ISSUE 49, SUMMER 2004

XCELL JOURNAL

XILINX, INC.

Issue 49
Summer 2004

Xcell journal

THE AUTHORITATIVE JOURNAL FOR PROGRAMMABLE LOGIC USERS

Designing High-Speed Serial Links

SIGNAL INTEGRITY

Issues, Tools, and Methodologies

High-Speed Interconnect

Debugging MGT Designs

Power Distribution Networks

BACKPLANES

Next-Generation Serial Backplanes

Create ATCA-Compliant Designs

Mesh Fabric Switching

DSP

XtremeDSP Kit-II

Increase Image Processing System Performance

Low-Cost Co-Processing

COVER STORY

Celebrating 20 Years of Leadership

XILINX®

Only FPGAs on 90nm
Shipping Since March 2003

The New SPARTAN™-3 Make It Your ASIC

SPARTAN™-3

The world's lowest-cost FPGAs

Spartan-3 Platform FPGAs deliver everything you need at the price you want. Leading the way in 90nm process technology, the new Spartan-3 devices are driving down costs in a huge range of high-capability, cost-sensitive applications. With the industry's widest density range in its class — 50K to 5 Million gates — the Spartan-3 family gives you unbeatable value and flexibility.

Lots of features ... without compromising on price

Check it out. You get 18x18 embedded multipliers for XtremeDSP™ processing in a low-cost FPGA. Our unique staggered pad technology delivers a ton of I/Os for total connectivity solutions. Plus our XCITE technology improves signal integrity, while eliminating hundreds of resistors to simplify board layout and reduce your bill of materials.

Easiest to use Software

With the lowest cost per I/O and lowest cost per logic cell, Spartan-3 Platform FPGAs are the perfect fit for any design … and any budget.

MAKE IT YOUR ASIC

XILINX®
The Programmable Logic Company℠

For more information visit
www.xilinx.com/spartan3

Pb-free devices available now

FORTUNE® 2004 100 BEST COMPANIES TO WORK FOR

20 YEARS

Close Isn't Good Enough Anymore

As I was preparing to write this editorial, I asked myself: "What did I do in the past that was relatively simple then, but has gotten vastly more complicated now?" The answer was tuning up a car's engine.

I've always liked cars. As a teenager, I would get together with my buddies on weekends to extract the finest performance from our machines. We lived for the automotive trinity: high speed, loud sounds, and great looks.

I remember replacing the spark plugs, which were factory-set to a gap clearance specific to my car's engine. However, this factory setting was rarely correct. If the gap was too wide, I tapped the end of the spark plug on the garage floor and remeasured. If it was too tight, I used a screwdriver to spread open the electrode, widening the gap.

Tuning up a car's engine used to be quite easy. I wasn't concerned with tight tolerances – close was good enough. But advances in automotive technology have made it virtually impossible for me to work on my car anymore.

Similarly, advances in PCB technologies pose far more difficult engineering challenges today than they did just a short time ago. Feature size reduction, market demands, and the need for reduced power consumption have driven core voltages down and operating frequencies up. These changes in signal voltage and frequency require new design practices that take into account electrical effects that could previously be ignored.

This issue features a section on signal integrity issues, tools, and methodologies pertaining to high-speed PCB design. We also have a section on end-to-end programmable solutions for line cards and high-speed serial backplanes. Together with many of our partners, Xilinx is addressing these issues to help you resolve the technical difficulties that affect performance, system development, and product introduction schedules.

As the new Managing Editor for *Xcell*, I'd like your feedback on the signal integrity series in this issue, as I endeavor to continually improve the magazine. Please visit our website at *www.xilinx.com/si_xcell.htm*, where you will find a short survey form.

Forrest Couch
Managing Editor

Xcell journal

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

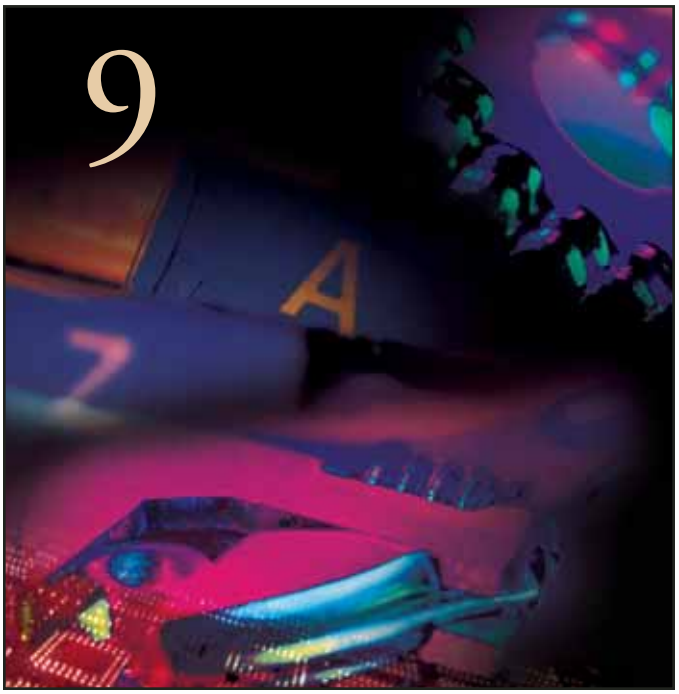

High-Speed PCB Design: Issues, Tools, and Methodologies

In this series on signal integrity, *Xcell* explores tools and methods you can use to combat signal and power integrity distortions throughout product development.

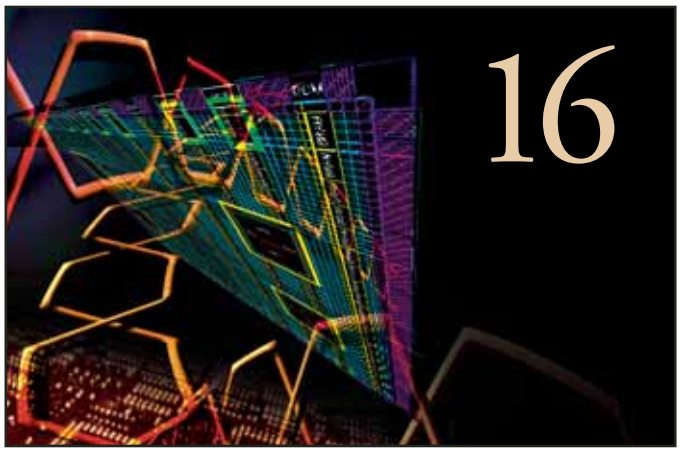

For Synchronous Signals, Timing Is Everything

Mentor Graphics highlights a proven methodology for implementing pre-layout Tco correction and flight time simulation with Virtex-II and Virtex-II Pro FPGAs.

The Next Gold Standard

The Advanced Telecom Compute Architecture standard has great potential for widespread adoption in next-generation infrastructure applications.

COVER STORY

Celebrating 20 Years of Leadership 6

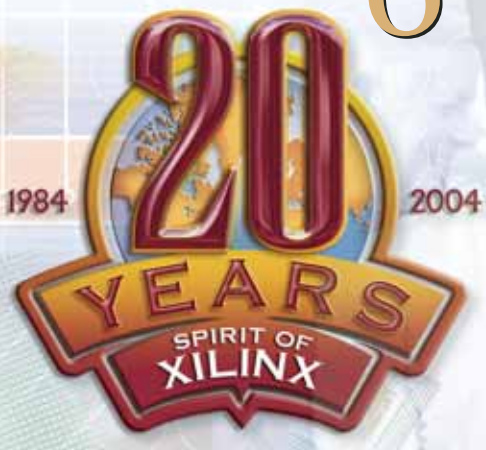

When the Xilinx founders created their first business plan in 1984, they agreed on a lofty goal: *"To be the leading company designing, manufacturing, marketing, and supporting user-configurable logic arrays for the application-specific market."*

SUMMER 2004, ISSUE 49

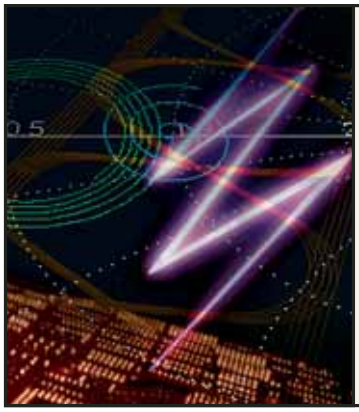

Backplane Characterization Techniques

High-bandwidth measurements of backplane differential channels are critically important for all high-speed serial links.

31

Better... Stronger... Faster

Virtex-II Pro FPGAs offer marked performance advantages over a competing device.

53

Create ATCA-Compliant Designs

Xilinx and Avnet have released a new design kit that reduces time to market for a wide range of serial backplane applications.

65

Secure Your Consumer Design with CoolRunner-II CPLDs

CoolRunner-II CPLDs offer unique features to ensure a more secure design and reduce the risk of reverse engineering.

92

Xcell journal

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

Celebrating 20 Years of Leadership

When the Xilinx founders created their first business plan in 1984, they agreed on a lofty goal: "To be the leading company designing, manufacturing, marketing, and supporting user-configurable logic arrays for the application-specific market."

Xilinx CEO Wim Roelandts in the Xilinx Hall of Patents.

By Xilinx Staff

In 1984 when Xilinx was founded, configurable logic arrays were viewed as exotic curiosities, the semiconductor industry was mired in a slump, and the personal computer – destined to become the driving force in silicon consumption – had just been introduced to skeptical reviews. That's why many people thought that Xilinx founders Ross Freeman, Bernie Vonderschmitt, and Jim Barnett were overly ambitious with their written missive. But the driving force in their plans was the goal of leadership – that sometimes vague, often elusive goal that all high technology companies seek but few ever attain.

Today, everyone in the industry knows that the Xilinx founders made good on their promise. As we enter our third decade as the preeminent supplier of programmable logic devices (owning more than 50 percent of the market), we increasingly find that our technology is the preferred choice for most digital logic designs. By almost any definition, Xilinx is setting a new standard for success.

Indeed, today's stated vision makes our founding fathers' objective seem comparatively tame. As Xilinx celebrates its 20th year in business, our market leadership is unquestioned and our current goal stretches far into uncharted territory: "To put a programmable device in every piece of electronic equipment within the next 10 years." This guiding principle is etched into the mind of every Xilinx

Since inventing the FPGA in 1984, Xilinx has progressively achieved new technological milestones ahead of its competition.

employee around the world, and is the emotional force behind the steady stream of innovation and operational excellence for which Xilinx is known.

Leadership Starts from Within

Talk to our CEO Wim Roelandts about leadership, and you won't hear a lot about market share dominance, a litany of industry firsts, or impressive statistics that typify most companies' definitions of what it means to be a leader. Instead, Wim speaks passionately about core values, management philosophy, corporate culture, and building a legacy. That's why the second Xilinx company goal is: "To build a company that sets a new standard for managing high-tech companies." Xilinx was named The Best Managed Semiconductor Company by *Forbes* magazine in 2004, just one indicator that this goal is now a reality.

Wim's own style draws upon his years of experience at Hewlett-Packard, something of a high-tech pioneer itself in terms of corporate culture with its legendary "HP Way." But he makes it clear that his team's goal for Xilinx is a new, unique style of management: one that combines the best of traditional hard-driving, top-down, win-at-all-costs approaches with "softer," consensus-oriented, people-centric models. And he insists you can have the best of both worlds. "We have a culture where people are treated with respect, where there is consensus management, and still we are a leader. How do we do it? Through innovation! We have a process that fosters innovation, and with innovation comes leadership."

VP of Human Resources Peg Wynn describes the competitive attitude at Xilinx like this: "We're fierce competitors with hearts of gold." That competitive attitude has led to no shortage of innovation and industry firsts at Xilinx during its first 20 years, as more than 900 patents attest. Such a record of achievement is the result of a well-thought-out process to inspire employees to greatness, with a business model that allows the company to focus on what it does best.

A Holistic Management Philosophy

Xilinx leadership is based on its ability to continuously innovate. Therefore, its management philosophy is based on simple tenets:

- People want to do a good job and they come to Xilinx to do their best work
- Work has to have meaning and value
- The company must provide a sense of community
- There must be an opportunity for personal growth
- Everyone should be an owner.

Because of this, a rare team attitude exists at Xilinx that is not often found in the hallways and meeting rooms of other high-technology companies. It meshes with a sense of quiet confidence that pervades the company. In fact, about the only "leadership" statistic that Wim likes to spend any time discussing is an employee retention percentage that is the envy of the industry. "We have set the standard for employee turnover in our industry. It's something like five or six percent, compared to an average in the mid 20s in our business."

Wim talks a lot about the importance of walking the talk, or as he puts it, "maintaining consistency and credibility" with the employee base as well as with the company's other stakeholders: partners, customers, and shareholders. It's one reason why he is fanatical about returning e-mails from employees, and moves his office every year to a new location "to get a different perspective on the company." Such an attitude underpins a sense of values and integrity that has led Xilinx to be voted the "Most Respected Public Company" by its peers in the Fabless Semiconductor Organization (FSA) two years in a row, as well as earn us a top-10 rating in *Fortune* magazine's "Best Places To Work" for the last four years.

Innovation and Leadership

Xilinx has put the structure in place to make all employees and partners successful. It begins with focus. From our inception in 1984, Xilinx strategy has relied on a partnership model through which we develop mutually beneficial relationships with experts in manufacturing, sales, and other activities that are impractical for us to do ourselves. For example, company founder Bernie Vonderschmitt essentially invented the fabless semiconductor model on a handshake agreement with Seiko in 1984. That agreement saw the first Xilinx-designed chips roll off the manufacturing lines at Seiko's plants. Today, Xilinx relies – and in fact, drives forward – our manufacturing partners as we reach new milestones together.

Since 1984, Xilinx has developed an extensive and growing "ecosystem" of partnerships for a wide variety of needs. We partner with experts in sales, design tools, intellectual property cores, and chip design services – a strategy that has allowed an unwavering focus on our own areas of core competence: designing, marketing, and supporting our programmable chips. "You can only be a leader in a few areas so you have to define where you want to be a leader and use partners to complement what you do," says Wim. "We want to be a leader in technology and in innovation. To do that, we need partners and there always has to be something in it for the partner – it has to make them better. Our philosophy on partnerships is that it should minimally be a ratio of 51 to 49, in favor of our partner."

The Xilinx track record of innovation is impressive. Since inventing the FPGA in 1984, Xilinx has progressively achieved new technological milestones ahead of its competition, and set new standards for

semiconductor design. Most recently, we were the first to produce production devices in 90 nm process geometries. Along with Intel™, we are also producing the most chips on state-of-the-art 300 mm wafers – both testaments to the design prowess of our engineering teams. Not content to rest on our laurels or follow trends, Xilinx management proudly points to the ratio of employees working on future business activities: about three-quarters of the company.

"Being a leader means taking risks," says VP of Marketing Sandeep Vij. "And the culture here at Xilinx rewards risk-taking. The whole concept behind our technology – programmability – was based on a giant risk by the founders. That's what inspires innovation. Because of the way we are set up, every employee feels like an owner, people feel like they are part of a team; they're part of something beyond an individual contribution." And with innovation comes leadership, although it's not always an overnight effect.

In fact, Sandeep looks at the first 20 years of Xilinx in two distinct phases. First was the decade that saw the first few generations of products take shape; market adoption of programmable technology happened on a gradual basis. Next came the decade when Xilinx products became more mainstream; new milestones were reached – including one million devices shipped, one billion transistors on a chip, and $1 billion in revenue. "Leadership is different than being a winner," Sandeep notes. "In our view of the world, there can be more than one winner – in fact, that's required because we want our partners to win too. There are a lot of intangibles in being a leader. A leader evokes respect. A leader inspires people to follow. A leader has to look at what's happening today and see its impact on the future. That's what the founders of Xilinx did 20 years ago, and that's what we must continuously do now."

Leading the Way to the Future

Wim likes to call Xilinx a reconfigurable company, a tribute not just to the innovative technology the company delivers to a wide range of electronics companies, but also to the flexible management style that he sees as essential to survival in high technology. "The challenge is in keeping Xilinx nimble and responsive. Every day we change. Whether it's the technology, a business process, or our geographical focus, we have to be comfortable with change. And we have to continue to re-innovate from within." What else would you expect from the company that invented programmable chips?

Original Xilinx Mission Statement 1984

To be the leading company designing, manufacturing, marketing, and supporting user-configurable logic arrays for the application-specific market.

Xilinx's strategies to be the leading company are:

1. Maximize our strengths in product architecture and design
2. Complement our strengths with a long-term fab partner who has high quality, high volume, competitive cost capability, and state-of-the-art process technology
3. Provide a logic solution that is easier to design-in and more cost-effective than SSI/MSI, PALS, and gate arrays with densities of 4,000 to 5,000 unit cells
4. Provide support for all user volumes with both softwired and hardwired products
5. Develop and support design tools to minimize the customer's design efforts

High-Speed PCB Design: Issues, Tools, and Methodologies

In this series on signal integrity, *Xcell* explores tools and methods you can use to combat signal and power integrity distortions throughout product development.

Table of Contents:

by Suresh Sivasubramaniam
Senior Design Engineer
Xilinx, Inc.
suresh.subramaniam@xilinx.com

Philippe Garrault
Technical Marketing Engineer
Xilinx, Inc.
philippe.garrault@xilinx.com

RocketIO™ transceivers are a standard feature on the most advanced Xilinx FPGAs. Our first-generation transceivers operate in the 1-3.125 Gbps bandwidth, while the latest generation transceivers have an operating bandwidth of up to 10 Gbps.

These data rates mean that bit period and signal rise and fall times are extremely small. Designing physical links/channels on a PCB for these high-speed devices with small amplitude and timing budgets necessitates a careful analysis of the possible signal integrity (SI) and power integrity (PI) distortions.

SI/PI issues and reduced amplitude and timing budgets are not limited to just high-speed serial links. In recent years, the amount of logic cells inside Xilinx devices has grown tremendously. Additionally, the pin count on device packages has gone from a few to more than a thousand. Increased I/O performance, in conjunction with the large number of I/Os available, means that for each of your new designs, a lot more transistors are switching more often.

The common denominator in these problems is poor management of the three "bad boys" of electric circuits: resistance, inductance, and capacitance.

Additional requirements for the efficient design of high-speed buses may be dictated by the needs of a specific application. For example, a 266 MHz, 64-bit DDR RAM interface will be sensitive to skew between the different byte lanes. Large parallel buses also have the potential to generate simultaneous switching output (SSO) noise and voltage droop. All of these factors translate into the need to manage the transient current demands of a particular application through proper design of the power distribution system (PDS).

Resistance, Inductance, and Capacitance Pull the Strings

In general, SI and PI issues arise when designers pay inadequate attention to these broad categories:

- Termination schemes
- Skin effect (frequency-dependent attenuation)
- Dielectric losses
- Impedance discontinuities/reflections
- Data coding (DC balanced codes, run length, channel memory)
- Equalization/pre-emphasis
- Inter-symbol interference
- Crosstalk
- Decoupling/bypassing in power distribution
- Board stack-up
- Signal edge rates.

The common denominator in these problems is poor management of the three "bad boys" of electric circuits: resistance, inductance, and capacitance (Figure 1). In addition, you must understand and employ the right measurement techniques in the lab to accurately measure and validate designs against simulations or design specifications.

The objective is to build systems right the first time.

Minimize SI/PI Effects

In this special series on signal integrity, we have assembled articles that will provide you with practical and technical resources towards achieving that goal. From characterization and model extraction techniques in the lab to methods for simulating signal degradations of synchronous parallel/asynchronous serial systems to case studies, this series covers many aspects of SI.

In a sidebar to this article, Xilinx Principal Engineer Austin Lesea lists "Ten Reasons Why Performing SI Simulations is a Good Idea." Although this may sound very familiar to some of you, understanding the benefits of performing SI analysis throughout the design cycle can help you achieve your performance, reliability, and time-to-market goals.

"Interfacing SMA Connectors to Virtex-II Pro MGTs" details Warren Miller and Vince Gavagan's experience designing the interface between Virtex-II Pro™ multi-gigabit transceivers (MGTs) and Sub Miniature version A (SMA) connectors for the Virtex-II Pro Aurora Design Kit. Through prototyping and time domain reflectometry (TDR) measurements, they illustrate how SMA connector choice influences signal quality.

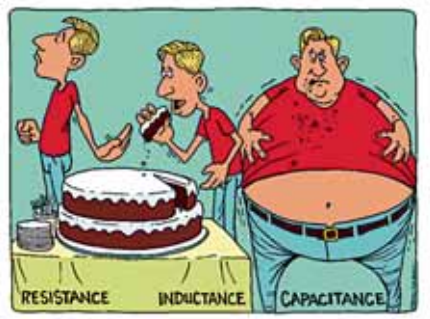

Figure 1 – The three bad boys of electric circuits (Courtesy of Educator's Corner/Agilent Technologies)

Bill Hargin believes that "For Synchronous Signals, Timing Is Everything." His article outlines a method for extracting correction values that can be applied to the clock-to-out and flight time numbers. The resulting timing values in the datasheet are representative of the actual load and topology of your design. This technique specifically applies to source-synchronous links.

Predicting the interconnect performance of high-speed links made of complex via, connector, and trace structures is no easy task. However, as Ansoft's Lawrence Williams explains in "Designing High-Speed Interconnects for High-Bandwidth FPGAs," combining electromagnetic, circuit, and system simulations greatly helps in the design of reliable and fast data transmission channels.

When designing multi-gigabit asynchronous channels, you must carefully analyze the link's physical and electrical properties. In his article, "Accurate Multi-Gigabit Link Simulation with HSPICE," Dr. Scott Wedge explains how the combination of an EM solver, coupled transmission lines, S-parameter support, and SPICE and IBIS modeling to the HSPICE® circuit simulator helps accurately account for high-speed signal distortions.

With "Eyes Wide Open," Steve Baker shows you how to use the RocketIO Design Kit for ICX™ to evaluate pre-layout options (such as placement, connectors, or stackup) as well as post-layout options (such as detailed routing structures) to achieve high-speed serial link performance.

As much as Xilinx recommends SI simulation and analysis before manufacturing a PCB, there are two very valuable lab instruments that you can use on prototype/exploration boards. With these instruments, you can characterize interconnect properties and high-speed signal behavior, explore different topology performances, or extract simulation models. In his article, "Backplane

Characterization Techniques," Eric Bogatin explains the need for making measurements, illustrating the concept of measurement and model bandwidth. He also discusses SMA launches, information contained in TDR traces, and differential S-parameters.

In "A Low-Cost Solution for Debugging MGT Designs," Joel Tan presents a solution comprising a bit-error rate testing module connected to a flexible on-chip logic analyzer core, both implemented in FPGA fabric. Together with the ChipScope™ Pro software suite, these two components allow you to perform diagnostic testing, debugging, and development of an MGT system without the use of expensive lab equipment such as logic analyzers and BERT testers.

And in "Tolerance Calculations in Power Distribution Networks," Sun Microsystems' Istvan Novak walks you through different scenarios of bypass capacitor configurations to demonstrate the importance and influence of the capacitors' technology, value, and number in designing a decoupling/power distribution network.

Conclusion

We hope you will find in this series instructive material on the sources of SI/PI effects, along with practical information about the resources and tools available to you. Our experience tells us that careful simulations, analysis, and measurements of PI and SI effects early in the design process guarantees first-time success more often than not.

If you'd like to send us feedback about the topics discussed, please e-mail us at *si_xcell@xilinx.com* Σ

Ten Reasons Why Performing SI Simulations is a Good Idea

by Austin Lesea
Principal Engineer, Advanced Products Group
Xilinx, Inc.

Not so long ago, the rise and fall times of signals, the coupling from one trace to another, and the de-coupling of power distribution on a PCB were tasks that were routinely handled by a few simple rules. Occasionally, you might use the back of an envelope, scribbling down a few equations to make sure that the design would work.

Those days are gone forever. Sub-nanosecond, single-ended I/O rise and fall times, 3 to 10 Gb transceivers, and tens of ampere power needs at around 1V have all led to increased engineering requirements.

Your choice is simple: simulate now and have a working result on the first PCB, or simulate later after a series of failed boards. The cost of signal integrity tools more than outweighs the cost of making the board over and over with successive failures.

In keeping with the theme of this special issue, here are my 10 best reasons why signal integrity engineering is a good idea:

1. You're tired of making PCBs over and over and still not having them work.

Seriously, without simulating all signals, as well as power and ground, you risk making a PCB that will just not work. IR (voltage) drop, inadequate bypassing or de-coupling, crosstalk, and ground bounce are just a few of the possible problems.

2. You're tired of being late to market and watching your competition succeed.

Every time you have to fix a problem with a PCB, it necessitates a new or changed layout, a new fabrication, and another assembly cycle. It also requires the re-verification of all parameters. Taking the time to do these things right has both monetary and competitive advantages.

3. You're tired of spending all this money, only to scrap the first three versions of PCBs and all of the components that went with them.

See reason number two.

4. Your eye pattern is winking at you.

If the eye pattern of a high-speed serial link is closing, or closed, it's likely that the link has a serious problem and will have dribbling errors – or worse, will be unable to synchronize at all. You must simulate every element of the design to assure an error-free channel.

5. All 1s or all 0s suddenly breaks the system.

Unfortunately, many systems do not have a choice of what data may be processed. Often the data pattern will create conditions that, if not simulated a priori, will cause errors in the system.

6. Hot and cold, fast and slow, and high and low voltages cause failures.

Without simulating the "corners" of the silicon used as well as the environmental factors, you're playing Russian Roulette with five of the six chambers loaded.

7. You cannot meet timing, and you are unable to find out why.

Poor signal integrity is the primary cause of adding jitter to all signals in a design. Ground bounce, crosstalk, and reflections all conspire to add jitter. And once added, jitter is virtually impossible to remove.

8. The FCC Part 15 or VDE EMI/RFI test fails every time you test a board.

Radiated and conducted radio frequency emissions, as well as susceptibility to radio frequency sources, is a sign of poor SI practices. Fixing the problem by shielding increases the system cost substantially, and may not even be possible in extreme cases.

9. Customers complain, but when you get the boards back, you don't find any problems.

One of the biggest problems with SI is that the errors and failures observed are difficult to correlate and sometimes impossible to find. Was it a problem with voltage, temperature, or with the data pattern itself? It might have been someone turning lights on and off (ground disturbance). Don't risk a return that cannot be fixed.

And last, but certainly not the least:

10. Your manager has suggested that you look for other employment.

Do not let this happen to you. Stay current, educated, and productive. Get the right tools to do the job. Realize that signal integrity engineering is a valuable and irreplaceable skill in great demand in today's design environments.

Interfacing SMA Connectors to Virtex-II Pro MGTs

SMA connector choice has a surprising effect on signal integrity.

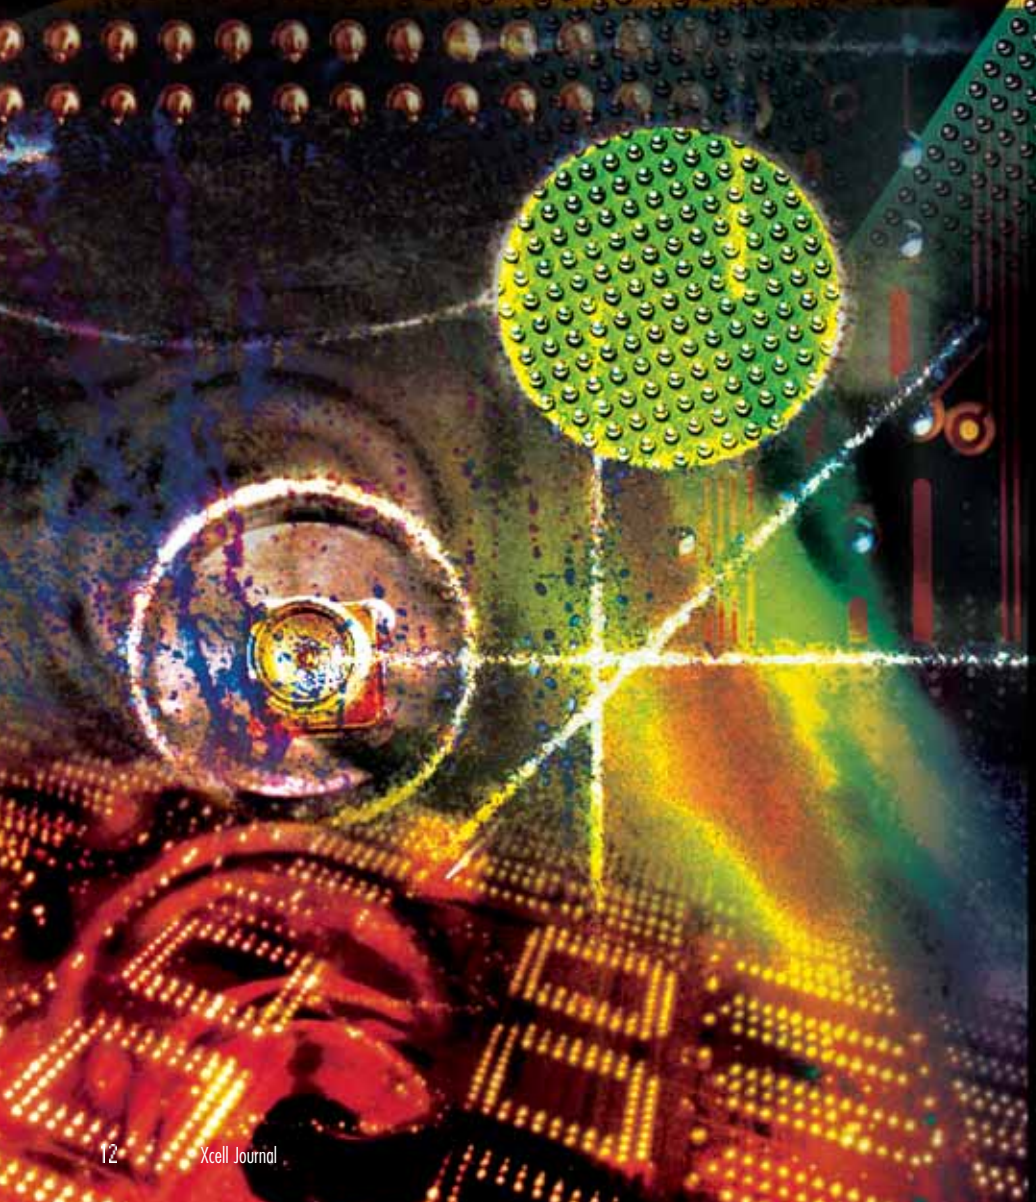

by Warren Miller
VP of Marketing
Avnet
warren.miller@avnet.com

Vince Gavagan
Design Engineer
Avnet
vince.gavagan@avnet.com

While designing the Avnet Design Services Virtex-II Pro™ 3.125 Gbps Aurora Design Kit, we found that there were a variety of options for the interface between the Virtex-II Pro multi-gigabit transceivers (MGTs) and the SubMiniature version A (SMA) connectors on the board. After trying a few of these options and measuring the results on the prototype board, we discovered that the signal integrity performance of the interface varied widely depending on the type of SMA connector used, the location, and characteristics of the board traces.

In this article, we'll review several specific design options and their impact on signal integrity. Our test results will show why the final "optimal" design was selected. We'll also provide detailed measurements on the signal integrity of the final design.

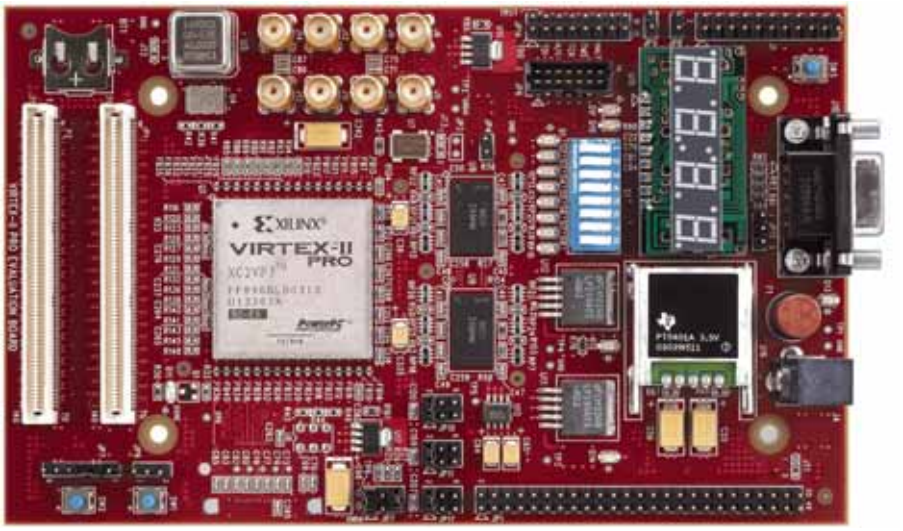

Figure 1 – The Avnet Virtex-II Pro evaluation board. The eight SMA connectors are located in the upper middle of the board, very close to the FPGA.

Avnet Virtex-II Pro Aurora Design Kit

The Aurora Design Kit includes the Avnet Design Services Virtex-II Pro evaluation board and its associated documentation, board support package, applications code, and example designs. The Aurora reference design has been ported to the board; an example design communicating between serial ports at 3.125 Gbps demonstrates the features and capabilities of the design kit. The circuit board used in the design kit is shown in Figure 1.

The hardware features of the evaluation board include the user FPGA, on-board memory, on-board communications, expansion, and configuration. A complete block diagram of the board's hardware components is shown in Figure 2.

The user FPGA is a Virtex-II Pro XC2VP7-FF896 device that includes an embedded PowerPC™ processor, eight high-speed serial I/O channels, and RocketIO™ MGTs.

The high-speed communications functions of the board include eight SMA connectors (TX/RX pairs for two RocketIO ports) with board-configurable loop-back for two RocketIO transceivers and pads for four additional RocketIO ports.

The connectors featured on board include two 140-pin general-purpose I/O expansion connectors (AvBus), up to 30 LVDS pairs, and a standard 50-pin 0.1-inch header for custom expansion.

Memory includes Micron™ DDR SDRAM (64 MB) for use as code storage space for the PowerPC and packet storage for serial I/O ports. Communication can also occur over a standard RS-232 serial port for simple monitor or debug functions.

An included 5.0V AC/DC power supply provides up to 22.5W for the on-board Texas Instruments™ 3.3V 6A module and National Semiconductor™ linear regulators.

You can configure the FPGA via two Xilinx XC18V04-VQ44 PROMs, a Parallel IV cable for JTAG, and fly-wire support for Parallel-III and Multilinx™ configuration cables.

Mictor™ connectors are available to access the remaining high-speed I/O signals (MGTs) on the device for test or characterization.

Prototype Design

When we first received the prototype design from the manufacturing house, we tested the MGT-to-SMA connections. In preliminary testing, we used a loop-back connection over an SMA cable from one MGT to another. No FR4 of significant length was inserted.

Test packets were simple 00 to FF data words (256 bytes) with idle sequences (k28.5, d21.4, d21.5, d21.5) between packets. We captured eye diagrams using the repeating idle sequence (Figure 3). Although the eye opening is fairly clean, the pattern is a repeating idle sequence and thus doesn't provide a very extensive test.

During our initial prototype testing, we sent several thousand packets successfully. However, testing additional boards

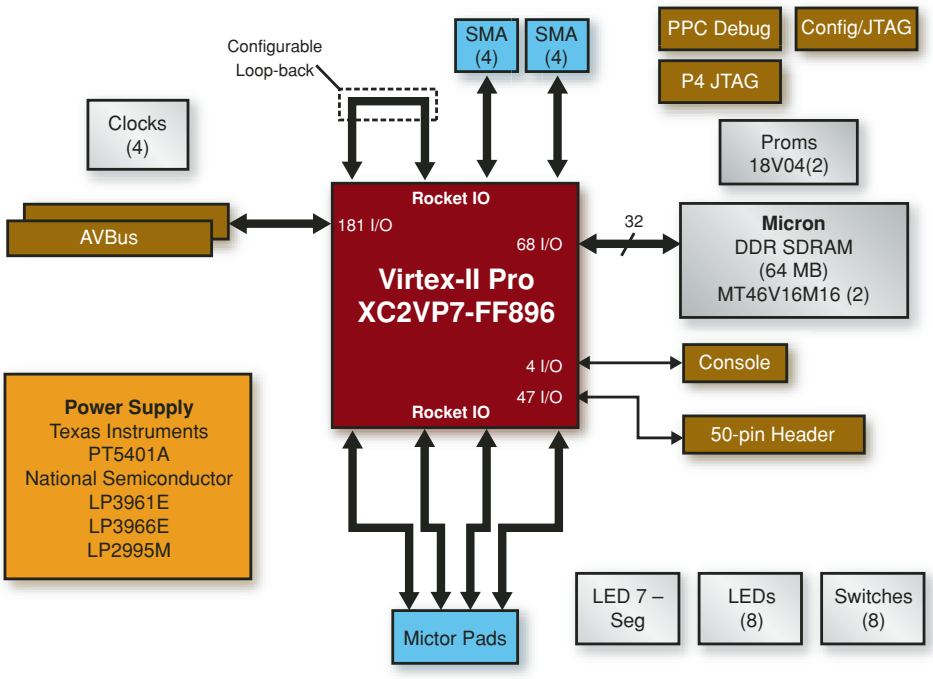

Figure 2 – Block diagram of the Aurora Design Kit circuit board

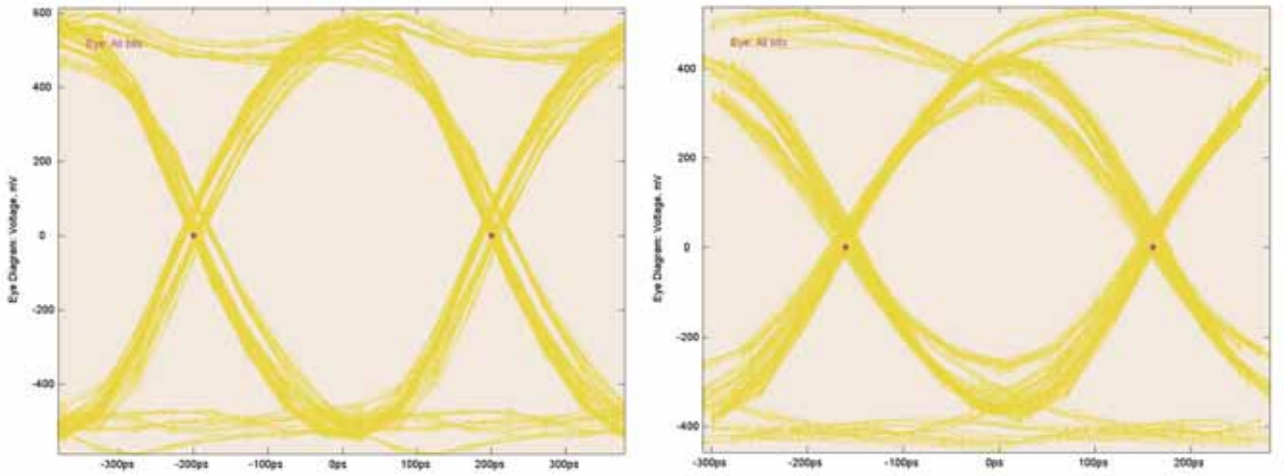

Figure 3 – Eye diagrams of 2.5 Gbps (left) and 3.125 Gbps (right) ▶

revealed that performance was not repeatable, and in fact was much worse for the majority of boards. Because these initial boards used -5 speed grade parts, we surmised that the errors could be attributed to the 2.0 Gbps limitation of the -5 speed grade devices.

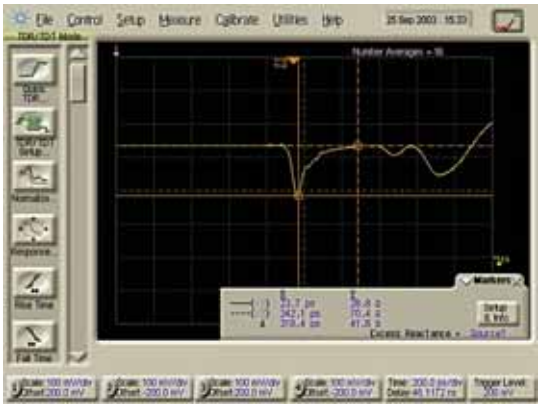

Figure 4 – Initial TDR results

Yet procurement of -6 speed grade parts revealed that this was not the case; a substantial number of errors continued to occur. Because the design includes two RocketIOs that are looped back via FR4 on the board in addition to the SMA breakout, we repeated the test using the non-SMA loop-back. These tests yielded substantially better results. In fact, the non-SMA loop-back was capable of 3.125 Gbps with identical payload and zero errors.

During these tests, we performed Time Domain Reflectometry (TDR) on the board as well, with results shown in Figure 4. After reviewing these results, we concluded that although the board impedance was matched very well, a severe impedance mismatch existed at the SMA connectors.

This information suggested we should look more closely at the layout, and we determined that during the design phase, a stub was overlooked. As the FPGA and SMAs both reside on the top (component) side of the PCB, and the traces are also on layer one (a 100 Ohm differential impedance micro-strip), a stub is created at the through-hole SMA.

Prior to a board spin, two boards were used to test the theory. Testing unmodified boards with 15 inches of FR4 (using an FR4 characterization board) yielded the following results:

Board #1:		
	# of packets	# byte errors
Test 1	0x10000 (65,536)	136
Test 2	0x10000 (65,536)	306

Board #2:		
	# of packets	# byte errors
Test 1	0x10000 (65,536)	5527
Test 2	0x10000 (65,536)	8270

We modified the poorer performing board by cutting the through-hole stub protruding from the backside of the board and grinding the stub flush with the backside of the board.

With stubs cut, we observed the following results:

Board #2:		
	# of packets	# byte errors
Test 1	0x10000 (65,536)	168
Test 2	0x10000 (65,536)	273

With stubs filed flush, we obtained the following results:

Board #2:		
	# of packets	# byte errors
Test 1	0x10000 (65,536)	115
Test 2	0x10000 (65,536)	134

This seemed to be a promising approach, so we decided to try another option to eliminate the stub before doing a re-layout of the board. The SMAs were removed from board #2 and placed on the backside of the board. This effectively removed the stub, since the via and center SMA conductor became part of the intended transmission line.

