs. The conventional digital signal processor (DSP)-based solution cannot meet this need. For example, today's cutting edge analog to digital converters deliver data rates of over 100 Mega Samples per Second (MSPS) with 12 to 14-bit resolution. The performance of a traditional DSP device falls short by an order of magnitude – a new way of processing data is needed if you are to take advantage of the SDR concept and meet the challenges created by future wireless communication standards.

Virtex-II FPGAs – Ideal for SDR

SDR solutions are required to perform many sophisticated signal processing tasks, including advanced compression algorithms, power control, channel estimation, equalization, forward error correction, adaptive antennas, rake processing, and protocol management. While there is a plethora of silicon alternatives available for implementing these functions, such as traditional DSPs and Application Specific Integrated Circuits (ASICs), FPGAs are ideal solutions and are often the only option.

Advanced process technology has enabled us to develop high-density FPGA devices that are extremely well suited to the needs of high-performance real-time signal processing. The Virtex-II family provides a valuable combination of high performance and configurability – both required for SDR systems.

Example SDR Applications

An example of an SDR system is shown in Figure 1. Digital filters are at the heart of the system. On the transmit side these filters help shape and translate a signal from the baseband to the Intermediate Frequency (IF) before it is converted to an analog signal and sent by the antenna. On the receive side, filters are used in the digital down converter, in the channel equalizer, and for the digital re-sampling of a signal in the timing recovery and acquisition loop. With recent FPGA technology, these essential filter functions are easily designed using the highly parallel structure of a Virtex-II FPGA. Figure 2 shows how the Virtex-II platform can be used to implement a "farm" of forward error correction (FEC) algorithms.

Extreme DSP Performance

The digital filter used most often within an SDR is the Finite Impulse Response (FIR) filter. There are multiple ways to implement a FIR filter within an FPGA.

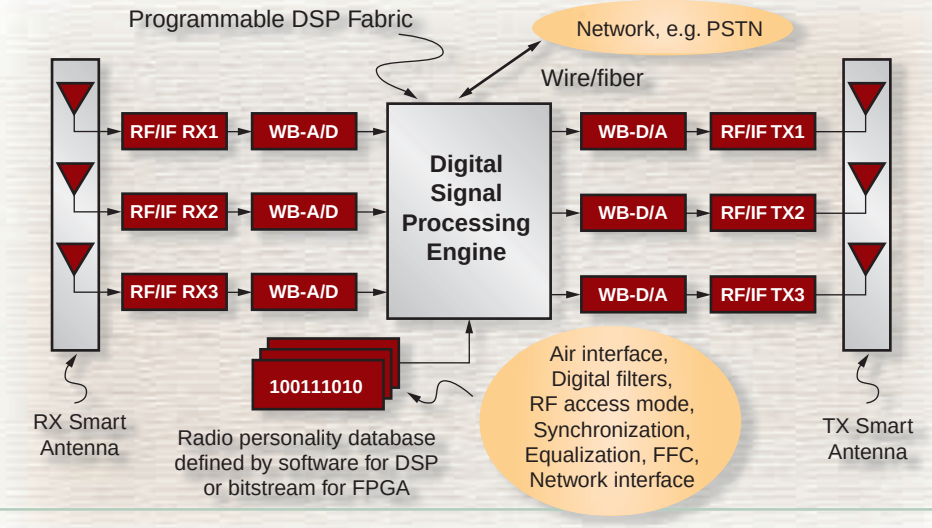

Figure 1 - SDR system diagram

However, the most common approach is to use a multiply-accumulate (MAC) unit. Depending on the performance required, you may need multiple MAC units for your filter design. The Virtex-II architecture is ideal for implementing multiple MACs. The performance of Virtex-II-based DSPs already exceeds 128 billion MACs per second. This is significantly higher than that of conventional DSPs available from mainstream DSP suppliers. Virtex-II FPGAs can extend this performance capability to 600 billion MACs per second because of the highly parallel architecture within the device.

In addition to implementing billions of MACs per second, Virtex-II FPGAs give you virtually complete control of the silicon. This enables you to decide how much real estate to allocate for the MAC units and determine the necessary performance/area trade-offs that all designs require. You can combine the MAC units, used to create the digital filter within the SDR solution, with additional Virtex-II resources to create a sophisticated, high-performance DSP engine.

Additional DSP-enhancing features include configurable dual-port block memories, distributed RAM, and multiplier arrays. The block and distributed memories are ideal for storing large amounts of data, required in the calculations for wireless standards. The multiplier array allows the system to have anywhere from 4 to 192 multipliers. Therefore, you can implement complex SDR designs in very compact solutions. The flexibility of the Virtex fabric enables you to replace multiple DSP processors, often referred to as a DSP farm, with a few Virtex-II devices. Not only does this reduce board complexity, but it also reduces power consumption.

Flexibility

One of the key aspects of an SDR system is that the same hardware can support multiple standards. Additionally, the continuing evolution of communication standards and the competitive pressures in the marketplace dictate that you must start your design and development while standards are still fluid. The flexible Virtex-II fabric enables you to control, integrate, and adapt critical DSP algorithms within the signal-processing engine.

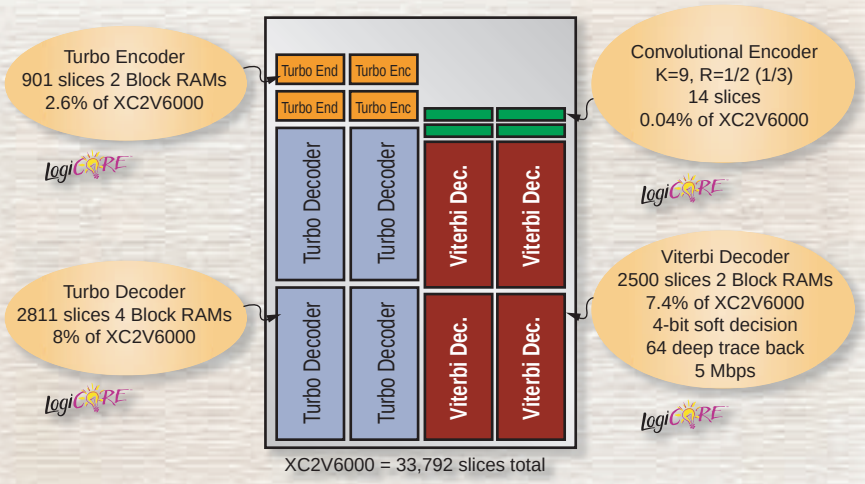

Figure 2. Virtex-II Integrates a 3G FEC Farm

As the SDR standard develops, more and more of the design will become digital. Using an FPGA, you can optimize and incorporate more and more of the digital functionality without adding more hardware – your system architecture grows and develops as the standards grow and develop. Additionally, you aren't locked into one design and you don't need go through the risks, the costs, and the long delays of respinning an ASIC solution.

The Complete SDR Solution

In addition to the Virtex-II platform FPGA, Xilinx is expanding the XtremeDSP™ Initiative to support SDR. As part of the initiative, Xilinx and its partners are providing Virtex-II development boards and a wide range of pre-engineered DSP-related algorithms including a Viterbi Decoder and a Turbo Codec.

Xilinx also offers a complete solution for general DSP development and has made significant advances to dramatically shorten the FPGA DSP design cycle. To enhance productivity, Xilinx has over 70 DSP algorithms (cores) that you can use to shorten development time and accelerate time-to-market. These solutions enable you to get to market quickly using proven technology. A complete list can be found on the Xilinx IP Center at *www.xilinx.com/ipcenter*.

Two examples of productivity enhancing DSP tools are The MathWorks System Generator for Simulink and the Filter Generator. The System Generator bridges the gap between FPGA and conventional DSP design flows. It enables you to develop DSP designs using the familiar Simulink/MATLAB tools. The Filter Generator is unlike any other filter tool in the market. Not only does it speed up the otherwise lengthy process of filter design, but it also has the ability to generate the most optimal area for a given performance level or degree of parallelism.