Testing this modified board yielded zero errors. This confirmed our finding that the

stub was the cause of the unacceptable error rate and that a layout change was required to remove the stub.

Because we wanted to keep the SMAs on the top side and minimize modification to the existing microstrip, a board spin was required. Noting the performance of surface-mount SMAs in simulation, we decided to proceed with a similar SMA configuration for the board spin.

Revision of the Prototype Design

The prototype board was redesigned to eliminate the stub by using a surface-mount SMA connector. More extensive testing would also be necessary to further verify the operation of the new board design.

In addition to FR4 characterization, we chose to use a more exhaustive test pattern to validate the performance of the new board. Furthermore, partial reconfiguration would allow on-the-fly adjustments to MGT parameters such as pre-emphasis and differential swing. The results are shown in Figure 5.

```
Device Type is 2VP7

Read Frame MGT_X = 0
           MGT_Y = 1
Left MGT Found

**** COMP ENABLED
    TX_PREEMPHASIS = 2 (25 Percent)
    TX_DIFF_CTRL = 600 mV
-------------------------------------


Read Frame MGT_X = 1
           MGT_Y = 1
Right MGT Found

**** COMP ENABLED
    TX_PREEMPHASIS = 2 (25 Percent)
    TX_DIFF_CTRL = 600 mV
-------------------------------------
```

Figure 5 – MGT settings

```
+-------------------------------------------------------+
|MGT# | RocketIO Transceiver ACTIVE MGT Status          |
|-----+-------------------------------------------------|
|     | Received  Frames # : 4,788,376 H + 421,617      |
|     | Dropped   Frames # : 0                          |
|     | Total Bit Errors # : 0                          |
|     | Error     Factor # : Infinite                   |
|     | Line Rate        # : 3,127 MBps                  |
|MGT4 | Bit Error Rate   # : 0                          |
|     | Data Pattern     # : 12-PRBS32: 1+x^10+x^30+x^31+x^32
|     | Link-1 Abort-0 PowrDun-0 TxInhibit-0 LoopBk-00  |
|     |-------------------------------------------------|
|     | [1] Toggle TxInhibit       [2] Toggle PowerDown |
|     | [3] Toggle Serial Loopback [7] Toggle ErrInsert |
|     | [d/a] Change Data Pattern                       |
|-----+-------------------------------------------------|
|     | Received  Frames # : 4,788,376 H + 422,548      |
|     | Dropped   Frames # : 0                          |
|     | Total Bit Errors # : 0                          |
|     | Error     Factor # : Infinite                   |
|     | Line Rate        # : 3,127 MBps                  |
|MGT6 | Bit Error Rate   # : 0                          |
|     | Data Pattern     # : 12-PRBS32: 1+x^10+x^30+x^31+x^32
|     | Link-1 Abort-0 PowrDun-0 TxInhibit-0 LoopBk-00  |
|     |-------------------------------------------------|
|     | [4] Toggle TxInhibit       [5] Toggle PowerDown |
|     | [6] Toggle Serial Loopback [0] Toggle ErrInsert |
|     | [c/z] Change Data Pattern                       |
|-----+-------------------------------------------------|
| [+/-] Change Data Pattern on Both MGTs                |
| [9] Switch MGT Clock Source (Current=BOARD OSC)       |
| [0] Reset Both MGTs                                   |
| [ESC] Back to main menu                               |
+-------------------------------------------------------+
```

Figure 6 – Test results

We found it was easy to control the length and impedance of the traces, but we overlooked the impact SMA connector choice had on the signal integrity.

The test setup parameters were:

- MGT 4 connected through a 12-inch RG 316 cable (Johnson 415-0029-012) to a 20-inch FR4 trace on the Xilinx MGT characterization board (Xilinx)
- Pre-emphasis at 25% (setting 2 of 0-3) and differential swing of 600 mV (setting 2 of 0-4)
- Test pattern is PRBS32 (using the Xilinx BERT design)
- MGT 6 connected through 24-inch RG316 (manufacturer unknown) to 15-inch FR4 characterization board (Xilinx)
- Pre-emphasis at 25% (setting 2 of 0-3) and differential swing of 600 mV (setting 2 of 0-4).
- Test pattern is PRBS32.

We ran the test until the 16550 UART timed out (an evaluation-licensed core); the total frames are shown in Figure 6. Note that there are no errors; hence the bit error rate is zero. Also, the error factor is defined on page nine and table 2 in Xilinx Application Note XAPP661 as the shortest gap between errors, expressed in frames. Because there are no errors, it makes sense that this would be infinite.

The final TDR test results are shown in Figure 7. A careful reading of the results shows a 50% improvement over the previous reading and confirms the improvement of the new design.

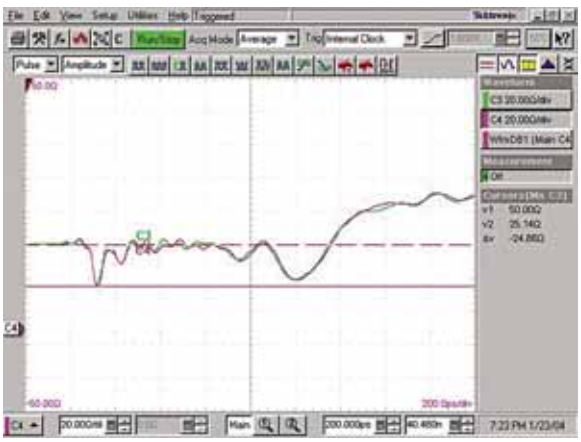

Figure 7 – TDR measurement results of the revised design show a 50% improvement, supporting the bit error rate test results. ▶

Conclusion

During the design phase, you must be very careful to identify all sources of trace and stub length. In particular, watch for a mismatch between your choice for through-hole or surface-mount SMA devices and the layout of your traces between the SMA connector and the Virtex-II Pro FPGA. We found it was easy to control the length and impedance of the traces, but we overlooked the impact SMA connector choice had on the signal integrity.

Detailed design files and test measurements are available with the purchase of an Avnet Design Services Virtex-II Pro 3.125 Gbps Aurora Design Kit. You can order the kit, part number ADS-XLX-V2PRO-EVLP7-6, for $599.

This design kit is just one of many available to speed the development cycle for complex processor, communications, FPGA, DSP, and networking applications. All of our design kits are modular and can accept matching add-on modules, applications software, design and debug tools, and compatible IP cores. For more information, visit *www.avnetavenue.com.*

For Synchronous Signals, Timing Is Everything

Mentor Graphics highlights a proven methodology for implementing pre-layout Tco correction and flight time simulation with Virtex-II and Virtex-II Pro FPGAs.

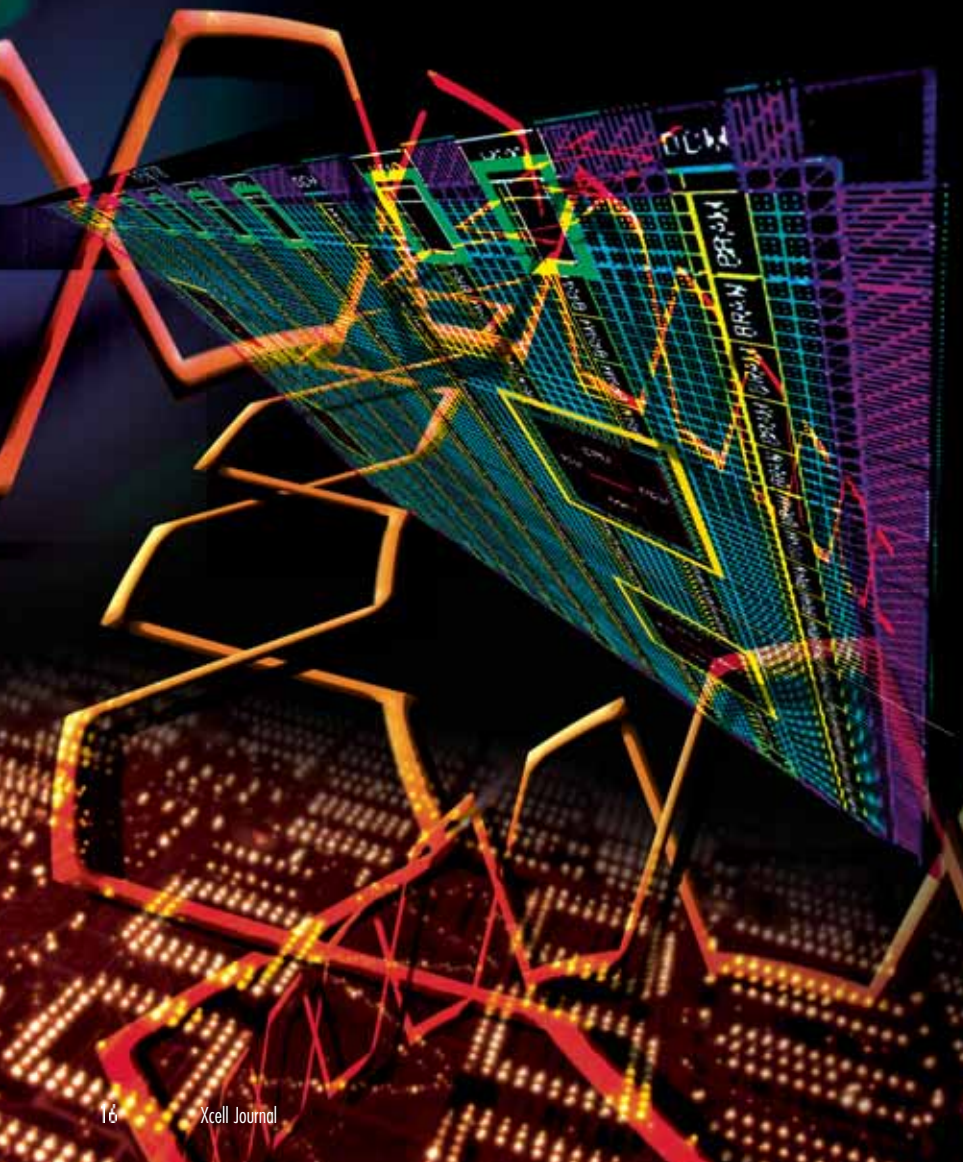

by Bill Hargin
Product Manager, HyperLynx
Mentor Graphics Corp.
Bill_Hargin@mentor.com

We've all heard the phrase "timing is everything," and this is certainly the case for the majority of digital outputs on modern FPGAs. Timing-calculation errors of 10 or 20 percent were fine at 20 MHz, but at 200 MHz and above, they're absolutely unacceptable.

As Xilinx Senior Field Applications Engineer Jerry Chuang points out, "The toughest case usually is a memory or processor bus interface. Most designers know that they have to account for Tco (clock-to-output) as it relates to flight time, but don't really know how."

Another signal integrity engineering manager who preferred to remain anonymous explains, "We've got lots of things that hang on the hairy edge of working. That's one of the reasons why they give you so many knobs to turn on newer memory interfaces."

To complicate matters, manufacturer datasheets and application notes use multiple, often-conflicting definitions of many of the variables and procedures involved, requiring you to investigate the conventions used by manufacturer A versus manufacturer B. Most of the recently published signal integrity books either gloss over the subject or avoid it altogether. We hope that this article will serve to blow away some of the fog and reinforce some standard definitions.

System Timing for Synchronous Signals

An FPGA team will typically place and route an FPGA according to their specific timing requirements, leaving system-level timing issues to be negotiated later with the system-design team. With the sub-nanosecond timing margins associated with many signals, it's common for the system side to be faced with PCB floor-planning changes, part rotation, and sometimes the need to negotiate pin swaps with the FPGA team to accommodate timing goals. Proactive, pre-layout timing analysis and some careful accounting can keep both the FPGA and system teams from spending a month or more chasing timing problems.

Two classes of signals pose problems for FPGA designers and their downstream counterparts at the system level: timing-sensitive synchronous signals and asynchronous, multi-gigabit serial I/Os. We'll concentrate on parallel, synchronous designs in this article.

Margins

The system-timing spreadsheet for synchronous designs is based on two "classic" timing equations:

Tco_test(Max) + Jitter + TFlight(Max) + TSetup < TCycle

Tco_test(Min) + TFlight(Min) > THold

Or, once Tco_test is corrected, becoming Tco_sys, as outlined in this article:

Tco_sys(Max) + Jitter + Tpcb_delay(Max) + TSetup < TCycle

Tco_sys(Min) + Tpcb_delay(Min) > THold

Each net's timing is initially set up with a small, positive timing margin. This margin is allocated to the TFlight(Max) and TFlight(Min) values (or Tpcb_delay[Max] and Tpcb_delay[Min], respectively) in the preceding equations; these are timing contributions of the PCB interconnect between each net's driver and receivers.

If there is insufficient margin left to design the interconnects, either the silicon numbers need to be retargeted and redesigned, or the system speed must be slowed. Figure 1 shows how timing margins shrink relative to frequency.

There are two ways to come up with the interconnect values for the timing spreadsheet. Some signal integrity tools automatically make calculations that produce a single "flight-time" value. However, especially for designers just learning about the timing challenges of high-speed systems, a two-step approach is more instructive. First, you learn how to correct a datasheet's driver Tco value to match the behavior in your real system; second, you add the additional delay between the driver and each of its receivers.

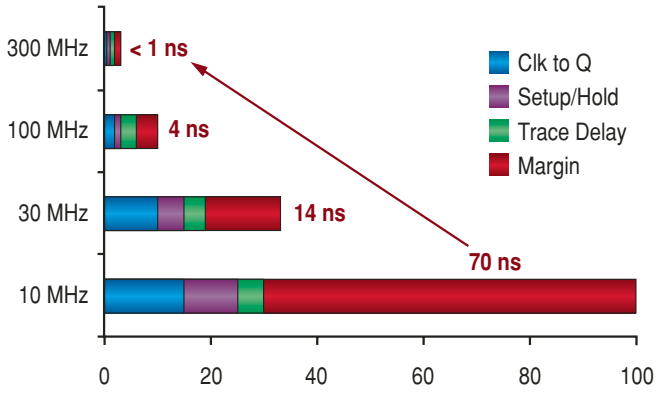

Figure 1 – Drastically narrowed system-timing margins, as clock frequency moves from 10 to 300 MHz, are shown in red.

Data Book Values

Initially, timing spreadsheets are populated with values from the silicon vendor's data book. You'll need first-order estimates from silicon designers on the values of Tco and setup and hold times for each system component. You can usually obtain this data from the component datasheet.

Test and Simulation Reference Loads

To arrive at the datasheet value for your drivers' Tco, standard simulation test loads (or reference loads) provide an artificial interface between the silicon designer and the system designer.

You'd prefer, of course, to have Tco specified into the actual transmission-line impedance you're driving on your PCB, but the silicon provider has no way of knowing what that will be. Knowing what loading the vendor assumed when publishing Tco is critical so that you can adjust for the difference between that load and your real one.

The Recipe for a Problem

As shown in Figure 2, if the reference load is significantly different from the actual load that the output buffer will see in your design, the sum of the datasheet and PCB-interconnect timing values will not represent actual system timing. Actual or total delay may be represented as:

Total Delay = Tco_sys + Tpcb_delay
≠ Tco_test + Tpcb_delay

where Tpcb_delay is the extra interconnect delay between the time at which the driver switches high or low until a given receiver switches.

Note that this "PCB delay" is not just the time it takes for a signal to travel along the trace (sometimes called "copper delay"

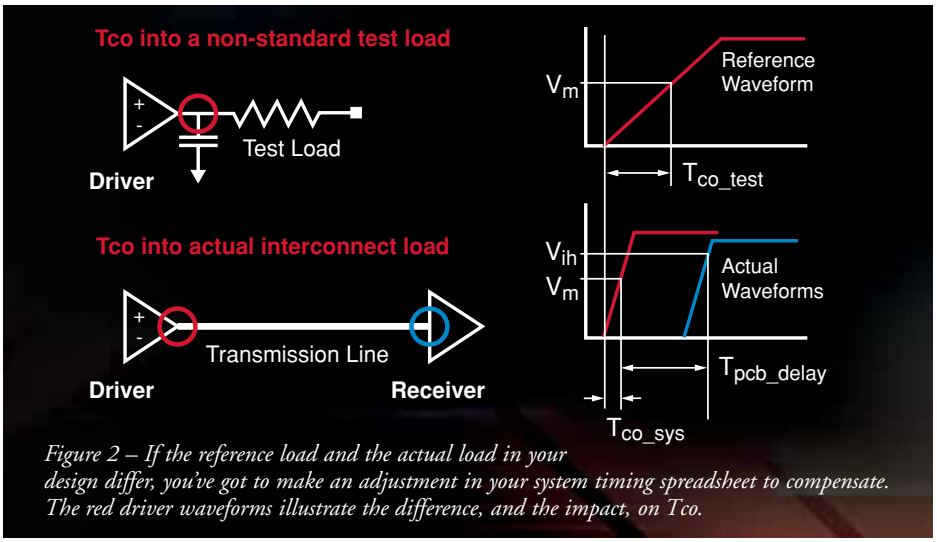

Figure 2 – If the reference load and the actual load in your design differ, you've got to make an adjustment in your system timing spreadsheet to compensate. The red driver waveforms illustrate the difference, and the impact, on Tco.

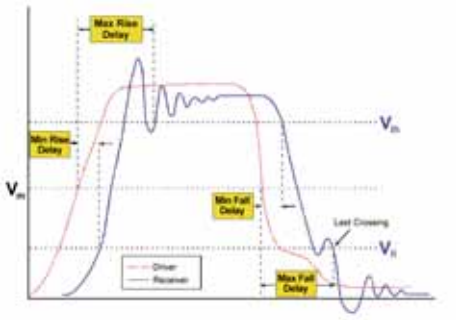

Figure 3 – "PCB delay" refers to the difference between the driver waveform switching through Vmeas and the waveform at the receiver as it switches through Vih (rising) or Vil (falling). Finding this value requires simulation, not just a simple "copper-delay" calculation.

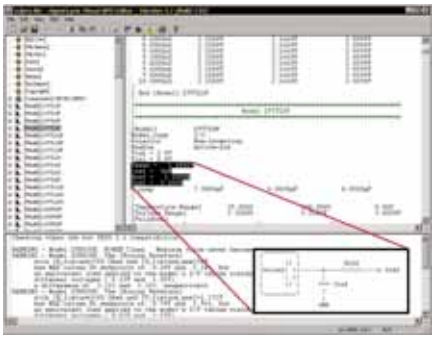

Figure 4 – Mentor Graphics' HyperLynx Visual IBIS Editor, a free tool for navigating the 50,000-plus lines of Xilinx Virtex-II Pro, Virtex-II, and Spartan IBIS models, shows reference load information for an LVTTL8F buffer as well as the assumed connections – from the IBIS specification – for Cref, Rref, and Vref in the insert.

or "propagation delay"). Here, Tpcb_delay accounts for effects such as ringing at the receiver, as shown in Figure 3. Its value could (on a poorly terminated net) easily be longer than the simple copper delay.

Calculating accurate timing involves more than finding Tpcb_delay. If the difference between Tco_sys and Tco_test is significant – even in the neighborhood of 100 ps – your board may not function properly if you don't account for the difference. But because Tco_test is a value created with an assumed test load, it almost never matches Tco_sys, the clock-to-output delay you'll see in your actual system.

For example, Lee Ritchey, author of "Get it Right the First Time" and founder of the consulting firm Speeding Edge, was hired to resolve a timing problem on a 200 MHz memory system. After digging into the design, he found that unadjusted datasheet values were used, based on Tco values that were measured on a 50 pF load rather than something resembling the design's 50 Ohm transmission-line load. As a result, this improper accounting "threw timing off by just over one nanosecond," he says. "That's 20 percent of the total timing budget, a major error."

In the following sections, we'll see how you can correct Tco_test to become Tco_sys, avoiding this type of error altogether.

The Process

Measuring Tco_test

To measure Tco_test, you need to set up a simulation with just the driver model and the datasheet test load. Though they're an optional sub-parameter in the IBIS specification, most IBIS models (including Xilinx IBIS models) contain a record of the test load (Cref, Rref, Vref) and the measurement voltage (Vmeas) to use with these values. Figure 4 shows these values for the LVTTL8F buffer in the Virtex-II Pro™ IBIS model, as well as a generic reference load diagram taken from the IBIS specification.

Once you've gathered these load values from the IBIS model, you simulate rising and falling edges, and for each, measure the time from the beginning of switching until the driver pin crosses the Vmeas threshold. These are the Tco_test values.

Obtaining "Tcomp," the Timing-Correction Value

Now you need to calculate a compensation value, Tcomp, that will convert the datasheet Tco value into the actual Tco you'll see in your system. Tcomp is the delay between the time the driving signal, probed at the output, crosses Vmeas into the silicon manufacturer's standard reference load, and the time it crosses Vmeas for your actual system load. Tcomp is then used as a modification to the Tco value from the vendor datasheet, as shown in Figure 5.

The revised computation of actual delay from the previous equation is then:

Total Delay = Tco_sys + Tpcb_delay
= (Tco_test + Tcomp) + Tpcb_delay

Note that Tcomp may be negative or positive, depending on whether the actual load in your system is smaller or larger than the standard test load. Traditionally, silicon vendors used capacitive test loads (like 35 pF) to measure Tco; almost all real PCB transmission lines do not present as heavy a load, so Tcomp is usually negative in this situation.

Xilinx, for its current generation of FPGAs, uses a 0 pF test load for output driver wave shape accuracy. Real transmission lines will represent a different load – some mixture of inductance, capacitance, and resistance. Because the transmission-line load is heavier than a 0 pF "open load," Tcomp will be positive. Simulation is the only way to accurately predict the exact value of Tcomp.

Simulating Tpcb_delay

At this point in the process, you've completed the first step in finding accurate delays for your timing spreadsheet, and you've compensated the datasheet Tco to match your real system load. Next, you need to determine Tpcb_delay, the additional delay caused by the interconnect from driver to receiver.

A signal integrity simulator is the only way to accurately do this, because only a simulator can account for subtle effects like reflections, receiver input capacitance, line loss, and so forth.

From here, we'll explore some detailed examples based on Xilinx-provided IBIS models – the process of calculating Tcomp and then using the HyperLynx™ simulator to determine an interconnect's Tpcb_delay through pre-layout topology analysis. You could enter the values that we come up with directly into your system-timing spreadsheet.

The process using Mentor Graphics' HyperLynx product is straightforward. You look up the manufacturer's test load in the IBIS model (see Figure 4), enter it in the LineSim schematic, set up your actual interconnect topology just below the reference load, and begin a simulation, probing at both drivers so that you can measure Tcomp and Tpcb_delay, as shown in Figure 6.

Running the Numbers on a Real Problem

An important design for an electronic equipment manufacturer had a Xilinx FPGA talking to a bank of SRAMs at 125 MHz, meaning the cycle time (Tcycle) was 8 ns.

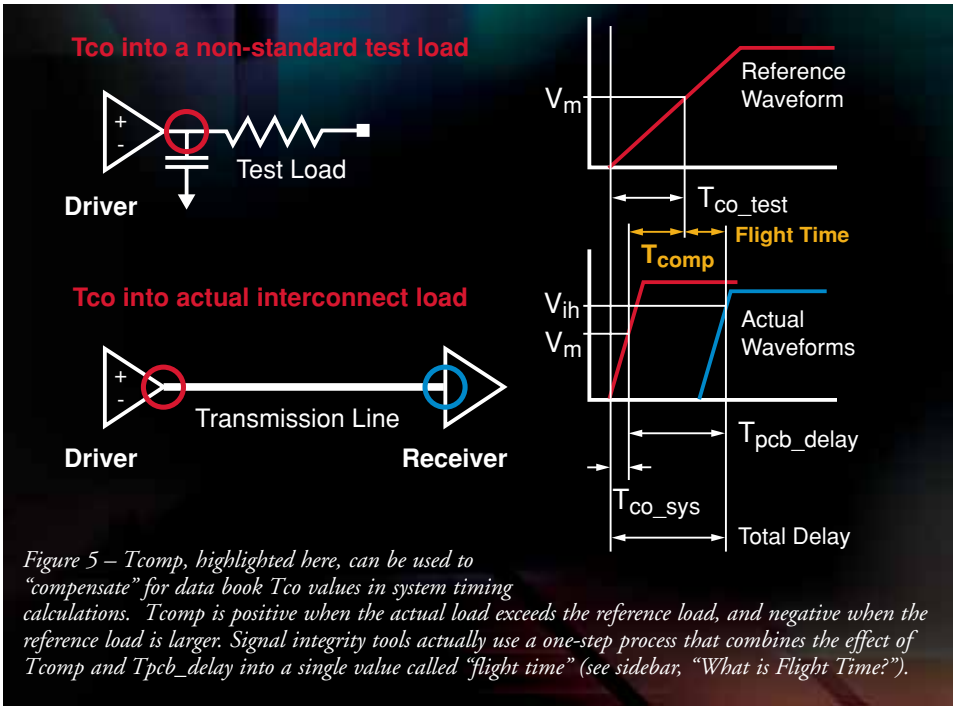

Figure 5 – Tcomp, highlighted here, can be used to "compensate" for data book Tco values in system timing calculations. Tcomp is positive when the actual load exceeds the reference load, and negative when the reference load is larger. Signal integrity tools actually use a one-step process that combines the effect of Tcomp and Tpcb_delay into a single value called "flight time" (see sidebar, "What is Flight Time?").

The Xilinx datasheet specified Tco as 4 ns (i.e., Tco_test). The SRAM's setup time was 2 ns.

Some of the traces connecting the FPGA to an SRAM were six inches long; a signal integrity simulation showed a worst-case maximum PCB delay (to the receiver's "far" threshold) of 2.5 ns. This yielded in the design's timing spreadsheet a total time of 4 + 2.5 + 2 = 8.5 ns (Tco_test + Tpcb_delay + Tsetup), violating the 8 ns cycle time.

However, the Tco value, when corrected for the actual design load, was 4-1.2 = 2.8 ns (Tco_sys = Tco_test + Tcomp), meaning that the actual total delay value was 2.8 + 2.5 + 2 = 7.3 ns (Tco_sys + Tpcb_delay + Tsetup), leaving an acceptable timing margin of 700 ps.

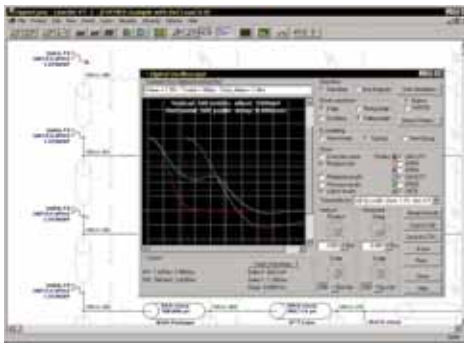

Figure 6 – Total Delay, Tco_test, Tcomp, Tco_sys, and Tpcb_delay, as well as flight time, are all measurable for this falling-edge waveform using Mentor Graphics' HyperLynx software.

Note that in this calculation, we measured to the time at which the receiver signal crossed the farthest-away threshold to get the worst-case, longest possible Tpcb_delay. For a rising edge, we measured to the last crossing of Vih; for a falling edge, to the last crossing of Vil.

Conclusion

For seamless interaction between the FPGA designer and the system designer, it's prudent to do as much pre-layout, "what-if" analysis as possible. And, though not covered explicitly in this article, you can also verify that your laid-out printed circuit boards meet your timing requirements using a post-layout simulator with batch analysis capabilities.

Some Mentor products that perform this type of analysis are HyperLynx, ICX, and XTK. Running these simulations, you're revising simulated representations of interconnect circuits in minutes as compared to the weeks required to spin actual PCB prototypes.

The new HyperLynx Tco simulator is available on Mentor Graphics' website, *www.mentor.com/hyperlynx/tco/*. Included with the Tco simulator are the Virtex-II Pro, Virtex-II™, and Spartan™ IBIS models; boilerplate schematics that will help you make adjustments to data book Tco values; and a detailed tutorial on Tco and flight-time correction that parallels this article.

What is "Flight Time"?

In this article, we've shown conceptually how Tco values specified into a silicon vendor's test load can be corrected on a per-net basis to give the actual clock-to-output (Tco) timing you'll see on your PCB, and then added to the additional trace delays between drivers and receivers to give accurate timing values. However, signal integrity (SI) tools actually deal with corrected timing values in a different (but equal) way.

The most convenient output from an SI tool is a single number – called "flight time" – shown in Figure 5 as (Total Delay - Tco_test) or (Tpcb_delay - Tcomp). You can add this value to the standard data book Tco values in your timing spreadsheet to give the same effect as the two-step process described in this article.

When an SI tool calculates timing values, it 1) simulates each driver model into the vendor's test load, measures the time for the output to cross the Vmeas threshold, and stores the value (Tco_test); 2) simulates the actual nets in the design and measures the time at which each receiver switches (Total Delay); and 3) for each receiver, subtracts the driver-switching-into-test-load time from the receiver time (Total Delay – Tco_test). The resulting flight time is a single number that can be added to each net's row in a timing spreadsheet, and that both compensates Tco_test for actual system loading and accounts for the interconnect delay between driver and receiver.

The term "flight time" is somewhat unfortunate, although it's become the industry standard. The name suggests the total propagation delay between driver and receiver, but the value calculated is actually the delay derated to compensate for the reference load. For old-style capacitive reference loads (e.g., 50 pF), flight time can even be negative.

Designing High-Speed Interconnects for High-Bandwidth FPGAs

Commercial EM software combines with circuit and system simulation to achieve reliable data transmission.

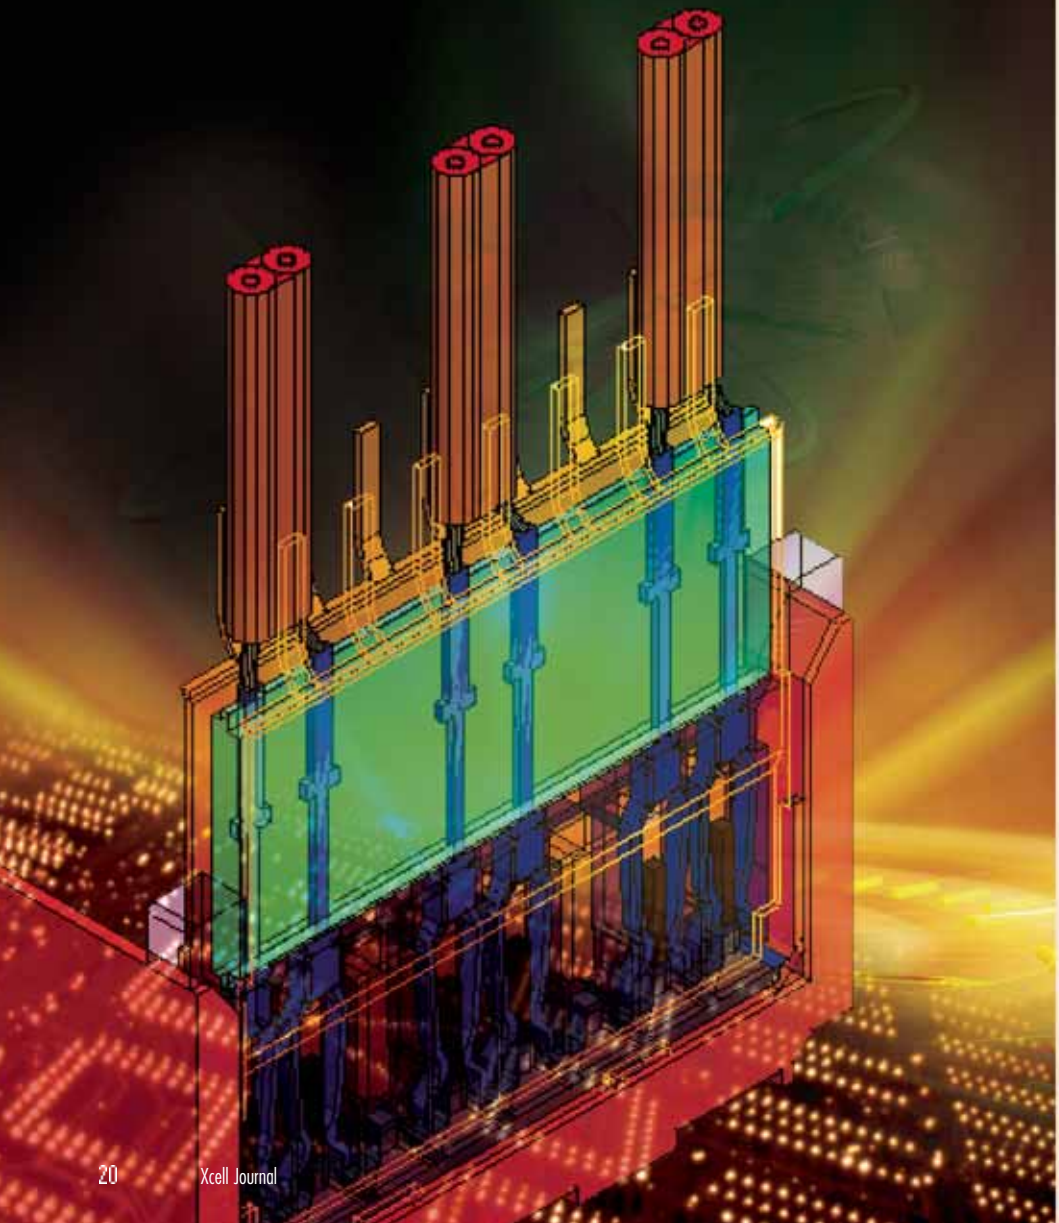

by Lawrence Williams, Ph.D.
Director of Business Development
Ansoft Corporation
williams@ansoft.com

The push toward FPGA platform solutions with high-bandwidth DSP and gigahertz-speed I/O functionality has led to devices that place greater demands on PCB design. The high serial data rates of Xilinx Virtex-II Pro™ FPGAs (3.125 Gbps) and Virtex-II Pro X™ FPGAs (10 Gbps) require careful signal integrity design for proper system operation.

In this article, we'll explain how to combine commercial electromagnetic (EM) software with circuit and system simulation to characterize transmission lines, vias, and connectors for systems that incorporate high-bandwidth FPGAs [1]. We used two-dimensional EM simulation to extract quasistatic circuit models for the PCB transmission lines and three-dimensional EM simulation to extract models for vias and connectors. For end-to-end simulations, we applied a convolution simulator. Thus, it's possible to achieve reliable data transmission with proper use of modern design tools.

High-Performance PCB Design

PCB designers aim to create interconnects that reliably transmit high-speed serial signals. Transmission lines, via structures, and connectors are the building blocks of the design – and all have their particular challenges. These structures are designed individually to meet particular metrics and are then assembled into a system-level interconnect to evaluate end-to-end performance.

The most common PCB transmission structures are the microstrip and stripline transmission line; they are easy to construct, and you can use both for signaling at gigabit speeds. Designers have also used single-ended lines successfully for lower speed designs; modern gigabit designs use differential signaling because of the advantages of noise immunity and reliable current return paths. The key parameters associated with PCB transmission lines are the characteristic impedance, delay, insertion loss, and crosstalk.

Via structures allow you to route circuit traces between layers of a multilayer board. Vias are particularly useful for transitioning from the pins of a ball grid array or connector down to stripline traces within the board. The most common and inexpensive via structure is the "through-hole" via.

Alternatives to the through-hole via are the blind via and the back-drilled via. Although these alternatives generally provide higher performance, most high-volume designs continue to use the lower cost through-hole via. Key issues in the design of through-hole via structures are unterminated via stubs and antipad radii.

Connectors provide an electrical and mechanical interface between circuit boards, or between boards and cabling. Connector performance is highly dependent on the escape-routing PCB interface. Designs can succeed or fail depending on the choice of route layer and resultant via stub length, antipad dimensions, board materials, and escape-routing layout. Additionally, transmission bends within connectors skew the transmission path and can lead to mode conversion.

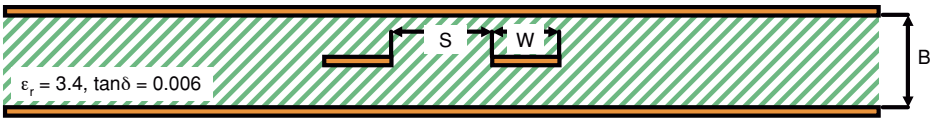

Layer	B	W	S	Zse	Zd	Zcom
S10	0.272	0.125	0.250	49.15	96.05	25.13

All dimensions are in millimeters

Figure 1 – Two-dimensional quasistatic simulations performed on stripline transmission structures using Ansoft Q3D. The table lists single-ended, differential, and common impedances.

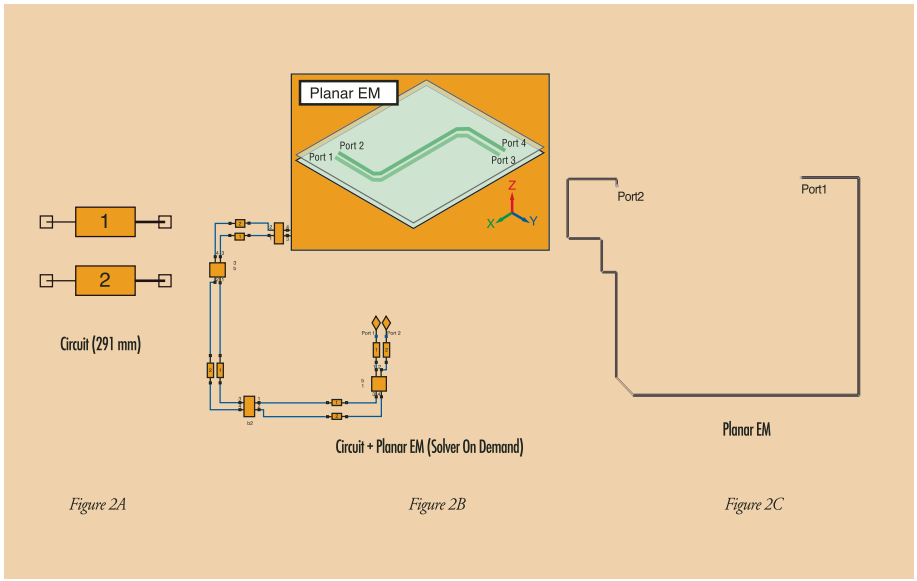

Figure 2 – You can model PCB interconnects using various methods. Circuit models (A) are the simplest and least expensive computationally; planar EM (MoM) simulations (C) are most expensive computationally but also the most accurate; a combined circuit + planar EM (B) provides accurate results with relatively low computational effort.

Electromagnetic Model Extraction

The most common printed circuit board material is FR4. Although inexpensive for circuit fabrication, FR4 suffers significant dielectric losses at high frequencies. Typical material properties for FR4 are ε_r = 4.2 and loss tangent tand = 0.022.

An alternative to FR4 is to use a lower loss Getek™ material. Getek II's material properties are ε_r = 3.4 and loss tangent tand = 0.006. Figure 1 depicts a layer within a typical backplane board. The layer height is 0.272 mm (10.7 mils); trace width is 0.125 mm; trace separation is 0.250 mm. Half-ounce copper plating for the traces provides a trace thickness of 0.7 mils.

We performed simulations using the two-dimensional, quasistatic finite element simulator within the Ansoft Q3D software suite. The stripline geometries were designed to provide nominally 100 Ohms of differential impedance, and simulations confirmed that the impedance was within 4% of the nominal value.

Figure 2 depicts three methods by which you can model the PCB interconnects. The simplest is to use a coupled-line circuit model (Figure 2A), found in popular high-frequency circuit simulators like Ansoft Designer™. In this instance, the interconnect is modeled with a uniform differential coupled transmission line without any discontinuities.

On the other end of the modeling spectrum is a full-wave planar EM field simulator based on the method of moments (MoM) (Figure 2C). The Ansoft Designer Planar EM simulator separates the traces into thousands of triangular elements. Numerical simulations compute the current flow on all triangles based on the EM coupling between them. As such, these computations completely characterize signal transmission and reflection on the interconnect.

Although accurate, MoM simulations are also the most computationally expensive. A compromise that offers the accuracy of planar EM simulations and some of the speed of circuit simulation is to use a combination of the two (Figure 2B). Ansoft Designer allows you to subdivide interconnects into a model with circuit elements and EM elements. Circuit elements are used for long, uniform sections of the coupled transmission line. EM simulation is used for all coupled line bends, as shown in Figure 2. This "solver on demand" approach automatically calls the planar EM solver whenever a bend is encountered.

Figure 3 plots the results of the three simulation methods outlined in Figure 2. All methods accurately predict the insertion loss. The circuit model cannot provide meaning-

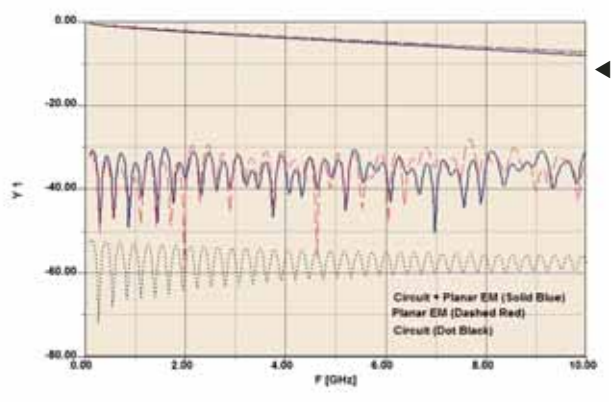

Figure 3 – PCB interconnect simulation results show that all methods outlined in Figure 2 accurately predict insertion loss. Return loss cannot be predicted with the circuit model alone.

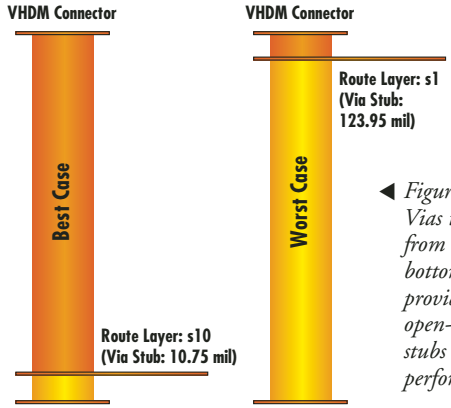

Figure 4 – Vias that transition from the top to the bottom of a board provide minimal open-circuited via stubs and "best-case" performance.