Conclusion

The Xilinx Virtex-II solution is uniquely positioned to meet the digital signal processing requirements demanded by SDR. Through the XtremeDSP Initiative, Xilinx provides the performance, flexibility, and productivity you need. The Virtex-II architecture delivers the platform for creating high performance, flexible, and upgradeable SDR systems. Parameterizable algorithms, partnerships with industry leaders, and third party development boards enable you to get to market quickly by using proven technologies. Additional details are available on the Xilinx DSP website at: *www.xilinx.com/dsp/*.

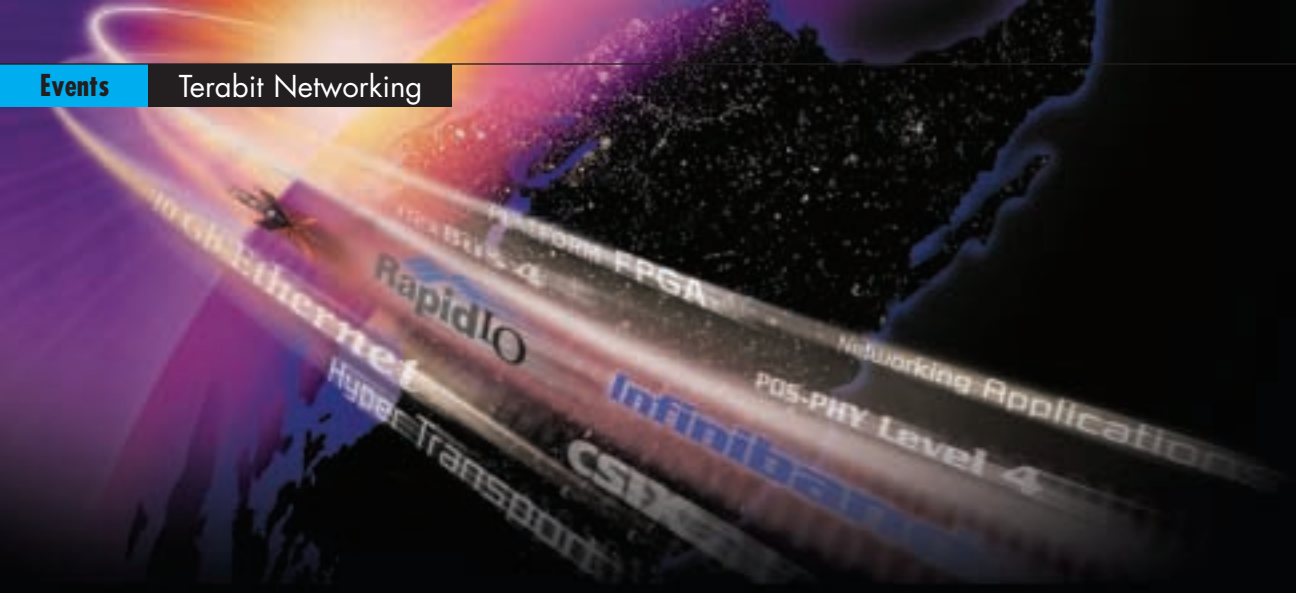

The Terabit Networking Forum

Xilinx helps you meet the challenges of building next generation communication systems.

Terabit Networking Forum

by Abhijit Athavale
System Interface Solutions Marketing Manager
abhijit.athavale@xilinx.com

In response to the networking industry's need for the latest news and technical information, Xilinx and Cahners Electronics Group, in association with other networking industry leaders, recently hosted the Terabit Networking Forum – a premier gathering of visionaries and experts who addressed the challenges and opportunities associated with today's evolving interconnect technologies. Over 2500 engineers and system architects attended, as industry experts provided insight into the future of interface standards such as InfiniBand™, RapidIO™, HyperTransport™, 3GIO, PCI-X, CSIX, 10 Gb Ethernet, and Packet-over-SONET technologies.

Here's an overview of what was discussed.

Defining the Challenge

Consumers' increasing appetite for continuous Internet access and multimedia-rich Web applications are creating a huge demand for fast and efficient processing of data across wireless or wired media. In a typical terabit system, the bottleneck is at the system interconnect level – the speed at which various components inside the box communicate with each other, as well as with the outside world. Clearly, faster CPUs and optimized network packet processors, along with an efficient, fast data transfer mechanism, are required to solve this problem. System architects are also looking for the ability to support multiple standards, to "future proof" their products against constantly evolving interconnect standards, and to integrate many of the functions performed by various chips on a board.

With this rapidly increasing emphasis on connectivity, integration, flexibility, and re-

HyperTransport or RapidIO – who will win this battle of emerging system interfaces?

Gabriele Sartori of AMD explains the HyperTransport technology.

Sam Fuller of Motorola gives his views on system interconnect technologies and RapidIO in particular.

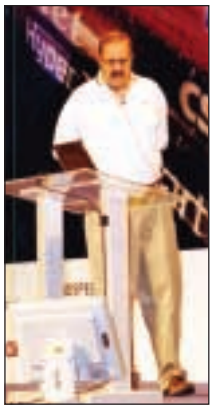

Jim Pappas

Jim Pappas of Intel explains the role of InfiniBand in Terabit Networks

Solving the Connectivity Problem

The PCI local bus has been the ubiquitous system interconnect standard for the past few years. From PCs to networking to communication systems, PCI has supported systems requirements quite satisfactorily and will still be used quite heavily in applications that do not require large quantities of data to be transferred quickly. However, PCI technology – with its centralized arbitration model, limited reliability, and limited scalability – cannot effectively meet the requirements of terabit applications.

Tracy Vanik of RHK explains why 10 Gb Ethernet is the hottest thing in town.

programmability, Xilinx FPGAs are now a mainstream solution and the central component in many new communication systems. Our Virtex-II series platform FPGAs include 3.125 Gbps serial I/O capability and a portfolio of high-performance IP cores that help you quickly implement complex functions such as PCI-X, RapidIO, 10 Gb Ethernet, CSIX, and POS PHY4, HyperTransport, and many more. In short, Virtex FPGAs are ideal for interconnecting communications systems.

Attendees viewing the technology demonstrations offered by participating companies.

Crowd enjoying the afternoon break.

Lauri Vickers, the panel moderator, listens as Elie Massabki from Mindspeed explains how multi-gigabit serial I/O will be used to build terabit systems.

Lauri Vickers

As Moore's law states, processor speed doubles every 18 months, but the PC bus performance doubles every three years. Clearly, this presents a large gap in performance that can not be overcome just by continually increasing the CPU speeds. Various system interconnect standards, notably InfiniBand, RapidIO, and HyperTransport (in addition to enhancements to the trusted PCI bus) have been proposed in the last couple of years to solve this communication bottleneck. New ones such as the recently announced third-generation I/O standard are still being proposed. It is very early to say which standard will win this battle of system standards and emerge as the new PCI.

Taking the Next Step

All the standards mentioned above and more were discussed at the Terabit Networking Forum. Various presentations and product demonstrations made on-site were recorded and are now available via Video-on-Demand at: *www.xilinx.com/terabit.*

If you are trying to create high-bandwidth systems using next generation interfaces, you should view the video on demand featuring technology overviews and detailed interface discussions by 12 industry experts, including:

- Siva Ananmalay, VP at Nortel, discussing industry requirements for High-bandwidth Network Systems
- Jim Pappas, Director at Intel, on InfiniBand architecture, discussing the serialization of server I/O.
- Gabriele Sartori, Director at AMD, discussing the role of HyperTransport in next-generation high-speed I/O.
- Sam Fuller, Director at Motorola, discussing RapidIO Interconnect Architecture.