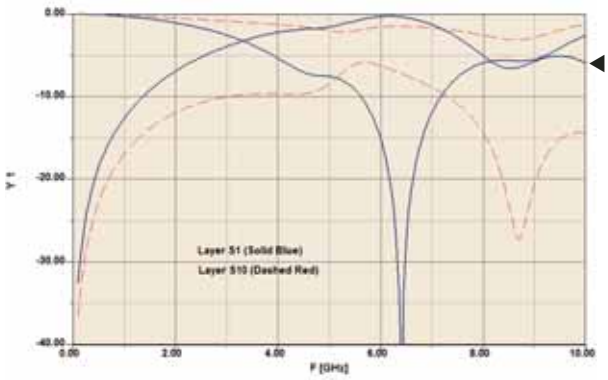

Figure 5 – "Through-hole" via performance as simulated using Ansoft HFSS. Note the sharp resonance in the insertion loss for the worst-case via routed to layer S1.

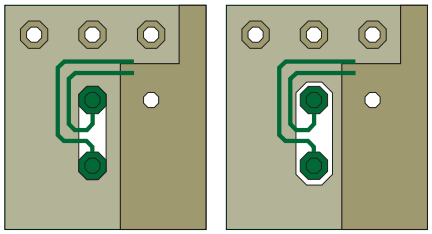

Figure 6 – Antipad radii should be sufficiently large to avoid capacitive coupling to power and ground.

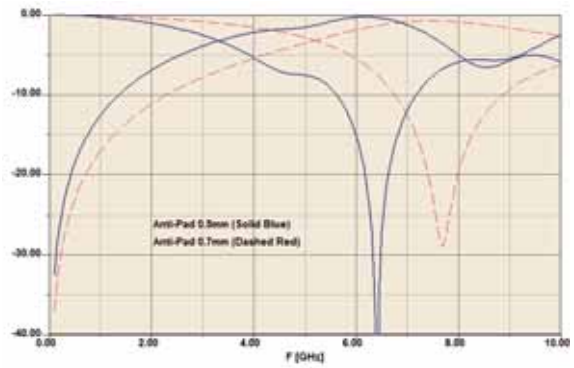

Figure 7 – Differential via performance for layer S1 (worst-case) routing for two antipad radii

ful return loss, as it does not contain any of the coupled line bends. The circuit plus planar EM method (solver on demand) provides return loss results that are in close agreement with the planar EM results. This method provides accurate results with a greatly reduced computational expense.

Vias

A common signal integrity design practice is to place high-speed route layers on opposite sides of the board in order to avoid open-circuit via stubs [2]. Figure 4 depicts two via structures: a best case and worst case. The best case occurs when routing from the top layer to layer S10, as this results in a very short (10.75 mil) via stub. The worst case occurs when routing to layer S1, leaving a very long (123.95 mil) via stub.

Figure 5 plots the insertion and return loss of an isolated differential via computed using the three-dimensional full-wave field solver Ansoft HFSS. The solid blue curve represents the via that transitions to layer S1. This is considered the worst case, as it has a very significant open-circuited via stub and an associated resonance in the insertion loss near 6.5 GHz. The dashed red curve represents the via that transitions to layer S10. This is considered the best case, as it provides a very flat insertion loss response to 10 GHz, and return loss is good to roughly 4.5 GHz.

Another consideration when designing vias are the antipads that exist on all power and ground layers. Figure 6 depicts two differential via structures with antipad radii of 0.5 mm and 0.7 mm. You can improve performance by using the larger antipad radius [2]. We performed simulations using Ansoft HFSS to predict the performance of each.

Figure 7 shows the swept frequency results for both via antipad radii for differential vias routing to layer S1 (worst case). As you can see in the plot, a significant increase in bandwidth is possible with this simple modification. The resonance in the

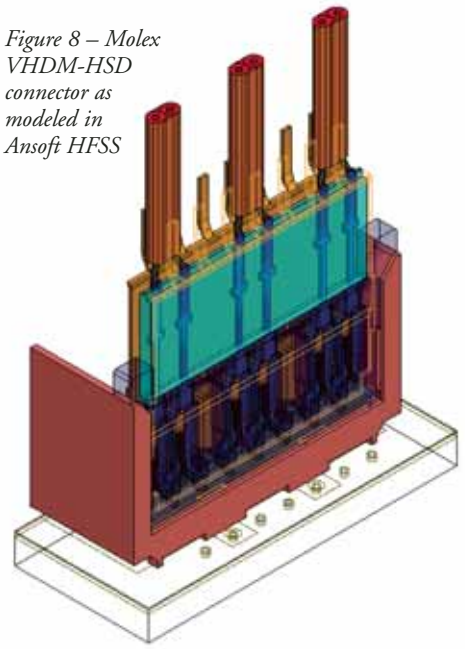

Figure 8 – Molex VHDM-HSD connector as modeled in Ansoft HFSS

insertion loss has been pushed up from 6.5 GHz to roughly 7.75 GHz. This is one of the simplest modifications that can be made to a PCB board file and should be considered for all high-performance designs.

Connectors

A common connector used to transition between boards and differential coaxial cables is the Molex™ very high density metric-high-speed differential (VHDM-HSD). Ansoft HFSS performed simulations of such a connector (Figure 8). On one side of the connector are three twinax cables; on the other side is a backplane board with its associated escape routing.

Figure 9 plots the insertion and return loss versus frequency for the VHDM connector without the escape routing. This connector provides a very flat insertion loss across the band. Return loss is below 10 dB up to 3 GHz.

Results for the connector (including all escape routing) are computed by cascading S-parameters from the individual HFSS models for the connector and the backplane escape routing. Including the backplane board, escape routing to the model has a significant effect.

Figure 10 plots the differential S-parameters for a channel containing a worst case via transition that leaves a long unterminated via stub. The performance of the VHDM connector is dominated by the sharp resonance of the via stub that manifests itself at 6.5 GHz.

System Simulation

It is possible to cascade results generated from EM and circuit simulations to get a full system simulation. Figure 11 plots circuit simulation results displaying the insertion and return loss up to 10 GHz. As expected, the channel has a response similar to a low pass filter.

We performed time domain simulation using the system simulator in Ansoft Designer. This simulator uses a convolution algorithm to process the frequency domain channel data with user-defined input bitstreams. Insertion and return loss are included

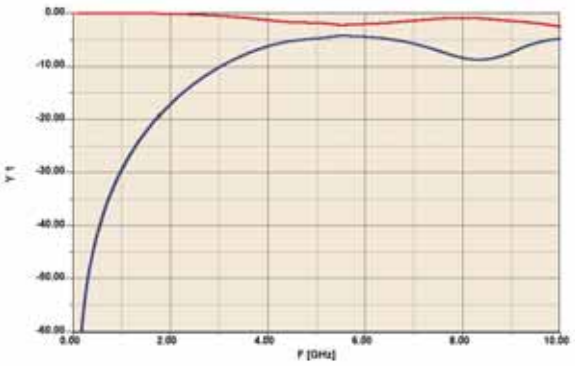

Figure 9 – Differential S-parameters for the Molex VHDM-HSD connector in isolation

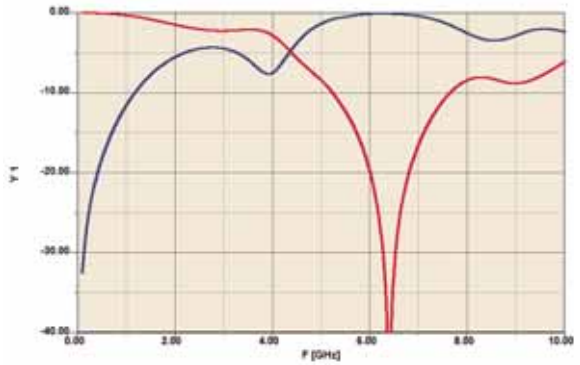

Figure 10 – Differential S-parameters for the Molex VHDM-HSD connector with backplane escape routing. This worst-case channel with large via stub shows signature resonance at 6.5 GHz.

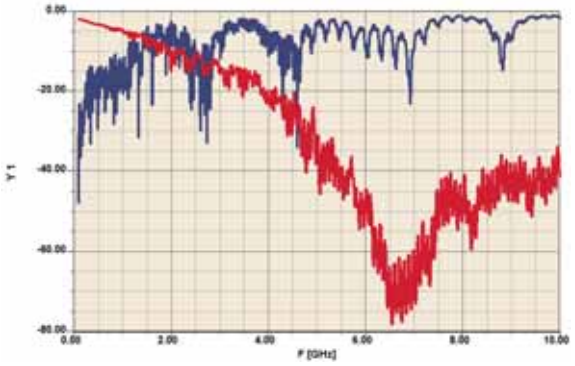

Figure 11 – Full-channel cascaded performance using the models developed from EM simulations up to 10 GHz

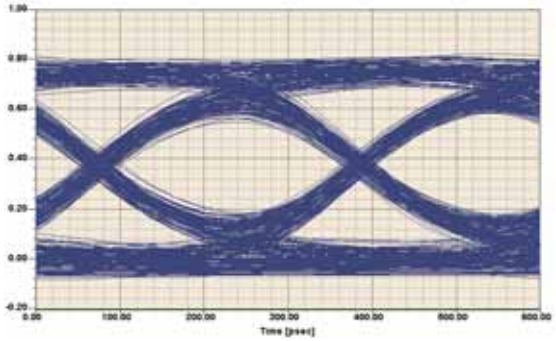

Figure 12 – Full-channel eye diagram using convolution system simulator in Ansoft Designer for the cascaded model with a 125 ps risetime

in the simulation. A 3.2 Gbps pseudo-random bit source with a 1V peak-to-peak amplitude and 125 ps risetime was applied to the channel. The channel was terminated in single-ended 50 Ohm resistors.

Figure 12 shows the resulting eye diagram as very clear and open, despite the significant channel impairments in the frequency domain results. We did not apply any pre-emphasis in the simulation. You should anticipate that some pre-emphasis would sharpen the time-domain response.

Conclusion

Modern platform FPGA devices provide wide bandwidth processing and high-speed I/O. Serial I/O with speeds in the gigabit realm creates new challenges for PCB designers.

You can solve the high-speed I/O challenges posed by modern platform FPGA devices using EM, circuit and system simulators. Although we focused our attention on the passive interconnect in this article, it is possible to include nonlinear I/O drivers and receivers in the simulation to obtain additional insight to system performance. Indeed, you can use a new tool from Ansoft called Nexxim™ to simulate all circuit behavior for systems including EM-based models, linear, and nonlinear circuits. Visit *www.ansoft.com* for more information about Ansoft Designer and Nexxim.

References

[1] Williams, L., S. Rousselle, and B. Boots, "Cray Supercomputer 3.2 Gb/s Serial Interconnect Simulation Using Full-wave Electromagnetics," in DesignCon 2004 Conference Proceedings, *Santa Clara, CA, Feb. 2-5, 2004.*

[2] Williams, L., S. Rousselle, and B. Boots, "Circuit board design for 10Gbit XFP optical modules." EDN, *May 29, 2002, pp. 63-70.*

Accurate Multi-Gigabit Link Simulation with HSPICE

With a built-in EM solver, coupled transmission lines, S-parameter support, and IBIS I/O buffer models, HSPICE provides a comprehensive multi-gigabit signal integrity simulation solution.

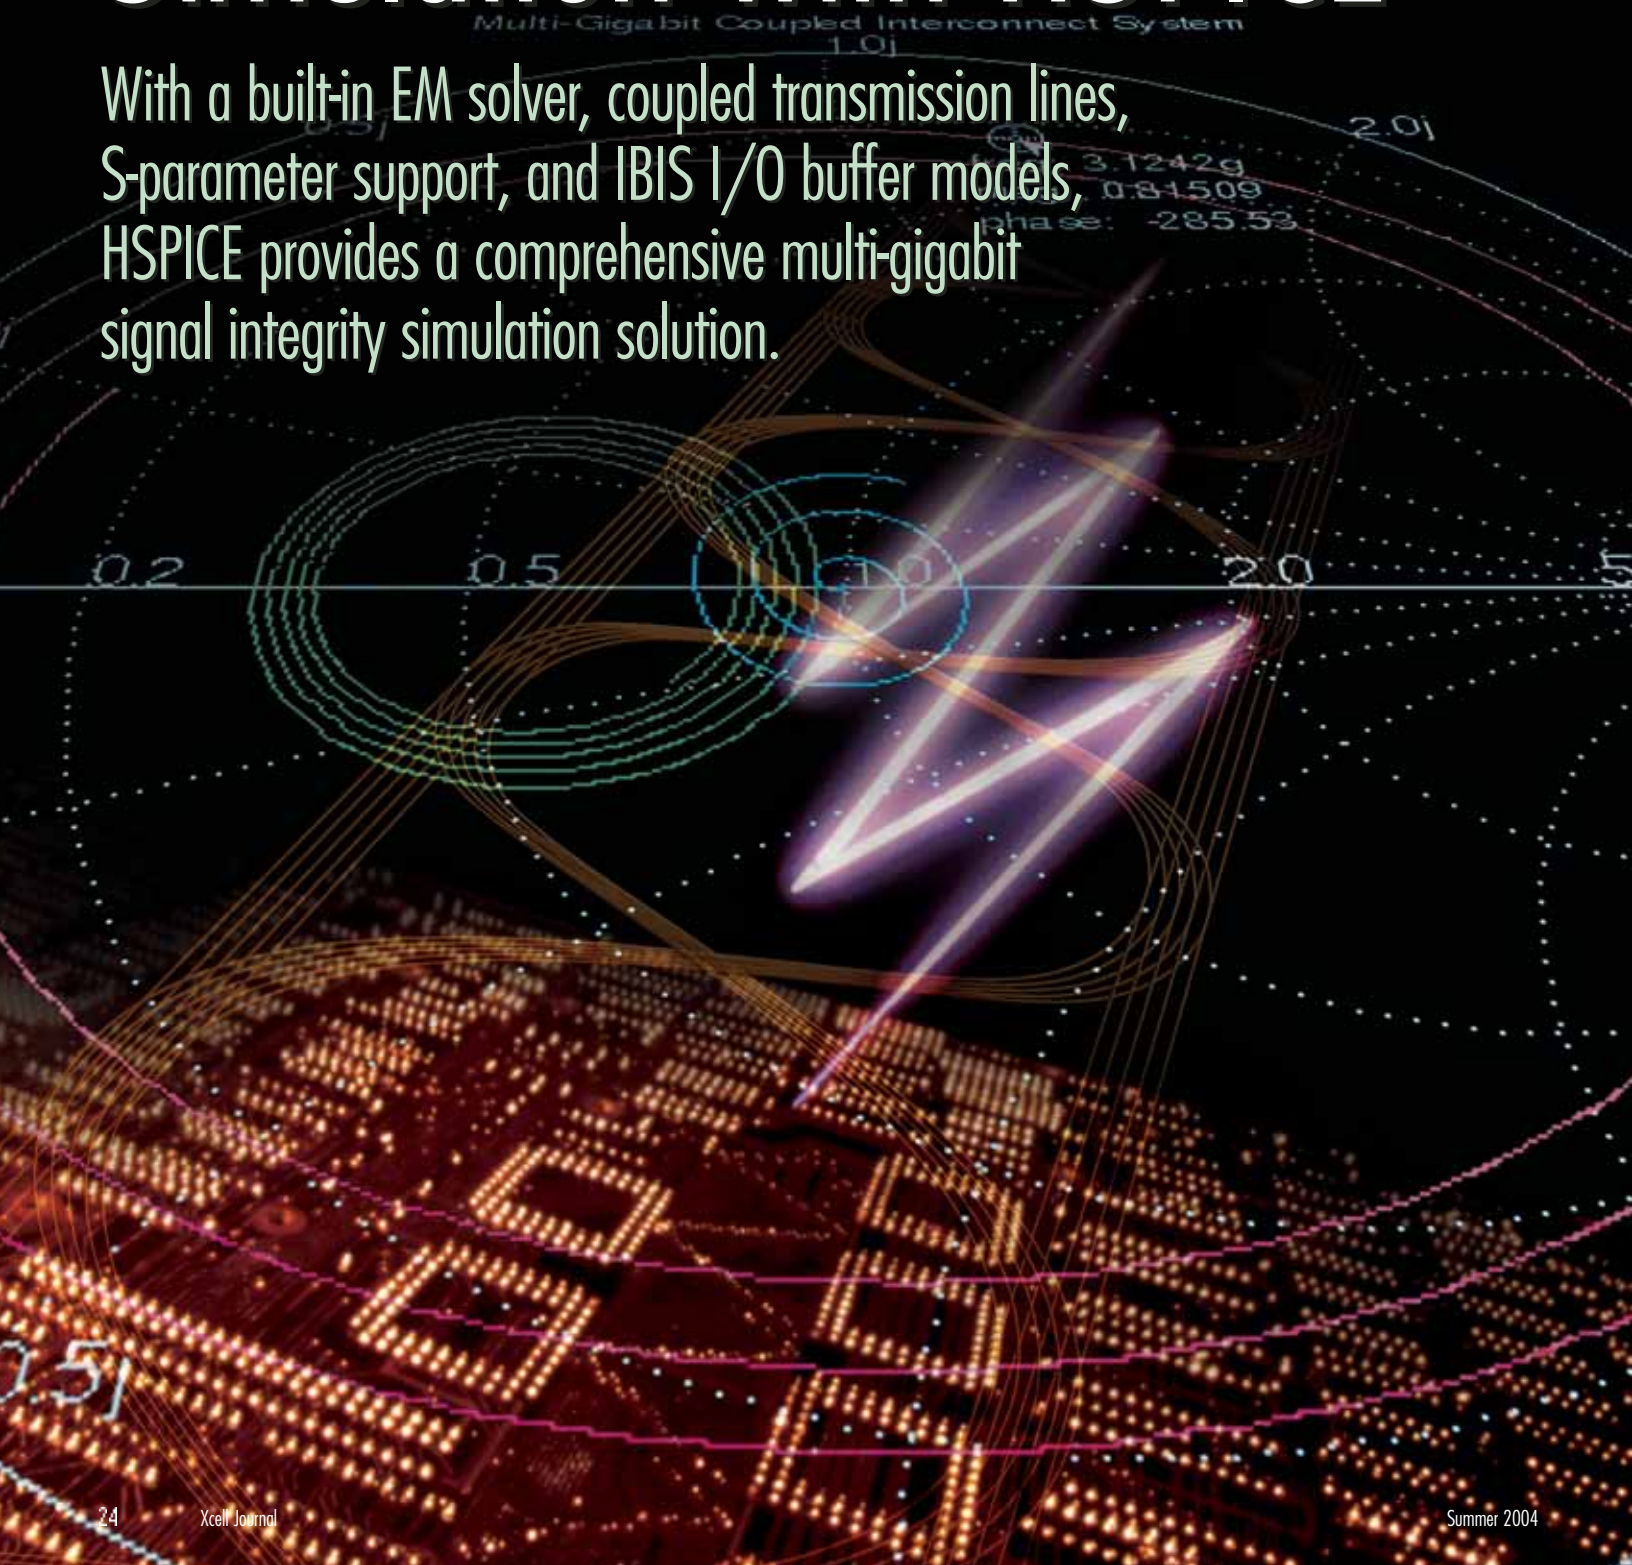

by Scott Wedge, Ph.D.
Sr. Staff Engineer
Synopsys, Inc.
wedge@synopsys.com

The Xilinx Serial Tsunami Initiative has resulted in a host of multi-gigabit serial I/O solutions that offer reduced costs, simpler system designs, and scalability to meet new bandwidth requirements. Serial solutions are now deployed in a variety of electronic products across a range of industries. Reduced pin count, reduced connector and package costs, and higher speeds have motivated the trend towards serialization of traditionally parallel interfaces.

RocketIO™ multi-gigabit transceivers (MGTs), for example, offer tremendous performance and functionality for connecting chips, boards, and backplanes at gigabit speeds. Whether your application is InfiniBand™, PCI Express™, or 10 Gigabit Application Unit Interface (XAUI), RocketIO MGTs offer ideal interface solutions.

However, the transition from slow, wide synchronous parallel buses to multi-lane, multi-gigabit asynchronous serial channels introduces new physical and electrical design challenges that traditionally fall more into the realm of radio frequency (RF) design than digital I/O design. The physical characteristics of the signal channel must be known and carefully controlled to ensure proper performance. At such high data rates, you must take into account a long list of analog, RF, and electromagnetic effects to guarantee a working design.

Life in the Fast Lane

Reliable operation of multiple transmit and receive lanes running up to 3.125 Gbps requires special attention to power conditioning, reference clock design, and to the design of the lanes themselves. You must match the differential signal trace lengths to tight tolerances. A length mismatch of 1.4 mm will produce a timing skew of roughly 10 ps, which is appreciable at these data rates. You must carefully control trace impedances and keep reference planes intact to avoid mismatches and signal reflections. Spacing between lanes must be adequate to avoid crosstalk, but remain space-efficient.

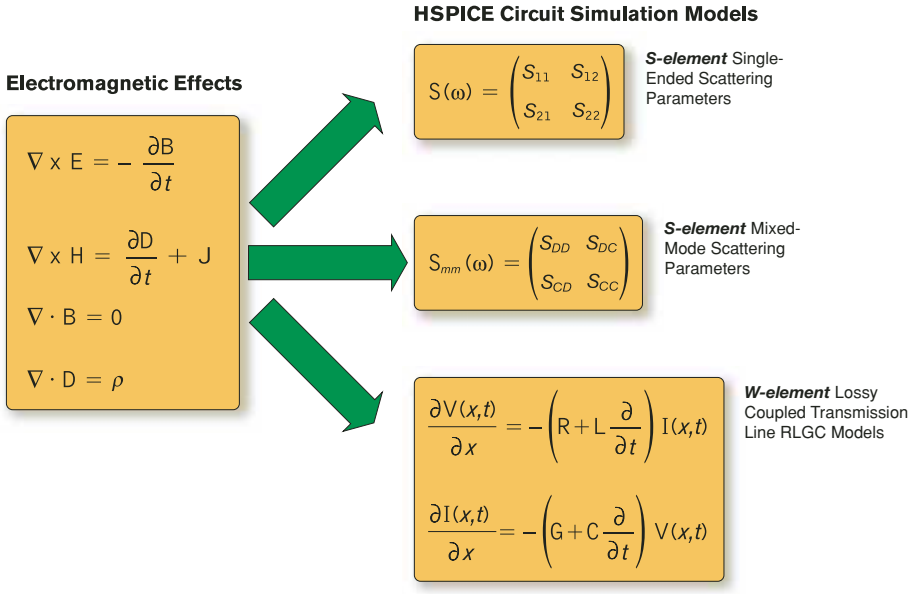

Figure 1 – Achieving accurate gigabit signaling channel simulations mandates the use of models that can take into account key electromagnetic effects.

Meeting these challenges requires using signal integrity (SI) simulations to uncover and help solve potential problems before fabrication. This is nothing new, but the trick is to now take into account several previously ignored factors that are detrimental to gigabit link design.

Consider the traces. Perhaps by now you've grown accustomed to using transmission lines in signal integrity simulations. But simple lossless, uncoupled transmission line models are just not good enough for MGT links. Frequency-dependent conductor and dielectric losses – especially in FR4 – are substantial and mandate a more sophisticated approach. Your basic gigabit trace is a differential coupled transmission line with considerable loss and must be treated as such to find optimal driver pre-emphasis settings.

To address these and other problems, HSPICE® provides a comprehensive set of SI simulation and modeling capabilities to help you achieve the necessary accuracy for multi-gigabit SI simulations. HSPICE includes:

- Built-in electromagnetic (EM) solver technology for trace geometries
- Lossy, coupled transmission line modeling with the W-element
- Single-ended and mixed-mode S-parameter modeling with the S-element
- I/O buffer modeling with I/O Buffer Information Specification (IBIS) models and encrypted netlists.

Getting from Maxwell to Models

According to electromagnetic theory, at high frequencies every millimeter of metal will influence electrical behavior. As depicted in Figure 1, one challenge in multi-gigabit SI is to reduce the significant aspects of EM theory into something useful for circuit-level simulation. Maxwell's equations must be reduced to something manageable; you must analyze the electromagnetic characteristics of the interconnect system to build an appropriate model for circuit simulation.

HSPICE includes a built-in electromagnetic field solver for computing the electrical characteristics of coupled transmission line systems. The solver is ideal for multi-lane, multi-gigabit applications. It uses a Green's function boundary element and filament method that yields very accurate resistance, inductance, conductance, and

capacitance (RLGC) matrices for the types of differential traces you'll need for gigabit design. You need only perform a field solver analysis for each unique cross-sectional geometry.

HSPICE field solver analysis will produce a characterization of the interconnect system in terms of distributed RLGC matrices. Frequency-dependent loss effects are included in the R_s and G_d matrix elements. Be sure to enable these field solver options; at gigabit data rates these losses can be substantial.

The conductor losses ($\propto \sqrt{f}$) and dielectric losses ($\propto f$) are both significant at 3.125 Gbps, and must be well modeled to determine your pre-emphasis needs for long lane lengths. Don't guess when specifying your material properties. The relative dielectric constant (4.2-4.7 for FR4) will influence line impedance (C matrix) values; electrical conductivity (5.8e7 for copper) will show up as skin effect (R matrix) losses; and dielectric loss tangent values (typically 0.015-0.03 for FR4) will show up as substrate (G matrix) losses.

Fortunately, board manufacturers are getting better at measuring and sharing such information. Many accurate W-element RLGC matrix models are available directly from vendors. Be sure to verify that frequency-dependent R_s and G_d values are included to ensure that loss modeling was taken into account. HSPICE's built-in EM solver is also well suited for copper cable geometries in cases where manufacturers do not have W-element models available.

Mixed-Mode Scattering Parameters

As shown in Figure 2, accurate SI simulation of multi-gigabit links involves a variety of models. For certain package, trace, connector, backplane, and cable sections, measured data or very accurate three-dimensional EM solver data is often available in the form of scattering parameters (Figure 3).

S-parameters represent complex ratios of forward and reflected voltage waves. Used as an alternative to other frequency domain representations (such as Y- or Z-parameters), S-parameters lack the dramatic magnitude variations that other representa-

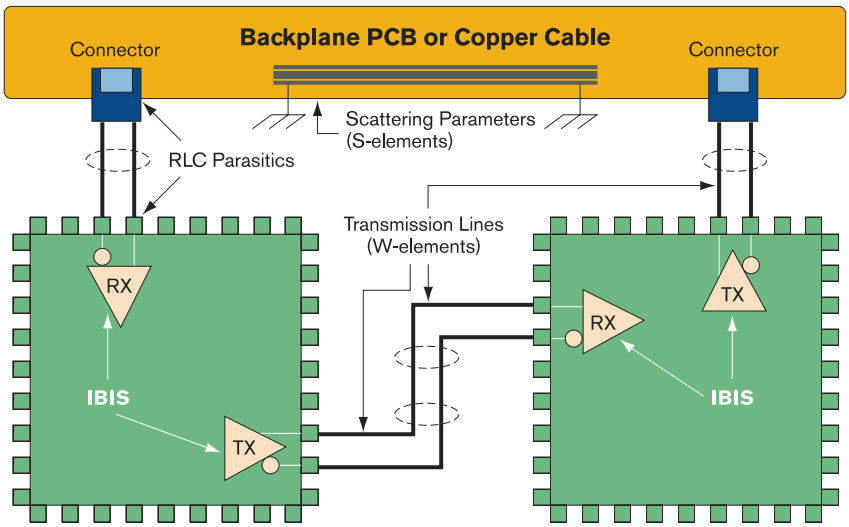

Figure 2 – Simulations for MGT chip-to-chip, backplane, and copper cable applications combine a diverse set of models for accurate signal integrity predictions.

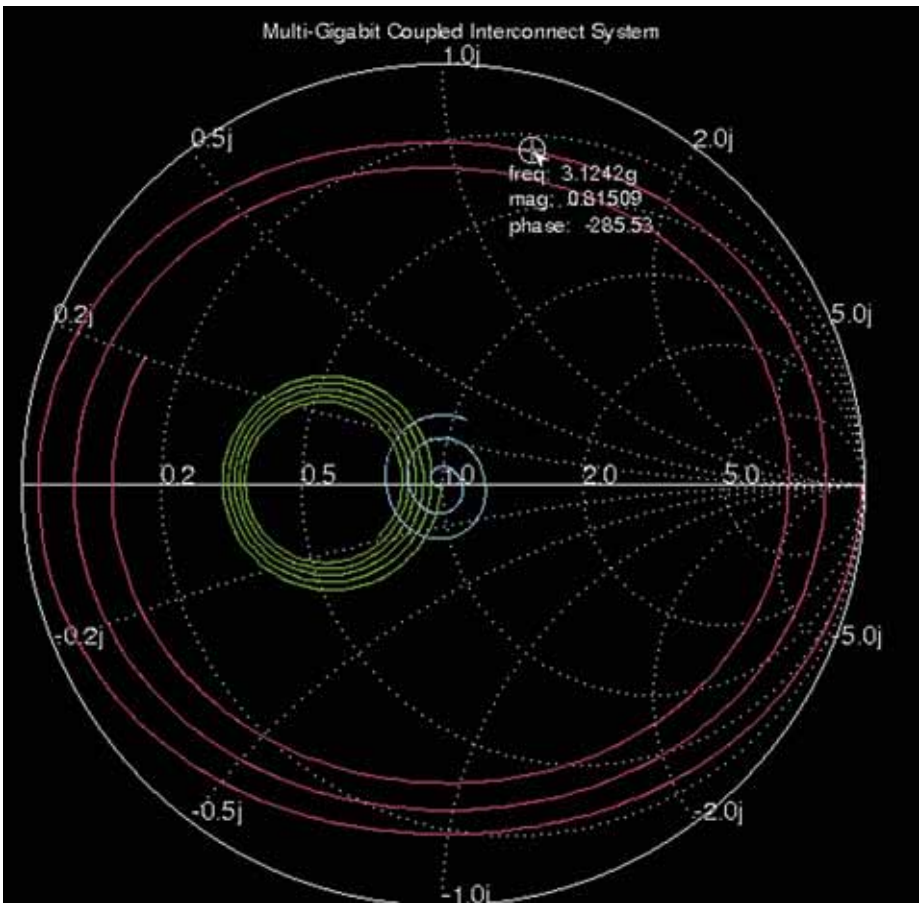

Figure 3 – Typical scattering parameters for an interconnect system showing the transmission coefficient (S21) for one interconnect (violet), the reflection coefficient (S11) for the same interconnect (green), and the coupling coefficient (S31) between adjacent interconnects (light blue) over a frequency sweep of 0-10 GHz.

tions have associated with high-frequency resonance. In addition, they can be measured directly with vector network analyzers. With differential traces the norm for XAUI and other links, mixed-mode S-parameters are particularly useful. They provide a means to characterize a differential trace in terms of its differential, common-mode, and cross-coupled behavior.

HSPICE provides single-ended and mixed-mode S-parameter modeling capability through the S-element. You can input S-parameter data in Touchstone™ file, CITI file, or table formats. Make sure your S-parameter data covers as broad a frequency range as possible with good sampling.

HSPICE will apply convolution calculations that need high-frequency values for crisp simulations of waveform rises and falls. If you have data up to 20 or 40 GHz, use it. A frequency range nine times your data rate (28 GHz for 3.125 Gbps) is considered optimal, although often hard to come by. Good low-frequency data (including DC) is also important for direct-coupled applications.

Beware of "measurement noise" with S-parameters. A poor network analyzer calibration can result in S-parameter data that will make your passive traces appear to have gain. HSPICE also supports S-parameter modeling for active devices, as is common with some RF/microwave designs. HSPICE uses a convolution algorithm for S-parameter modeling that is not limited to passive devices, avoiding the creation of intermediate, reduced-order models required by other time-domain simulation approaches. HSPICE uses the S-parameter response directly for maximum accuracy.

I/O Buffer Modeling

Ideally, you can perform SI simulations using transistor-level models and netlists for the input/output buffers. This level of detail may be unwieldy, but is sometimes necessary. The IBIS standard provides a means of encapsulating the key electrical characteristics of I/O buffers into accurate behavioral models. These models include data tables for buffer drive and switching ability, and package parasitic information. These models may or may not be appropriate for high-speed applications, depending on their intended use. Be sure to check the notes in the header of your IBIS model files so that you're not pushing the model outside its range of validity. There is also a new IBIS Interconnect Modeling Specification (ICM) for exchanging S-parameter and RLGC matrix data for connectors, cables, packages, and other types of interconnects.

HSPICE provides single-ended and mixed-mode S-parameter modeling capability through the S-element.

Another advantage of IBIS is that it allows vendors to deliver good buffer models to their customers without disclosing proprietary design information. This is also accomplished with encrypted HSPICE netlists. Multi-gigabit transceiver modeling is particularly difficult, so be prepared to see several buffer modeling approaches.

In the case of RocketIO transceivers, Xilinx provides special MGT models verified with HSPICE; visit the Xilinx Support SPICE Suite at *www.xilinx.com/support/software/spice/spice-request.htm* for more information. Whether you're using IBIS, SPICE netlist, or encrypted buffer models, HSPICE provides the most comprehensive and validated solution available.

Don't Skimp on the SPICE

So now you've got S-parameter models based on measured data, W-element trace models built from EM solvers, and accurate I/O buffer models. Are you ready to simulate? Maybe not. You may still be missing lumped R, L, and C values needed to capture all the parasitic effects in your design. Are you using AC coupling capacitors? At gigabit frequencies, no passive component behaves completely as expected. Even coupling capacitors must be modeled as lumped RLC circuits to capture resonance effects. Using off-chip terminations? The same is true with resistors. Are you leaving out any package lumped RLC or S-parameter models? Thankfully, manufacturers are getting better at providing accurate SPICE models for most of their components. You just need to ask.

Conclusion

Multi-gigabit signal integrity simulations must take into account a great deal of previously ignorable effects. Every trace is a transmission line, and you must account for every bump, bend, turn, and millimeter of metal with appropriate electrical models.

HSPICE is constantly being improved to better address these accuracy needs for multi-gigabit SI simulation. The W-element has been enhanced for faster and more accurate modeling of frequency-dependent losses in coupled transmission lines. HSPICE's built-in EM solvers can build accurate W-element models based on trace geometries (Table 1). The S-element has been enhanced to support both single-ended and mixed-mode S-parameter data sets. This, combined with HSPICE's trustworthy device and IBIS models, provides a comprehensive signal integrity simulation and modeling solution.

For more information about the latest capabilities of HSPICE and the integration of HSPICE into overall design processes, visit the HSPICE Update page at *www.hspice.com.*

Use the Following Command:	To Specify Trace:
.MATERIAL	Conductor and dielectric properties
.SHAPE	Conductor geometries
.LAYERSTACK	Ground planes and dielectric thicknesses
.MODEL	W-element model derived from the field solver analysis

Table 1 – Use HSPICE's built-in EM solver to turn material properties and trace geometry specifications into accurate lossy, coupled transmission line models.

Eyes Wide Open

The RocketIO Design Kit for ICX reduces the burden of implementing working multi-gigabit channels.

by Steve Baker
High Speed Architect, Systems Design Division
Mentor Graphics Corporation
steve_baker@mentor.com

If you're migrating from traditional bus standards such as PCI and ATA to serialized asynchronous architectures such as PCI Express™ and ATA-2, you've probably discovered that the tools for simulating the designs and models for the various buffers, connectors, transmission lines, and vias have become more complex.

Although setup and hold, crosstalk and single-ended delay are well understood, accurately modeling these new parts and their various complex behaviors adds to the job's complexity. To reduce the complexity of interacting with model and design parameters, Mentor Graphics and Xilinx have jointly developed the RocketIO™ Design Kit for ICX™ software, producing a design environment that allows you to fully confirm what's required to satisfy your design specifications.

The Design Kit

The RocketIO Design Kit for ICX is a companion to the standard Xilinx Signal Integrity Simulation (SIS) Kit and comprises a set of designs that match various Xilinx-supplied SPICE transmission line implementations. The kit is hierarchical, so all of the different elements – such as documentation, system configuration, simulation models, and ICX databases – are stored in different, relative location folders. These folders are located within the ICX kit in the same parent directory as the Xilinx SIS kit.

The design kit enables easy simulation analysis through the RocketIO menu and through existing features of ICX products, including eye-diagram, jitter, and inter-symbol interference analysis using predefined and custom multi-bit stimuli with lossy transmission line modeling.

Additionally, the IBIS 4.1 models, which ICX uses for simulation, reference the encrypted models supplied by Xilinx. You can progress from design to design through the kit's environment, learning more about the behavior of the RocketIO buffers with each design or simulation, such as what is achievable with these buffers in a multi-gigabit channel and what settings are required to maximize system performance.

The custom menu is more full-featured, allowing direct simulation and eye diagram display of any of the 10 pairs from a single menu selection.

Standard Designs

The three standard designs supplied with the RocketIO Design Kit include:

- Correlation
- Example
- Evaluation.

You can also verify your own design, either in pre- or post-route states, in the kit's design area.

The Correlation Design

In a correlation design, the ICX database reproduces the interconnect scheme (Figure 1) from the Xilinx backplane example and uses the same drivers and receiver buffer models and parameters. The ICX database provides virtual "push button" operation so that you can run a signal integrity simulation and compare the resulting waveform with that provided in the Xilinx Rocket IO Design Kit in eye-diagram form. You can also verify that simulation results match those supplied by Xilinx with either the ICX self-contained simulation environment using ADMS SI or with HSPICE® as an external simulator called from within ICX.

The Example Design

The example design has an expanded set of transmission line examples to match the 10 examples that Xilinx supplies. Each of the 10 paths comprises a RocketIO transmitter connected to a Teradyne™ HSD five-row connector through two inches of differential board traces; 16 inches of differential board traces to a second Teradyne HSD five-row connector; and finally two inches of differential board traces from the second Teradyne HSD five-row connector to a RocketIO receiver.

The custom menu allows direct simulation and eye diagram display of any of the 10 pairs from a single menu selection. The menu also includes additional configuration and pulse train dialogs that you can use to change the simulation parameters, thus allowing investigations of RocketIO buffer behavior with these different settings and stimuli.

In the example design, because the transmission lines are fixed, you modify the various settings of the buffer itself and then conduct a simulation on whichever differential channel you want to investigate.

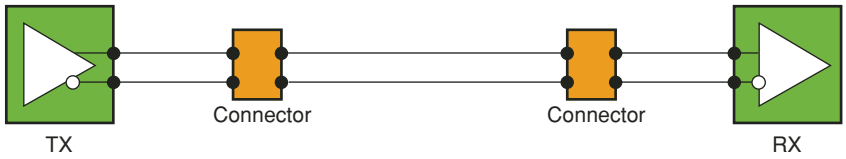

Figure 1 – Generic schematic of the design under simulation

The built-in RocketIO configuration utility allows changes to the temperature and bit duration settings when using the models directly from the Xilinx IBIS writer utility. It also gives you additional freedom to set the pre-emphasis level, driver/receiver termination values, and differential voltage swing when evaluating other possible solutions.

To enable different bit-patterns and speeds, you can also change the pulse train from the standard 3.125 GHz to your own specified pulse train using the pulse train generator. This utility allows you to specify bit patterns that can be used directly in ICX or exported as an ASCII file, in either SPICE PWL format or VHDL-AMS time vectors, toggling between state transitions.

The bit-patterns have an underlying pulse duration over which you can add jitter, where the peak-to-peak value specifies the six sigma points in picoseconds of this Gaussian random number. The pattern can be a user-defined set of ones and zeros, automatically defined as a random number of user-defined pattern length or as a pre-defined pattern. Pre-defined pattern styles include several pseudo-random bit sequences and Fibre Channel pulse trains (Figure 2).

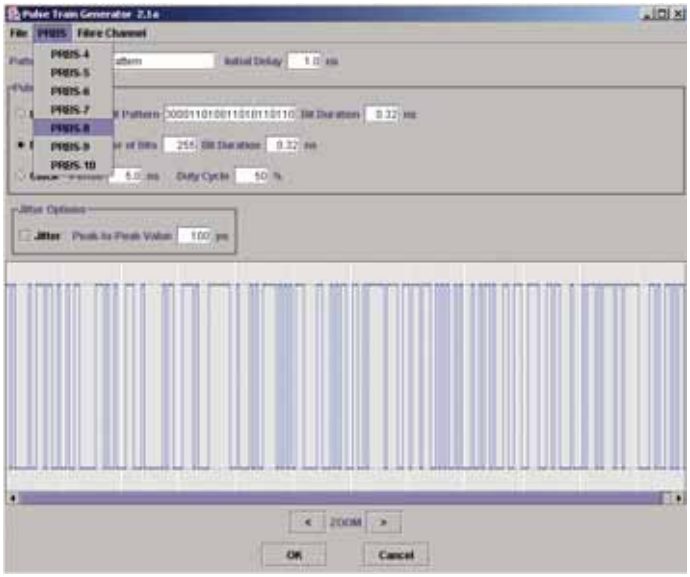

Figure 2 – Pulse train dialog showing pseudo-random bit pattern

The Evaluation Design

The evaluation design allows you to load a pre-defined cross section that matches one of the cross sections from the example design. In this virtual prototype environment, you can place actual parts, try "what-if" routing, and see the results in an eye diagram. As the IBIS part models include other buffers for Virtex-II Pro™ devices, you can simulate the whole of the FPGA rather than just the RocketIO channel.

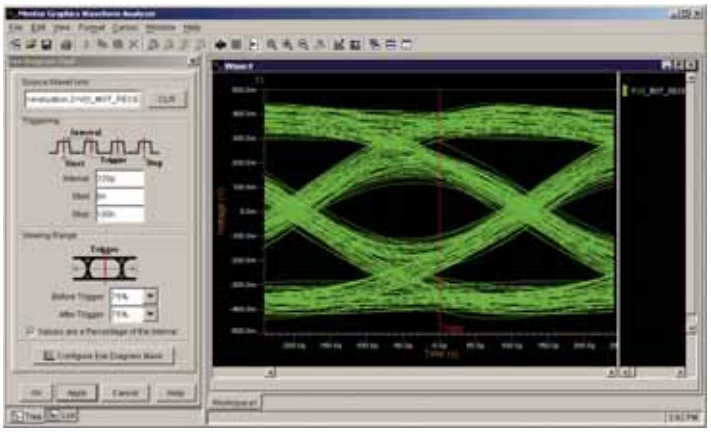

Figure 3 – A 3.125 GHz eye diagram from the evaluation design

This is where the channel's design is investigated in greater detail, as you initially place the devices to match your expected end design rather than using a fixed set of transmission lines. Using the electrical editor functionality of the IS floorplanner tool, you can add additional parts such as connectors or terminators and evaluate the impact of these on the resulting eye diagram. When working with these items, you can quickly determine the result of the different pre-emphasis settings. Additionally, you can see the impact of different routing strategies, including the fan-out pattern and tightly or loosely coupled differential pairs.