You'll find that Xilinx programmable logic technology can satisfy the requirements of today's systems by reducing risk, improving system integration, and improving time-to-market.

Year 2001-2002 Worldwide Xilinx Events Schedule

Year 2001 - 2002 North American Event Schedule

Nov 13-15	EDA: Front-to-Back	San Jose, CA
Dec 10-11	Wind River's Worldwide Developers Conference and Exhibition	Santa Clara, CA
Jan 16-17	Portable Design	Santa Clara, CA
Feb 26-28	Wireless/Portable Symposium	San Jose, CA
March 13-15	Embedded Systems Conference	San Francisco, CA
May 14-16	ICASSP	Orlando, FL
June 10-12	Design Automation Conference	New Orleans, LA
July 15-17	NSREC	Phoenix, AZ
Sept 17-18	Applied Computing Conference	Santa Clara, CA

Year 2001 - 2002 Japanese Event Schedule

Nov 2001	MST Fair 2001	Tokyo, Japan
Jan 24-25	Electronic Design and Solution Fair	Yokohama, Japan

For more information about Xilinx Worldwide Events, please contact one of the following Xilinx team members, or see our Web site at: *www.xilinx.com/company/events.htm*

- **North American Shows:** Jennifer Waibel at: *jennifer.waibel@xilinx.com*
- **European Shows:** Andrea Fionda at: *andrea.fionda@xilinx.com* or Andrew Stock at: *andrew.stock@xilinx.com*
- **Japanese Shows:** Yumi Homura at: *yumi.homura@xilinx.com*
- **Asia Pacific Shows:** Mary Leung at: *mary.leung@xilinx.com*

Terabit Networking Forum

August 21, 2001

Santa Clara Convention Center

DID YOU MISS IT?

The most widely viewed high performance interface discussion.

See it all now with Video on Demand

www.xilinx.com/terabit

Brought to you by the whole industry.

…ts joined together

…plied to Optical Networks,

…port, Network Servers, Gigabit and Terabit

…rs, and more.

…evolution of these

…eration system overviews,

…s, and silicon solutions required for the

…opment of the high-speed Internet.

Simulcast LIVE across North America

The Terabit Networking Forum will be simulcast live to multiple cities across North America including Los Angeles, Dallas, Raleigh, Boston, and Ottawa. Register today to reserve your seat in the city nearest you!

Seating is limited. Register now at www.xilinx.com/terabit

XILINX®
The Programmable Logic Company℠

Cahners ELECTRONICS GROUP

MINDSPEED Skyrail · MOTOROLA digital dna · PMC PMC-SIERRA · RHK · intel · AMD · Infineon · Mentor Graphics · AMCC · COMPAQ

Improve Your Time to Market with Online Solutions from Xilinx

You've got the Xilinx silicon, you've got the Xilinx software – now get the technical support you need online from Xilinx Education Services.

by Kristi Gloppen
Manager, eBusiness Marketing
kristi.gloppen@xilinx.com

Over the past several years, Xilinx has pioneered a number of online solutions. Now, we have taken the leadership role in providing you with complete online design solutions. Xilinx online solutions improve your time to market by offering an easy-to-use streamlined process that takes you all the way from design concept to production.

With our online solutions, Xilinx is focused on eliminating downtime and enhancing productivity as you progress through your design cycle. Online solutions from Xilinx give you the means to attain education and support, device selection tools, the fastest programmable logic software on the market, and online purchase.

Let Your Fingertips Do the Learning

The Xilinx desktop solution begins at your desktop – in the discovery stage of your research. In addition to providing detailed device and software information, the product web pages on *www.xilinx.com* educate you and your customers on the benefits of using PLDs. The recently launched Design Tools Center – *www.xilinx.com/ise/xcell/* – focuses on accelerating time to market by providing a consolidated location for education and research on the complete design process. From design entry to verification, the Xilinx Design Tools Center presents links to pertinent sites and walks you through the design flow process.

Xilinx Education Services offers you the most comprehensive suite of online training modules in the industry. Classes range from introductory to advanced levels. The course descriptions and schedules can be found by going to *support.xilinx.com* and selecting the "Education" tab. There, you have the flexibility to set your own training schedule and to begin your education immediately.

Lukose Ninan, hardware group manager of ConSonics Inc., recently took advantage of

our online training: "e-Series I is exactly what I wanted. I get the information I need without having to travel."

Our *support.xilinx.com* website supplies additional resources with "techXclusives," Problem Solvers, an Answer Browser, and Forums, all enabling you to reduce your time to market. "Both the hotline and *support.xilinx.com* are very impressive in providing quick resolutions for technical problems," says Rod Neal, electrical engineer at Harris Corp.

Fastest Access to Design Tools

For instant productivity, Xilinx gives you two free software tools you can access from the Xilinx website. The WebFITTER™ online tool not only has the ability to fit your design into the appropriate Xilinx CPLDs, but it also allows you to convert an existing design for use in a Xilinx solution.

The WebPACK™ software package is the entry-level configuration of the Xilinx ISE (Integrated Software Environment) suite of design tools. It is the fastest free software on the market, and you can download WebPACK modules right onto your desktop PC. The WebPACK modules support all CPLDs and selected FPGA families with up to 300,000 logic gates.

If, after "test driving" the free ISE package, you determine you need additional software features, you can easily purchase all ISE upgrade configurations – including BaseX, Foundation™, and Alliance Series™ software – by clicking the "Buy Online" button on the top navigation bar of *xilinx.com* web pages.

Xilinx software solutions are recognized for helping companies get their products to market faster than the competition. For instance, Robel Borja, a senior electrical engineer at Advanced Bionics Corp., explains: "There were a couple of reasons for going with the Xilinx software tools. First, I was impressed with the software's ability to take a design from another tool and easily convert it to fit a Xilinx device. The software provided a means for conversion and ultimately, saved me a great deal of time. The second reason for going down this path was the technical support I was receiving. The learning curve was short and issues were addressed promptly."

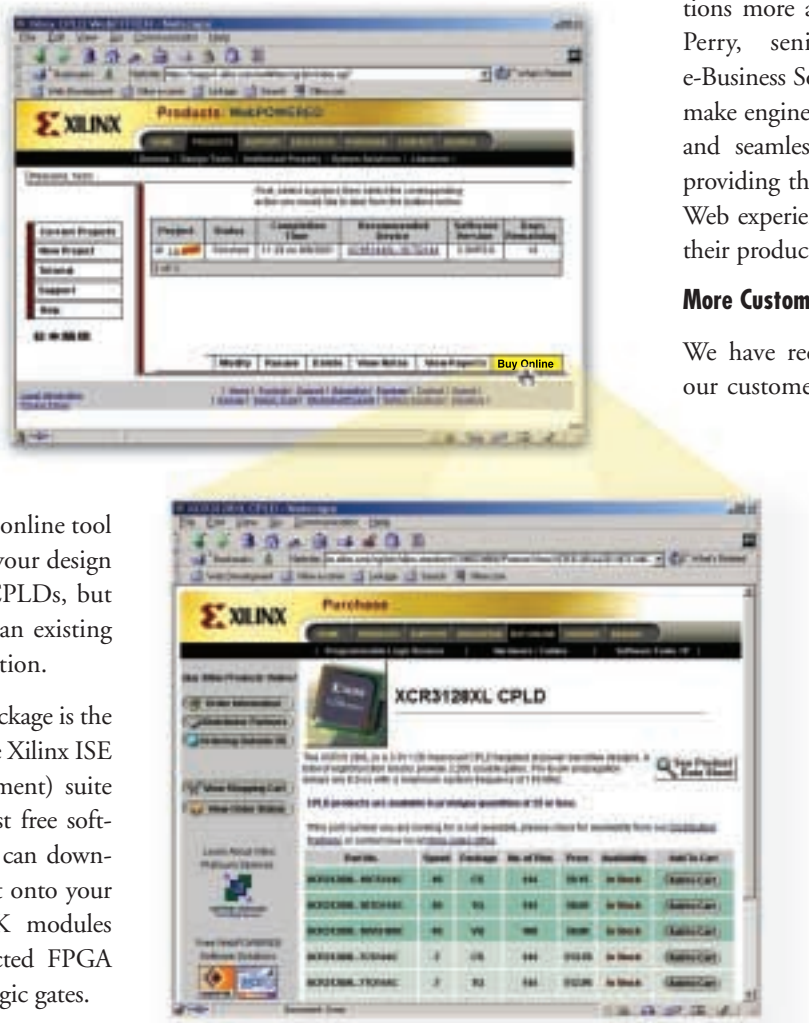

Figure 1 - Online purchasing is only a mouse-click away.

Purchase Products Online

To further accelerate your time to market, Xilinx now sells an increasing variety of products online, including design tools, programming hardware, intellectual property (IP) cores, and programmable logic devices themselves. Most products are shipped the same day, allowing you to get to work on your design, program, or prototype with minimal delay.

As shown in Figure 1, Xilinx has developed an integrated shopping experience that makes moving from discovery to transaction seamless. "The purpose of our eBusiness initiatives is to make our customers more successful by making our products and solutions more accessible to them," says Kenn Perry, senior managing director of e-Business Solutions at Xilinx. "We want to make engineers' design experience as simple and seamless as possible, which includes providing them with personalized, adaptive Web experiences that assist them in getting their products to market faster."

More Customer Endorsements

We have received positive feedback from our customers on this personalized, value-added shopping service:

- *"I am very pleased with the quality and speed of service from the Xilinx e-commerce site,"* reports David C. Brown of the Jodrell Bank Observatory at Manchester University in the United Kingdom.
- *"Great service!"* agrees Kurt Mahan of Silicon Whisper, *"I wish every e-commerce site was as responsive as you guys are, very painless."*
- Chris van Wijhe of Apeldoorn, Netherlands, says, *"A compliment for the Xilinx e-commerce site and the initiative for it. Beautiful!"*

Conclusion

Xilinx can help you get your products to market faster if you utilize our online solutions. Regardless of what phase of your design you are in, Xilinx has online services to help you. Log on to *www.xilinx.com* today and let us assist you in achieving your design goals.

Virtex Series FPGAs

Virtex™-II FPGAs offer you unprecedented capability and flexibility. The first of the new Platform FPGAs, the Virtex-II Series delivers enhanced system memory and lightning-fast DSP through a flexible IP-immersion fabric. Virtex-II devices – with densities ranging from 40,000 up to 10 million system gates – offer new capabilities, including SystemI/O, XCITE, triple DES encryption, and digital clock managers to address system-level design issues.

Additional, mainstream members of the Virtex family are also available.

The Virtex-E family offers high logic gate count, ranging from 50,000 up to 3.2 million system gates. The Virtex-E family supports 20 I/O standards, including LVPECL, LVDS, and Bus LVDS.

The Virtex-EM (Extended Memory) family consists of two devices that have high RAM-to-logic ratios that target specific applications, such as gigabit-per-second network switches and high definition graphics.

The original Virtex family consists of devices that range from 50,000 up to one million logic gates. This family supports 17 I/O standards and offers 5 V PCI compliance.

See www.xilinx.com for more information.

FPGA Product Selection Matrix

		DENSITY							FEATURES							
DEVICES	KEY FEATURES	Logic Cells	Maximum Logic Gates	Typical System Gate Range	Max. RAM Bits	CLB Matrix	CLBs	Flip-Flops	Max. I/Os	Output Drive (mA)	PCI Compliant	1.5 Volt	1.8 Volt	2.5 Volt	3.3 Volt	5.0 Volt
XC2V40	Virtex-II Family: Highest Density/ Ultra Fast Block RAM Distributed RAM System I/O XCITE Up to 12 DCMs	576	6.9K	40K	80K	8x8	64	512	88	2/24	Y	X	–	I/O	I/O	–
XC2V80		1152	13.8K	80K	160K	16x8	128	1024	120	2/24	Y	X	–	I/O	I/O	–
XC2V250		3456	41.5K	250K	480K	24x16	384	3072	200	2/24	Y	X	–	I/O	I/O	–
XC2V500		6912	82.9K	500K	672K	32x24	768	6144	264	2/24	Y	X	–	I/O	I/O	–
XC2V1000		11520	138K	1M	880K	40x32	1280	10240	432	2/24	Y	X	–	I/O	I/O	–
XC2V1500		17280	207K	1.5M	1104K	48x40	1920	15360	528	2/24	Y	X	–	I/O	I/O	–
XC2V2000		24192	290K	2M	1344K	56x48	2688	21504	624	2/24	Y	X	–	I/O	I/O	–
XC2V3000		32256	387K	3M	2176K	64x56	3584	28672	720	2/24	Y	X	–	I/O	I/O	–
XC2V4000		51840	622K	4M	2880K	80x72	5760	46080	912	2/24	Y	X	–	I/O	I/O	–
XC2V6000		76032	912K	6M	3648K	96x88	8448	67584	1104	2/24	Y	X	–	I/O	I/O	–
XC2V8000		104832	1.26M	8M	4480K	112x104	11648	93184	1108	2/24	Y	X	–	I/O	I/O	–
XCV50E	Virtex-E Family: Density Block RAM Distributed RAM SelectI/O 8 DLLs LVDS, BLVDS, LVPECL	1728	21K	47K-72K	88K	16x24	384	1536	176	2/24	Y	–	X	I/O	I/O	**
XCV100E		2700	32K	105K-128K	118K	20x30	600	2400	196	2/24	Y	–	X	I/O	I/O	**
XCV200E		5292	64K	215K-306K	186K	28x42	1176	4704	284	2/24	Y	–	X	I/O	I/O	**
XCV300E		6912	83K	254K-412K	224K	32x48	1536	6144	316	2/24	Y	–	X	I/O	I/O	**
XCV400E		10800	130K	413K-570K	310K	40x60	2400	9600	404	2/24	Y	–	X	I/O	I/O	**
XCV600E		15552	187K	679K-986K	504K	48x72	3456	13824	512	2/24	Y	–	X	I/O	I/O	**
XCV1000E		27648	332K	1,146K-1,569K	768K	64x96	6144	24576	660	2/24	Y	–	X	I/O	I/O	**
XCV1600E		34992	420K	1,628K-2,189K	1062K	72x108	7776	31104	724	2/24	Y	–	X	I/O	I/O	**
XCV2000E		43200	518K	1,857K-2,542K	1240K	80x120	9600	38400	804	2/24	Y	–	X	I/O	I/O	**
XCV2600E		57132	686K	2,221K-3,264K	1530K	92x138	12696	50784	804	2/24	Y	–	X	I/O	I/O	**
XCV3200E		73008	876K	2,608K-4,074K	1846K	104x156	16224	64896	804	2/24	Y	–	X	I/O	I/O	**
XCV405E		10800	130K	1,068K-1,307K	710K	40x60	2400	9600	404	2/24	Y	–	X	I/O	I/O	**
XCV812E	Virtex Extended Memory Capabilities	21168	254K	2,569K-3,062K	1414K	56x84	4704	18816	556	2/24	Y	–	X	I/O	I/O	**
XCV50		1728	21K	34K-58K	56K	16x24	384	1536	180	2/24	Y	–	–	–	X	*
XCV100	Virtex Family: Density Block RAM Distributed RAM SelectI/O 4 DLLs	2700	32K	72K-109K	78K	20x30	600	2400	180	2/24	Y	–	–	–	X	*
XCV150		3888	47K	93K-165K	102K	24x36	864	3456	260	2/24	Y	–	–	X	I/O	*
XCV200		5292	64K	146K-237K	130K	28x42	1176	4704	284	2/24	Y	–	–	X	I/O	*
XCV300		6912	83K	176K-323K	160K	32x48	1536	6144	316	2/24	Y	–	–	X	I/O	*
XCV400		10800	130K	282K-468K	230K	40x60	2400	9600	404	2/24	Y	–	–	X	I/O	*
XCV600		15552	187K	365K-661K	312K	48x72	3456	13824	512	2/24	Y	–	–	X	I/O	*
XCV800		21168	254K	511K-888K	406K	56x84	4704	18816	512	2/24	Y	–	–	–	X	*
XCV1000		27648	332K	622K-1,124K	512K	64x96	6144	24576	512	2/24	Y	–	–	–	X	*