In the evaluation design, you can determine how much pre-emphasis is required to create the desired eye, as well as what level of noise is introduced on adjacent signals, on the board, or through the connector due to that level of pre-emphasis. The results of this virtual prototyping, as seen in the eye diagram in Figure 3, can be passed forward in the flow as constraints to drive the electrical design, as well as placement and routing examples.

Verification

The most advanced part of the kit allows you to simulate your design or system. The various parts of the system, backplane and plug-in cards, or just a single card with on-board channel, can be run through verification using the same complex pulse trains and model settings as before.

If required, you can modify settings to improve channel performance as measured by the eye. You can also define additional corner cases to evaluate best- and worst-case scenarios, including the impact of one pair on the other in terms of crosstalk; its impact on the shape and size of the eye; and the impact of other signals on the channel.

Conclusion

Iteration happens in any design process. The quicker decisions can be made in those iterations and the smaller the impact on existing design implementations, the happier we all are.

The RocketIO Design Kit for ICX allows you to make initial evaluations of the technology before any of the actual design implementation has occurred. As the design progresses forward from initial evaluations to the virtual prototype environment, you can confirm, in a pseudo-physical implementation, that the specifications can still be achieved, or use the kit to determine what changes are required to achieve the desired performance.

Finally, by verifying the placement, the routing of the multi-gigabit channels, or the whole design, you can confirm that you are within specification. For more information about the RocketIO Design Kit for ICX, visit *www.mentor.com/highspeed/resource/design_kits/icx-rocketio_designkit.html.*

Backplane Characterization Techniques

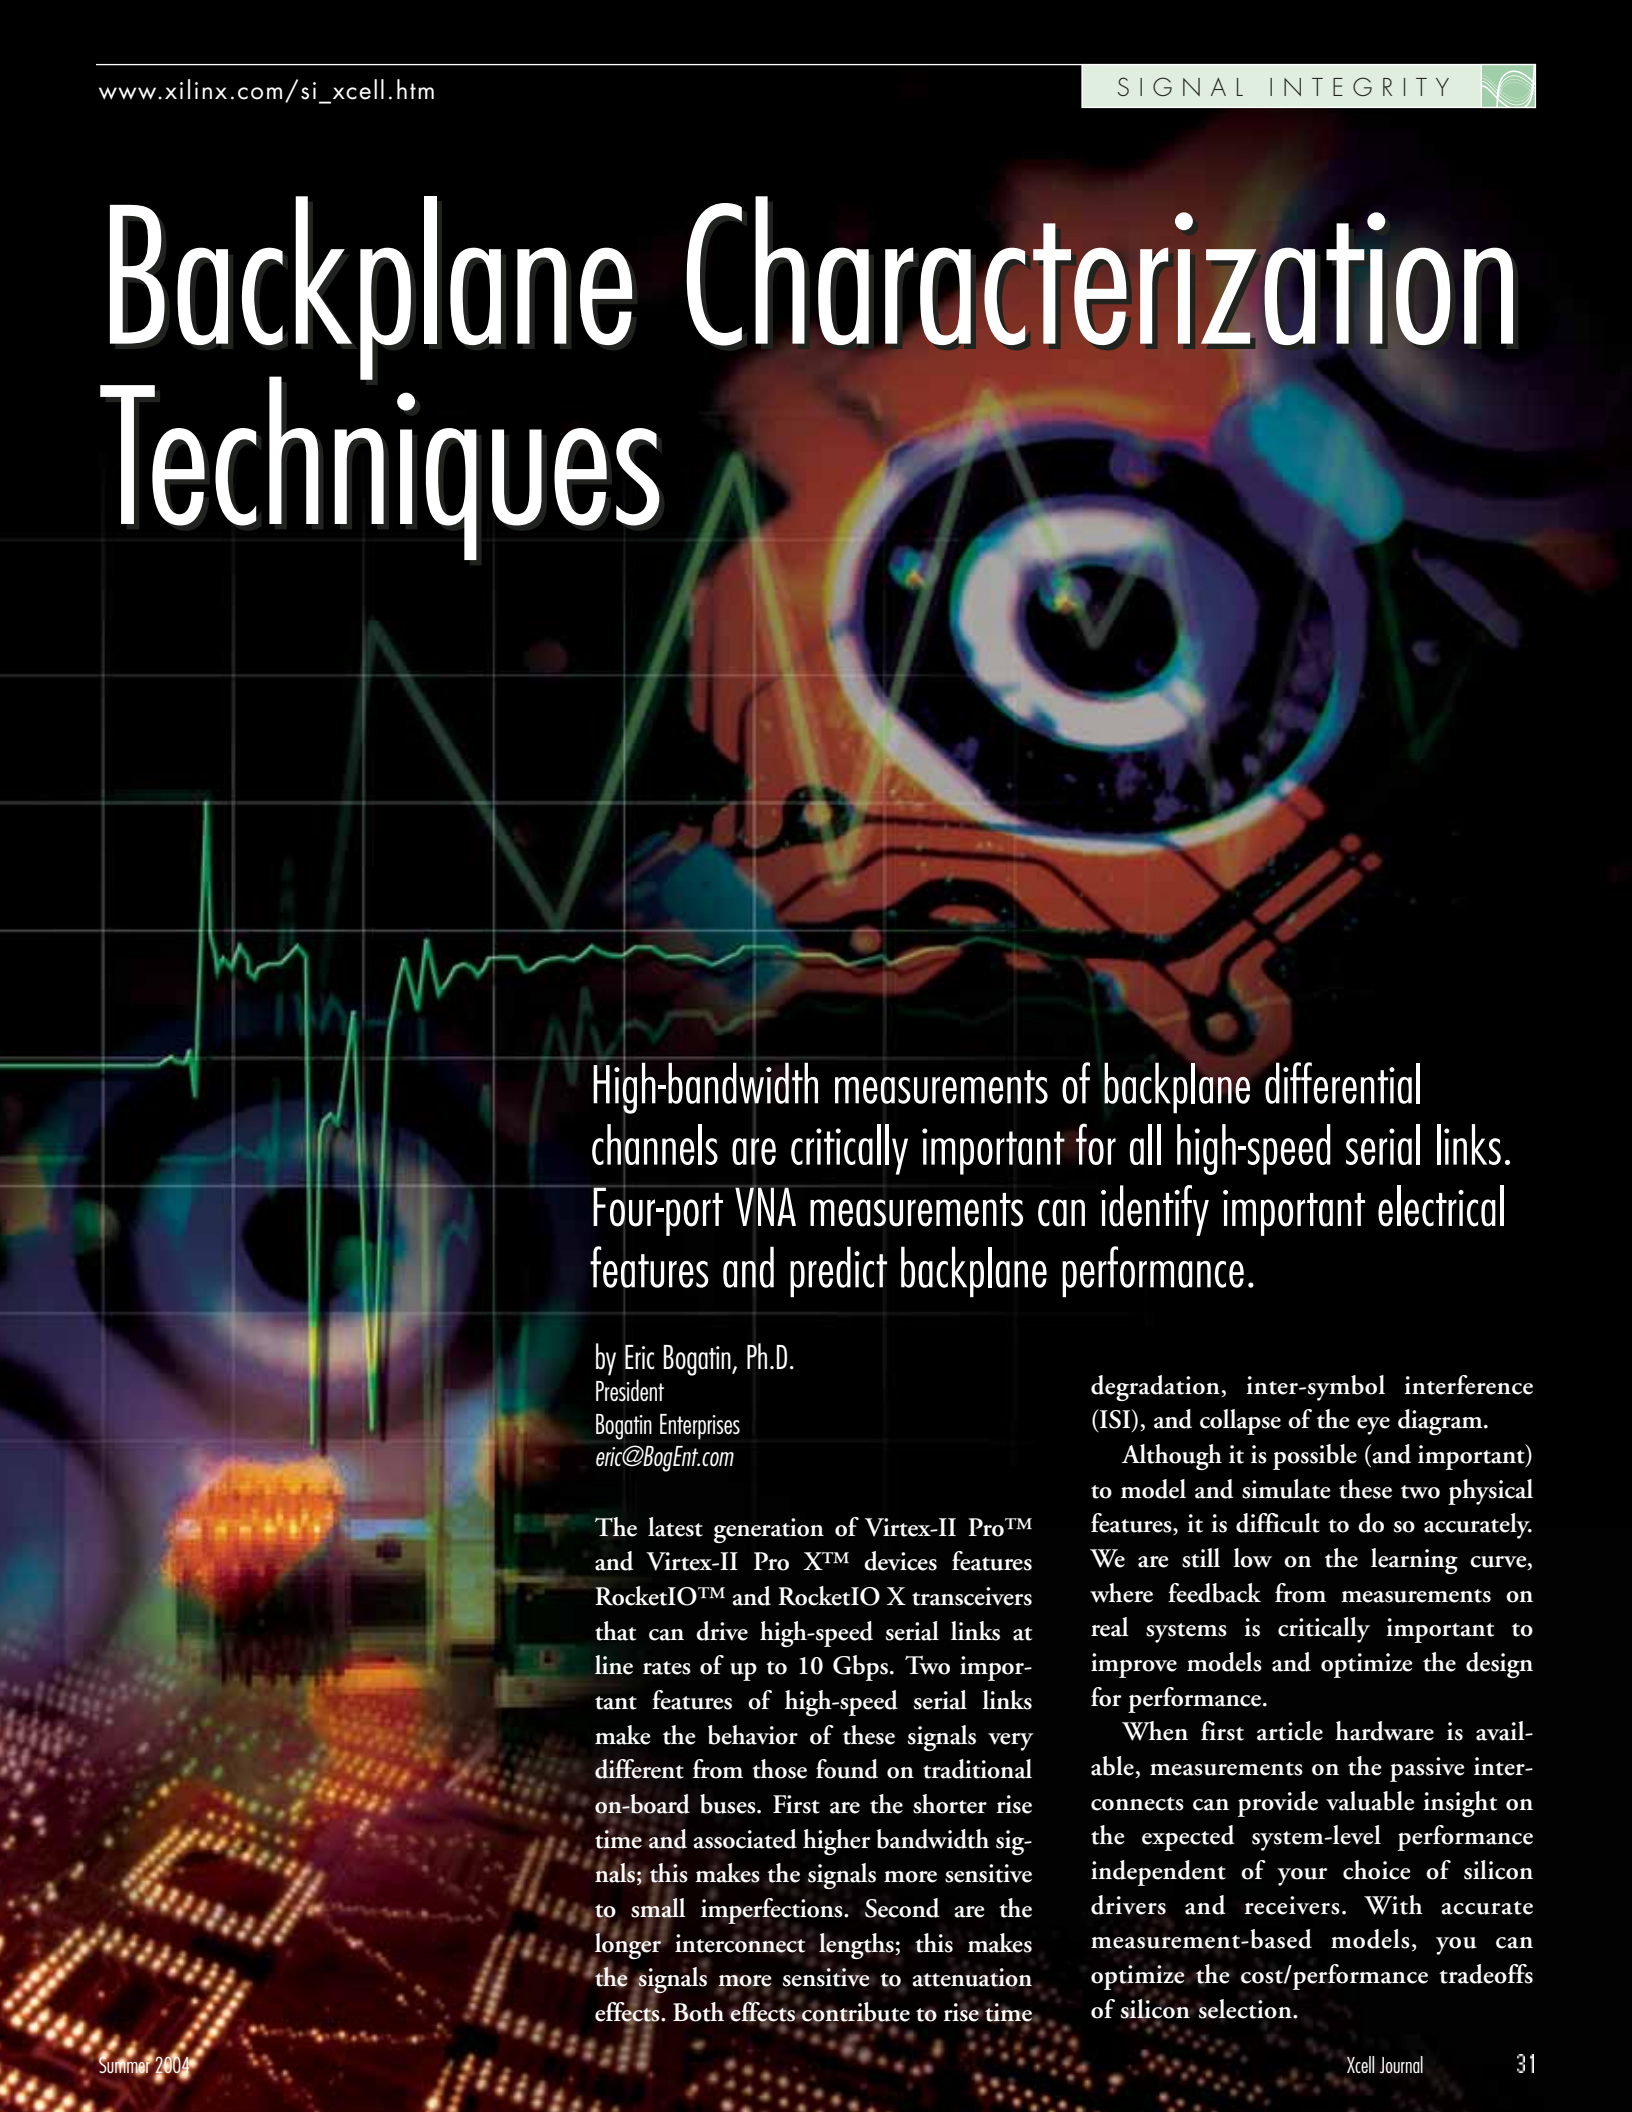

High-bandwidth measurements of backplane differential channels are critically important for all high-speed serial links. Four-port VNA measurements can identify important electrical features and predict backplane performance.

by Eric Bogatin, Ph.D.
President
Bogatin Enterprises
eric@BogEnt.com

The latest generation of Virtex-II Pro™ and Virtex-II Pro X™ devices features RocketIO™ and RocketIO X transceivers that can drive high-speed serial links at line rates of up to 10 Gbps. Two important features of high-speed serial links make the behavior of these signals very different from those found on traditional on-board buses. First are the shorter rise time and associated higher bandwidth signals; this makes the signals more sensitive to small imperfections. Second are the longer interconnect lengths; this makes the signals more sensitive to attenuation effects. Both effects contribute to rise time degradation, inter-symbol interference (ISI), and collapse of the eye diagram.

Although it is possible (and important) to model and simulate these two physical features, it is difficult to do so accurately. We are still low on the learning curve, where feedback from measurements on real systems is critically important to improve models and optimize the design for performance.

When first article hardware is available, measurements on the passive interconnects can provide valuable insight on the expected system-level performance independent of your choice of silicon drivers and receivers. With accurate measurement-based models, you can optimize the cost/performance tradeoffs of silicon selection.

The Bandwidth of the Measurement

Bandwidth is the highest sine wave frequency component that is significant. "Significant" means the frequency at which a harmonic of the signal is greater than -3 dB of the amplitude the same harmonic an ideal square wave at the same clock frequency would have.

If the signal edge is roughly Gaussian with a 10-90% rise time (RT), the bandwidth (BW) is approximately:

$$BW = \frac{0.35}{RT}$$

For example, a rise time of 0.1 ns has a bandwidth of about 0.35/0.1 ~ 3.5 GHz.

Usually, the bit rate is specified in a high-speed serial link. To estimate the bandwidth of the signal, we need to have an estimate of the rise time. Assuming that the rise time is 25% of the bit period, then the bandwidth of the signal is approximately:

$$BW_{signal} = \frac{0.35}{0.25} BR \sim 1.4 \times BR$$

As a general rule of thumb, the highest sine wave frequency component in a high-speed serial link is about 1.4 times the bit rate. For a 2.5 Gbps signal, the bandwidth is about 3.5 GHz. If it is important to know whether the bandwidth is really 3.5 GHz or 4 GHz, the term "bandwidth" is misused, as it is not accurate enough to make this fine a distinction. Rather, you should use the entire spectrum.

To have confidence in the accuracy of a model, the bandwidth of that model – the highest sine wave frequency at which the simulated electrical performance still matches the measured performance of the real structure – should be at least twice the bandwidth of the signal to allow for a reasonable margin. Likewise, the bandwidth of the measurement should be at least twice the bandwidth of the signal. This rule of thumb suggests that the bandwidth of the measurement should be at least:

$$BW_{measurement} = 3 \times BR$$

If the bit rate is 10 Gbps, the bandwidth of any model used (or the bandwidth of the measurement of the interconnect) should be at least 30 GHz. Of course, if the rise time of the bit pattern is longer than 25% of the bit period, the measurement bandwidth might be reduced from this rule of thumb.

Unfortunately, the higher the bandwidth required, the more expensive it is (both in resources, time, and money) to perform a measurement or create a model of an interconnect. That is why it is so important to have a rough idea of the bandwidth requirements so as to minimize the cost. As high-speed serial links approach the 10 Gbps rate, measurement bandwidths need to be at least 30 GHz. Accurate measurements in this regime get increasingly more difficult with each generation of bit rate.

No Such Thing as a Free Launch

Credit that clever turn of phrase to Scott McMorrow, president of Teraspeed Consulting. Probing a channel on a board or a backplane introduces errors that might not be there, or be of a different magnitude, than in the actual product when signals are launched from chips in packages.

All high-performance measurement instruments, such as a time domain reflectometer (TDR) or a vector network analyzer (VNA), have a standard connector on the front face, typically APC-7 or 3.5 mm. High-performance cables are used to get from the instrument to the device under test. However, the interface from the cable to the board traces under test can introduce impedance discontinuities which degrade the signal getting onto the trace.

The larger the discontinuity, the more high-frequency components reflect back to the source, and the fewer that get launched into the transmission line. If characterizing a path for 5 Gbps signals, the connection method may limit the measured system performance. To increase the bandwidth of the characterization, you must consider the launch before designing and building the board.

A key ingredient in the design for test for high-bandwidth characterization is to use a pad and via design transparent to the signal. This typically means using a small diameter via with a surface-mount connector and optimizing the clearance holes in the planes. Alternatively, you could use a copper fill adjacent to the signal via being probed, with the copper fill connected to return path vias adjacent to the signal via so you could use microprobes.

Figure 1 shows the TDR response for different connection designs. The top curve is the TDR response (with a roughly 35 ps rise time) for a conventional through-hole Sub Miniature version A (SMA) connec-

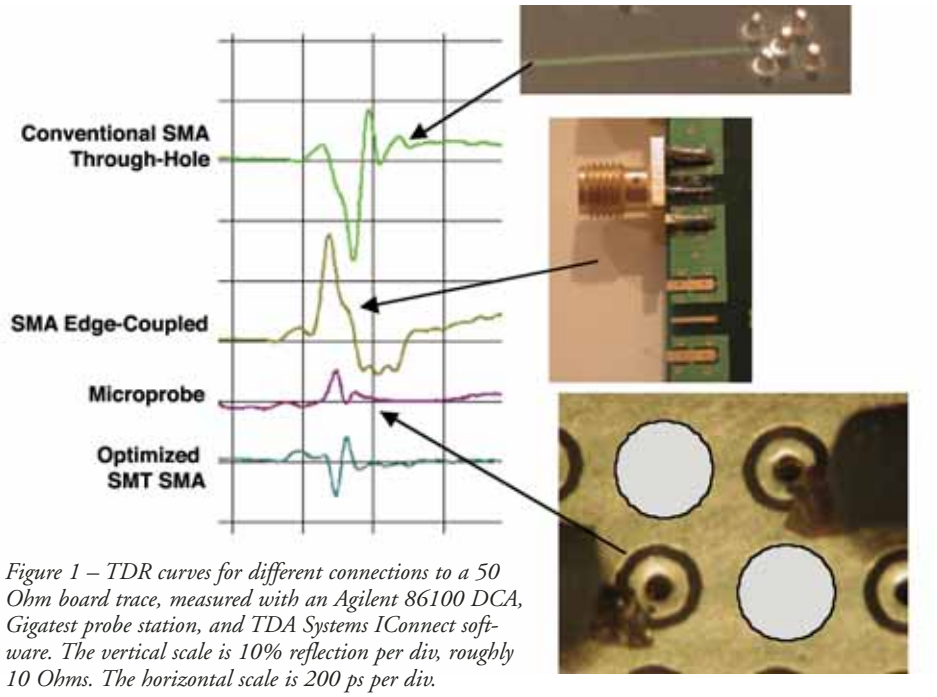

Figure 1 – TDR curves for different connections to a 50 Ohm board trace, measured with an Agilent 86100 DCA, Gigatest probe station, and TDA Systems IConnect software. The vertical scale is 10% reflection per div, roughly 10 Ohms. The horizontal scale is 200 ps per div.

You might think that avoiding the vias will prevent the impedance discontinuity, but just as many problems can be generated by an edge-coupled SMA attached directly to a surface trace.

tion to a bottom trace. On this scale, one division is a reflection coefficient of 10% and corresponds to an impedance change of about 10 Ohms. At this rise time, the impedance discontinuity is more than 18 Ohms, and is predominately capacitive.

You might think that avoiding the vias will prevent the impedance discontinuity, but just as many problems can be generated by an edge-coupled SMA attached directly to a surface trace. The second curve in Figure 1 shows the measured TDR response of an edge-coupled launch using an SMA. The impedance discontinuity is more than 18 Ohms at this rise time and is inductive.

One way to avoid this problem is to use microprobes and design the surface pads for probing. The key feature is to use a copper fill shorted to all adjacent ground vias. In Figure 1, the gray vias have been shorted to the copper fill. With this configuration, you can probe every signal.

The third TDR curve in Figure 1 shows the response of a microprobe launch into an optimized 50 Ohm stripline. The impedance discontinuity at this rise time is less than 5 Ohms and is inductive.

Finally, it is possible to use an SMA connection to a circuit board trace if it is optimized. The bottom curve in Figure 1 shows such a connection. Its impedance discontinuity, less than 5 Ohms, compares to a microprobe launch.

High-Bandwidth Measurements

All high-bandwidth measurements take advantage of what is normally a problem encountered by high-bandwidth signals: reflections from impedance discontinuities. As a signal propagates down an interconnect, if the instantaneous impedance the signal sees ever changes, a reflection will occur and the transmitted signal will be distorted. The magnitude of the reflected signal will depend on the change in impedance.

By using a calibrated reference signal – a sine wave in the frequency domain and a Gaussian step edge in the time domain – and measuring the amount of signal reflected back from an interconnect as well as transmitted through it, you can extract the electrical properties of the interconnect. All of the electrical properties of the interconnect path are contained in these two basic measurements.

When displaying data in the frequency domain, the reflected signal is called the return loss and the transmitted signal is called the insertion loss. These two metrics have become the universal standard to characterize the fundamental properties of an interconnect, such as a channel path in a backplane. Many of the important physical layer properties of a backplane can be read directly from the return and insertion loss of both single-ended and differential channels.

When displaying data in the time domain, the reflected signal gives direct insight into how the physical structure contributes to electrical impedance discontinuities. The transmitted signal in the time domain gives a direct measure of the propagation delay and rise time degradation. From this result, an eye diagram can be synthesized.

Whether you've measured the data in the time or frequency domain, it can be transformed into either one. A VNA will measure the response in the frequency domain, while a TDR will measure the response in the time domain. With appropriate software, you can convert the data from either instrument into both domains.

All high-speed serial links today use differential signaling and backplane channels routed on differential pairs. For these structures, the same metrics of return and insertion loss are used, but there are additional terms. Both differential and common signals will have a return and insertion loss, with mode conversion terms of differential signal in, common signal out and common signal in, and differential signal out.

Differential S-Parameters

The description of return and insertion loss measurements borrows from a formalism heavily used in the RF world based on scattering or S-parameters. It's just a shorthand way of keeping track of all the different measurements.

In a differential channel, the interconnect is a single, differential pair, with the two ends labeled port 1 and port 2. The ratio of the reflected sine wave signal coming out of port 1 to the incident sine wave signal going into port 1 is labeled S11. This is the return loss.

The ratio of the transmitted sine wave signal coming out of port 2 to the incident sine wave signal going into port 1 is labeled S21. This is the insertion loss.

A complication arises in a differential pair, where you must consider not only the port at which signals appear but also the nature of the signal (differential or common). There are four choices:

- A differential signal going in and coming out, which would be the differential return and insertion loss, SDD11 and SDD21
- A common signal going in and coming out, which would be the common return and insertion loss, SCC11 and SCC21
- A differential signal going in and a common signal coming out, a type of mode conversion, SCD11 and SCD21
- A common signal going in and a differential signal coming out, a type of mode conversion, SDC11 and SDC21.

Don't forget the case of the signal going in from port 2 rather than port 1. All of these combinations result in 16 differential S-parameters, which are arrayed in a matrix. Each set of terms has significance, but the most important are the differential return and insertion loss and the differential to common mode conversion.

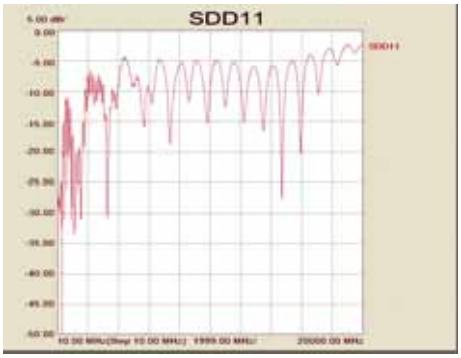

Figure 2 – SDD11 in the frequency domain for a backplane channel, measured with an Agilent PNA N4421b four-port VNA and PLTS software.

Figure 3 – SDD11 in the time domain for a backplane channel, measured with an Agilent PNA N4421b four-port VNA and PLTS software.

Figure 4 – SDD21 in the frequency domain for two different length backplane channels, measured with an Agilent PNA N4421b four-port VNA and PLTS software. The red line is about 26 inches and the green is about 40 inches.

Differential Return Loss

SDD11 is a direct measure of the impedance discontinuities encountered by the differential signal propagating through the channel. Figure 2 is an example of the measured differential return loss of a backplane trace in the frequency domain up to 20 GHz. The more negative the decibel value, the less reflected signal and the better the impedance match.

It's a little difficult to interpret the measurement in the frequency domain. This is a case where transforming the data to the time domain gives immediate insight.

Figure 3 is the same data displayed in the time domain. In this display, you can identify the discontinuity from the SMA launch, the high impedance of the daughtercard, and the capacitive discontinuity of the vias in the backplane.

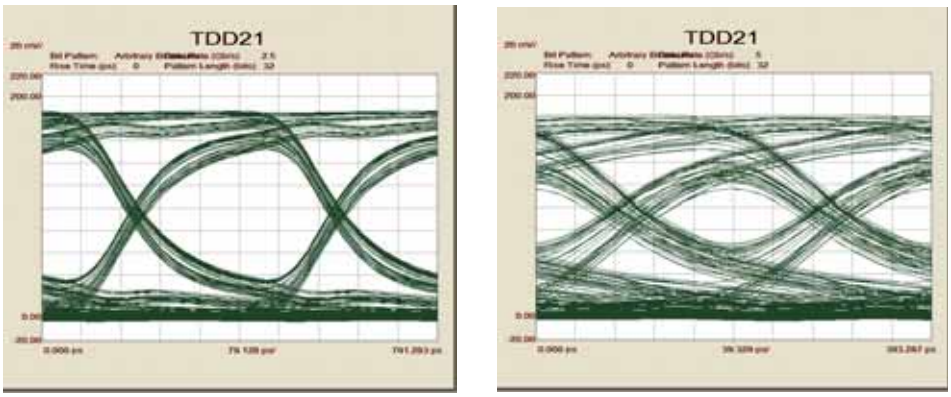

Figure 5 – Eye diagram calculated from SDD21 in the frequency domain for a backplane differential channel, measured with an Agilent PNA N4421b four-port VNA and PLTS software. Left is 2.5 Gbps and right is 5 Gbps.

Differential Insertion Loss

SDD21 is a direct measure of the quality of the transmitted differential signal through the channel. In the frequency domain we can read the bandwidth of the interconnect directly off the screen. The maximum useable bandwidth of the channel is set by the frequency at which the attenuation is below the usable value, typically about -15 dB of loss, depending on the SerDes. The more discontinuities and losses, the higher the attenuation, and the lower the bandwidth.

Figure 4 shows the measured SDD21 for two different length channels, including the higher bandwidth of the shorter channel.

Using the limiting attenuation as -15 dB, the short channel has a usable bandwidth of about 4 GHz, and the long channel has a usable bandwidth of about 3 GHz. This would correspond to a usable bit rate of roughly 2.5 Gbps and 2 Gbps. However, it is more than just the attenuation that determines the maximum usable bit rate.

A better estimator for the maximum usable bit rate is the eye diagram. Even though this differential insertion loss was measured in the frequency domain, it can be translated into the time domain, and as a response function can be used to calculate an eye diagram.

Figure 5 shows the calculated eye diagram for a 25-inch channel with 2.5 Gbps and 5 Gbps signals. Based on this measured response, this channel might be useful for even 5 Gbps data rates, with an appropriate receiver.

Mode Conversion

Any asymmetry between the two lines that make up the differential pair will convert some of the transmitted differential signal into common signal. This will create two problems. If any of this created common signal gets out of the channel onto external twisted pairs, it will potentially contribute to electromagnetic interference. Of course, every good design should have integrated common signal chokes in all external twisted pair connectors. However, it is always good practice to try to reduce the source of the noise before filtering.

The second problem isn't so much from the common signal created but from the impact on the differential signal from what caused the conversion. One of the most common sources of mode conversion is a difference in the time delay of each channel. This line-to-line skew within a channel

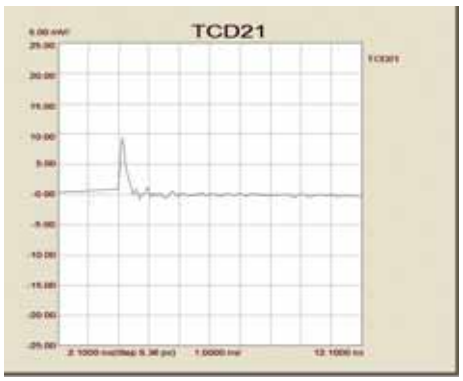

Figure 6 – SCD21 displayed in the time domain, showing the converted common signal when the incident differential signal is 400 mV. The conversion is about 2.5%.

will convert differential signals to common signals and result in increased rise time degradation of the differential signal and larger deterministic jitter.

The total amount of common signal coming out of port 2, based on a pure differential signal going into port 1, is described by the SCD21 term. Figure 6 shows the response for this channel.

Looking at the time evolution of the creation of the converted common signal coming out of port 1, we can gain insight into where the conversion might be occurring. Figure 7 shows the SCD11 term, displayed in the time domain, compared with the SDD11 term, which has information about the physical features of the channel.

It appears as though most of the mode conversion occurs in the via field of the backplane side of the connector to the daughtercard. Additional mode conversion exists at each of the connector locations in the backplane. This might be caused by the via fields or an asymmetry between the two lines in the differential pair, such as a spatial difference in the dielectric constant each trace sees.

Conclusion

Everything you ever wanted to know about the electrical characteristics of a differential channel is contained in the differential S-parameters. They can be measured in the time domain or the frequency domain and displayed in either, and each one offers a different insight.

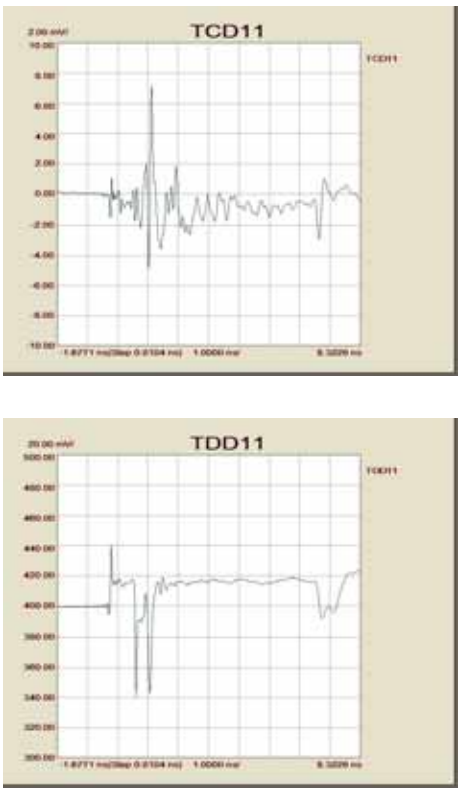

Figure 7 – Comparing SCD11 (top) with SDD11 (bottom) displayed in the time domain, showing the converted common signal coming out of port 1 coincident with the reflected differential signal out of port 1. This helps identify the location of the mode conversion.

Measurements play an important role in risk reduction when designing systems incorporating Rocket IO or RocketIO X transceivers. Although it is important to integrate simulation tools into the design process to perform cost/benefit analyses of technology and design tradeoffs, it is also important to use measurements to verify the accuracy of the simulation process.

Measurements can also offer immediate insight into the behavior of first article hardware to evaluate whether they meet specifications, and how well the interconnects will interact with the silicon. ☒

Additional Resources

For more information about this and other signal integrity topics, visit www.BogEnt.com.

Acknowledgments

The data in this paper was graciously provided by Maria Brown of Agilent Technologies and Al Neves and Dima Smolyansky of TDA Systems Inc.

A Low-Cost Solution for Debugging MGT Designs

Choose serial I/O technology for your designs without relying on expensive high-speed lab equipment.

by Joel Tan
Applications Engineer
Xilinx Global Services Division – Asia Pacific
joel.tan@xilinx.com

Xilinx Virtex-II Pro X™ devices contain RocketIO™ X multi-gigabit transceivers (MGTs) capable of 10 Gbps line rates, representing the leading edge of serial I/O performance. In Virtex-II Pro™ devices, up to 3.125 Gbps are available from each RocketIO transceiver, with the largest device in the family possessing 20 MGTs. When channel bonded together, they yield a single aggregated data channel with 62.5 Gbps of bandwidth.

At line rates as much as two orders of magnitude higher than single-ended I/O, lab and test equipment used in the development environment must keep up. Unfortunately, equipment designed for use with high-speed serial I/O systems may consume a large portion of program budgets.

Should limited access to high-speed equipment stop you from reaping the benefits of serial I/O? In this article, we'll present a solution that can lift this barrier to entry and make serial technology more accessible. It can also maximize the availability of expensive lab equipment for other projects.

The solution comprises a bit-error rate (BER) testing module connected to a flexible on-chip logic analyzer core, both implemented in FPGA fabric. Together with ChipScope™ Pro software tools, these two components can replace the diagnostic functions of a high-speed BER tester and logic analyzer, which together could cost more than $50,000.

RocketIO Design Flow Overview

Designing a RocketIO system requires you to simulate the system's digital and analog portions. Figure 1 shows the typical flow for an MGT design.

To ensure a reliable link, SPICE simulation of the analog system is mandatory.

An accurate setup must include all of the physical connections between transmitter to receiver, using accurate models for each of the vias, traces, connectors, and transmission media. (The importance of SPICE simulation is highlighted elsewhere within this series of signal integrity articles.)

At the same time, you must also simulate MGT functionality together with user logic; Xilinx provides MGT SmartModels for this purpose. Please refer to Answer Record #14596 in the Xilinx Answers Database for HDL simulator requirements.

Using the simulation results, you can then design and build the prototype board for further testing. It is during this hardware test, debug, and development phase that you can realize the benefits of this complete, low-cost debugging solution.

Debugging Challenges

The RocketIO MGT functional block diagram shown in Figure 2 is divided into two layers. Functions in the physical media attachment (PMA) layer are implemented digitally, while those in the physical coding sublayer (PCS) are predominantly analog.

Diagnosing a serial link issue is also split along the same divide: analog and digital.

Locating errors in digital logic is a familiar process because symptoms are easily reproducible and isolated. You can detect and fix deterministic errors in hardware by comparing captured data from a logic analyzer against expected data from simulation.

Problems are more difficult to diagnose for the analog portion, especially if errors seem to occur infrequently and randomly. Results vary from trial to

trial because of the random nature in which errors occur. However, over a number of repeated trials, it is possible to reproduce them reliably. The BER test does just this, and provides a useful metric for link performance.

Why Use BER Measurements?

BER equals the number of bit errors divided by the total number of bits transmitted. To measure the BER, test patterns are sent over the serial link and then compared to the original pattern at the receiver. Because the occurrence of errors is modeled as a stochastic process, a calculated minimum number of bits are transmitted before the BER is statistically valid. Xilinx Application Note XAPP661 discusses the method for calculating the confidence and precision of the BER measurement in detail.

Although many factors affect link performance, the final figure of merit for link reliability is the BER. These factors include signal trace design, clock quality, power integrity, and even impedance mismatches due to loose manufacturing tolerances. The BER metric has a systemic scope that covers all these factors, such that an anomaly in any part of the link (or its associated subsystem) will manifest as a higher than expected BER.

One assumption inherent to the BER measure is that the errors follow a Gaussian distribution. You should always test this by examining the distribution of errors in the data stream. If you observe bursts of errors, then the errors are non-random. This should prompt you to check if they are related to any noise sources, or even to the data pattern itself.

To simplify MGT designs, Xilinx provides a comprehensive list of power supply and oscillator recommendations within the RocketIO User's Guide. Power integrity is virtually eliminated as a potential cause for a high BER if these recommendations are strictly followed. Similarly, clock quality is addressed by the oscillator recommendations. To date, the majority of signal integrity issues have been traced to non-recommended power supply and oscillator configurations.

BER testing also verifies that your SPICE simulations resulted in a physical connection that delivers all the performance of which the silicon is capable. With power and clock quality taken care of, any difference between measured and simulated results comes down to the accuracy of the models and manufacturing processes. To differentiate between these, use time-domain reflectometry (TDR) measurements of the high-speed traces to check impedance deviations from the PCB specification.

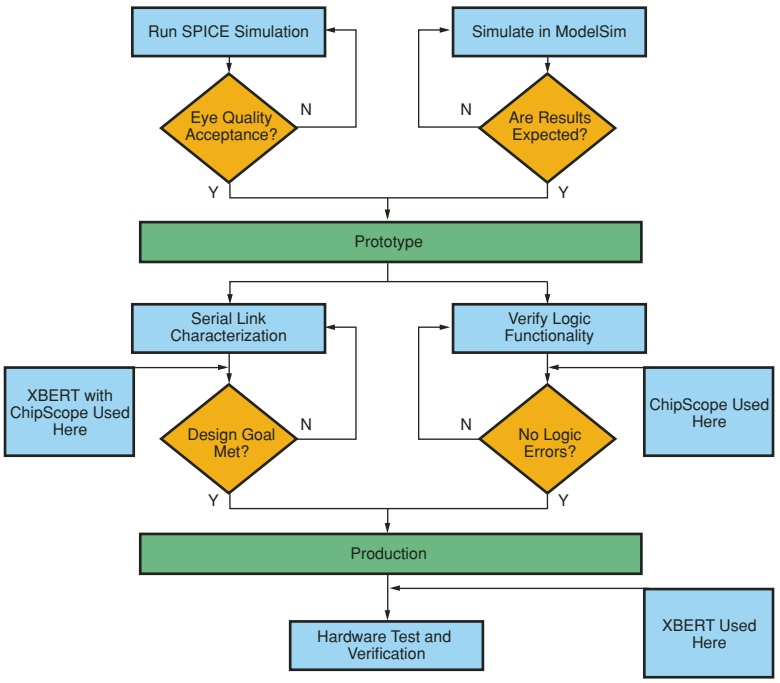

Figure 1 – Typical RocketIO design flow

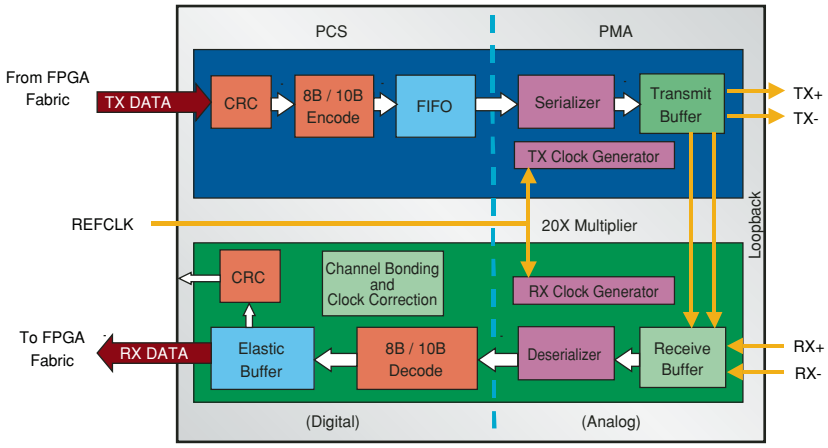

Figure 2 – MGT block diagram

Determining the root cause of poor BER is not straightforward, since multiple factors interact to produce the measured effect. However, you can observe how incremental changes affect link performance by comparing BERs before and after each change.

This is useful for quick what-if scenario testing of changes made to any part of the link, such as the PCB, power supply, clock source, connectors, and cables. An example of this is during a cost-down effort, where cost reductions are traded off with performance based on how each component change influences BER.

XBERT – The "Soft" BER Tester

The XBERT module pictured in Figure 4 measures BER and is delivered as a reference design with XAPP661. It uses an MGT to transmit serial data constructed

by a pattern generator, while a pattern follower and compare logic detects bit errors at the receiver.

Control signals into the module toggle resets and select between various pseudo-random bit sequence (PRBS) and clock patterns, while the outputs provide statistics for BER calculation.

An idle MGT, placed close to the active MGT, provides a simulated noise source that is useful when diagnosing interference from nearby MGTs. Such active noise is often coupled to other MGTs through the power supply or through poorly designed traces.

An appropriate test pattern must stress the link sufficiently to accurately simulate the data-dependent stresses that it will encounter with real traffic. The patterns in XBERT are International Telecommunication Union (ITU) recommended test patterns used in standards such as SONET and 10 Gigabit Ethernet.

By stepping through the various stress levels and running each for a short time, you can obtain a coarse measure of link performance quickly. As the link reliability improves, patterns should get harder and you will need to run tests for longer periods.

On its own, XBERT is by no means a complete replacement. (For example, jitter tolerance testing is required by some standards, which XBERT cannot perform.) But it can perform many of the more time- and resource-intensive measurements than BERT test equipment can. XBERT frees up lab equipment for other measurements and makes more lab resources available.

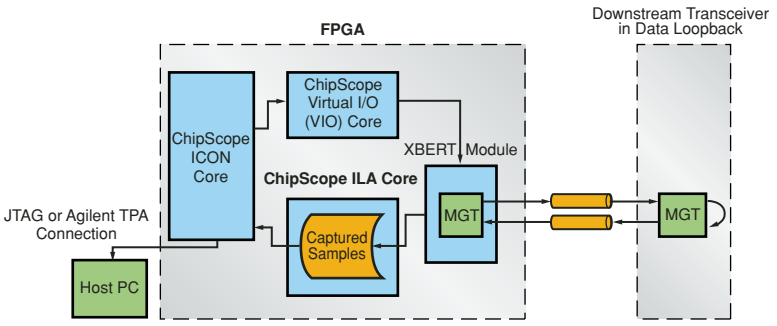

Figure 3 – XBERT with ChipScope software

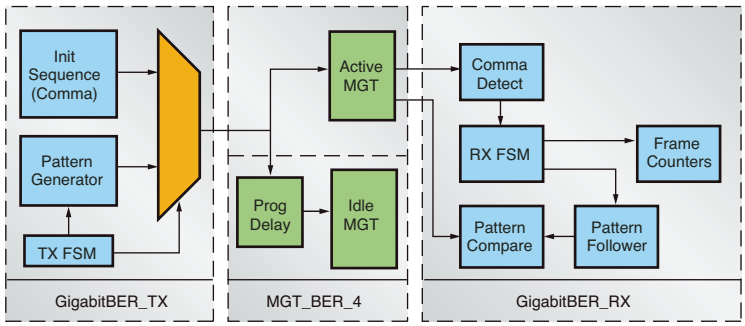

Figure 4 – Single-channel XBERT block diagram

Solution Overview

When implemented in Virtex-II Pro devices, the combination of XBERT and ChipScope software takes a form similar to the block diagram in Figure 3. In this particular test setup, the data is looped back at the far end so that both links are tested in the same trial. Alternatively, you can have another XBERT at the far end to test each link independently.

The inputs to XBERT are connected to ChipScope virtual I/Os (as shown) or to user logic. XBERT outputs such as the frame error count and bit error count are read by the ChipScope integrated logic analyzer (ILA) core and used as trigger conditions.

Together, the pair provides powerful diagnostic functionality as a data analyzer. You can trigger on a bit error or a combination of conditions to isolate certain types of errors. At the same time, you can sample the received data to examine the data pattern around an error condition.

This provides useful clues to identify the root cause of a bit error, especially if it is data-dependent. For example, if DC balance is disrupted, then bit errors will probably occur after long run lengths.

The ChipScope Pro tool implements a logic analyzer within the FPGA without additional hardware. It is a real-time debugging solution that lets you look at signals in a design as it is running. You can examine more ports simultaneously with ChipScope Pro software than with any other logic analyzer equipment available today.

Each FPGA requires a ChipScope ICON core to enable this JTAG connection to the host PC. In turn, the ICON core supports as many as 15 ILA, ILA/ATC, IBA/OPB, IBA/PLB, and VIO cores. The maximum number of signals possible per ILA core is limited by the amount of logic resources available up to a maximum of 16 trigger ports, each with a maximum width of 256 bits.

The ChipScope Pro analyzer GUI has a convenient waveform viewer that formats the sampled data in the same way as common HDL simulators. You can view MGT data and status signals as they appear in simulation, thus speeding up the verification process.

Typical Debugging Flow

Let's consider a scenario where you are debugging a new prototype board and bit errors are reported by the user logic.

ChipScope software can monitor any bus or signal in the design. By manipulating the ChipScope probe locations in the design hierarchy, you can narrow in on the problem by comparing the data in hardware against simulation results at various checkpoints. When the digital logic has been eliminated as a possible cause, you can then proceed to debug the analog portion.

Here are some debugging steps to take when using the solution:

1. Double-check the power supply and oscillator choices against Xilinx recommendations.
2. Using ChipScope software, examine the received data and status signals directly from the MGT outputs before any user logic. If all is as expected, then the user logic is at fault.
3. Use parallel and serial loopback modes to check transceiver settings and verify correct MGT operation.

4. Use ChipScope software to check the associated status signals for each of the MGT functions in the following order:
   a) Clock and Data Recovery
   b) Comma alignment
   c) 8b/10b
   d) Clock correction
   e) Channel bonding
   f) Cyclic redundancy check (CRC).
5. Run BER tests on the PCB traces to see if the physical link itself can operate reliably at the target line rate. Try progressively more challenging patterns if no errors are detected with easier test patterns.
6. Using XBERT with ChipScope software as a data analyzer, examine the distribution of bit errors and check if these errors are related to any noise sources.
7. Measure TDR and analyze trace and via construction.
8. If possible, gather more information using other lab equipment.

Debug Faster

The ChipScope tool speeds up debugging. When using ChipScope software, changing the trigger signal or data signal source does not require changes to the HDL code or re-synthesis, so you can change probe points to any signal within the same clock domain very quickly. To effect these changes, you need only rerun post-synthesis implementation, resulting in significantly shorter implementation iterations.

The ChipScope cores can be quickly and easily removed and inserted via the core inserter GUI. You can also place signal probes much faster than with a conventional logic analyzer, especially with wide signal buses.

The XBERT with ChipScope solution operates independently of user logic, software, and system-level control. Before measuring BER, the FPGA is simply configured using an image containing XBERT and ChipScope software. You can modify that same image to fit different devices and easily reuse the same design and techniques.

Crowded Boards and Remote Control

With increasing FPGA device densities, high pin counts make attaching test equipment probes a real challenge. Given the bus widths common today, numerous external test points are necessary; this greatly reduces the number of remaining I/Os.

In applications where board space is a concern, connectors for these test points consume precious real estate. The problem is further complicated by having to route these bus traces in tight places.

ChipScope software addresses this by requiring only a four-pin JTAG connection to the host PC. Because this connection is often provided for Boundary Scan testing during production, in most cases no additional pins are needed for the ChipScope tool.

Another advantage of the solution is that ChipScope virtual I/Os are used to toggle ports on the MGT and other control signals, when board space restrictions do not allow push buttons or DIP switches. In addition, they can also replace manual controls in an environmental testing context, giving full control over any net in the design.

If the selected device is too fully utilized for ChipScope software, try using the next larger footprint-compatible device during development. You can keep costs low by switching back to the smaller device for production.

The additional logic resources available through the use of footprint compatibility are freed up when ChipScope software and XBERT are not in use. Should the need arise, these resources can accommodate new features and design revisions that outgrow the original device. This eliminates the need for a board redesign, as the footprint can fit a range of FPGA densities.

Even without the option of a footprint-compatible device, you can employ a divide-and-conquer strategy to debug parts of user logic at a time, leaving sufficient resources to implement the two solution components.

Conclusion

The Xilinx XBERT with ChipScope solution enables faster diagnostic testing, debugging, and development of an MGT system without the use of expensive lab equipment such as logic analyzers and BERT testers. These significant cost savings reduce total serial system development costs, allowing even more budgets to benefit from multi-gigabit serial technology.

Xilinx will be offering a signal integrity course in the coming months. In the meantime, to find out more about the complete serial connectivity solution from Xilinx, please contact your local FAE for more information, or visit the following web resources:

- XAPP661 – *http://direct.xilinx.com/bvdocs/appnotes/xapp660.pdf*
- ChipScope Pro – *www.xilinx.com/ise/verification/chipscope_pro.htm*
- "Designing with Multi-Gigabit Serial I/O" Course – *www.xilinx.com/support/training/abstracts/rocketio.htm*
- Serial Tsunami Solutions – *www.xilinx.com/xlnx/xil_prodcat_product.jsp?title=hsd_high_speed.*

A Success Story

"My application uses four channel-bonded MGTs to communicate between processor boards in a universal mobile telecommunications system. The 128-bit wide channel-bonded data and numerous status signals made it very difficult to debug using a traditional logic analyzer.

ChipScope Pro™ enabled me to easily and accurately examine even the widest data paths and internal signals. XBERT also proved useful in verifying my PCB and backplane design. This solution enabled me to locate and fix a particularly elusive bug and is a great debugging tool.

With the assistance of a Xilinx Engineer on-site via the Xilinx Titanium Technical Service program, we very quickly started debugging using the advanced capabilities of ChipScope Pro. The Xilinx AE also introducted us to the use of XBERT as described in this article. The use of Xilinx Titanium Technical Service saved us many weeks of debug time!"

Hyung-Rak Kim
Hardware Engineer, UMTS Wireless Systems
LG Electronics

Tolerance Calculations in Power Distribution Networks

The impedance gradient of power planes around bypass capacitors depends on the impedance of planes and the loss of bypass capacitor.

by Istvan Novak, Ph.D.
Senior Signal Integrity Staff Engineer
Sun Microsystems
istvan.novak@sun.com

More designers are determining the requirements and completing the design of power distribution networks (PDN) for FPGAs and CPUs in the frequency domain. Although the ultimate goal is to keep the time-domain voltage fluctuation (noise) on the PDN under a pre-determined maximum level, the transient noise current that creates the noise fluctuations may have many independent and highly uncertain components, which in a complex system are hard to predict or measure.

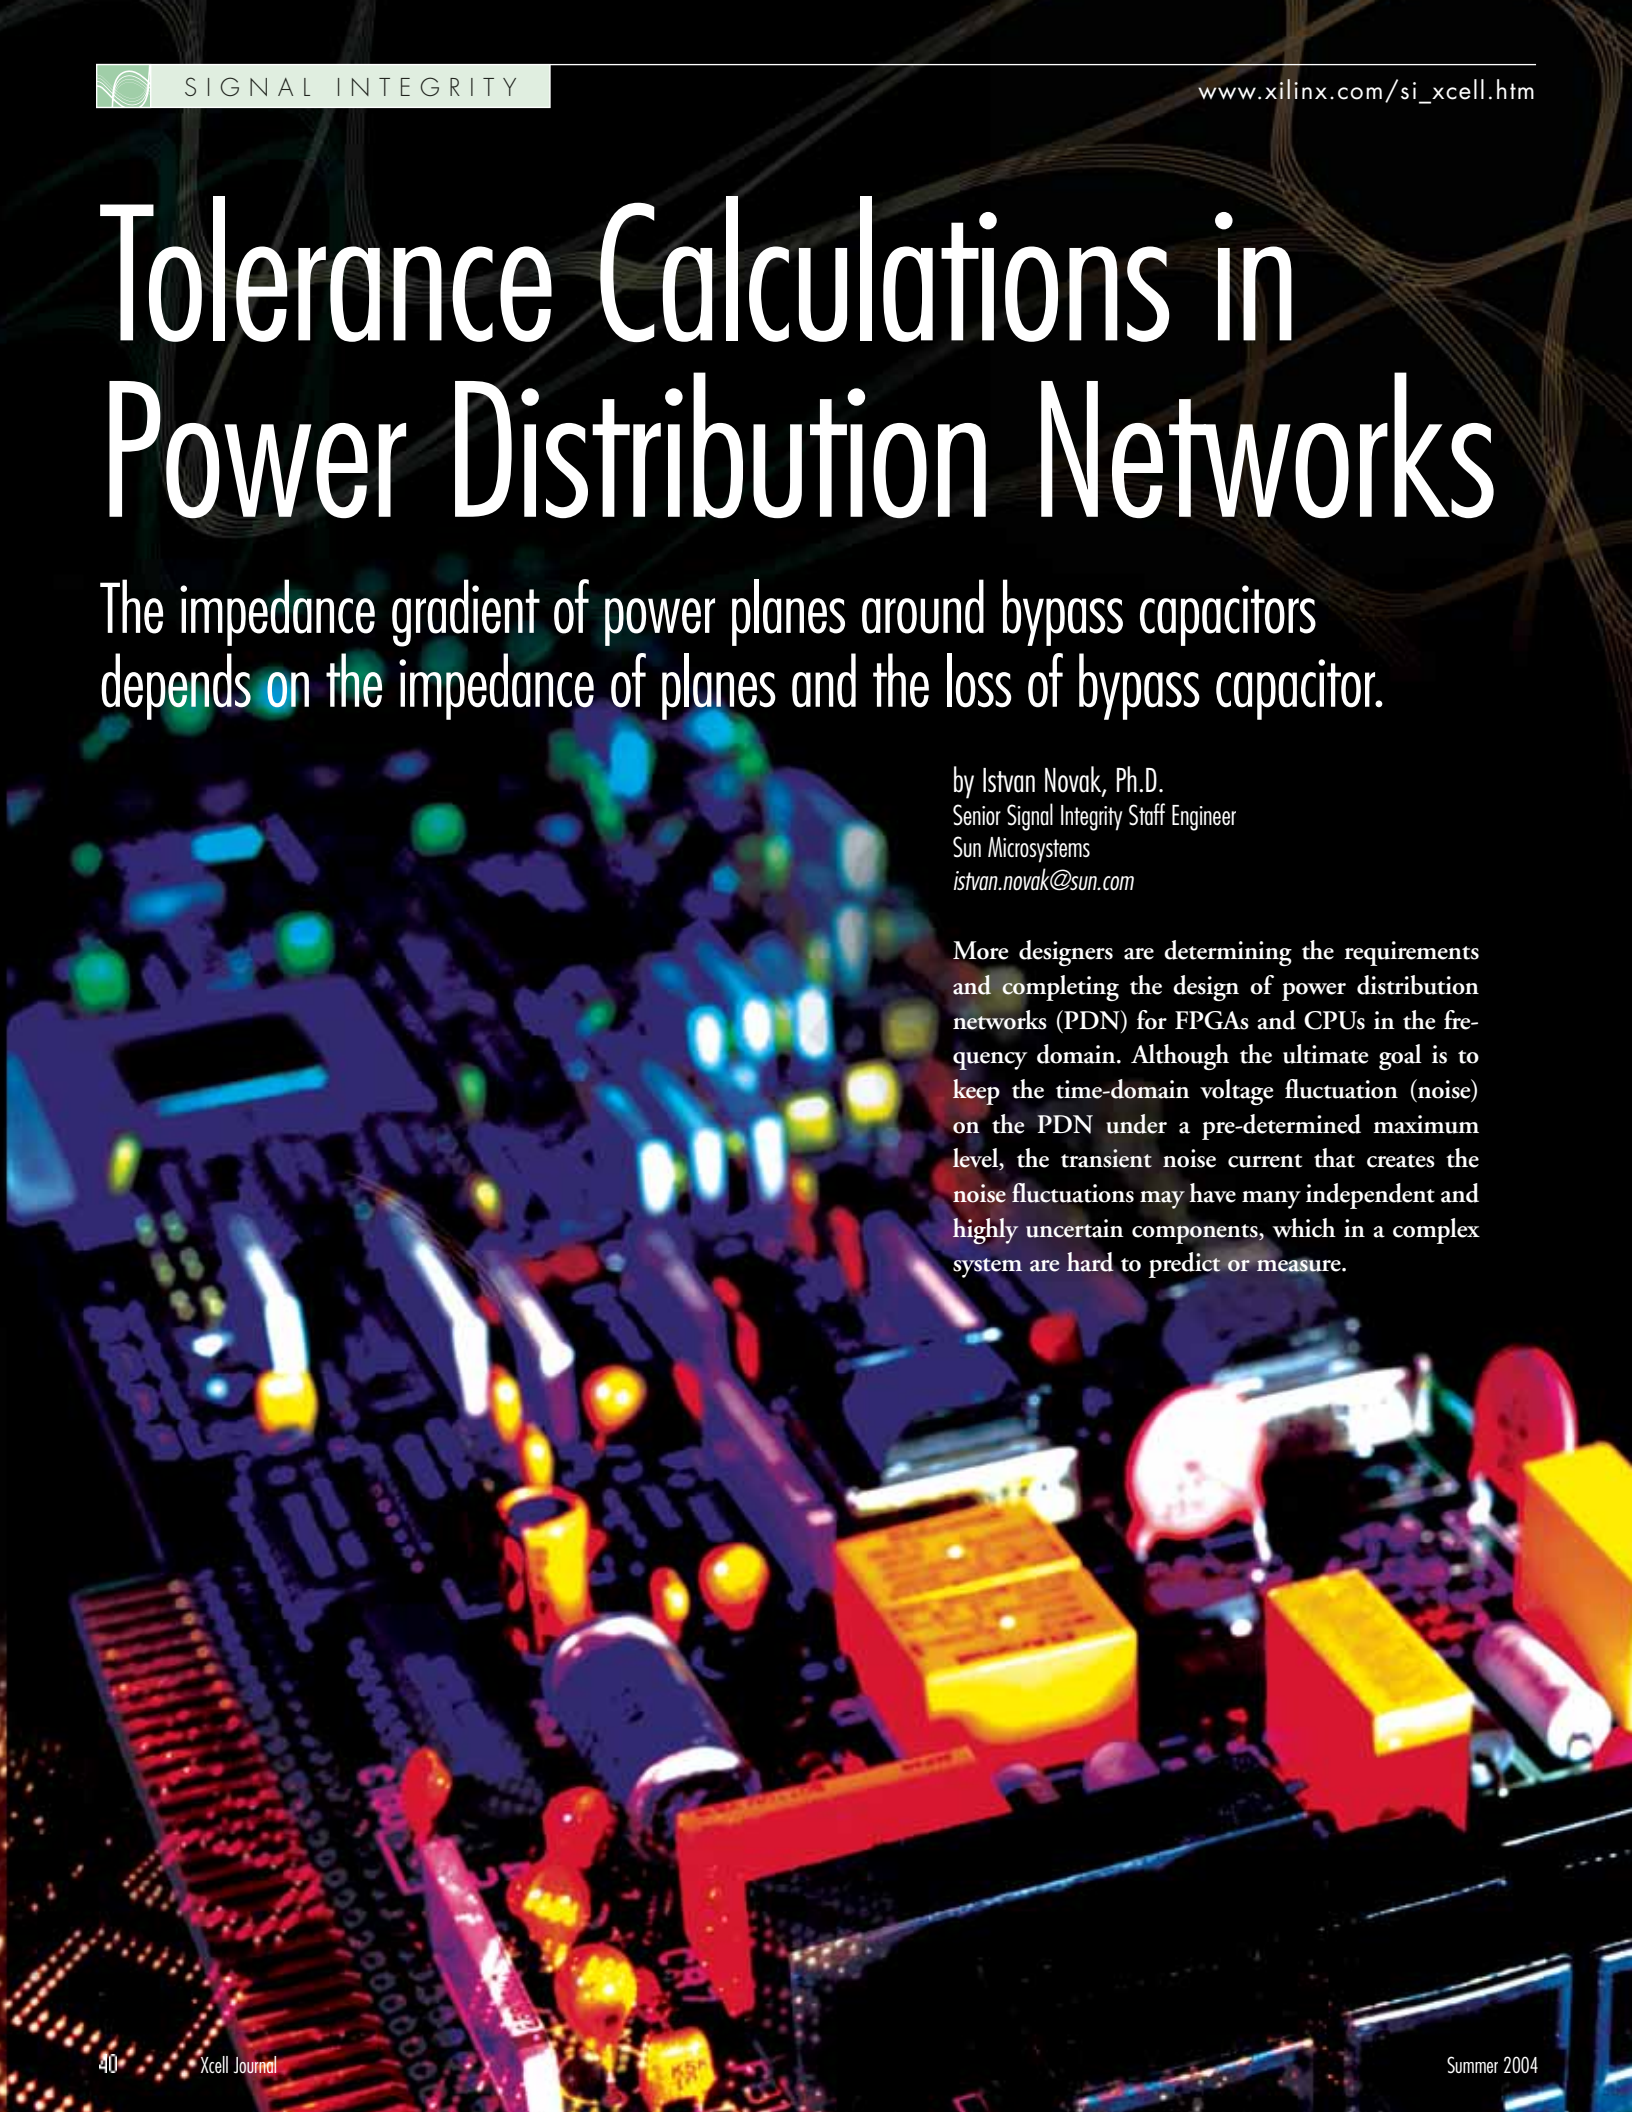

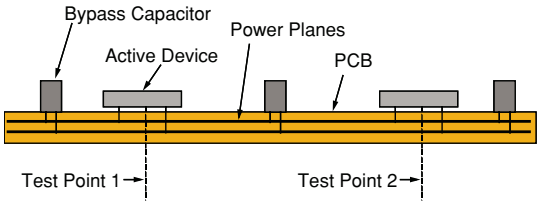

Figure 1 – Simple sketch of a PDN with two active devices, three capacitors, and one pair of power planes

Figure 2 – Three-element equivalent circuit of bypass capacitors

Figure 1 is a simple sketch of a PDN [1] with two test points. In the frequency domain, you can describe this network with a two-by-two impedance matrix, where the indices refer to the test points. Z_{11} and Z_{22} are the self impedances at test points 1 and 2, respectively, and Z_{12} and Z_{21} are the transfer impedances between test points 1 and 2.

With very few exceptions, the PDN components are electrically reciprocal; therefore the two transfer impedances are identical, and can be replaced with a mutual impedance term:

$$Z_{12} = Z_{21} = Z_M.$$

You cannot assume electrical symmetry, however, so Z_{11} and Z_{22} are, in general, different. You can calculate the noise voltages at test points 1 and 2 generated by the noise currents of $I_1(t)$ and $I_2(t)$ of the two active devices with the following formula:

$$V_1(t) = Z_{11}I_1(t) + Z_MI_2(t)$$
$$V_2(t) = Z_MI_1(t) + Z_{22}I_2(t)$$

A PDN comprises power sources (DC/DC, AC/DC converters, batteries); low- and medium-frequency bypass capacitors; PCB planes or other metal structures (a collection of traces or patches); packages with their PDN components; and the PDN elements of the silicon [2]. When dealing with board-level PDN, its impedance contributions to the overall PDN performance are much more stable and predictable, so much so that we often forget to analyze our PDN designs against component tolerances. In this article, we'll show how tolerances of bypass-capacitor parameters, such as capacitance (C), effective series resistance (ESR), effective series inductance (ESL), and capacitor location impact the impedance of PDNs.

C, ESR, and ESL Tolerance Effects

Figure 2 shows the simple equivalent circuit of a bypass capacitor when neglecting the parallel leakage of the capacitor. The series capacitor-resistor-inductor circuit shows a resonance frequency with a given quality factor (Q). You can calculate the series resonance frequency (SRF) and Q from the equations below:

$$SRF = \frac{1}{2\pi\sqrt{C * ESL}};\quad Q = \frac{\sqrt{\frac{ESL}{C}}}{ESR}$$

Although in a general case all three elements in the equivalent circuit are frequency-dependent [3], for the sake of simplicity, and because it would not change the conclusions of this article, we'll use frequency-independent constant parameters.

Figure 3 shows the impedance magnitudes of three different capacitors you could use in a PDN. Each curve has a label, giving the C, ESR, and ESL values assumed for the part. The SRF and Q values are also shown for each part. With these numbers, the 100 uF part could be a tantalum brick; the 1 uF and 0.1 uF parts could be multi-layer ceramic capacitors (MLCC).

When connecting capacitors with different SRFs in parallel, they may create anti-resonance peaks where the impedance magnitude exceeds the lower boundary of the composing capacitors' impedance magnitude values [4] [5]. The impedance penalty gets bigger as the Q of capacitors gets bigger, or as their SRFs are farther apart in frequency.

The anti-resonance peaks get even bigger when you consider the possible tolerances associated with the capacitor parameters. We illustrate this in Figure 4, which shows what happens in typical, best, and worst cases when you connect the three capacitors from Figure 3 in parallel. The plot assumes no connection impedance or delay between the capacitors. You can use this assumption as long as the distance between the capacitors is much less than the wavelength of higher frequency of interest, and the connecting series plane impedance is much less than the impedance of capacitors.

The frequency plot extends up to 100

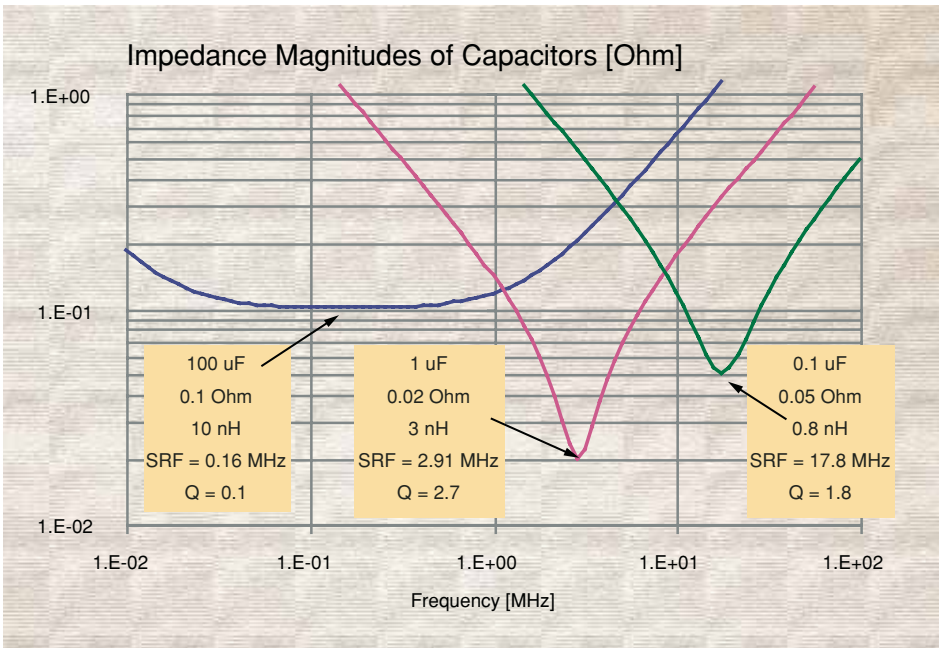

Figure 3 – Impedance magnitudes of three stand-alone bypass capacitors

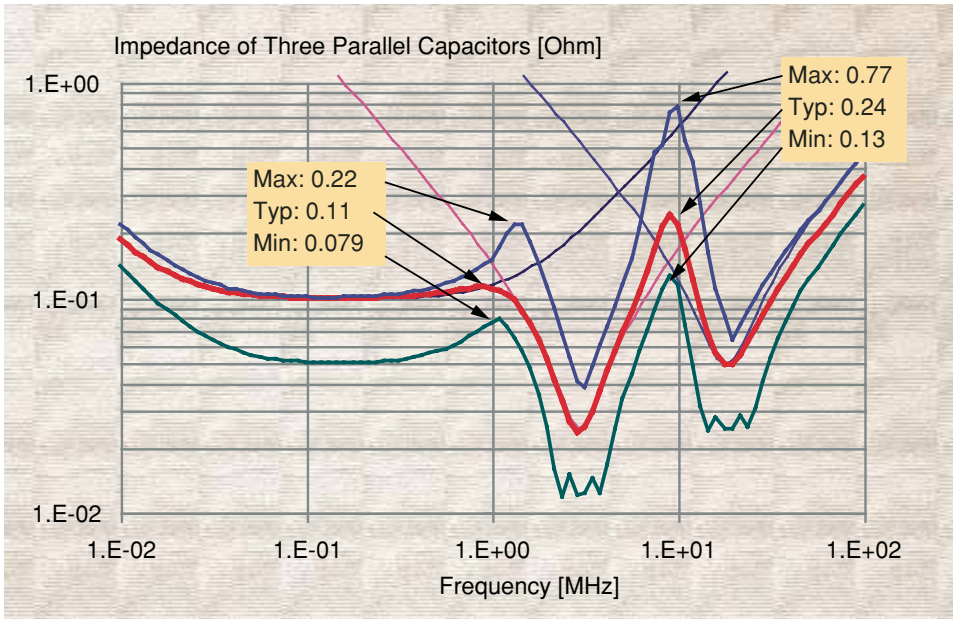

Figure 4 – Typical, highest, and lowest impedance curves of the three parallel connected capacitors shown in Figure 3

MHz, which represents a wavelength of 15 meters in FR4 PCB dielectrics. This tells us that the lumped approximation is valid in this entire frequency range, no matter where we place these capacitors on a typical-size PCB.

Table 1 lists the percentage tolerance ranges for the C, ESR, and ESL values used in Figure 3. We calculated the impedance curves and tolerance analysis with a simple spreadsheet [6]. The spreadsheet calculates the complex impedance resulting from the three parallel connected impedances. During tolerance analysis, the spreadsheet steps each parameter systematically though their minimum and maximum values – specified by the tolerance percentage entered – and accumulates the lowest and highest magnitudes at each frequency point.

For Figure 4, we assume a capacitance tolerance of +-20% for all three capacitors. For ESR, datasheets usually state the maximum value but no minimum, so we can assume a +0 to -50% tolerance around the nominal value. ESL strongly depends on both the capacitor's construction and its mounting geometry. For this example, we assume +-25% inductance variation.

Figure 4 also shows the impedance magnitudes of the individual capacitors with thin lines. The three heavy lines in the figure represent the maximum, typical, and minimum values from all possible tolerance permutations. All three curves exhibit two peaks: the first around 1 MHz and a second around 10 MHz.

The trace representing the typical case has an impedance magnitude of 0.11 Ohms and 0.24 Ohms at these peak frequencies, respectively. Impedance at and around the first peak is mostly below the impedance curves of the 100 uF and 1 uF capacitors. The second peak, however, exceeds the lower boundary of the impedance curves of the 1 uF and 0.1 uF capacitors by about a factor of two. This is a typical anti-resonance scenario.

In a worst-case combination of component tolerances, the second anti-resonance peak increases from 0.24 Ohms to 0.77 Ohms, a 220% increase. The contributors to the second anti-resonance peak are the ESR and ESL of the 1 uF capacitor, and the C and ESR of the 0.1 uF capacitor. The sum of the tolerances of these four parameters is 145%, but they increase the impedance at the peak by 220%. This illustrates that the resonance magnifies the tolerance window.

	C1	tol. [%]	C2	tol. [%]	C3	tol. [%]
Capacitance [uF]:	100	20 -20	1	20 -20	0.1	20 -20
ESR [Ohms]:	0.1	0 -50	0.02	0 -50	0.05	0 -50
ESL [nH]:	10	25 -25	3	25 -25	0.8	25 -25

Table 1 – Parameters used for Figure 4

Bypass Capacitor Range

Bypass capacitors are considered to be charge reservoir components, and common wisdom tells you to put them close to the active device they need to feed. We will show here that when the capacitor and the active device are connected with planes, the ratio of plane impedance and ESR of capacitor will determine the spatial gradient of impedance around the capacitor. Even at low frequencies, the impedance gradient can be significant.

Let's look at the self-impedance distribution over a 2" x 2" plane pair with 50 mil plane separation. You will get this plane separation if you have just a few layers in the board and if they are not placed next to each other in the stack-up. The characteristic impedance of these planes is approximately 1.7 Ohms. You can calculate the approximate plane impedance from our third equation [7]:

$$Z_p = \frac{532}{\sqrt{\varepsilon_r}} \frac{h}{P}$$

where Z_p is the approximate plane impedance in Ohms and h and P are the plane separation and plane periphery, respectively, in the same but arbitrary units.

We assume one piece of capacitor located in the middle of the planes. MLCC capacitors are available with as much as a few hundred uF capacitance in the 1210 case style, and their ESR can be as low as one milliohm. For this example, we use C = 100 uF, ESR = 0.001 Ohm, ESL = 1 nH. The SRF of this part is 0.5 MHz.

The surface plot of Figure 5 shows the variation of self-impedance magnitude over the plane at 0.5 MHz. The gray bottom area of the graph represents the top view of the planes. The grid on the bottom area shows the locations where the impedance was calculated: the granularity was 0.2 inches. The logarithmic vertical scale shows the impedance magnitude between 1 and 10 milliohms.

We calculated the surface impedance with a spreadsheet [8]. The macro in the spreadsheet calculates the impedance matrix by evaluating the double series of cavity resonances. It then combines the complex impedance of plane pair with the complex impedance of the bypass capacitor.

The impedance surface at 0.5 MHz has a sharp minimum in the middle; here the capacitor forces its ESR value over the plane impedance. However, as we move away from the capacitor, the impedance rises very sharply. At 0.2 inches away, the impedance is approximately 50% higher; 0.4 inches away, the impedance magnitude doubles. At the corners of the 2" x 2" plane pair, the impedance magnitude is almost 10 milliohms.

When changing either the plane impedance or the ESR of capacitor so that their values are closer, the variation of impedance over the plane shape gets smaller. Figure 6 shows the impedance surface of the same plane shape and same capacitor in the middle, except we increased ESR from 1 to 7 milliohms and decreased the plane separation from 50 to 20 mils. Now the impedance surface at SRF varies only about 10% over the plane area.

For Figures 5 and 6, you can see the same characteristic behavior if you sweep the frequency over a wider frequency range in the spreadsheet. The impedance surface of Figure 5 changes and fluctuates significantly, while the impedance surface of Figure 6 changes less with frequency.

Note that this trend does not change if we have more capacitors on the board. If we have significantly different plane impedance and cumulative ESR of capacitors, the impedance gradient will be big, and we must use many capacitors to hold the impedance uniformly down over a bigger area even at low frequencies.

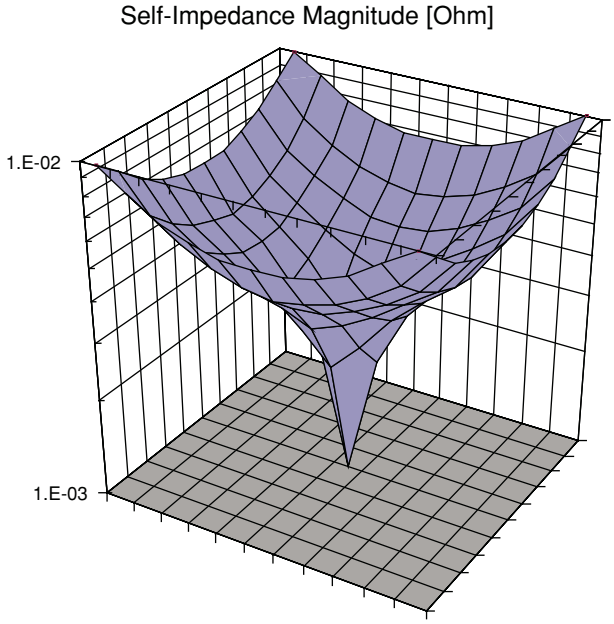

Figure 5 – Self-impedance at 0.5 MHz on a 2" x 2" plane pair with 50 mils dielectric separation, with a 100 uF, 0.001 Ohm, 1 nH capacitor located in the middle

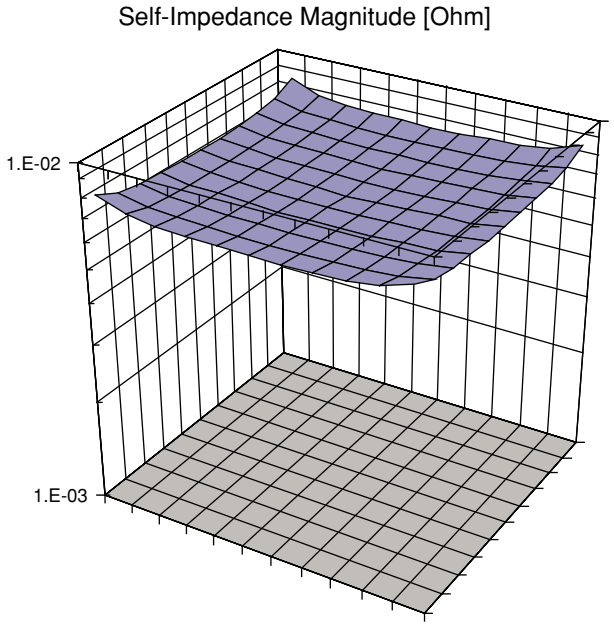

Figure 6 – Self-impedance at 0.5 MHz on a 2" x 2" plane pair with 20 mils dielectric separation, and a 100 uF, 0.007 Ohm, 1 nH capacitor located in the middle

Conclusion

The impedance tolerance window at the anti-resonance peak of paralleled discrete bypass capacitors widens with higher Q capacitors. To keep the impedance window due to tolerances small, you need either many different SRF values tightly spaced on the frequency axis, or the Qs of capacitors must be low.

Contrary to popular belief, the service range of low-ESR capacitors is severely limited when connected to planes of much higher impedance. But you can achieve the lowest spatial impedance gradient if the cumulative ESR of bypass capacitors is close to the characteristic impedance of planes. Σ

References

[1] Novak, I. "Frequency-Domain Power-Distribution Measurements – An Overview, Part I" in HP-TF2, *Measurement of Power Distribution Networks and their Elements.* DesignCon East, June 23, 2003, Boston.

[2] Smith, L.D., R.E. Anderson, D.W. Forehand, T.J. Pelc, and T. Roy. 1999. Power Distribution System Methodology and Capacitor Selection for Modern CMOS Technology. *IEEE Transactions on Advanced Packaging* 22(3): 284-290.

[3] Novak, I., and J. R. Miller. "Frequency-Dependent Characterization of Bulk and Ceramic Bypass Capacitors" in *Proceedings of EPEP*, October 2003, Princeton, NJ.

[4] Brooks, Douglas. 2003. *Signal Integrity Issues and Printed Circuit Board Design.* Upper Saddle River: Prentice Hall.

[5] Ritchey, Lee W. 2003. *Right the First Time, A Practical Handbook on High Speed PCB and System Design, Volume 1.* Glen Ellen: Speeding Edge.

[6] Download Microsoft™ Excel spreadsheet at *http://home.att.net/~istvan.novak/tools/bypass49.xls*

[7] Novak, I., L. Noujeim, V. St. Cyr, N. Biunno, A. Patel, G. Korony, and A. Ritter. 2002. Distributed Matched Bypasssing for Board-Level Power Distribution Networks. *IEEE Transactions on Advanced Packaging* 25(2):230-243.

[8] Download Microsoft Excel spreadsheet at *http://home.att.net/~istvan.novak/tools/Caprange_rev10.xls*

High-Speed PCB Design Resources

If you've read up to this point in the series, you're probably thirsty for additional information about the tools and methods discussed. Thus, we've compiled different sources of information available from the Xilinx website and design and education services. We've also included references to partner software development tools, hardware development platforms, and literature from renowned personalities.

by Suresh Sivasubramaniam
Senior Design Engineer
Xilinx, Inc.
suresh.subramaniam@xilinx.com

Philippe Garrault
Technical Marketing Engineer
Xilinx, Inc.
philippe.garrault@xilinx.com

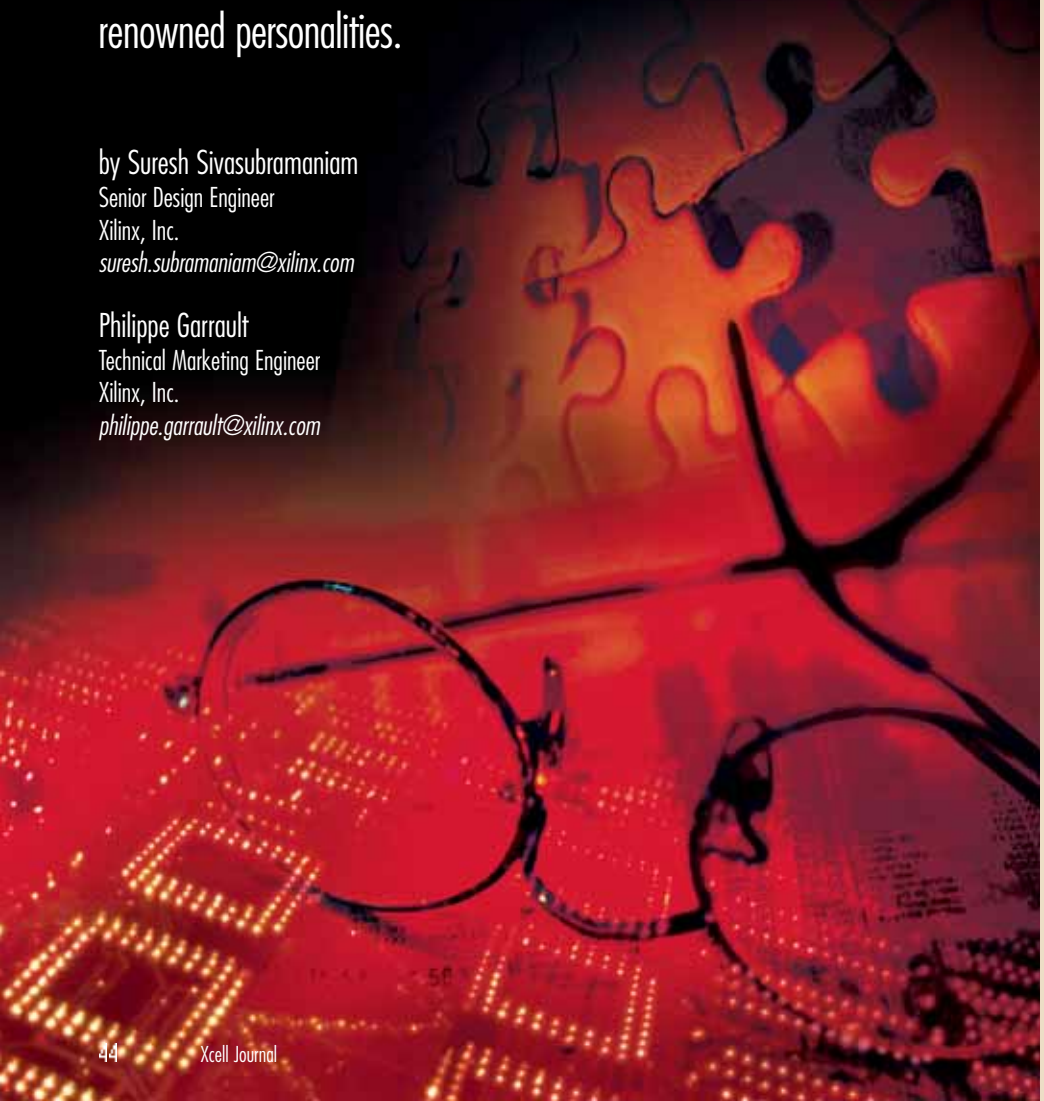

Xilinx Website Resources

Signal Integrity Central

www.xilinx.com/xlnx/xil_prodcat_landingpage.jsp?title=Signal+Integrity

This Xilinx portal has everything you need to achieve reliable PCB designs on the first pass. It covers signal integrity fundamentals, power supply and bypassing, simulation tools, PCB design and thermal considerations, and multi-gigabit signaling, with a variety of documents, FAQs, and links.