* I/Os are voltage tolerant
** 5 V tolerant I/Os with external resistor
X = Core and I/O voltage
I/Os = I/O voltage supported

Say hello to a new level of performance: the Spartan™-II family now includes devices with more than 200,000 system gates. You get 100,000 system gates for under $10, at speeds of 200 MHz and beyond, giving you design flexibility that's hard to beat. These low-powered, 2.5V devices feature I/Os that operate at up to 3.3V with full 5V tolerance. Spartan-II devices also feature multiple delay locked loops, on-chip RAM (block and distributed), and versatile I/O technology that supports over 16 high-performance interface standards. You get all this in an FPGA that offers unlimited reprogrammability, and can even be upgraded in the field, remotely, over any network.

Robust Feature Set

- Flexible on-chip distributed and block memory
- Four digital delay-locked loops for efficient chip-level/board-level clock management
- SelectI/O™ technology for interfacing with all major bus standards such as HSTL, GTL, SSTL, and so on
- Full PCI compliance
- System speeds over 200 MHz

Extensive Design Support

- Complete suite of design tools
- Extensive core support
- Compile designs in minutes

Advantages over ASICs

- No costly NRE charges
- No time consuming vector generation needed
- 100% tested by Xilinx
- Field upgradable (remotely upgradable, using Xilinx Online technology)
- No lengthy prototype or production lead times
- Priced aggressively against comparable ASICs

See www.xilinx.com for more information.

FPGA Product Selection Matrix

		DENSITY							FEATURES					
DEVICES	KEY FEATURES	Logic Cells	Maximum Logic Gates	Typical System Gate Range	Max. RAM Bits	CLB Matrix	CLBs	Flip-Flops (1)	Max. I/O	Output Drive (mA)	PCI Compliant	2.5 Volt	3.3 Volt	5.0 Volt
XCS05	Spartan Family: High Volume ASIC Replacement/ High Performance/ SelectRAM Memory	238	3K	2K-5K	3K	10x10	100	360	77	12	Y	–	–	X
XCS10		466	5K	3K-10K	6K	14x14	196	616	112	12	Y	–	–	X
XCS20		950	10K	7K-20K	13K	20x20	400	1120	160	12	Y	–	–	X
XCS30		1368	13K	10K-30K	18K	24x24	576	1536	192	12	Y	–	–	X
XCS40		1862	20K	13K-40K	25K	28x28	784	2016	205	12	Y	–	–	X
XCS05XL	Spartan-XL Family: High Volume ASIC Replacement/ High Performance/ SelectRAM Memory	238	3K	2K-5K	3K	10x10	100	360	77	12/24	Y	–	X	*
XCS10XL		466	5K	3K-10K	6K	14x14	196	616	112	12/24	Y	–	X	*
XCS20XL		950	10K	7K-20K	13K	20x20	400	1120	160	12/24	Y	–	X	*
XCS30XL		1368	13K	10K-30K	18K	24x24	576	1536	192	12/24	Y	–	X	*
XCS40XL		1862	20K	13K-40K	25K	28x28	784	2016	224	12/24	Y	–	X	*
XC2S15	Spartan-II Family: High Volume Block RAM Distributed RAM SelectI/O 4 DLLs	432	5K	5K-15K	22K	8x12	96	384	86	2/24	Y	X	I/O	*
XC2S30		972	12K	12K-30K	36K	12x18	216	864	132	2/24	Y	X	I/O	*
XC2S50		1728	21K	21K-50K	56K	16x24	384	1536	176	2/24	Y	X	I/O	*
XC2S100		2700	32K	32K-100K	78K	20x30	600	2400	196	2/24	Y	X	I/O	*
XC2S150		3888	47K	47K-150K	102K	24x36	864	3456	260	2/24	Y	X	I/O	*
XC2S200		5292	64K	64K-200K	130K	28x42	1,176	4704	284	2/24	Y	X	I/O	*

* I/Os are tolerant
X = Core and I/O voltage
I/O = I/O voltage supported

(1) Does not include I/O flip-flops for Spartan-II devices.

XC9500 and CoolRunner CPLDs

From high-speed networking to power-conscious portable designs, Xilinx CPLDs give you a complete range of value oriented products.

The **XC9500™** Families deliver industry-leading speeds, while giving you the flexibility of an enhanced customer-proven pin-locking architecture along with extensive IEEE 1149.1 JTAG Boundary-Scan support. In addition, the XC9500 family easily complements our higher density FPGAs, providing a total, seamless logic solution.

- **XC9500XV** – Offers 30% lower power consumption than 3.3V CPLDs, resulting in lower heat dissipation and increased long-term device reliability.
- **XC9500XL** – With state-of-the-art pin-locking architecture and 5ns pin to pin speed, the XC9500XL is targeted for leading-edge systems that require rapid design development, longer system life, and robust field upgrade capability.
- **XC9500** – The architecture is rich, including individual p-term output enables, three global clocks, and more p-terms per output than any other CPLD.

The **CoolRunner™** Family is targeted for low power applications that include portable, handheld, and power sensitive designs. Each member of the XPA3 family includes Fast Zero Power™ (FZP) design technology that combines low power and high speed. With this design technique, the XPLA3 family offers true pin-to-pin speeds of 5.0ns, while simultaneously delivering power that is <100µA (standby).

- **WebPOWERED™ Software Solutions** – Offers you the flexibility to target Xilinx CPLD and FPGA products online or on the desktop, including:
- **WebFITTER™** – Offers you an online device-fitting and evaluation tool that accepts HDL, ABEL, or netlist files, and provides all reports, simulation models, and programming files, along with price quotes. Available to support all Xilinx CPLD products.
- **WebPACK™ ISE** – Offers downloadable desktop solutions that offer free CPLD and FPGA software modules for ABEL/HDL synthesis and simulation, device-fitting, and JTAG programming.

Through leading performance, free Internet-based WebPOWERED software, and the industry's lowest power consumption, Xilinx has the right CPLD solution for every designer's need.

See www.xilinx.com/products/cpldsolutions for more information.

CPLD Product Selection Matrix

Core Voltage	CPLD Family	Device	Key Features	Density: Macrocells	Density: Max. I/O	Features: Pin-to-Pin Delay (ns)	Features: System Frequency	Features: Individual OE Ctrl	Features: JTAG	Features: Ultra Low Power
2.5 VOLT ISP	XC9500XV	XC9536XV	Best Pin-Locking JTAG w/Clamp High Performance High Endurance	36	36	3.5	278	√	√	–
		XC9572XV		72	72	4	250	√	√	–
		XC95144XV		144	117	4	250	√	√	–
		XC95288XV		288	192	5	222	√	√	–
3.3 Volt ISP	XC9500XL	XC9536XL	Best Pin-Locking JTAG w/Clamp High Performance High Endurance	36	36	5	178	√	√	–
		XC9572XL		72	72	5	178	√	√	–
		XC95144XL		144	117	5	178	√	√	–
		XC95288XL		288	192	6	208	√	√	–
	CoolRunner XPLA3	XCR3032XL	Ultra Low Power JTAG Increased Logic Flexibility	32	36	5	200	–	√	√
		XCR3064XL		64	68	6	145	–	√	√
		XCR3128XL		128	108	6	145	–	√	√
		XCR3256XL		256	164	7.5	140	–	√	√
		XCR3384XL		384	220	7.5	127	–	√	√
		XCR3512XL		512	260	7.5	127	–	√	√
5 Volt ISP	XC9500	XC9536	Best Pin-Locking JTAG High Endurance	36	34	5	100	√	√	–
		XC9572		72	72	7.5	83.3	√	√	–
		XC95108		108	108	7.5	83.3	√	√	–
		XC95144		144	133	7.5	83.3	√	√	–
		XC95216		216	166	10	66.7	√	√	–
		XC95288		288	192	10	66.7	√	√	–

XC18V
XC17V
XC17S

FPGA Configurations PROM

Xilinx offers a full range of configuration memory devices optimized for use with Xilinx FPGAs. Our PROM product lines are designed to meet the same stringent demands as our high-performance FPGAs, taking full advantage of the same advanced processing technologies. In addition, they were developed in close cooperation with Xilinx FPGA designers for optimal performance and reliability.