White Papers and Application Notes

http://direct.xilinx.com/bvdocs/whitepapers/wp174.pdf

White Paper WP174, "Methodologies for Efficient FPGA Integration into PCBs," describes how PCB design considerations play a major role in obtaining the expected performance from FPGAs. It then focuses on early analysis and simulation methodologies as a way of performing a guided implementation.

www.xilinx.com/bvdocs/appnotes/xapp623.pdf

Application Note XAPP623, "Power Distribution System (PDS) Design: Using Bypass/Decoupling Capacitors," details Virtex™ power supply requirements and techniques for designing power distribution systems using bypass/decoupling capacitors.

www.xilinx.com/bvdocs/appnotes/xapp689.pdf

Application Note XAPP689, "Managing Ground Bounce in Large FPGAs," explains the ground bounce effect, with calculations to help you derive an FPGA pinout that meets input undershoot and logic-low voltage requirements for devices receiving signals from an FPGA.

www.xilinx.com/bvdocs/appnotes/xapp609.pdf

Application Note XAPP609, "Local Clocking Resources in Virtex-II™ Devices," describes the different local clocking resources available in the Virtex-II architecture. Along with a reference design, this application note explores local clocking resources in source-synchronous applications.

Xilinx and Partner Software Resources

www.support.xilinx.com/support/software_manuals.htm

The ISE software tool suite includes helpful tools like the pin assignment constraints editor (PACE) to help you during the I/O planning and pinout assignment phases. Another helpful tool is Xpower, which allows you to plan and estimate your FPGA power requirements.

AllianceEDA

www.xilinx.com/xlnx/xil_prodcat_landingpage.jsp?title=Alliance+EDA+Program

Visit this website to learn more about the tools mentioned in the *Xcell* SI series. This page provides information on all third-party tools interfacing with Xilinx software.

Education Courses

www.xilinx.com/support/education-home.htm

Xilinx offers courses about PCB considerations when designing with high-performance FPGAs. The "High-Speed Signal Integrity Design" course, for example, teaches how signal integrity techniques are applicable to high-speed interfaces between Xilinx FPGAs and semiconductor memories.

The course covers high-speed bus and clock design, including transmission line termination, loading, and jitter. For additional details, visit the education services website.

Design Services

www.xilinx.com/xds/

Xilinx provides a comprehensive service offering that includes education services, support services, and design services to accelerate time to knowledge and time to market. To effectively design new multi-gigabit serial systems, it is imperative to understand the complete channel. Xilinx services leverage new techniques, relying on in-house expertise and state-of-the-art tools to create accurate models, evaluation platforms, and production backplanes. Xilinx helps companies to design, simulate, and characterize every aspect of the channel from 1 Gbps to more than 10 Gbps.

RocketIO Multi-Gigabit Serial Transceivers

RocketIO Resources

www.xilinx.com/serialsolution/

The Serial Tsunami Solutions portal provides you with a wealth of information on RocketIO™ transceivers, with access to transceiver data sheets, characterization data, protocols, articles, white papers, and a gateway to the Serial Backplane Simulator tool.

The Serial Backplane Simulator

www.xilinx.com/products/xaw/hsd/simulator.htm

The Serial Backplane Simulator provides signal integrity simulations for more than 300 situations when using Virtex-II Pro™ devices to drive signals across backplanes and line cards. The analyzer covers several different backplane materials, speed, pre-emphasis, lengths, peak-to-peak differential swings, number of connectors, and connector types.

RocketLabs

www.xilinx.com/rocketlabs/

With 15 locations around the world, RocketLabs is the first network of labs to provide system designers with free access to high-speed equipment, multiple hardware evaluation boards, application expertise, signal integrity simulation tools, and presentations and specialized training.

RocketIO Design Kits

www.xilinx.com/support/software/spice/spice-request.htm

From the SPICE suite, you can download RocketIO SPICE models and design kits. The kits have comprehensive documentation and examples that will jumpstart the simulations process and get you to the simulation results analysis phase faster and easier.

Reference Books

"Computer Circuits Electrical Design" by Ron K. Poon

"Digital Systems Engineering" by William J. Daly and John W. Poulton

"High-Frequency Characterization of Electronic Packaging" by Luc Martens

"Signal Integrity – Simplified" by Eric Bogatin

"High-Speed Digital Design: A Handbook of Black Magic" by Howard Johnson

"High Speed Signal Propagation: Advanced Black Magic" by Howard W. Johnson

"Digital Signal Integrity: Modeling and Simulation with Interconnects and Packages" by Brian Young

"MECL System Design Handbook," Motorola Semiconductor Products, Inc.

"High-Speed Digital System Design: A Handbook of Interconnect Theory and Design Practices" by Stephen H. Hall, Garrett W. Hall, and James A. McCall

Other Resources

www.signalintegrity.com

www.speedingedge.com

www.gigatest.com

www.nesa.com

www.qsl.net/wb6tpu/si_documents/docs.html

www.ultracad.com

www.teraspeed.com

www.tdasystems.com

Conclusion

We hope that you enjoyed reading this special SI series and that you feel better informed in dealing with signal and power integrity issues in your current and future high-speed PCB designs.

If you have any comments or feedback about the topics discussed, please e-mail us at *si_xcell@xilinx.com.*

Helping To Launch Your Designs...

3

Spartan-3. The World's Lowest-Cost FPGAs.

2

Unbeatable Deals on Development Kits and Software.

1

Great Opportunity for Industry-leading On-site Workshops.

Nu Horizons WORKSHOPS

Go!

www.nuhorizons.com/S3workshops
1-888-747-NUHO

XILINX®
The Programmable Logic Company℠

NU HORIZONS ELECTRONICS CORP.

The FPGA Dynamic Probe

Innovative technology significantly increases in-circuit debug productivity.

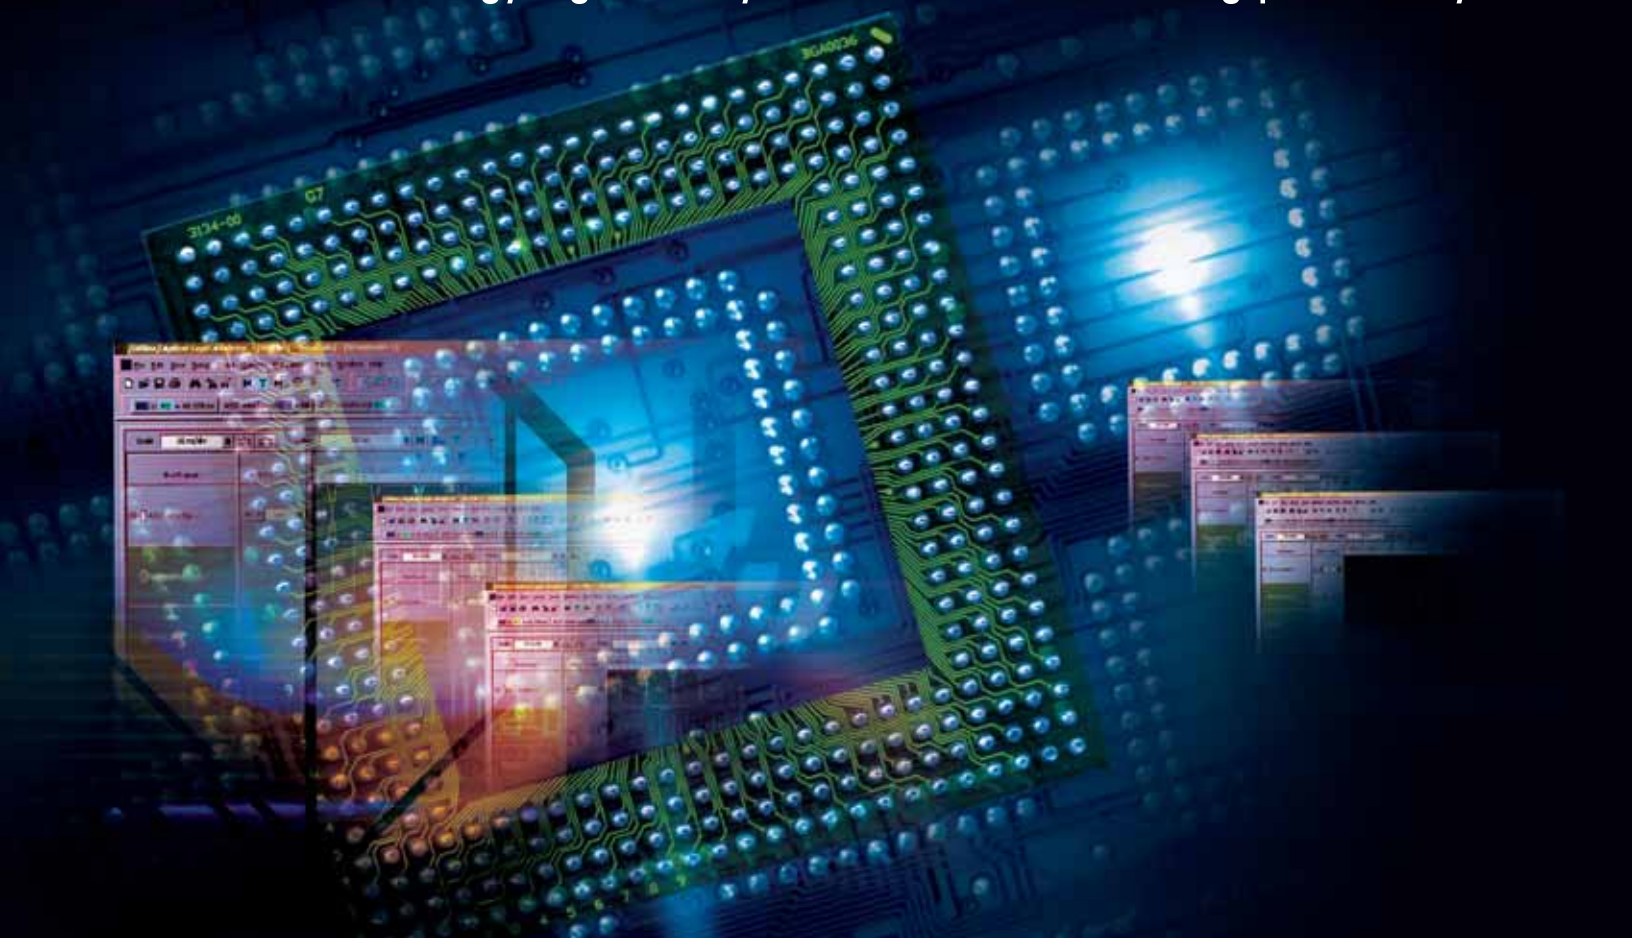

by Joel Woodward
Logic Analysis Project Manager
Agilent Technologies
joel_woodward@agilent.com

FPGAs play an increasingly important role in project development, where the need for high-performance designs with flexible architectures collides with lean engineering teams, constrained budgets, and rapid development schedules. Yet traditional in-circuit debug methodologies limit how quickly designers can uncover design problems.

Often, design defects in increasingly complex systems may occur exclusively in real time, when multiple subsystems and software interact. Using FPGAs, design teams can move quickly to system integration, increasing the importance of effective debug and validation. With sufficient visibility, in-circuit debug of FPGA designs can uncover in just a few minutes problems that might have required hours, days, or weeks to simulate.

Logic analyzer measurements are particularly effective in the debug of FPGAs and surrounding systems. A typical measurement approach is to take advantage of the programmability of the FPGA to route internal signals to a small number of pins. Although this is a very useful approach, it has limitations that inhibit productivity.

Because pins on the FPGA are typically an expensive resource, there are a relatively small number available for debug. One pin is required for each internal signal to be probed, thereby limiting the visibility of internal nodes to the same small number of signals. Design teams rarely find this width of visibility adequate.

When different internal signals are measured, new signals are routed out to pins; sometimes, a recompile of the design is required. In either case, the change consumes valuable engineering resources and

can change the timing of the FPGA. To make sense of the measuring, engineers must manually update logic analyzer label names and probe locations to match the new configuration of the measurement every time new signals are routed to pins.

New technology from Agilent Technologies and Xilinx mitigates the issues described above by combining ChipScope™ Pro technology with Agilent's FPGA dynamic probe logic analysis application. Figure 1 shows the key components of the application.

You can use the Xilinx Core Inserter or CORE Generator™ tool to insert an Agilent Trace Core 2 (ATC2) into an FPGA, thus facilitating a more productive debug session. The core is controlled by Agilent's FPGA dynamic probe logic analysis application software. The application runs on Agilent's 1680, 1690, or 16900 series logic analyzers.

Figure 1 – The FPGA dynamic probe application software can change virtual probe points inside Xilinx FPGAs in less than a second. The logic analysis application communicates to a debug core via JTAG.

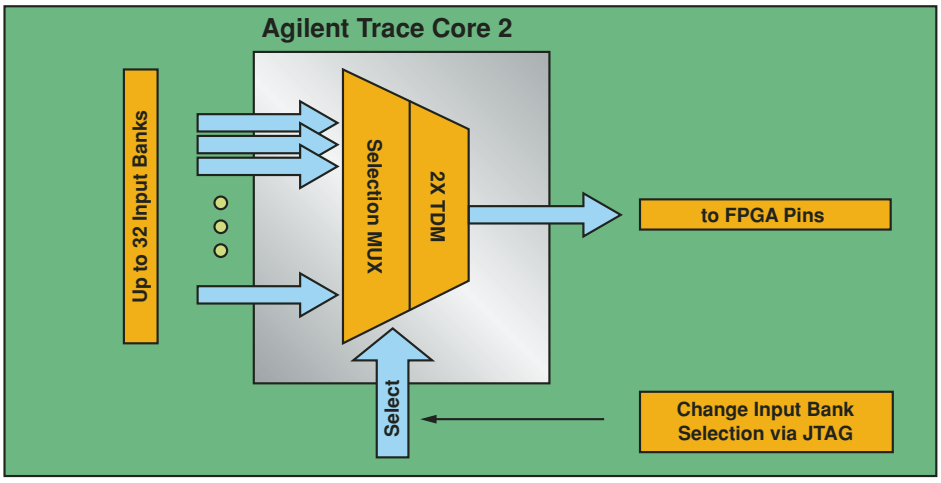

Figure 2 – Agilent's second-generation configurable trace core provides visibility to as many as 64 internal FPGA signals for each pin dedicated to debug.

Time-to-Market Advantages

The FPGA dynamic probe delivers four primary benefits:

1. The ATC2 core allows a dynamic approach to choose internal signals for logic analysis – without incurring the limitations (such as potential recompiles and the associated timing impact) of the traditional "route out signals to pins" approach.

   Using Core Inserter, you can specify groups of internal FPGA signals that might need measurement. Each group of signals represents an input to the ATC2 core. The core allows one group of input signals to be routed to pins. With a mouse-click in the logic analysis application software, the analyzer changes which group of internal FPGA signals are routed through the core. This capability eliminates the need to recompile to change signal probing, saving days of development time per FPGA design. In addition, this method keeps timing constant.

2. Although a 1:1 internal signal-to-pin ratio normally exists for debug, the FPGA dynamic probe increases this visibility ratio to 64:1. With 32 input groups into the ATC2 core, a single pin can sequentially gain access to 32 internal signals. With an optional 2X compression mode, each pin accesses two signals on each of the 32 input groups, for a total visibility of 64 signals per pin.

   This means that for each pin dedicated to debug, you can access as many as 64 internal signals (Figure 2). With this increased visibility width for certain types of validation requirements, you can bypass the time-consuming process of creating test benches and perform the validation more quickly in-circuit.

3. The FPGA dynamic probe automates the process of label naming when a new set of internal signals is selected. Logic analyzers with this application read a file from a .cdc file that the Core Inserter generates. This file con-

tains all node names of signals that may be eventually selected.

Because the tool tracks which signals are currently routed through the ATC2 core, the software application running on the logic analyzer automatically enters signal names and channel locations on the logic analysis setup menu each time a new set of internal signals is probed (Figure 3). This additionally saves time and eliminates errors.

4. The FPGA dynamic probe helps you make more accurate state measurements. The core invokes test stimulus that is acquired by the logic analyzer. The logic analyzer samples the test pattern and automatically determines when to best sample each signal relative to the clock. This calibration capability compensates for path length variances, ensuring accurate state measurements. This is particularly beneficial on high-speed circuits with narrow data valid windows.

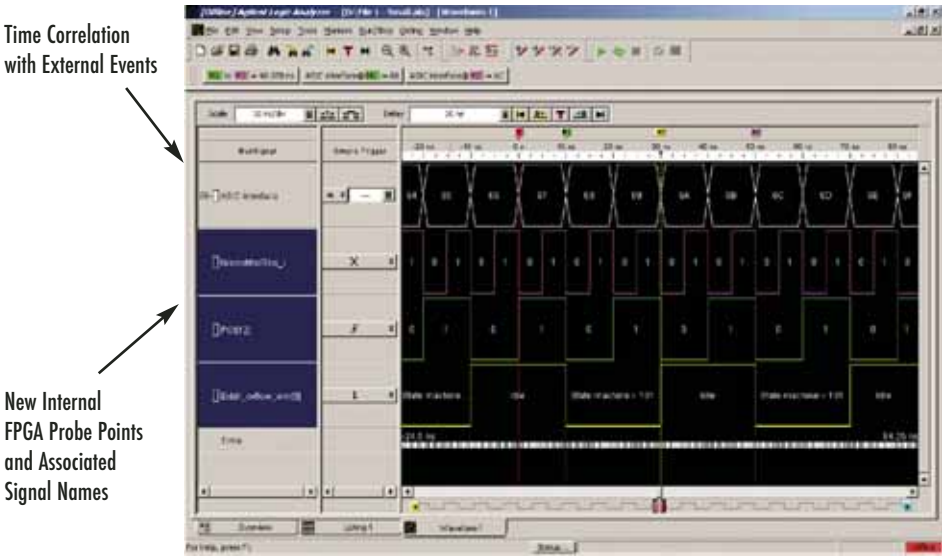

Figure 3 – The FPGA dynamic probe automatically extracts internal signal names and updates the logic analyzer each time new probe points are selected.

Configure the Core to Match Debug Needs

The ATC2 core is configurable to match your design requirements. Number of pins, number of input banks, and sampling mode (timing or state) are some of the configurable parameters. The ATC2 core has been crafted to take minimal space inside the FPGA.

As an example, an ATC2 core with eight bits of visibility on each of 32 input banks consumes only about 2% of the slices in a XC2V3000 device. This core offers access to 256 signals using just eight pins.

A smaller number of input banks and fewer pins allows the core to consume fewer FPGA resources. A higher number of input banks and more pins increases visibility. You can make tradeoffs depending on the specific device and visibility requirements.

The ATC2 core runs as fast as the device runs, so measurement speeds are limited only by the acquisition capabilities of the logic analyzer. With state speeds well in excess of 200 MHz and timing speeds of 4 GHz, most new logic analyzers contain enough headroom to make accurate FPGA measurements for the next several years.

When you have significant debug issues, you can create multiple ATC2 cores that coexist peacefully within a single device. The FPGA dynamic probe application software can also control ATC2 cores in multiple FPGAs, as long as the FPGAs are on the same scan chain.

The new technology allows you to more easily correlate internal FPGAs to external events, thus isolating problems more quickly. When the ATC2 core facilitates measurements internal to the FPGA, the logic analyzer can time-correlate these measurements with measurements elsewhere on the target system. This capability allows you to gain insight into your system designs more quickly.

The FPGA dynamic probe virtual probing technology, combined with a logic analyzer, blurs the boundary between internal FPGA measurements and external measurements.

Conclusion

The joint collaboration between Xilinx and Agilent in producing the royalty-free ATC2 core and the FPGA dynamic probe will enable more productive in-circuit debug. Agilent has already used this technology internally to shave weeks of development time from a critical project that used multiple Xilinx FPGAs. The lead hardware engineer found that the solution allowed him to uncover in a few minutes problems that would have traditionally required hours or days to reveal.

As FPGA sizes increase and bigger designs take advantage of increased densities, successful design teams will adapt by employing innovative debug methodologies. The FPGA dynamic probe and ATC2 core provide critical capabilities for effective debugging. With these new tools, you can plan for debug early in the development process. Employing a design-for-debug methodology will allow you to keep pace with ever-increasing design sophistication.

ChipScope Pro software makes it possible for Xilinx FPGA users to easily debug designs. ChipScope Pro cores are integrated into the FPGA to provide real-time debug and verification capabilities via a standard JTAG port. ChipScope Pro is available from Xilinx for $695. A 30-day downloadable evaluation version is available for free. For more information, visit *www.xilinx.com/chipscopepro/*.

You can purchase Agilent's FPGA dynamic probe logic analysis application for an introductory price of $995 through the end of 2004. For more information on the Agilent FPGA dynamic probe, the ATC2 core, and supported logic analyzers, visit *www.agilent.com/find/FPGA/*.

Xilinx 6.2i Design Tools

The latest releases of ISE and ChipScope Pro design tools slash design and verification times while delivering the fastest performance available in PLD-based designs.

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

Xilinx Integrated Software Environment (ISE) 6.2i, the newest version of industry-leading Xilinx logic design tools, is focused on delivering you the highest performance available in PLD design. With ISE 6.2i, Virtex-II Pro™ FPGAs are now on average 40% faster than the nearest delivering competitive FPGA offering. That's up to three speed grades faster, and on silicon and software delivering today.

Spartan-3™ designers will also benefit significantly from using ISE 6.2i. You can improve performance by as much as 50% when using ISE 6.2i over our last release through a series of Spartan-3 enhancements:

- The Spartan-3 -4 speed grade has been enhanced to deliver higher performance
- The new, faster Spartan-3 -5 speed grade
- The clock-to-output performance has improved by 35-40%
- Embedded multiplier performance is as much as 50% faster – greater than 225 MHz
- ISE 6.2i now supports automatic local clock placement for Spartan-3 designs, delivering quicker and more accurate off-chip memory interface designs.

ISE 6.2i also continues to deliver 15% better logic utilization over competing solutions; you can get more design into a Xilinx FPGA using ISE. These performance improvements, combined with industry-leading cost advantages, are fueling the rapid replacement of ASICs and ASSPs with Spartan-3 FPGAs in numerous high-volume applications.

But faster performance has implications to all Xilinx customers, whether or not you're currently attacking a high-speed project. High performance means that ISE will hit your design targets first, with fewer costly design iterations requiring you to tweak your code to meet timing.

This is the first time a quantitative timing optimality study has been reported on any FPGA placement and routing tools.

Nearly Optimal Place and Route Results

Many design tools claim leadership, but ISE place and route (PAR) algorithms were recently tested by researchers from the University of California, Los Angeles (UCLA). These independent benchmark tests presented at the International Conference on Computer-Aided Design (ICCAD) showed that ISE PAR tools produce near-optimal timing-driven results. In an ICCAD paper titled "Optimality and Stability in Timing-Driven Placement Algorithms," Microelectronics Center of North Carolina (MCNC) benchmarks demonstrated that ISE came between 8.3 and 4.1% of the optimal PAR solution.

"As part of our placement optimality study, we generated a set of placement benchmark examples with known optimal solutions. Our study showed the Xilinx place and route tools produced consistently near-optimal timing results on Virtex-II™ series devices," said Dr. Jason Cong, a professor at the UCLA Computer Science Department and the faculty member directing the research. "We believe the excellent placement timing results were achieved by employing advanced timing-driven placement algorithms with efficient exploitation of the segmented routing architecture used in the Virtex-II series FPGAs." This is the first time a quantitative timing optimality study has been reported on any FPGA placement and routing tools.

A Unique New Approach to Logic Debug

If you're looking for a way to slash your verification cycle, you'll want to see what's new in the ChipScope™ Pro 6.2i release. The industry standard for real-time debug, the ChipScope Pro tool (along with Agilent Technologies' new FPGA Dynamic Probe) combines to create a logic debug solution that can't be matched by ASICs or competing FPGA solutions. ChipScope Pro can slash your verification cycle by as much as 50%, saving you significant time and money.

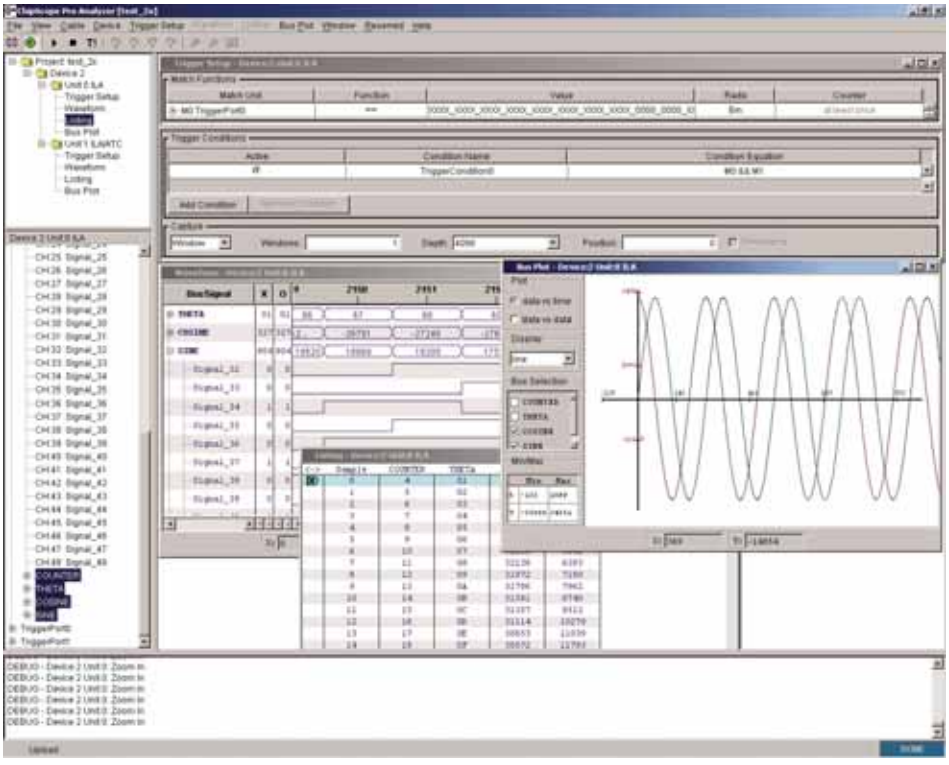

Figure 1 – ChipScope Pro logic analyzer

ChipScope Pro software lets you insert low-profile logic analyzer (ILA), bus analyzer (IBA), and Virtual I/O (VIO) software cores into your design or post-synthesis netlist. These cores allow you to view any internal signal or node within your FPGA, including the IBM™ CoreConnect processor local bus, on-chip peripheral bus for the IBM PowerPC™ 405 inside Virtex-II Pro Platform FGPAs, or the MicroBlaze™ soft processor core. Signals are captured at or near operating system speed, and brought out through the programming interface, freeing up pin assignments for your design. The ChipScope Pro logic analyzer can then analyze the captured signals (Figure 1).

ChipScope Pro and FPGA Dynamic Probe

ChipScope Pro software also links internal FPGA debug to your Agilent™ 16900, 1690, or 1680 series logic analyzer through the new ATC2 core. ATC2 synchronizes ChipScope Pro software to Agilent's new FPGA Dynamic Probe technology, delivering the first integrated application for FPGA debug with logic analyzers.

This unique partnership between Xilinx and Agilent gives you deeper trace memory, faster clock speeds, and more trigger options, all using fewer pins on the FPGA. For more details on ATC2 and the FPGA Dynamic Probe, see Joel Woodward's article, "The FPGA Dynamic Probe," also in this issue of *Xcell*.

Conclusion

ISE 6.2i and ChipScope Pro 6.2i tools can help you realize lower project costs immediately. Release to release, Xilinx is committed to delivering higher performance and shorter implementation and verification cycles, helping you slash design times and lower your costs. With a performance advantage of as many as three speed grades, the slowest Virtex-II Pro device is still faster than the fastest competing FPGA in production, helping you save in device costs with the added potential to get more design into your target device.

Download the free 60-day ISE 6.2i evaluation at *www.xilinx.com/ise_eval* or the free 60-day ChipScope Pro 6.2i evaluation at *www.xilinx.com/chipscope* today.

At your Service.
From start to finish.

www.xilinx.com/services

Xilinx offers the industry's broadest portfolio of education, support and design services to extend your technical capabilities, accelerate your time to market, and build a competitive advantage.

- Reduce the learning curve
- Speed your design time
- Jump-start your product development cycle
- Lower your overall development costs

THE ASSOCIATION OF SUPPORT PROFESSIONALS
Award Winner 2003
The Year's Ten Best Web Support Sites

Finish Faster

Xilinx delivers industry-leading service and support through every step of the design process—around the clock and around the world—to help you get to market faster. From technical support, consultative services, and dedicated design assistance to online support and training, you can find it all at **www.xilinx.com/services**.

XILINX®
The Programmable Logic Company℠

Better... Stronger... Faster

Virtex-II Pro FPGAs offer marked performance advantages over a competing device.

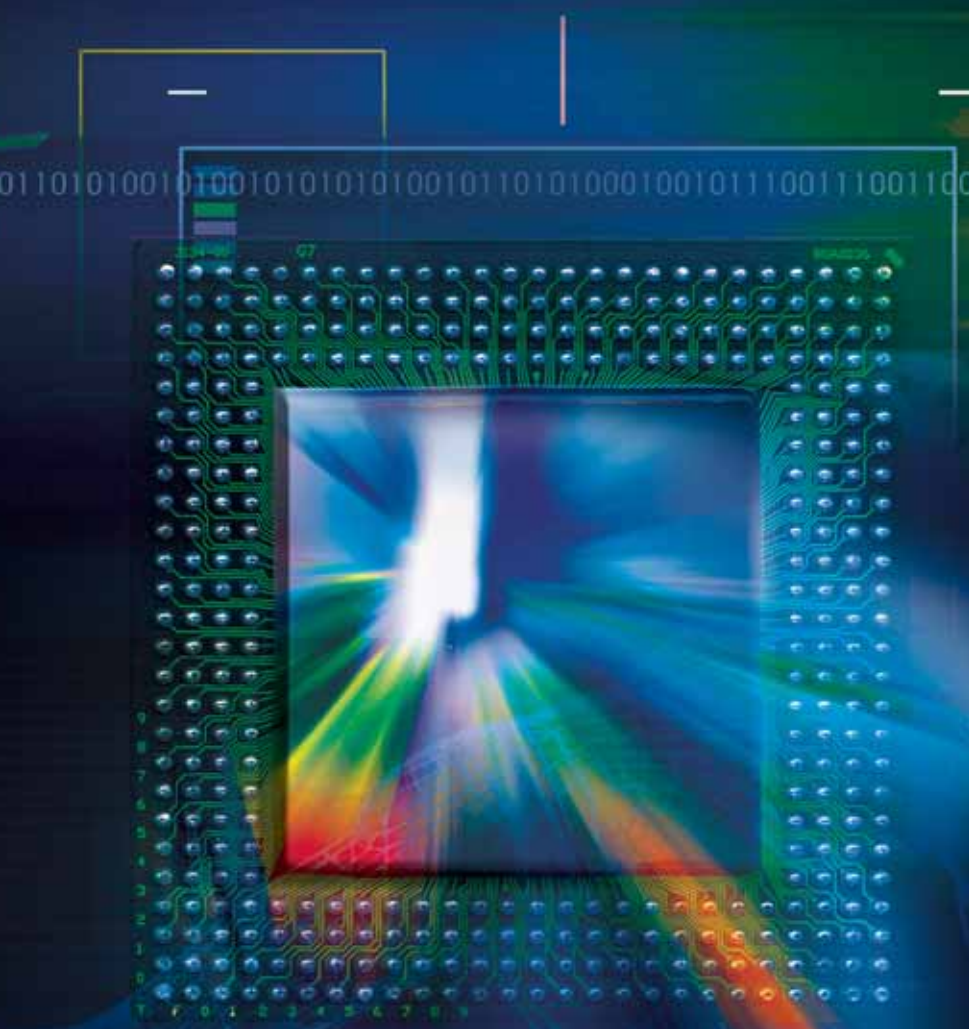

by Hitesh Patel
Sr. Manager, Software Product Marketing
Xilinx, Inc.
hitesh.patel@xilinx.com

As programmable logic devices increase in density and complexity, the combination of a feature-rich fabric and sophisticated design tools enables users to realize their performance goals faster. Shorter design cycle times also enable users to lower overall design costs and meet time-to-market requirements.

From analyzing 50 customer designs, we determined that Xilinx Virtex-II Pro™ FPGAs enjoy a 40% performance advantage over their nearest competitor, Altera™ Stratix™ FPGAs, to further realize the advantages of FPGAs. With densities ranging from 200,000 to 6 million system gates, the Virtex-II Pro device was as much as 123% faster than the Stratix device. Figure 1 shows the performance advantage distribution.

This article highlights how Virtex-II Pro FPGAs, along with ISE 6 design tools, provide a 40% performance advantage when compared to Stratix FPGAs.

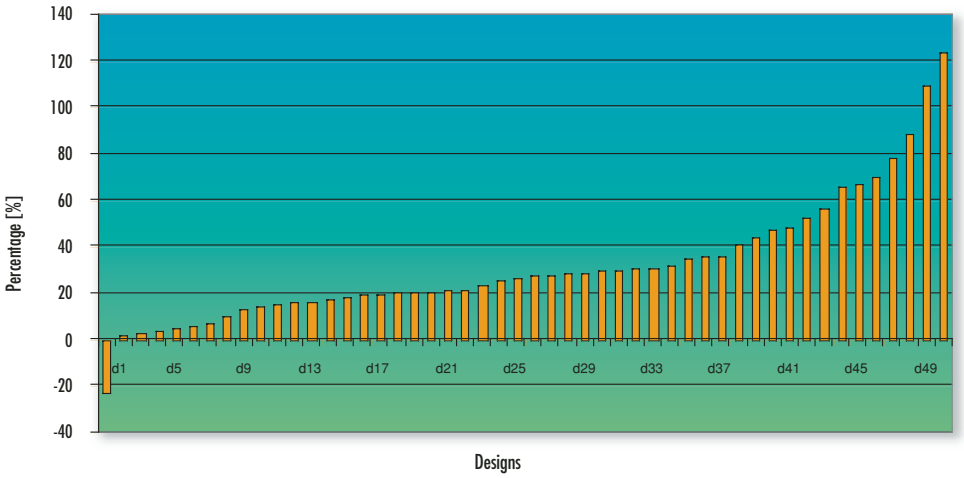

Figure 1 – Virtex-II Pro performance advantage versus Stratix FPGAs for 50 customer designs

Architectural Features

The basic building block in the Stratix architecture is called a logic element (LE). An LE contains three functional structures: a four-input look-up table (LUT), a register, and a carry chain.

Virtex-II Pro architecture not only includes the structures found in an LE, but also additional functionality, such as a function expander (MUXF), a MULT_AND arithmetic cell, and a more logic-rich carry structure.

Furthermore, the Virtex LUT can be used as a 16-bit shift register or as a single- or dual-port RAM element. These additional features in the Virtex-II Pro architecture enable users to realize higher design performance, as we'll describe in the next section.

MUXF Function Expander

One of the primary factors impacting circuit performance in FPGAs are logic levels in the signal path. The function expander cell represents a 2:1 MUX, which can be used to build functions wider than four inputs without the need for additional LUT logic levels.

For example, using the MUXF, only four LUTs are required to implement an 8:1 mux in a single LUT logic level. That same 8:1 mux in the Stratix PLD is implemented using five LUTs – and the implementation is two LUT logic levels. The additional LUT logic level adds delay to the signal path.

The function expander is not limited to multiplexers; it can be used for many other logic functions. For example, a MUXF combined with two LUTs can implement any function of five inputs, thereby implementing a full five-input LUT in a single LUT logic level. A Stratix implementation would require two or three LUTs, depending on the function, and would be implemented in two LUT logic levels.

Figure 2 shows a nine-input function mapped onto two LUTs (plus one function expander for the Virtex-II Pro architecture). The same function requires three LUTs for the Stratix device and two LUT logic levels, as opposed to a single LUT logic level for a Virtex-II Pro device.

The MUXFx component is like having a five- or six-input LUT. This leads to fewer logic levels and also far fewer LUTs consumed (10% on average) than for the same function in Stratix FPGAs. This results in higher performance for Virtex-II Pro designs because fewer logic levels are generally required for critical paths. At the same time, less placement and routing congestion occurs because 10% fewer resources (LUTs) are necessary to build the same functionality.

Shift Register LUT

A LUT in shift register mode (SRL) can implement a selectable 16-bit shift register in a single LUT. The same shift register in a Stratix device would be implemented using 16 flip-flops and as many as 10 LUTs or a memory block, a much less flexible manner.

In a Stratix PLD, if the shift register cannot be implemented in a memory block, a 16-bit shift register implemented using 16 LEs creates added routing congestion that may impact design performance. If the shift register requires variable tap selection, this will add logic levels on the output path, resulting in much slower operation.

MULT_AND

The MULT_AND arithmetic cell is commonly used in soft multiplication applications. However, the flexibility of the FPGA fabric allows some five-input functions to be mapped onto a single LUT. For example, loadable up and down counters implemented using the MULT_AND function utilize only one LUT per bit instead of two LUTs per bit, as in Stratix PLDs. This

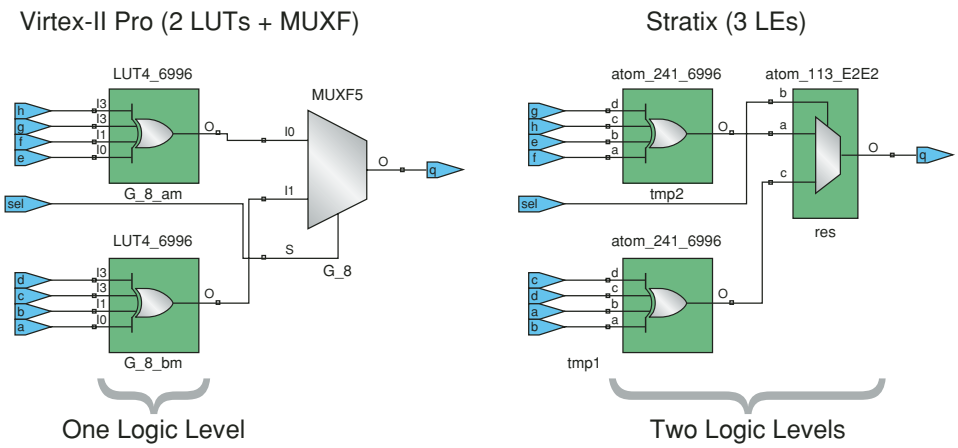

Figure 2 – Nine-input function mapped to Virtex-II Pro and Stratix devices

implementation can result in as much as 30% faster performance in Virtex-II Pro FPGAs because of the fewer logic levels and fewer required LUTs.

LUT-based RAMs

A LUT may also be configured as a single- or dual-port RAM, resulting in very fast read and write access for smaller data storing and buffering applications. In Stratix devices, the smallest RAM configuration (the M512 blocks) offers much slower RAM operation and less flexible dual-port access, while at the same time requiring greater latency for reads.

The maximum read speeds for the M512 RAMs are 266 MHz for one-clock cycle reads and 320 MHz for two-clock cycle latency, while the Virtex-II Pro SelectRAM™ memory allows 360 MHz read operation with a single clock latency, as well as asynchronous read capability for low-latency design requirements.

Because small RAMs are often used as data storage for small FIFOs, coefficient storage for DSP filters, buffers for packet processing, and other applications, having maximum performance in this structure can often enable designers to meet their system performance requirements.

Block RAMs

As most designs typically use a majority of the RAM memory available on the device, Stratix users are forced to use the MegaRAM memory blocks to create their desired functionality. For the wide (4k x 144) and deep (64k x 8) configuration of the MegaRAM, we evaluated the read/write performance of Virtex-II Pro block RAM configured to the same width and depths as the Stratix MegaRAM memory. The results, as presented in Table 1, show that for the deep and wide configuration with one clock delay, the memory read time performance in Virtex-II Pro FPGAs is approximately 40% and 95% faster than Stratix FPGAs, respectively.

The wide MegaRAM configuration has approximately 300 signals that need to be connected to the relatively small footprint of the memory block. This leads to registers and logic competing for optimal placement locations of a few sites in the array closest to these memory pins. The additional routing congestion of these signals impacts overall memory performance.

Because the Virtex-II Pro configuration was created using smaller RAMs spread out over a greater area of the chip, a more optimal placement and routing could be realized, resulting in higher performance.

Multiply and Accumulate

Stratix devices contain a dedicated DSP block; it is often assumed that it can outperform that same function created in a Virtex-II Pro device. Figure 3 highlights the maximum performance, with latency, for the two popular sizes of implementation for a multiply and accumulate (MAC): 9 x 9 and 18 x 18. This analysis shows that Virtex-II Pro devices have faster performance than Stratix devices for the MAC function.

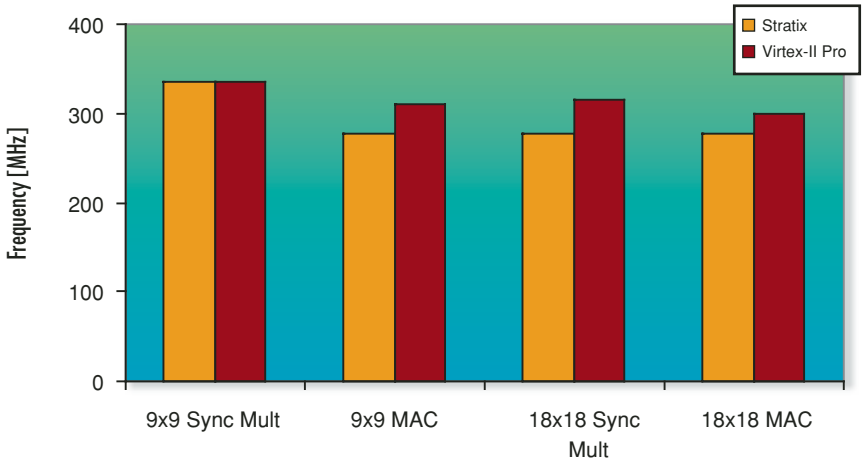

Figure 3 – Virtex-II Pro(-7) and Stratix(-5) MAC performance

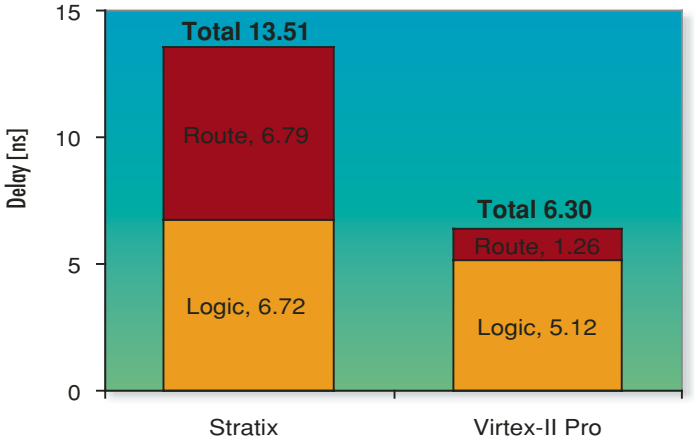

Figure 4 – Logic versus route delay on the critical path for "blowfish" design

Software Features

The FPGA fabric feature set continues to offer capabilities that improve design performance and reduce area. For users to realize these benefits, the software tools – both synthesis and place and route – need to use these architecture capabilities.