XC18V00 – Our in-system reprogrammable family provides a feature-rich, fast configuration solution available today, and provides a cost-effective method for reprogramming and storing large Xilinx FPGA bitstreams. This family is JTAG ready and Boundry-Scan enabled for exceptional ease-of-use, system integration, and flexibility.

XC17V00/XC17S00 – Our low cost XC17V and XC17S families are an ideal configuration solution for cost-sensitive applications. XC17V PROMs are pin-compatible with our XC18V family to allow for a cost-reduction migration path as your production volumes increase. The XC17S family is specially designed to provide a low cost, integrated solution for our Spartan families of FPGAs.

Configuration PROMs for Virtex-E/Virtex-EM

Device	Configuration Bits	XC17xx Solution	XC18Vxx Solution	8-pin TSOP	20-pin PLCC	20-pin SOIC	44-pin PLCC	44-pin VQFP
XCV50E	630,048	17V01	18V01	X*	X	X	–	X**
XCV100E	863,840	17V01	18V01	X*	X	X	–	X**
XCV200E	1,442,106	17V02	18V02	X*	X*	X*	X**	X**
XCV300E	1,875,648	17V02	18V02	–	X*	–	X	X
XCV400E	2,693,440	17V04	18V04	–	X*	–	X	X
XCV405E	3,430,400	17V04	18V04	–	X*	–	X	X
XCV600E	3,961,632	17V04	18V04	–	X*	–	X	X
XCV812E	6,519,648	17V08	2 of 18V04	–	–	–	X	X
XCV1000E	6,587,520	17V08	2 of 18V04	–	–	–	X	X
XCV1600E	8,308,992	17V08	2 of 18V04	–	–	–	X	X
XCV2000E	10,159,648	17V16	2 of 18V04	–	–	–	X	X
XCV2600E	12,922,336	17V16	3 of 18V04 + 18V512	–	X***	–	X**	X
XCV3200E	16,283,712	17V16	4 of 18V04	–	–	–	X	X

* Available in XC17Vxx only.
** Available in XC18Vxx only.
*** Available in XC18V512 only.

Configuration PROMs for Virtex

Device	Configuration Bits	XC17xx Solution	XC18Vxx Solution	8-pin TSOP	20-pin PLCC	20-pin SOIC	44-pin PLCC	44-pin VQFP
XCV50	559,200	17V01	18V01	X*	X	X	–	X**
XCV100	781,216	17V01	18V01	X*	X	X	–	X**
XCV150	1,041,096	17V01	18V01	X*	X	X	–	X**
XCV200	1,335,840	17V01	18V02	X*	X*	X*	X**	X**
XCV300	1,751,808	17V02	18V02	–	X*	–	X	X
XCV400	2,546,048	17V04	18V04	–	X*	–	X	X
XCV600	3,607,968	17V04	18V04	–	X*	–	X	X
XCV800	4,715,616	17V08	18V04 + 18V512	–	X***	–	X**	X
XCV1000	6,127,744	17V08	18V04 + 18V02	–	–	–	X	X

* Available in XC17Vxx only.
** Available in XC18Vxx only.
*** Available in XC18V512 only.

Configuration PROMs for Spartan-XL/Spartan-II

Device	PROM Solution	8-pin PDIP	8-pin VOIC	20-pin SOIC	44-pin VQFP
XCS05XL	XC17S05XL	X	X	–	–
XCS10XL	XC17S10XL	X	X	–	–
XCS20XL	XC17S20XL	X	X	–	–
XCS30XL	XC17S30XL	X	X	–	–
XCS40XL	XC17S40XL	X	X	X	–
XC2S15	XC17S15A	X	X	X	–
XC2S30	XC17S30A	X	X	X	–
XC2S50	XC17S50A	X	X	X	–
XC2S100	XC17S100A	X	X	X	–
XC2S150	XC17S150A	X	X	X	–
XC2S200	XC17S200A	X	X	–	X

QML-Certified FPGAs and PROMs

The Xilinx QPro™ family of radiation hardened FPGAs and PROMs are finding homes in many new satellite and space applications. Both the XQR4000XL and XQVR Virtex™ products are being designed into space systems that will use reconfigurable technology. Numerous communications and GPS satellites, space probes, and shuttle missions are included on the growing list of programs that will be flying these devices.

The Virtex QPro family of high reliability products is experiencing a high degree of success in the defense market. As designers find it more and more difficult to find components suitable for the harsh environments seen by defense systems, they are discovering that they can incorporate the functions of obsolete parts into Virtex QPro products. This has the added long term advantage of significantly reducing the costs of future requalifications, because their systems can retain consistent form, fit, and function through the use of Virtex QPro FPGAs. This cannot be achieved with costly and inflexible ASICs or custom logic.

Please visit www.xilinx.com/products/hirel_qml.htm for all the latest information about these products, including some new applications notes.

FPGA Product Selection Matrix

		DENSITY							FEATURES						
Device	Key Features	Logic Cells	Maximum Logic Gates	Typical System Gate Range	Max. RAM Bits	CLB Matrix	CLBs	Flip-Flops	Max. I/O	Output Drive (mA)	PCI Compliant	1.8 Volt	2.5 Volt	3 Volt	5 Volt
**XQR/XQ4013XL	XC4000 Series: Density Leadership/ High Performance/ SelectRAM Memory	1,368	13K	10K-30K	18K	24x24	576	1,536	192	12/24	Y	–	–	X	*
**XQR/XQ4036XL		3,078	36K	22K-65K	42K	36x36	1,296	3,168	288	12/24	Y	–	–	X	*
**XQR/XQ4062XL		5,472	62K	40K-130K	74K	48x48	2,304	5,376	384	12/24	Y	–	–	X	*
XQ4085XL		7,448	85K	55K-180K	100K	56x56	3,136	7,168	448	12/24	Y	–	–	X	*
XQV100	Virtex Family: Density/ Performance Leadership BlockRAM Distributed RAM SelectI/O 4 DLLs	2,700	32K	72K-109K	78K	20x30	600	2,400	180	2/24	Y	–	X	I/O	*
**XQVR/XQV300		6,912	83K	176K-323K	160K	32x48	1,536	6,144	316	2/24	Y	–	X	I/O	*
**XQVR/XQV600		15,552	187K	365K-661K	312K	48x72	3,456	13,824	512	2/24	Y	–	X	I/O	*
**XQVR/XQV1000		27,648	332K	622K-1,124K	512K	64x96	6,144	24,576	512	2/24	Y	–	X	I/O	*

* I/Os are tolerant
** XQR and XQVR devices are radiation hardened
X = Core and I/O voltage
I/O = I/O voltage supported

(1) Selected XQ4000E/EX devices also available

FPGA Product Selection Matrix

Device	System Gates	Logic Gates	CLB Array	Logic Cells	Differential I/O Pairs	User I/O	BlockRAM Bits	Distributed RAM Bits
XQV600E	985,882	186,624	48x72	15,552	247	512	294,912	221,184
XQV1000E	1,569,178	331,776	64x96	27,648	281	660	393,216	393,216
XQV2000E	2,541,952	518,400	80x120	43,200	344	804	655,360	614,400

QPro QML-Certified PROMs

Device	Density	Package			
		DD8	SO20	CC44	VQ44
XC1736D	36Kb	X			
XC1765D	64Kb	X			
XC17128D	128Kb	X			
XC17256D	256Kb	X			
XQR/XQ1701L*	1Mb		X	X	
XQR/XQ18V04*	4Mb			X	X**

* XQR devices are radiation hardened.
** XQ devices only.