Synthesis

FPGA-centric synthesis tools constantly look for new optimization techniques that go beyond mere LUT mapping. These synthesis tools can extract known functions such as arithmetic functions, memories, and multiplexers by parsing the RTL code, automatically mapping these functions to features on the target architecture.

Synthesis mapping to the MUXF, MULT_AND, and SRL are examples of synthesis tools providing architecture-specific

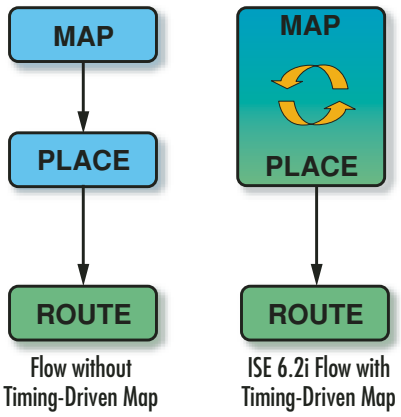

Figure 5 – ISE 6.2i timing-driven map flow

mapping to reduce logic levels on the critical paths, as well as reducing placement and routing congestion, thereby improving overall design performance. Synthesis tools will also automatically infer either the LUT RAM or block RAM based on the coding style and the size of memory being used. For example, the Synplicity® Synplify® software tool may infer fast LUT RAMs for as much as 2k of memory.

As FPGAs go deeper into sub-micron technologies, routing delays become more predominant, and design performance is highly influenced by cell placement. Thus, Xilinx provides detailed timing estimates to enable synthesis tools to not only select the best architecture element for the implementation, but also to improve timing predictability between post-synthesis and post-layout. This close technical collaboration ensures that synthesis optimization is focused on the path that is critical to place and route.

Place and Route

A study done by researchers at UCLA showed that timing-driven placement algorithms for FPGAs can average 30% off from optimal results. The study also found that Xilinx tools do much better than other tools in the industry. For instance, the delay generated by the Xilinx ISE placer was only 8.3% worse than optimal and only 4.1% worse after routing.

To illustrate this advantage, we compiled the "blowfish" encryption algorithm, an open source design, using ISE 6.2i and Altera Quartus™ 3.0 targeting Virtex-II Pro(-7) and Stratix(-5) devices, respectively. Figure 4 represents the breakdown of logic and route delay for the critical path.

This analysis shows that ISE placement technology is able to provide near-optimal placement, resulting in a 80:20 logic:route delay ratio for Virtex-II Pro FPGAs, whereas the Stratix implementation using Quartus leads to a 50:50 logic:route delay ratio. As a result, the design is two times faster when implemented in a Virtex-II Pro device.

Timing-driven map technology, new in ISE 6 software, is just one example of years of Xilinx expertise in place and route for segmented architectures. This technology enables the mapper to iterate between map and place, as shown in Figure 5, such that the placer can provide the mapper with suggested slice-level primitive mapping. This iterative loop leads to near-optimal slice mapping and placement, resulting in improved timing, because the router can now pick the best route with fewer conflicts for the same routing resources.

		Write Speed		Read Speed	
Configuration	Clock Delays	Stratix [MHz]	Virtex-II Pro [MHz]	Stratix [MHz]	Virtex-II Pro [MHz]
Deep Single-Port Memory 64k x 8	1	287	282	199	282
	2	287	282	287	282
Wide Single-Port Memory 4k x 144	1	255	284	145	282
	2	255	287	255	287

Table 1 – Virtex-II Pro(-7) and Stratix(-5) block RAM performance

Critical Settings

The performance graphs in Figure 1 show that the Stratix device outperformed the Virtex-II Pro device in one design. This is because our analysis uses default settings in synthesis, with pipelining "off." Because the design had a multiply function on the critical path, the Stratix design had an instantiated pipelined lpm (library of parameterized modules) multiplier, a black-box function generated by the Quartus MegaWizard. For the Virtex-II Pro design, synthesis inferred the MULT18x18 primitive.

By changing pipelining to "on," the synthesis tool inferred a MULT18x18S primitive for Virtex-II Pro FPGAs, resulting in an implementation with faster performance compared to Stratix FPGAs. So, in real-world designs, you'll see that Virtex-II Pro devices almost always outperform Stratix devices.

Conclusion

Advanced architecture features, such as MUXFs, SRLs, MULT_ANDs, fast SelectRAM and block RAM solutions, and fast dedicated multipliers contribute significantly to the performance advantage of Virtex-II Pro devices over Stratix devices.

The combination of an advanced architecture, the synthesis tool's capability to access architecture-specific features, and the place and route software's ability to deliver near optimal placement for a segmented architecture result in Virtex-II Pro FPGAs having a 40% average performance advantage over Stratix PLDs.

In most cases, the fastest Stratix speed grade must be used to realize the performance of the slowest Virtex-II Pro speed grade. A Stratix device in any speed grade cannot match the performance seen in the faster speed grades of Virtex-II Pro devices. Virtex-II Pro FPGAs reach a new level of performance not matched by any other FPGA in the industry today. Σ

The Ultimate Design Resource for Digital Consumer Electronics.

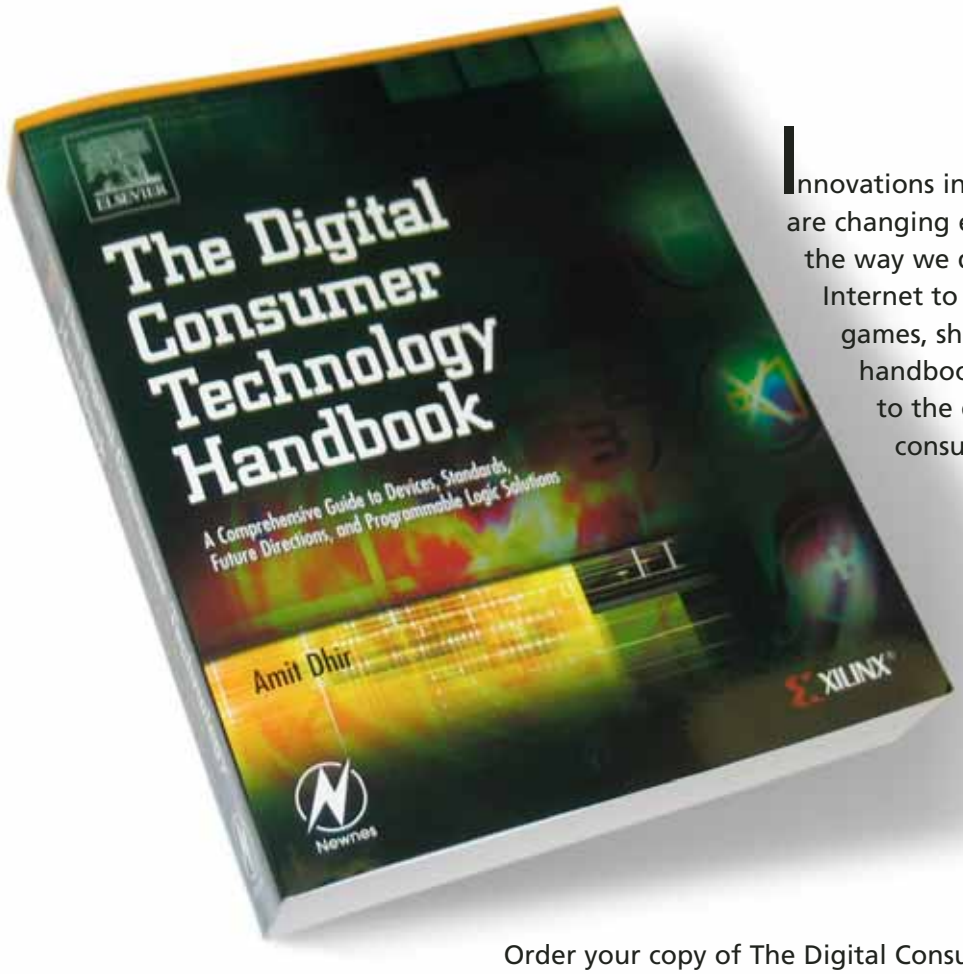

Innovations in the consumer electronics market are changing every aspect of our lives — from the way we drive, communicate, and access the Internet to the way we listen to music, play games, share images, and much more. This handbook provides a comprehensive guide to the dynamic, exciting world of digital consumer electronics, including:

- Thorough overview of all key categories — including digital TV, digital audio, mobile communications devices, PCs and peripherals, digital imaging devices, PDAs, and telematics
- Key enabling technologies
- Standards
- Delivery and reception systems
- Applications
- Networking systems

Order your copy of The Digital Consumer Technology Handbook today at **www.xilinx.com/esp** or **www.amazon.com**.

XILINX®
The Programmable Logic Company℠

The Next Gold Standard?

The Advanced Telecom Compute Architecture standard has great potential for widespread adoption in next-generation infrastructure applications.

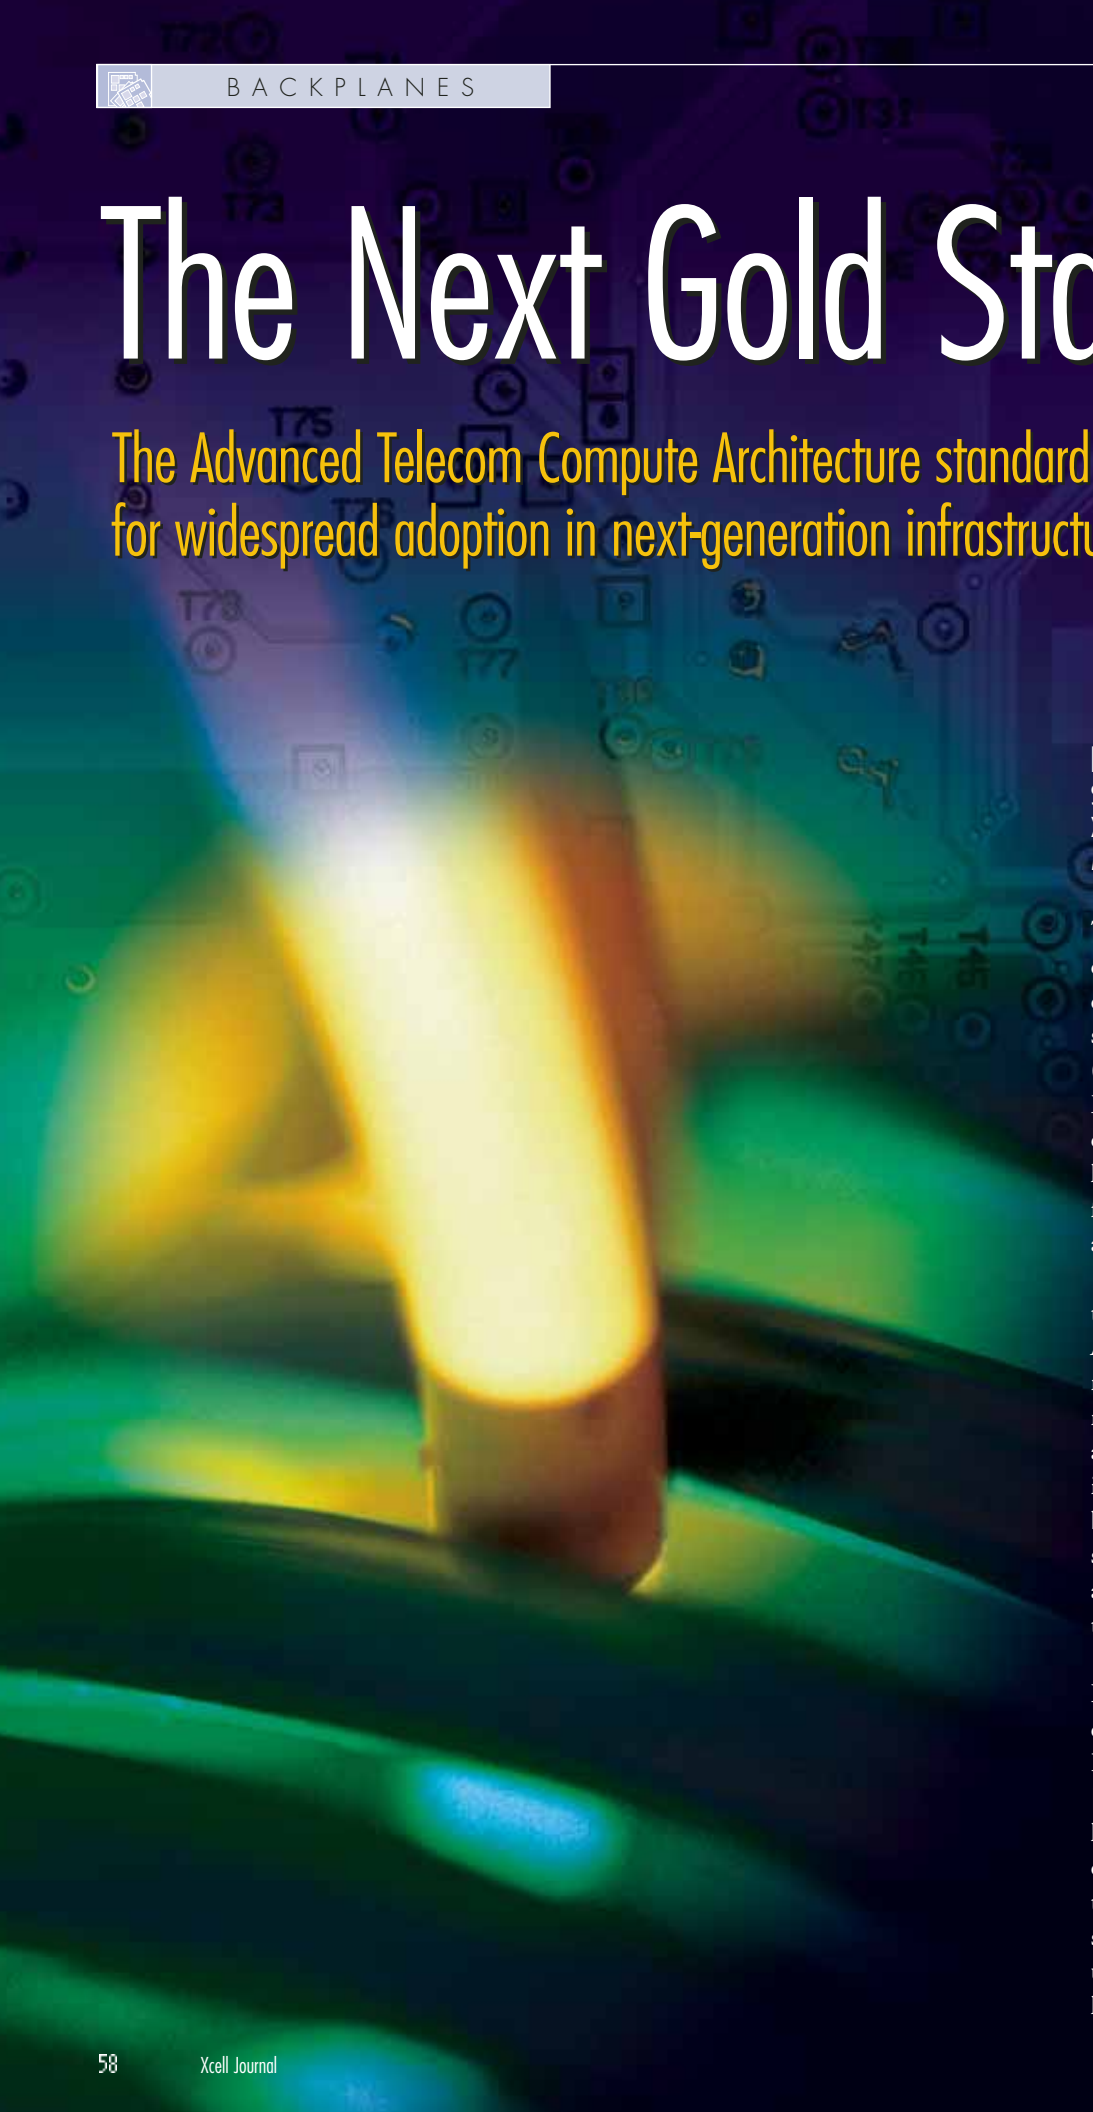

by Robert Bielby
Sr. Director of Strategic Solutions Marketing
Xilinx, Inc.
robert.bielby@xilinx.com

Technology developments and traffic demands are transforming the dynamics of the telecom market. The virtual explosion of bandwidth in local area networks (LANs), the deployment of Gigabit Ethernet, and the growth of dense wave division multiplexing (DWDM) in long-haul wide area networks (WAN) have all fueled the demand for servicing greater amounts of data traffic.

Today, it is believed that 80% of all telecommunications traffic is data traffic. Although this percentage is expected to rise, service providers continue to remain motivated to support legacy voice services, as this fundamental revenue-bearing service provides a significant base for carriers to build out their new service models. At the same time, service providers are deploying a wide range of new technologies to capitalize on new revenue opportunities.

Despite the focus on new or modified Layer 2 technologies (such as Ethernet over SONET [EOS], Resilient Packet Ring [RPR], Metro Ethernet Forum [MEF], and a host of others) that address legacy voice support as well as up-and-coming data services, challenges arise in the development of the platforms themselves. Aggressive business models continue to push for a continued model of lower-cost-per-megabit bandwidth.

The "data-friendly" Layer 2 technologies have come a long way in reducing data transport costs in some of the existing infrastructures. However, beyond those savings, achieving additional cost reductions has forced equipment providers to rethink their basic platform architectures.

A clear trend in the industry is the adoption of standard technologies over custom wherever possible. This trend is further exacerbated by the recent economic downturn – not only in the telecom market, but across almost every infrastructure market, forcing top-tier equipment providers to downsize and employ outsourced technologies. Furthermore, issues such as reduced margins, increased technology costs, rapid hardware obsolescence, and high competition have given even greater weight to a standards-based model.

Next-generation platform product development has been limited in the area of I/O signaling performance, more specifically at the point where the majority of traffic is aggregated in the backplane. The continuous scaling of system bandwidth is exceeding the capabilities of traditional backplane signaling technologies and architectures, in addition to challenging current power technologies and cooling systems.

The combination of these technical and economic factors has given rise to the definition of an industry standard for board and shelf, optimized to address the needs of next-generation infrastructure applications.

ATCA

In 2001, experts from more than 600 industries and companies collaborated to define a standardized platform that could address the challenges of future applications. This lead to the formation of a consortium under the PCI Industrial Computer Manufacturers Group (PICMG™). Previously, the consortium was responsible for the definition of PICMG 2, also known as the CompactPCI standard.

From the PICMG 3 specification, the next-generation platform dubbed ATCA™ (Advanced Telecom Compute Architecture) addresses the requirements of applications that could not be served by the CompactPCI (CPCI) standard or proprietary solutions.

Attribute	PICMG2 CPCI	PICMG3 ATCA
Board Size	57" sqr. + 2 Mez	140" sqr. + 4 Mez
Board Power	35-50W	150-200W
Backplane Bandwidth	~ 4 Gbps	~ 2.4 Tbps
Number of Active Boards	21	16
Power System	Central Converter 5,12, 3.3V Backplane	Distributed Converter Dual 48V Backplane
Management	OK	Advanced
I/O	Limited	Extensive
Clock, Update, Test Bus	No	Yes
Regulatory Conformance	Vendor-Specific	In Standard
Multi-Vendor Support	Extensive	Currently Limited
Base Cost of Shelf	Low	Moderate
Functional Shelf Density	Low	High
Lifecycle Cost Per Function	High	Low

Table 1 – PICMG2 versus PICMG3 features comparison

Finalized in January 2003, the ATCA standard has become one of the most rapidly adopted open specifications in the history of PICMG. ATCA's prime objective is to provide the benefits of a standardized yet scalable platform to address the key challenges of next-generation systems, with sufficient flexibility to be used across a broad class of applications without imposing constraints that might impact product differentiation. A key objective was that the platform could be employed in carrier-grade telecommunication applications, with support for such features as Network Equipment Building Specification (NEBS), European Telecommunications Standards Institute (ETSI), and 99.999% availability.

The ATCA platform was designed to be scalable to 2.5 Tbps; provide support for multi-protocol interfaces at rates as high as 40 Gbps; and provide high levels of modularity and configurability, allowing a range of vendors to drive competitive solutions to market.

ATCA architecture is optimized around connectivity requirements for media gateways, while providing scalability to address higher performance computing elements. ATCA was defined to support a scalable backplane environment that addresses a range of standard and proprietary fabric interfaces, primarily based on serial signaling technologies, robust system management, and support for higher performance power and cooling. Table 1 compares the key characteristics of the CPCI (PICMG 2) standard versus the ATCA (PICMG 3) standard.

The consortium employed a layered approach in the definition of the ATCA specification to accommodate support for new fabric technologies as they evolve. These layers are specified under the guidelines of the PICMG, and to date a number of them have already been defined. They include:

- PICMG 3.0 – the core specification defining architecture, mechanicals, power system management, and fabric connectors
- PICMG 3.1 – specification for Ethernet and Fibre Channel fabric interconnects
- PICMG 3.2 – specification for InfiniBand™ fabric interconnects
- PICMG 3.3 – specification for StarFabric™ interconnects
- PICMG 3.4 – specification for PCI Express™ fabric interconnects.

A report from Crystal Cube Consulting Inc. suggests that the ATCA equipment market will exceed $250 billion by 2007.

Many new layers are currently under proposal or in the process of being ratified.

In addition to supporting several fabric technologies, the backplane supports both star and full-mesh connectivity between boards in the system. System management is built on the Intelligent Platform Management Interface (IPMI) 1.5 specification. Each ATCA board supports up to 200W in a single slot, with power supplied via redundant 48V DC feeds. The result is a standard that enables solution providers to deliver products rapidly to market that support high availability and high performance, and at significantly lower costs than custom-developed or proprietary solutions.

The Market for ATCA

The confluence of a significant downturn in the infrastructure markets, competitive market pressures, and the need to address the complex and costly challenges associated with next-generation equipment platform development has caused many industries – including the telecom industry – to reconsider traditional business models. Thus, industry analysts expect the ATCA standard to achieve far greater adoption in the marketplace than previously introduced standards such as PICMG 2. A report from Crystal Cube Consulting Inc. suggests that the ATCA equipment market will exceed $250 billion by 2007.

The key benefits of the ATCA platform include lower materials costs, faster time to market, and lower development costs. Because the specification is modular in its definition, it is expected (and has already been seen through product introductions) to spawn an ecosystem of building blocks ranging from silicon solutions, boards, chassis, middleware, operating systems, and applications, among others.

The benefits to equipment manufacturers are many, as this standards-based ecosystem will allow for a lower cost of market entry/investment costs, more efficient inventory management, and a focus on higher value-added differential services while delivering cost-competitive products.

Industry analyst RHK expects shipments of more than 600,000 shelves based on the ATCA standard by the year 2007. Assuming that a shelf contains 16 cards, this translates to shipments of more than 9.6 million ATCA-based line cards.

Considering that this growth stems from an effective base of zero in January 2003, when the ATCA specification was first ratified, it's no surprise that ATCA has received a phenomenal amount of attention and press.

Industry analysts expect that the adoption of this standard will occur across various network segments at different rates – understandably so, as it provides different levels of benefits relative to where it is employed within the network. Table 2 lists the expected adoption of ATCA across various markets by 2007.

Conclusion

New business and technology paradigms continue to challenge existing business and product development models. The most recent downturn in the infrastructure markets and the introduction of many flawed business models have caused equipment suppliers to re-think their approaches to product development.

A new outsourced model based on industry standards that comprehends the requirements of specific needs for multiple markets appears to be the next major paradigm shift. Equipment suppliers need to embrace this shift to remain competitive for the next generation of platform solutions.

ATCA, which was developed, defined, and endorsed by experts from many industries, holds great promise in serving as the new disruptive technology to continue to drive down costs while increasing performance and features across a range of markets and applications.

The platform's inherent scalability and its sweeping applicability versus the significant investment costs required to develop proprietary platforms – further aggravated by the need to employ technically challenging serial signaling technologies to support next-generation backplanes – are causing equipment suppliers to seriously consider this new platform.

Once these suppliers begin to signal their intent to build products based on the ATCA standard, an entire ecosystem of modular component suppliers is expected to emerge to help further fuel the growth of this new outsourced model.

Segment	Equipment Types	ATCA System Units 2007
Wireless Access	BTS/Node B, BSC/RND, Transcoder	38%
Wireless Edge	MSC, HLR, GGSN, SGSN/PDSN, Billing Server, Multimedia Server	50%
Wireline Access	DSLAM, CMTS, MxU	1%
Edge	Edge Router, Multiservice Switch, Optical Edge Device	3%
New Access Edge	Media Gateway, Softswitch, Media Server	21%
Core Transport	Core Router, SONET/SDH, ADM, WDM	<1%
Signaling	Signaling Server, STP, SCP	5%

Table 2 – Estimated 2007 ATCA system unit shipments by equipment type (Source: RHK)

End-to-end Programmable Solutions—From the Line Card to the Backplane.

Xilinx delivers the complete, open standards-based, modular platform you need to develop designs for the line card, control plane, and high-speed serial backplane. Superior density, features, and performance make these solutions ideal for networking, telecom, data storage, and computing.

Unbeatable programmable solutions

The Xilinx Virtex-II Pro™ FPGA family offers advanced embedded features — including multi-gigabit transceivers and microprocessors — at prices that make a fully programmable line card practical, from the physical layer to the network layer.

Reduced latency, cost, and design time

Xilinx provides a proven framework for product development and deployment with the PICMG 3.0-compliant ATCA Development Platform — a 15-channel, 3.125 Gbps full mesh fabric interface, with headers for application-specific personality modules.

In addition, Xilinx offers comprehensive reference designs, IP cores and design services, making high-speed serial designs easy.

Visit **www.xilinx.com/esp/backplanes** for more information.

XILINX®

The Programmable Logic Company℠

Earth Friendly

Pb-free devices available now

FORTUNE® 2004 100 BEST COMPANIES TO WORK FOR

1984 20 YEARS 2004 LEADERSHIP

Programmable Logic Solutions for Next-Generation Serial Backplanes

Virtex-II Pro and Virtex-II Pro X FPGAs enable rapid development of flexible, serial backplane designs.

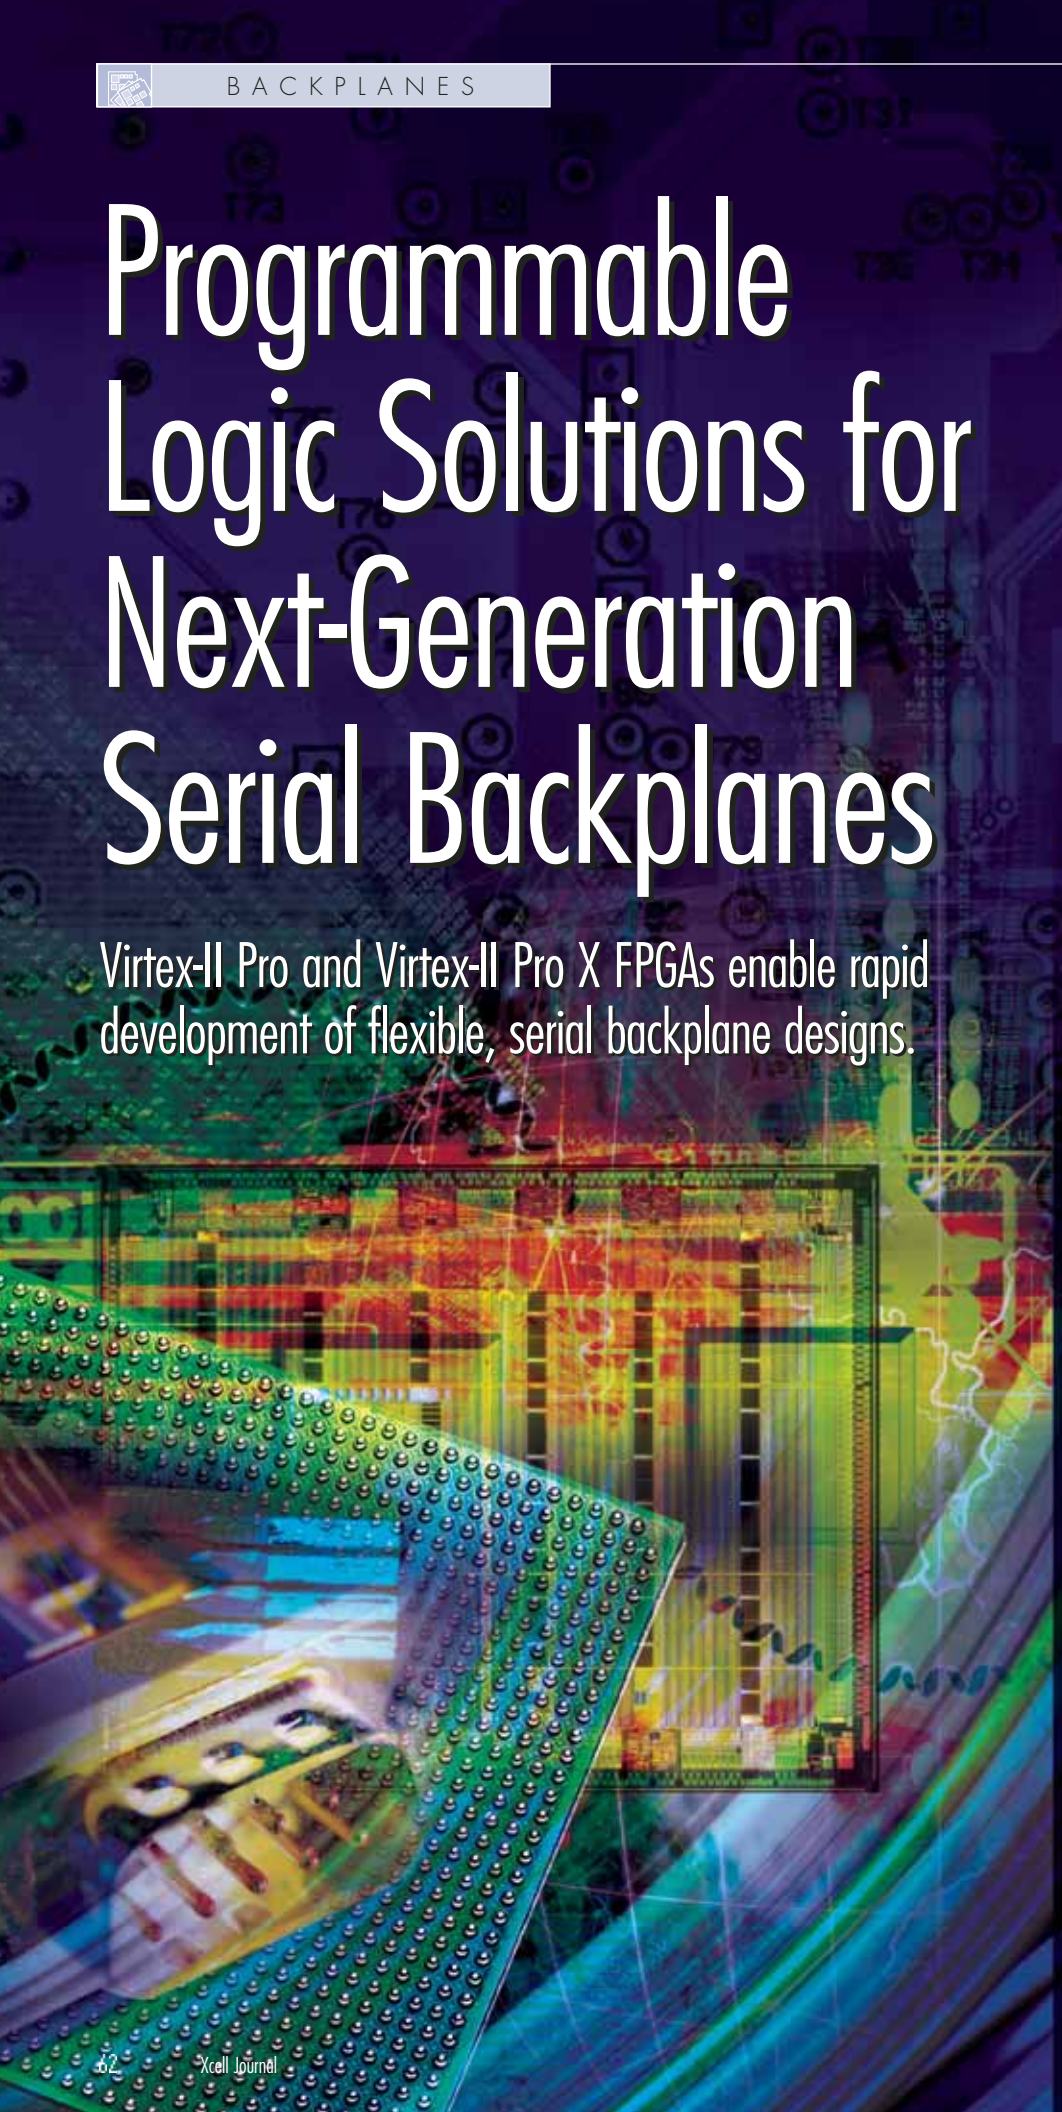

by Amit Dhir
Sr. Manager, Networking &
Telecom Markets, Strategic Solutions
Xilinx, Inc.
amit.dhir@xilinx.com

Delfin Rodillas
Manager, Networking &
Telecom Markets, Strategic Solutions
Xilinx, Inc.
delfin.rodillas@xilinx.com

Historically, designers improved bandwidth performance in telecom, datacom, and computing systems backplanes by widening buses and increasing signal clock rates. Now, with data rates exceeding 622 Mbps and reaching the 1 to 10 Gbps range across 20 inches or more of backplane trace, passing data reliably over parallel buses is a challenge. Characteristics such as signal skew and loading – non-issues before – are suddenly problematic. Consequently, designers have shifted from parallel buses to more advanced serial interconnects.

However, even serial technologies have limitations, especially at data rates beyond the 1 Gbps level, where new problems arise. These limitations include reflections due to impedance mismatches along the signal path; signal attenuation from backplane materials; and added noise due to crosstalk and inter-symbol interference.

Backplane designers should be aware of these issues and compensate accordingly to ensure that the bit error rate (BER), which is a measure of backplane robustness, is less than 10^{-12}. This challenging task becomes even more critical as system throughput requirements approach 40 Gbps.

Fortunately, you can reduce the effects of the signal degradation phenomena by several means, including:

- **Using better backplane material (FR4, Rogers)**
- **Using better connector types**
- **Improving layout trace to reduce the number of PCB layers and crosstalk**
- **Implementing different signaling schemes**
- **Using signal conditioning techniques.**

In addition, you can improve the signal integrity of a multi-gigabit serial link by selecting the appropriate serializer/deserializer (SerDes) device, comprising a transmitter, receiver, clock/data recovery (CDR), SerDes, integrated termination resistors, programmable output swing, transmit pre-emphasis, and receive equalization.

Virtex-II Pro's on-chip RocketIO MGTs allow all mesh cards on a full-mesh backplane to have direct, high-speed serial links to each other.

Standards for Serial Backplanes

The large investment required to develop a proprietary serial backplane subsystem led to the organization of the PCI Industrial Computer Manufacturers Group (PICMG™), which develops open specifications for high-performance telecommunications and industrial computing backplane architectures.

PICMG recently produced a series of specifications (PICMG 3.x) called the Advanced Telecom Computing Architecture (ATCA™) for next-generation carrier-grade telecommunications equipment. ATCA features a new form factor and is based on switched fabric architectures, including dual star, dual-dual star, and mesh topologies. The base specification, PICMG 3.0, was adopted at the end of 2002. Additional specifications in the series include PICMG 3.1 for Ethernet fabric, PICMG 3.2 for Infiniband™, PICMG 3.3 for StarFabric™ Interconnect, and PICMG 3.4 for the PCI Express™ architecture.

Xilinx Solutions for Serial Backplanes

Xilinx has made significant strides in making serial technology available in our FPGAs and developing solutions such as IP cores, reference designs, and tools to help our customers gain the benefits of serial technology easily and quickly. Let's take a look at the serial backplane solutions we offer.

Virtex-II Pro and Virtex-II Pro X

The Virtex-II Pro™ and Virtex-II Pro X family of FPGAs represent a high-end line of Xilinx FPGAs built on 130 nm, nine-layer copper and featuring an advanced fabric, embedded processors, and multi-gigabit SerDes devices. Both families are based on the same FPGA fabric, which provides abundant logic (as many as 125,000 logic cells), embedded memory (as much as 10 Mb block RAM), clock management, and DSP resources.

Standard SelectIO™ resources are also common, with as many as 1,200 user I/Os, 840 Mbps LVDS for interfaces such as 10 Gigabit Sixteen-Bit Interface (XSBI) and SerDes Framer Interface Level (SFI)-4, as well as XCITE (Xilinx Controlled Impedance Technology) on-chip termination.

Both devices also use the same embedded IBM™ PowerPC™ supporting 300 Mhz+ operation. The main difference is that the Virtex-II Pro FPGA has embedded RocketIO™ transceivers supporting speeds as high as 3.125 Gbps per channel, while the Virtex-II Pro X FPGA has embedded RocketIO X transceivers, providing up to 10.3125 Gbps per channel.

The largest of the 10-member Virtex-II Pro family of devices supports as many as 24 RocketIO transceivers. Each device can support operation from 622 Mbps to 3.125 Gbps, allowing up to 75 Gbps aggregate baud rate. Moreover, features such as programmable transmit pre-emphasis and output voltage enable the RocketIO transceiver to drive signals over 40" of FR4 material at 3.125 Gbps.

You can thus use RocketIO devices to address a number of emerging high-speed serial standards that fall within its range of operation, such as 1 Gigabit Ethernet, 10 Gigabit Ethernet (XAUI), PCI Express, Serial RapidIO, and Serial ATA.

RocketIO X transceivers found on Virtex-II Pro X FPGAs are capable of operating from 2.488 Gbps to 10.3125 Gbps. The larger of the two Virtex-II Pro X devices supports as many as 20 RocketIO X transceivers, providing an aggregate baud rate of more than 206 Gbps.

RocketIO X devices have the same features as RocketIO devices, as well as some additional features to improve signal integrity, such as receive equalization. With 10 Gbps capability, you can implement next-generation standard interfaces requiring serial 10 Gbps interfaces such as 10GBase-R Ethernet or SXI-5, or implement your own proprietary 10G interface.

Aurora

Aurora is a scalable, lightweight, link-layer protocol that you can use to move data across point-to-point serial links at baud rates as high as 75 Gbps. It is an open protocol that you can implement in any silicon device/technology. Aurora provides a transparent interface to the upper layers of proprietary or industry-standard protocols such as Ethernet or TCP/IP. This allows next-generation communication and computing system designers to achieve higher connectivity performance while preserving software infrastructure investments.

Mesh Technology on Xilinx

Mesh Technology on Xilinx (MTX) includes hardware and software reference designs and a bit error rate test (BERT) toolkit to enable rapid development of full-mesh serial backplane systems.

The PICMG 3.0 2.5G ATCA Development Platform is a reference board for PICMG 3.x line cards supporting port rates to 2.5 Gbps. The heart of the development platform is the Virtex-II Pro device, which serves as the interface to the full-mesh backplane.

Virtex-II Pro's on-chip RocketIO MGTs allow all mesh cards on a full-mesh backplane to have direct, high-speed serial links to each other. The full-mesh card also allows application flexibility by reserving an area of the board for a pluggable "personality module" (PM). You can use the PM to implement any application-specific line card and easily connect to the full-mesh card through the included headers.

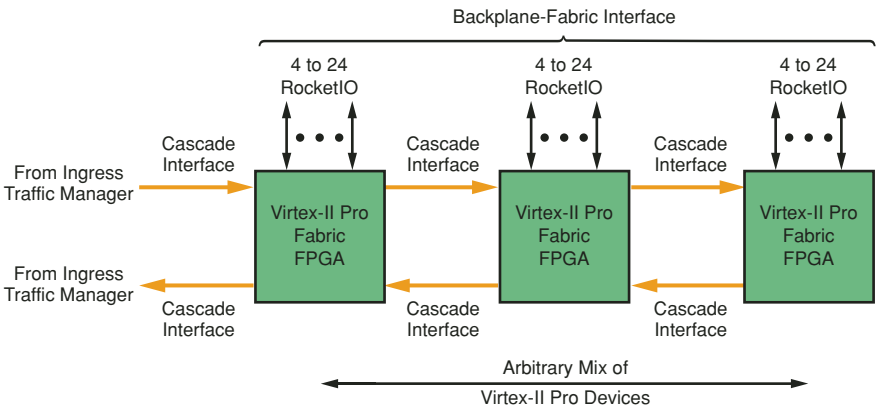

Figure 1 – You can implement the Mesh Fabric Reference Design in a single FPGA or in multiple, daisy-chained FPGAs.