Xilinx Intellectual Property Solutions

Cut Development Time and Reduce Design Risk with Xilinx IP Solutions

Xilinx offers hundreds of Intellectual Property (IP) Solutions (cores) and reference designs for a wide variety of industries and applications. By using our IP search engine on the Xilinx IP-Center (*www.xilinx.com/ipcenter* – the industry's most comprehensive Web portal for IP), you can easily find and often download, the right IP for your specific application. You'll find:

- **General purpose IP** such as parameterized memory and multiplier generators. Most of these IP blocks are included at no extra cost with the Xilinx software tools
- **Processor IP** such as the MicroBlaze processor, CoreConnect bus architecture, and peripherals
- **High-speed DSP IP** such as FFT/iFFTs, parameterized filter generators, MACs, and DCTs
- **Connectivity IP** such as PCI/PCI-X, POS PHY L3/L4, and FlexBus 4
- **Application-specific IP** such as 1Gb/10Gb Ethernet MACs, parameterized forward error correction, and voice codecs.

Complete System Solutions

Xilinx provides complete solutions that include IP products as well as software tools, development boards, reference designs, customer training, and special partnerships. These solutions give you access to the tools and industry-leading experts that can help you complete your design on time.

- **The Xilinx XtremeDSP™ solution** is the industry's fastest programmable DSP solution. It delivers Xtreme performance through tremendous parallel processing capabilities, Xtreme flexibility through an SRAM-based reconfigurable technology, and Xtreme productivity through exclusive FPGA partnership with The MathWorks that allows you to use familiar system-level DSP design tools.
- **The Xilinx EmPower! processor solution** offers you the capability to implement single- or multiple-processor solutions on Platform FPGAs so you get the highest possible computational power, bandwidth, and flexibility. The solution includes MicroBlaze™ – the industry's fastest soft processor – and the IBM PowerPC hard processor featured in the next generation Virtex-II FPGAs. Both processors are supported by common peripherals and the CoreConnect on-chip bus architecture.
- **The Xilinx SystemIO capability** is the industry's most complete interconnect solution, giving you a solution for most leading high-bandwidth interface standards. By using the Virtex-II SelectI/O-Ultra technology, a unique flexible physical interface, and the Xilinx pre-verified interface cores, you can quickly design for many different interface standards.

Access to Industry Leading Expertise

To support you through the entire design process, there are several partnership programs available:

- **The AllianceCORE Program** is designed to produce a broad selection of industry-standard solutions. Xilinx partners with independent third-party core developers and takes an active role to select the right cores, and raise the quality and usability of the cores that are offered.
- **The XPERTS™ Partner Program** is worldwide, providing more than 70 consultants certified in delivering turnkey system designs for the Xilinx architecture, including IP customization and integration.
- **The Reference Design Alliance Program** supports the development of third-party, system-level reference designs.
- **The SignOnce license agreement** is a set of common terms being accepted by Xilinx and the industry's leading providers of soft IP cores for use in programmable logic. With this agreement, you can review and execute a single license to ease access to FPGA cores from Xilinx and many of our IP partners.

Unique Features and Development Tools

The Xilinx IP Solution delivers significant benefits through a number of development tools and features:

- **Smart-IP™ technology** uses unique architectural features in the FPGA, along with effective implementation constraints in the Xilinx software tools, to automatically deliver the highest possible performance, predictability, and flexibility. Smart-IP technology gives you both reduced cost and faster time to market.
- **The CORE Generator™ software** enables easy cataloging and generation of parameterized cores that are high performance, predictable, and area optimized; it is included with the Xilinx ISE 4.1i software.
- **The System Generator for Simulink** allows you to use the popular MATLAB® and Simulink® tools from the MathWorks to model a DSP system, and then automatically generate an implementation for Xilinx FPGAs optimized for highest performance and best utilization.
- **Design Reuse** includes the IP Capture tools and the "FPGA Reuse Field Guide" that can be downloaded from the Xilinx IP Center website. The IP Capture tool packages your IP with simulation models, testbenches, and PDF or HTML files, allowing you to catalog and share your IP using the CORE Generator

Visit www.xilinx.com/ipcenter

Upgrade to the Fastest Performance Ever.

The ISE Development Systems Deliver All the Speed You Need …

- **Ultra-Fast Run Times**
- **Superior Device Performance**
- **Increased Productivity**

Xilinx High-Performance Development Systems

The ISE design environment brings you the fastest and easiest family of design tools you can get. ISE combines advanced technologies such as Proactive Timing Closure with a flexible, easy-to-use graphical interface to help you achieve the best possible designs with the least time and effort, regardless of your experience level.

Four Easy-to-Use Configurations

The Integrated Software Environment (ISE) development systems are available in the following configurations:

ISE *WebPACK*
ISE WebPACK™ is the easiest development system to get — it's free and it's on the Web! ISE WebPACK combines support for advanced HDL entry, synthesis, and verification capabilities for both CPLDs and select FPGAs. (Refer to the chart on the back)

ISE *BaseX*
ISE BaseX is the industry's most cost-effective, PC-based programmable logic design environment. The ISE BaseX configuration provides all of the capabilities contained within ISE WebPACK plus additional tools that will help you complete your programmable logic design even faster. (Refer to the chart on the right)

ISE *Alliance*
ISE Alliance is designed to fit into your existing design environment. The tools work seamlessly with those of our EDA partners.

ISE *Foundation*
ISE Foundation™ is a complete, ready-to-use design environment that integrates schematic, synthesis, and verification technology into an intuitive, yet highly advanced design solution.

Xilinx also offers these optional productivity tools:

ModelSim Xilinx Edition (XE)
Provides a complete HDL simulation environment.

ChipScope™ ILA
Enables in-system logic debugging and design verification.

Modular Design
Allows you to divide multi-million gate designs into smaller modules to increase your productivity.

Xilinx System Generator for Simulink®
Allows you to develop high performance DSP systems for Xilinx FPGAs.

See For Yourself

Upgrade now and take advantage of the fastest and easiest development environment in the industry — ISE will save you time and effort in your next design. To get the full details on ISE, go to ***www.xilinx.com/ise4.***

www.xilinx.com/ise4

Development Systems Feature Comparison Guide

	Feature	ISE *WebPACK*	ISE *BaseX*	ISE *Foundation*	ISE *Alliance*
	Platform	PC	PC	PC/UNIX	PC/UNIX
Device Support	Virtex™ Series	Up to 300K (Virtex-E and Virtex-II)	Up to 300K	ALL	ALL
	Spartan™ Series	Spartan-II Series	ALL	ALL	ALL
	XC4000™ Series	No	4KE & Newer	4KE & Newer	4KE & Newer
	CoolRunner™ Series	ALL	ALL	ALL	ALL
	XC9500™ Series	ALL	ALL	ALL	ALL
Design Planning	Modular Design	No	Sold as an Option	Sold as an Option	Sold as an Option
	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 (Gate & Block Level HDL)	Yes	Yes	PC Only	No
	HDL Editor	Yes	Yes	Yes	Yes
	State Diagram Editor	Yes	Yes	PC Only	No
	CORE Generator™	No	Yes	Yes	Yes
	Xilinx System Generator for Simulink	No	Sold as an Option	Sold as an Option	Sold as an Option
Synthesis	Xilinx Synthesis Technology (XST)	Yes	Yes	Yes	No
	Synplify/Pro Integration	No	Yes	PC Only	PC Only
	Leonardo Integration	Yes	Yes	Yes	Yes
	ABEL	CPLD	CPLD	CPLD (PC Only)	No
Implementation Tools	iMPACT	Yes	Yes	Yes	Yes
	FloorPlanner	Yes	Yes	Yes	Yes
	Xilinx Constraints Editor	Yes	Yes	Yes	Yes
	Timing Driven Place & Route	Yes	Yes	Yes	Yes
	Timing Improvement Wizard	No	Yes	Yes	Yes
Board Level Integration	IBIS Models	Yes	Yes	Yes	Yes
	STAMP Models	Yes	Yes	Yes	Yes
	LMG SmartModels	Yes	Yes	Yes	Yes
Verification	HDL Bencher™	Yes	Yes	PC Only	No
	Model*Sim* Xilinx Edition (XE)	Model*Sim* XE Starter Included	Model*Sim* XE Starter Included	Model*Sim* XE Starter Included	Model*Sim* XE Starter Included
	(See www.xilinx.com for more information)	(Model*Sim* XE Sold as an Option)	(Model*Sim* XE Sold as an Option)	(Model*Sim* XE Sold as an Option)	(Model*Sim* XE Sold as an Option)
	Static Timing Analyzer	Yes	Yes	Yes	Yes
	Chipscope ILA	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 Simulator Support	Yes	Yes	Yes	Yes
IP/CORE	For more information on the complete list of Xilinx IP products, visit the Xilinx IP Center at http://www.xilinx.com/ipcenter				

www.xilinx.com/ise4

"I have taken a number of courses offered by Xilinx Education Services and have benefited from their instructors' high levels of expertise, professional presentation and materials. The San José facilities are superb. I highly recommend Xilinx Education Services to anyone who chooses to design with Xilinx FPGAs."

- Jeffrey E. Journey,
Image Capture Development, IBM

Extend your technical capabilities and accelerate your time to market with Xilinx Global Services. Our portfolio of education, support, and design services, along with our award-winning website (support.xilinx.com), will give you the expert help you need to stay ahead of your competition and be the first to market with the most efficient and cost effective designs. We can reduce your learning curve, speed up your design time, jump-start your product development cycle, and get your products to market faster than ever before.

Education Services
Move to the head of the class

Xilinx Education Services keep your technical skills sharp. A broad range of classes are available for designers of all levels, from the novice to the most experienced. Hands-on training classes, led by instructors who are experienced designers themselves, are conducted at the Xilinx headquarters in San Jose and at other sites worldwide throughout the year. The classes cover a wide range of topics, including:

- Using cores
- Using high level design languages
- System and configuration design issues
- Migration from ASIC to FPGA.

support.xilinx.com
Knowledge at your fingertips

Support.xilinx.com is our online solution for resolving your design issues. Here, you'll find:

- Our Answers Database which contains more than 4,000 proven design solutions.
- Problem Solvers to help you troubleshoot device configuration, software installation, and JTAG issues.
- Discussion forums that let you share ideas and questions with other designers.
- A Web support interface that allows you to easily submit a problem, and can get your answer quickly.

Support.xilinx.com is always available to you, 24 hours a day.

Support Services
Experts on Call

Our Gold and Platinum Technical Services improve your productivity and accelerate your design process by reducing your design and troubleshooting time. As a Platinum customer, you receive access to a dedicated toll-free number (in North America only) so you can get quick assistance with any design problem; you have first priority to our dedicated team of skilled senior application engineers - the best in the business. You also get ten education credits for Xilinx training courses to improve your skills. As a Gold customer, you receive our standard level of service, at no charge. With either Platinum or Gold service, our team of professionals is ready to help you use the tools and techniques that have made Xilinx the market leader in FPGA technology.

Design Services
Your Virtual Project Team

With Xilinx Design Services you can take full advantage of our design expertise, so you're free to focus on what you do best. A Xilinx Design Services team - made up of best-in-class designers, silicon experts, and software specialists - becomes your virtual in-house project team, so you can immediately extend your technical staff's bandwidth and eliminate ramp-up time. The Xilinx Design Services team gives you our unique expertise with Xilinx tools and techniques, and the fruits of our extensive R&D investment.

See www.support.xilinx.com for more information on all of our services.

Features	Platinum	Gold**
Senior Application Engineers	✔	
Dedicated Toll-Free Number	✔	
Proactive Status Updates	✔	
Priority Case Resolution	✔	
Ten Education Credits	✔	
Service Packs and Software Updates	✔	✔
Application Engineer/Customer Ratio*	2X Gold Level	Standard

**Applicable in North America Only*
*** Gold is the standard level of service we provide.*

Do this. Do that.
Do this. Do that.

(this)

(that)

(this)

(that)

(this)

(that)

(this)

(that)

The new Spartan®-IIE FPGA series is the most advanced, cost-optimized technology for enabling consumer digital convergence. Supporting 19 I/O standards, including LVDS, the Spartan-IIE FPGAs also deliver nearly one billion MACs/s per dollar of DSP performance.

The Gateway to Consumer Digital Convergence

With devices ranging from 50,000 to 300,000 gates, consumer electronics designers can now enjoy the flexibility and time-to-market advantages of an FPGA with all the system-level integration of an ASIC.

Leading-Edge Software & IP . . . the Complete Solution

Our ISE 4.1i software is the fastest on the market, and works with all product families. Spartan-IIE FPGAs are also supported in our free online ISE WebPACK™. With the widest range of reference designs, soft processors and third-party support available, it's the complete solution only Xilinx can offer.

Visit ***www.xilinx.com/spartan2e2*** today to see how Xilinx is helping consumer electronics designers to do this . . . and do that . . . and do it right first time.

XILINX®

The Programmable Logic Company℠

www.xilinx.com/spartan2e2

FORTUNE® 2001
100 BEST COMPANIES TO WORK FOR

INTRODUCING ISE 4.1i

All the speed you need.

ISE

The fastest programmable design software just got faster

Get ready to fly. Our ISE 4.1i software release completes your design twice as fast as our previous generation, and up to 6 times faster than competitive tools. With an array of industry-unique features, ISE 4.1i gives you unbeatable performance advantages, increasing your design's clock performance by up to 75%, and enabling I/O speeds of 3.125 Gbits/sec.

Performance and Speed with ProActive™ Timing Closure

ISE 4.1i features Xilinx ProActive Timing Closure - a set of new technologies including physical synthesis integration, HDL analysis, timing wizards with cross-probing, and the most advanced timing-driven place and route algorithms available in the industry today. This means ISE 4.1i will intelligently identify, analyze, and solve any potential timing bottlenecks. Also, our new Directed Routing technology delivers unprecedented improvements in performance predictability for IP.

VIRTEX-II

SPARTAN-III

XPLA3 CoolRunner

XC9500XV

A single design environment for *all* Xilinx devices

Combine the superior performance of ISE 4.1i with the speed and density leadership of the Virtex®-II Platform FPGAs, the Spartan®-II ASIC replacement FPGAs, the ultra-low power CoolRunner™ CPLDs, plus our high-performance XC9500™ CPLDs. With the first and only FPGA support for industry-leading formal verification tools, and clock speeds over 300MHz, Xilinx provides the totally integrated design solution.

Visit *www.xilinx.com/ise* today for all the *speed* you need!

XILINX®
The Programmable Logic Company℠

www.xilinx.com/ise/xcell

FORTUNE® 2001
100 BEST COMPANIES TO WORK FOR