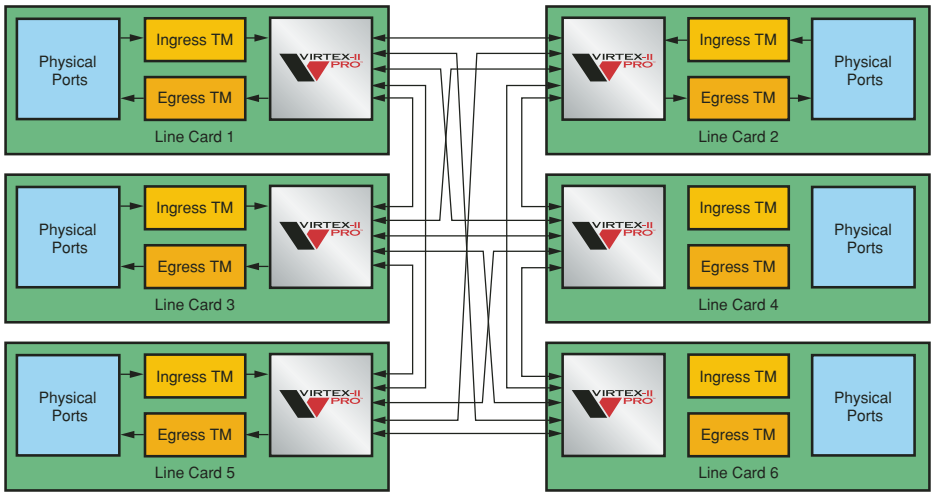

Figure 2 – In a Virtex-II Pro-based full-mesh backplane, you can implement serial channels and distributed switch functions using RocketIO transceivers and logic resources with the Mesh Fabric Reference Design.

PICMG 3.0 also specifies card and shelf management functionalities that are implemented in the development platform.

The Mesh Fabric Reference Design (MFRD) is a fully functional IP reference design that provides a building block for creating Virtex-II Pro-based mesh switch fabric interfaces. You can use the fabric reference design in a single Virtex-II Pro device or in several daisy-chained Virtex-II Pro devices, allowing up to 256 RocketIO serial channels. Figure 1 shows the concept of the mesh fabric interface and the daisy-chain scheme.

By having the flexibility to daisy-chain devices of different densities, you can choose an appropriate logic-to-RocketIO ratio. For example, more logic may be useful in a design where additional network processing functions are needed beyond those provided by an ASSP or ASIC. Flexible traffic scheduling is also possible with the support of as many as 16 priority levels and multiple scheduling algorithms on egress.

Figure 2 illustrates an example system with line cards that use the Virtex-II Pro FPGA and the full-mesh IP as the backplane interface.

GigaBERT is an IP toolkit that enables easy and comprehensive BERT of Virtex-II Pro-based, full-mesh fabric channels. Using GigaBERT, you can configure each RocketIO transceiver on each mesh fabric interface FPGA connected to a backplane as either a BERT tester or a far-end loopback. In effect, you can accomplish a scheme for simultaneous BERT testing of all links in a full-mesh fabric. Furthermore, GigaBERT's flexibility enables you to quickly and easily create a BERT stress test to check for signal integrity in specific configurations.

Legacy Backplanes Support

Xilinx FPGAs are also ideal for customers who still need to support their differential or single-ended legacy bus architectures as they transition to serial architectures. For the highest performance differential solution, you can use the Virtex-II Pro or Virtex-II Pro X FPGAs to achieve LVDS rates as high as 840 Mbps. For low-cost LVDS, Spartan™-3 FPGAs support rates as high as 622 Mbps. Together, these devices provide a complete differential I/O solution with coverage of popular standards such as LVDS, Extended LVDS, Bus LVDS, Ultra LVDS, LVPECL, LDT, and RSDS.

For legacy designs using older single-ended signaling standards, the Xilinx SelectIO technology available in Virtex-II Pro, Virtex-II Pro X, and Spartan-3 FPGAs allows the most comprehensive support for LVTTL, LVCMOS, PCI/PCI-X, GTL, HSTL, and SSTL signaling standards. As a result, Virtex-II Pro, Virtex-II Pro X, and Spartan-3 FPGAs provide you with everything you need to support legacy backplane interfaces.

Conclusion

Designers of high-end telecom, datacom, and computing platforms have looked towards serial I/O technologies to address the increasing performance requirements of next-generation systems. Additionally, consortia such as the PICMG have stepped up to the plate to define serial backplane standards.

Whether proprietary or standards-based, Virtex-II Pro and Virtex-II Pro X FPGAs with embedded multi-gigabit serial transceivers provide the technology to enable serial backplanes – including advanced, full-mesh architectures. Our growing portfolio of IP cores, reference designs, and toolkits for serial backplanes such as Aurora, Mesh Fabric IP, the PICMG ATCA Development Platform, and GigaBERT lead to shorter time to knowledge and ultimately shorter time to market. For more details about Xilinx solutions for serial backplanes, visit *www.xilinx.com/esp/backplanes/*.

Create ATCA-Compliant Designs

Xilinx and Avnet have released a new design kit that reduces time to market for a wide range of serial backplane applications.

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

Traditionally, designs for a variety of applications used high-speed backplanes to provide high-bandwidth communications between subsystem cards. Parallel bus implementations like PCI were popular because they offered the highest bandwidth in an industry-standard form factor.

However, for applications requiring very high bandwidth connectivity (such as telecom and networking), these parallel implementations ran into bandwidth and cost problems. Non-standard implementations were sometimes needed; these custom efforts slowed development and increased costs.

Technological advances now allow you to use high-speed serial interfaces cost-effectively in chassis-based, industry-standard designs. The new Advanced Telecom Compute Architecture (ATCA™) PICMG™ 3.1 specification creates a flexible, industry-standard platform that lets you cut-and-paste previously complex and expensive high-speed serial portions of your design. This improves time to market and significantly reduces the cost normally associated with creating high-speed backplane designs.

Figure 1 – ATCA card cage

We expect that this change will open the market to a wide range of new applications and companies that were historically shut out of these designs. Xilinx and Avnet have partnered to create a complete ATCA PICMG 3.1 Design Kit that can be used to quickly and easily implement the high-speed serial backplane portion of the ATCA PICMG 3.1 specification; it can also be used as a platform for a complete design.

PICMG 3.1 Design Kit

The card cage, shown in Figure 1, is a PICMG-standard 12U form factor sized for 16 slots in a 600 mm frame, with room for both front and rear fiber bend. The boards measure 8U x 280 mm x 1.2 in (140 in^2 + 4 mezzanine connectors), can run 150-200W of power, and can provide 2.4 Tbps of bandwidth. There are as many as 16 active boards per chassis.

The power is at 48V and sourced from the backplane.

The main component of the ATCA PICMG 3.1 Design Kit is the line card, which is a complete development platform for creating PICMG 3.1-compliant designs. Some of the key features are:

- 15-channel, one-port full mesh fabric interface
- Intelligent Platform Management Interface (IPMI)
- Base interface ShMC port
- Headers for an application-specific personality module
- Fully distributed system management
- Management firmware running on a PowerPC™ processor
- Linux™-based control plane software.

Figure 2 – PICMG 3.1 line card

Line Card

The Xilinx ATCA PICMG 3.1 full mesh line card (Figure 2) provides a baseline implementation of a PICMG 3.1 line card. It includes a Virtex-II Pro™ FPGA that implements both a full mesh fabric interface and a management subsystem.

The full mesh line card can serve as a development platform for PICMG 3.x line cards supporting port rates to 2.5 Gbps. It includes a Virtex-II Pro-based fabric interface that also includes all PICMG 3.0-defined card and shelf management functionalities. Management firmware executes on one of the Virtex-II Pro's PowerPC processors running an embedded Linux operating system.

The card also includes headers to interface to a user-defined personality module. This module is used to implement application-specific line card processing and external interfaces. I/O access for this module can be reached through the front panel or rear transition modules. The personality module also has full access to the PICMG 3.1 update channel interface.

Fabric Interface FPGA

The fabric interface FPGA implements not only the data plane functions needed to transfer data across the distributed fabric, but also all management functions defined in the PICMG 3.1 specification. When placed in slots one or two, the card is capable of acting as a shelf manager.

The control plane section of the fabric interface FPGA implements management functions for the card; a block diagram for the control plane implementation is shown in Figure 3. All of these functions are imple-

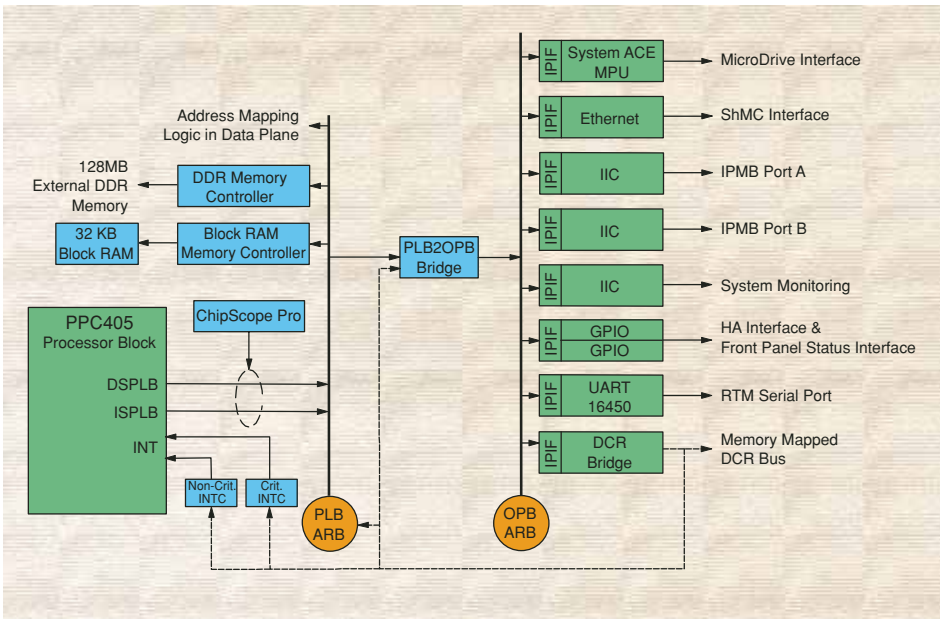

Figure 3 – Fabric FPGA control plane block diagram

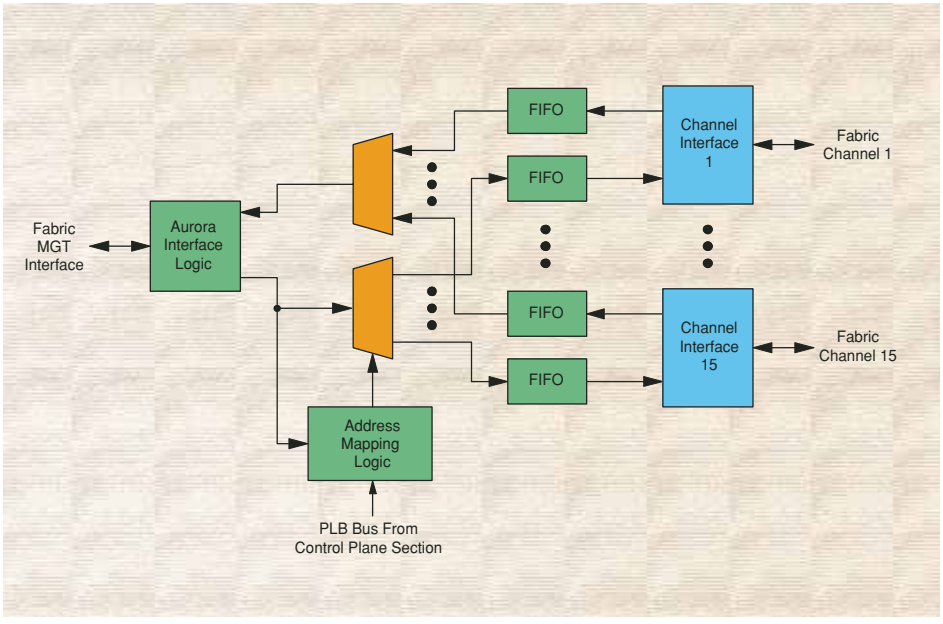

Figure 4 – Fabric FPGA data plane block diagram

mented as firmware running on an embedded Linux operating system. The functions provided include an IPMI agent, shelf manager, and hardware and software updates via an ShMC interface.

The Virtex-II Pro FPGA includes two 400 MHz PowerPC 405 processors. One processor is used to implement management functions. It interfaces to the rest of the management subsystem by way of a 64-bit CoreConnect processor local bus and a 32-bit on-chip peripheral bus. The second PowerPC processor is available for application-specific functions.

The data plane section implements a complete 15-channel distributed switch fabric interface. The configuration shipped with the card implements a PICMG 3.1 Ethernet transport, but it can also be customized to support other PICMG 3.x transports. Figure 4 shows a block diagram of the data plane section of the fabric interface FPGA.

The Aurora interface is used to transfer packets between user-defined logic on the prototyping module and the PICMG 3.x fabric. The Aurora interface uses the fabric interface multi-gigabit transceiver signals for connectivity, but you can substitute other interfaces. For example, if you used an alternative interface such as POS-PHY Level 3, the fabric interface GPIO signals would be used for connectivity.

Conclusion

Xilinx has certified Avnet Cilicon, via the Avnet Design Services Design Centers, to sell and support the ATCA PICMG 3.1 Design Kit. The kit includes detailed design files, a comprehensive board support package, and example designs, along with test results. Design Services can be bundled along with the Design Kit to help port a custom design to the line card FPGA.

To get the most up-to-date information on the ATCA PICMG 3.1 Design Kit, visit *www.avnetavenue.com* and select "ATCA Design Kit." To obtain pricing, delivery information, and a more complete description from an Avnet Cilicon representative, click on "To Register." Σ

APPLICATIONS

The ATCA PICMG 3.1 specification defines a flexible serial backplane development platform that is applicable to a wide variety of applications. In general, the specification targets Telco carrier-grade applications, but it is also applicable to data centers and other more computationally intensive applications.

Typical application areas include:

- Narrowband line units
- Narrowband local switch line or trunk unit
- Digital loop carrier local terminal/ONU
- PBX line unit
- Broadband line units
- DSLAM
- Cable modem termination system/head end
- FTTx line unit
- Wireless elements
- Base transceiver station
- Base station controller
- Wireless access gateway
- Radio network controller
- SGSN/GGSN
- Home location regulator
- Integrated mobile switching center
- Service nodes
- Echo canceller
- Network resource server/intelligent peripheral
- Remote access server/modem pool
- IVR/voicemail system
- Core data network elements
- Switched LAN hub
- IP switch/router
- ATM switch
- Optical transport terminal (DACS, WDM)
- Metro optical system
- Data network elements
- ASP server
- Storage area network element
- Compute server (thin client host, game host)
- Web server (e-commerce, web cache, firewall, filter)
- Database engine (RADIUS, LNP, billing)
- Video server
- Converged switch elements
- Softswitch
- Line access gateway
- Trunk access gateway
- Signaling gateway
- Internet telephony host
- Compression/vocoding/encryption gateway
- PSTN elements
- Universal AIN element (SCP, SCC, NCP, STP)
- DLC/GR-303 host terminal
- TDM switch core replacement
- PBX
- E.911, CALEA host
- Industrial applications
- Factory automation/robotics
- Multimedia studios
- Traffic control
- Military/avionics/shipboard

Ethernet Aggregation with GFP Framing in Virtex-II Pro

A new reference design from AMIRIX Systems and Xilinx allows aggregation of multiple Gigabit Ethernet ports to SPI-4.2, with frame-mapped GFP.

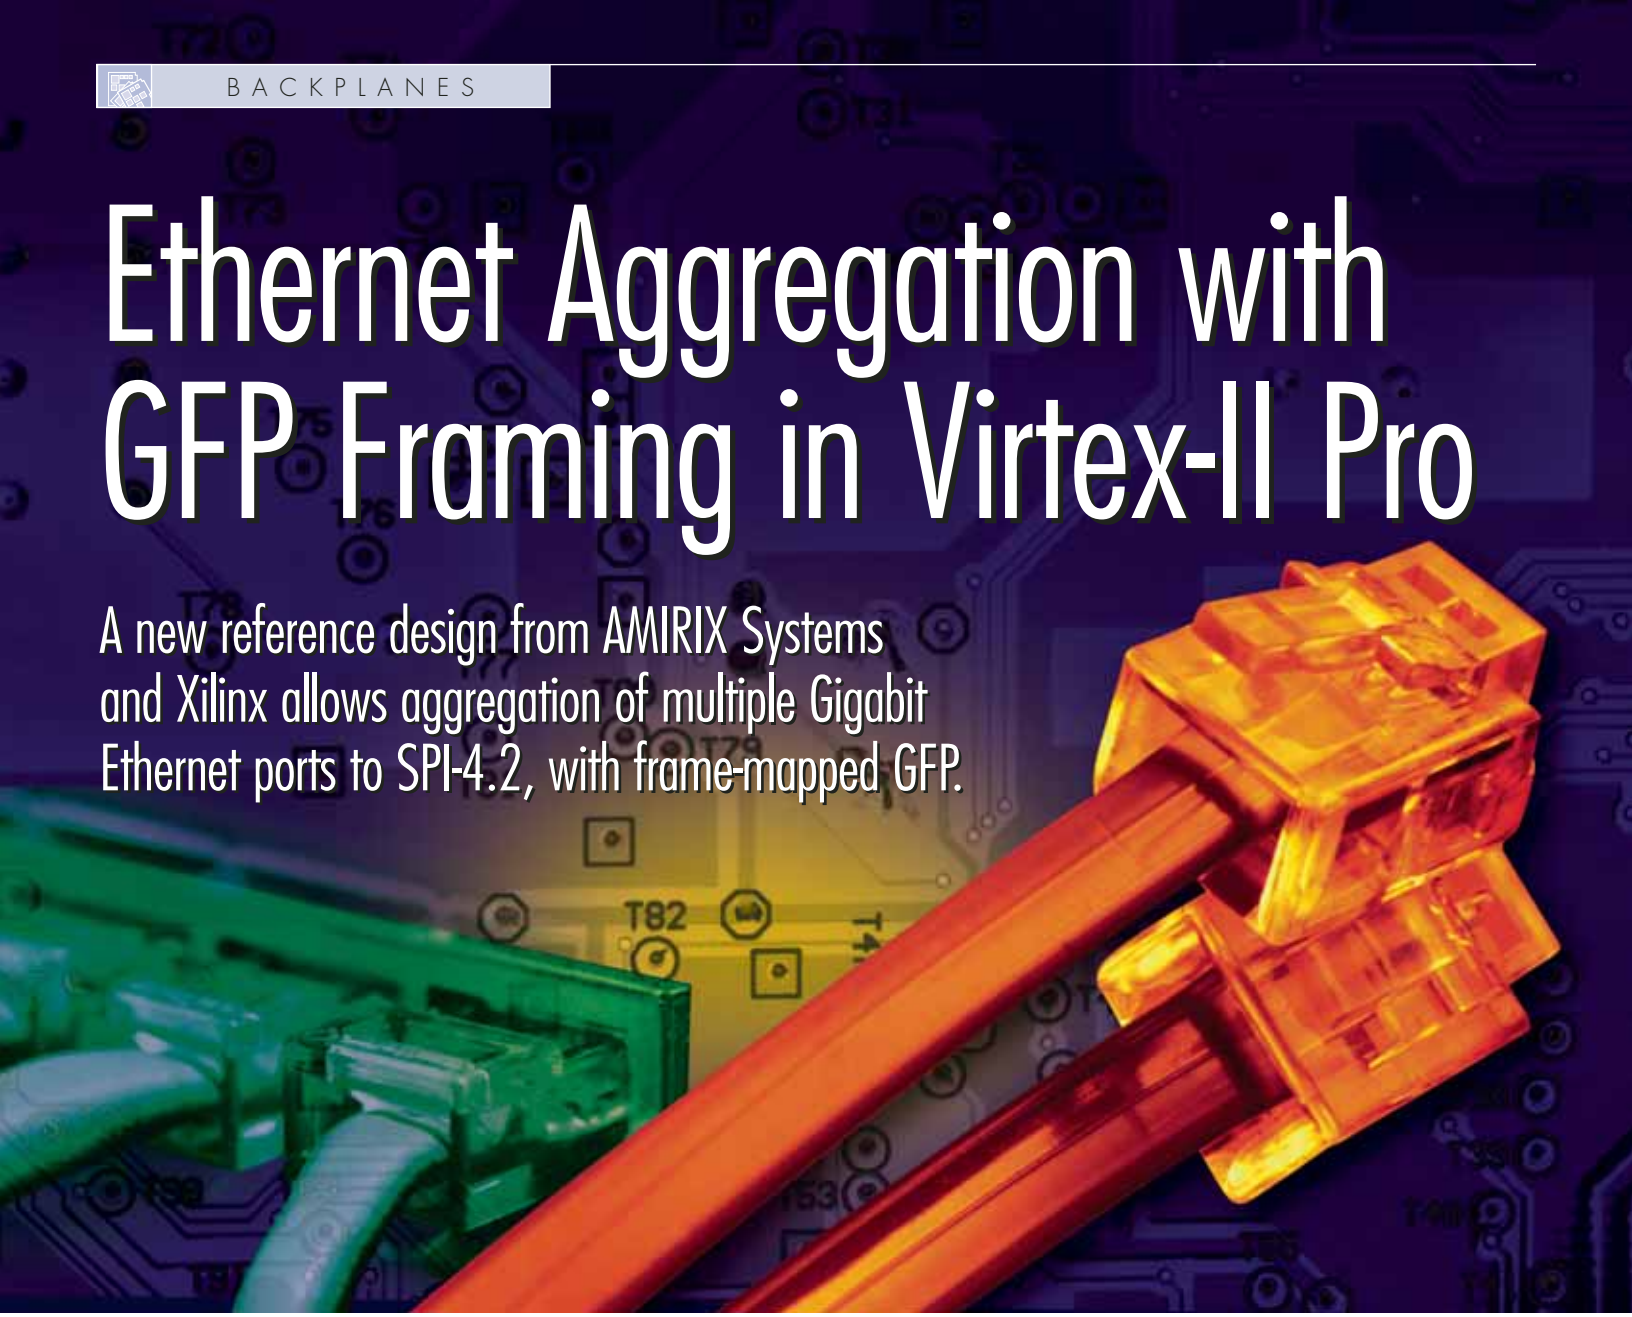

by Bruce Oakley
Director of Embedded Systems Design
AMIRIX Systems
bruce.oakley@amirix.com

Although the existing transport network infrastructure was built for carrying voice, it now carries other types of traffic, such as video, data, and storage. Network services like asynchronous transfer mode (ATM) carry this traffic with varying degrees of overhead and impact on performance.

The Generic Framing Procedure (GFP) – as defined by the International Telecommunication Union (ITU) Telecommunication Standardization Sector (ITU-T) Recommendation G.7041 – offers another solution. GFP defines framing methods for mapping different traffic types directly to the octet-synchronous optical network (SONET) infrastructure.

GFP comprises two stages: client-specific mapping of different protocols into frames, and common procedures to adapt frames to an octet stream. The flexibility of FPGAs makes them a natural solution for the first stage, in which supporting a variety of different interfaces is necessary. A multiplexer can then aggregate the resulting GFP frames and send them to a framer for adaptation to SONET. Framing and aggregation of Gigabit Ethernet frames to a SPI-4.2 interface, as shown in Figure 1, will be a very common building block.

The Xilinx Virtex-II Pro™ FPGA offers a very powerful platform on which to build such a system. The Gigabit Ethernet and SPI-4.2 interfaces can be driven directly by MGT and LVDS I/O, respectively, and proven IP cores for these functions exist.

Using the programmable array for framing and multiplexing allows a great deal of flexibility, which you can use to support different algorithms for application-specific functions such as mapping, scheduling, and flow control. You can also include a control plane subsystem in the same device using the embedded PowerPC™. Such a solution has been developed and tested by AMIRIX™ Systems, which Xilinx now offers as a free reference design.

Architecture and Data Flow

The basic architecture of the Ethernet Aggregation Reference Design (EARD) is shown in Figure 2. Although the architecture shown in the figure is for a four-port

system, the EARD can also be configured for eight ports. In the egress direction, frames arriving at the Ethernet ports are multiplexed and segmented into the SPI-4.2 interface. Segments are de-multiplexed and reassembled in the ingress direction.

Egress access to the SPI-4.2 interface is scheduled according to a simple round-robin algorithm, with a one-to-one mapping of Gigabit Ethernet ports to SPI-4.2 channels.

We should note that EARD traffic directions (ingress and egress) are defined from the point of view of a SONET framer. The opposite sense is used in the GFP standard.

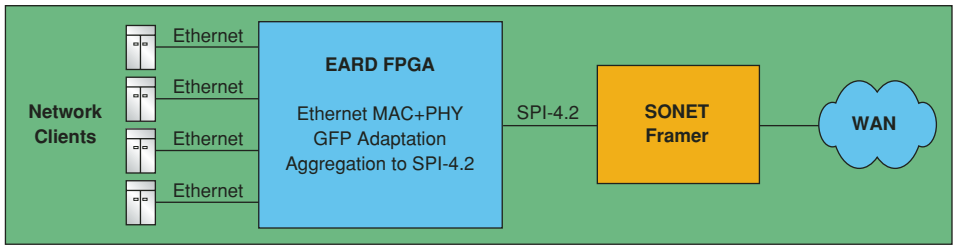

Figure 1 – EARD network context

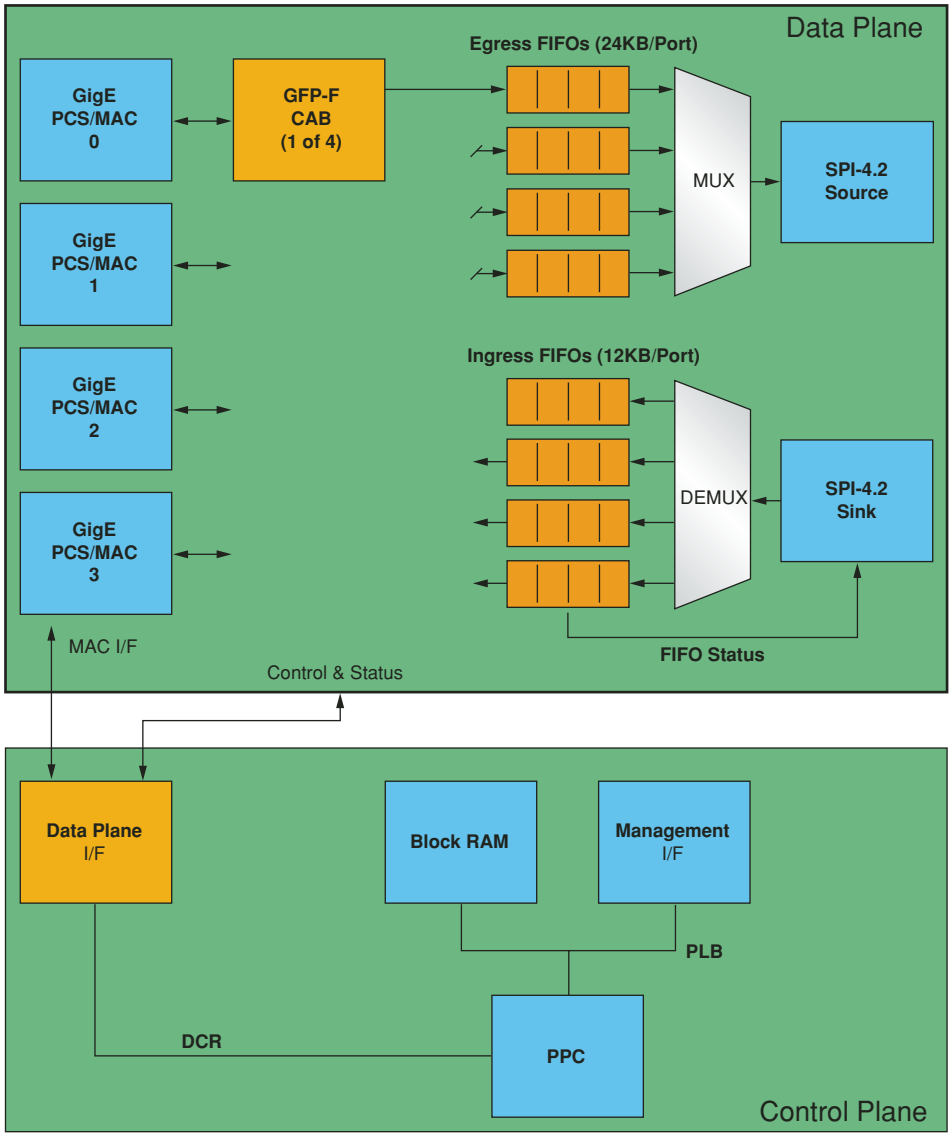

Figure 2 – EARD block diagram

GFP/Pass-Through/Loopback Modes

When in GFP mode, the EARD supports frame-mapped GFP for Ethernet medium access control (MAC) payloads, as defined in Section 7.1 of the GFP standard. The EARD adds headers in the egress path and strips them in the ingress path; header contents are set using compile parameters. To correctly encode the length during GFP encapsulation, an entire frame must be buffered before forwarding. This store-and-forward approach is used in both egress and ingress directions when GFP framing is enabled.

When GFP framing is disabled (pass-through mode), a lower latency approach is used. Forwarding begins upon receipt of an entire SPI-4.2 segment during egress, or when reaching a programmable threshold during ingress.

The EARD can also be configured in loopback mode, in which traffic at each port is fed back to itself. The SPI-4.2 sink client interface is connected directly to the SPI-4.2 source, and Gigabit Ethernet traffic is looped back through the egress and ingress FIFOs.

Interfaces

Both the Gigabit Ethernet and SPI-4.2 interfaces are implemented using Xilinx IP cores. Features of these interfaces include:

Gigabit Ethernet

- Core – Xilinx Gigabit Ethernet MAC revision 3.0
- Physical Interface – Built-in physical layer device (physical coding sublayer [PCS]/physical medium attachment [PMA]) using MGT
- Frame Size – Support for jumbo frames
- Flow Control – Support for incoming and outgoing pause frames; pause frames have fixed delay (compile parameters) triggered by a programmable egress FIFO threshold
- Statistics – Traffic statistics maintained by the MACs, accessible through the management interface.

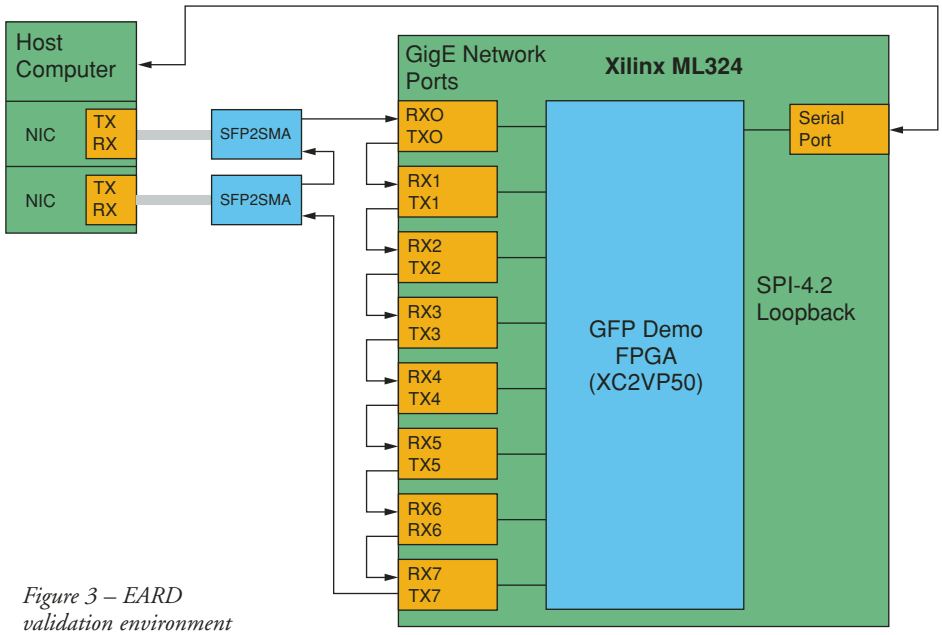

Figure 3 – EARD validation environment

SPI-4.2

- Core – Xilinx SPI-4.2 revision 6.0
- Phase Alignment – Dynamic phase alignment
- Flow Control – Sink status (ingress path) is reported based on programmable ingress FIFO thresholds; source status is not used.

Control Plane

EARD control plane software runs on an embedded PowerPC, clocked at 250 MHz. It manages system initialization and provides an external management interface. This management interface is based on a simple serial port, which is useful for demonstration purposes. In an actual application, a more sophisticated interface such as PCI or RapidIO would be more useful.

To facilitate porting to different physical interfaces, the management interface software is based on a generic message passing protocol. You can support changes to the physical interface by modifying a few low-level routines.

Because the control plane is a relatively small part of the system, we preferred an ISE-centric design flow. Thus, the control plane was built using EDK, but is exported as a sub-module and integrated into the EARD as a core.

Validation

We performed all EARD validation in a Xilinx XC2VP50 device. The four-port version should fit in an XC2VP30, and with some customization (such as running control plane software directly from cache), we expect that the eight-port version can fit in an XC2VP40. Table 1 shows the approximate resource usage.

The EARD was tested through a combination of simulation and hardware validation. The various configurations were subjected to extended heavy traffic, as well as tests focused on exercising segmentation and reassembly, scheduling, and error handling.

We performed hardware validation using a Xilinx ML324 board, as shown in Figure 3. LVDS headers were used to loop back the SPI-4.2 interface, and the Gigabit Ethernet ports were daisy-chained. We used an optical network interface card (NIC) in a Linux™ host computer, as well as laboratory network analysis equipment, to generate and check traffic. The EARD carried hundreds of millions of Ethernet frames of varying sizes at data rates well over 900 Mbps on all ports.

	Block RAM	4LUT	FF	DCM	GCLK	MGT	GPIO
Four-Port	123	18500	18100	5	11	4	96
Eight-Port	215	29800	27300	5	11	8	96

Table 1 – EARD resource usage

Conclusion

The Ethernet Aggregation Reference Design makes an excellent starting point for designs requiring Gigabit Ethernet aggregation, particularly those involving GFP framing. This reference design is described in Xilinx Application Note XAPP695 and can be downloaded from the Xilinx website at *www.xilinx.com/esp/networks_telecom/optical/xlnx_net/eard_download.htm.*

Using the EARD in real applications will likely involve some degree of customization to meet system needs. Examples include:

- Changing SPI-4.2 configuration options to comply with PCB requirements and SONET framer specifics
- Replacing the management interface with something more suitable for an embedded system, such as PCI or RapidIO
- Modifying the algorithms used for scheduling, mapping, and flow control.

Of course, you can make more extensive architectural changes as well, such as adding queue management or replacing the Ethernet ports with different interfaces. You can make changes yourself using the freely available source code, or leverage AMIRIX Systems' extensive experience with FPGA design and the EARD. We have applied our FPGA design capabilities to a number of communication systems involving queueing, classification, segmentation and assembly, and a variety of customized packet processing functions. For more information, visit *www.amirix.com,* or e-mail *info@amirix.com.*

Mesh Fabric Switching with Virtex-II Pro FPGAs

Implementing mesh fabric architectures has just gotten easier with the Xilinx Mesh Fabric Reference Design and ATCA Development Platform.

by Mike Nelson
Sr. Manager, Strategic Solutions
Xilinx, Inc.
mike.nelson@xilinx.com

The introduction of the Virtex-II Pro™ Platform FPGA with integrated multi-gigabit transceivers (MGTs) enabled a new era of system design. Specifically, Virtex-II Pro devices now enable designers to implement switched fabric system architectures efficiently, affordably, and entirely in programmable logic.

To illustrate this point and enable its rapid exploitation by our customers, Xilinx developed the Mesh Fabric Reference Design (MFRD), a modular, highly scalable, and configurable resource for building switched fabric system solutions, and the Advanced Telecom Compute Architecture (ATCA) Development Platform. In this article, we'll take a close look at both tools.

Switched Fabric Topologies

The classic switched fabric configuration is a star in which each node communicates with all of the other nodes through a central switch (Figure 1A). The obvious limitation of a star is that it is not fault tolerant. To address this limitation, you need a dual star (Figure 1B).

In a mesh fabric, the switching function is distributed across the system; every node connects directly to each and every other node. This configuration is inherently resilient, as shown in Figure 1C.

To compare the performance of these alternatives, let's consider two atypical 16-slot configurations: a dual star with 10 Gb links, and a mesh with 2.5 Gb links. Because these configurations require approximately the same number of MGT resources for implementation (224 for the star versus 240 for the mesh), they are essentially equal from a power and system cost perspective (i.e., connector and backplane routing resources).

The maximum theoretical system bandwidth for a dual star is equal to the number of nodes times the link rate times two (as all links are full duplex). In our 16-slot example, this works out to 14 nodes (two slots are required for the switches) x 10 Gb x 2 = 280 Gb.

The maximum theoretical system bandwidth for a mesh is equal to the number of nodes times the number of links per node (nodes minus 1) times the link rate. In our example, this works out to 16 (all slots are nodes in a mesh) x 15 x 2.5 Gb = 600 Gb.

The mesh configuration is able to achieve more than twice the system performance with essentially equal resources because half of the star is required simply for fault tolerance. Additionally, the star incurs a fractional performance hit because two slots must be dedicated to switching in its chassis, thus limiting the node count.

In fairness, we should note that a dual star can double its theoretical bandwidth to 560 Gb if it uses active-active load balancing, but not with fault tolerance. That would require the addition of a third switch for failover, increase the MGT count to 312, and reduce performance to 520 Gb in a 16-slot chassis, as the node count decreases to 13. Table 1 compares the performance of these configurations, along with additional examples.

			16-Slot Chassis Configuration	
Fabric Topology	MGT BW	Link BW	Aggregate System BW	MGTs Required
4X Star	2.5 Gb	10 Gb	300 Gb	120
4X Dual Star	2.5 Gb	10 Gb	280 Gb	224
Active-Active 4X Dual Star	2.5 Gb	10 Gb	560 Gb	224
A-A 4X Dual Star with HA*	2.5 Gb	10 Gb	520 Gb	312
1X Full Mesh	2.5 Gb	2.5 Gb	600 Gb	240
2X Full Mesh	2.5 Gb	5 Gb	1.2 Tb	480
4X Full Mesh	2.5 Gb	10 Gb	2.4 Tb	960

* Requires three switches

Table 1 – Performance comparison of various star and mesh fabric configurations

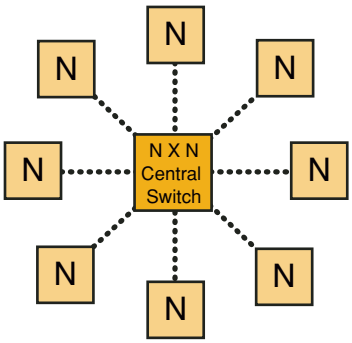

Figure 1A – Star fabric configuration

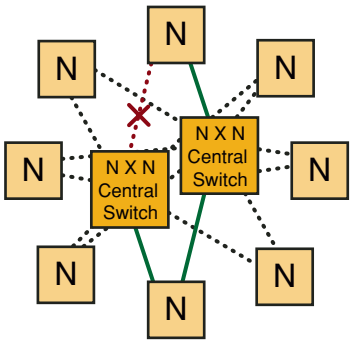

Figure 1B – Dual star fabric configuration

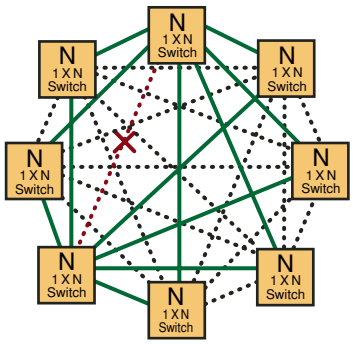

Figure 1C – Mesh fabric resiliency

Figure 1 – Switched fabric topologies

Mesh Fabrics Fit Virtex-II Pro FPGAs

Before the advent of abundant and affordable MGT resources, mesh fabrics were challenging to implement. Now, they're an emerging segment – historically an excellent home for programmable logic.

The *distributed* nature of switching in a mesh fabric enables a mesh to map extremely well to the resources available in Virtex-II Pro Platform FPGAs. These products have everything you need to build exceptional mesh fabric interconnects:

- Four to 24 MGTs per device for implementing serial links
- Block RAM for implementing queues
- Logic for implementing control and traffic management functions
- Embedded PowerPC™ processors that can be used to implement management functions.

Mesh fabrics will also scale well in next-generation Virtex-II Pro X™ Platform FPGAs. The Pro X family introduces 10 Gb MGTs that can quadruple the performance for our 16-slot mesh example to an incredible 2.4 Tb.

The Xilinx Mesh Fabric Reference Design

To enable Virtex-II Pro applications in mesh fabrics, Xilinx developed the Mesh Fabric Reference Design. The MFRD enables an extremely broad range of system configurations.

When designing the MFRD, Xilinx set out to address a number of key objectives:

1. Support system configurations from a few to hundreds of ports
2. Enable flexibility for implementing a chosen configuration and thus the ability to cost-optimize the solution
3. Provide configurable and competent queue management functionality
4. Enable efficient use of fabric bandwidth
5. Support standard Xilinx interfaces on modular boundaries
6. Enable processor-based switch management.

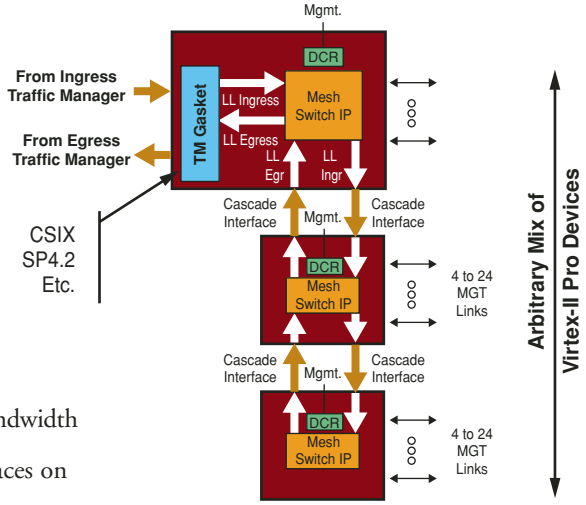

Figure 2 – MFRD architecture

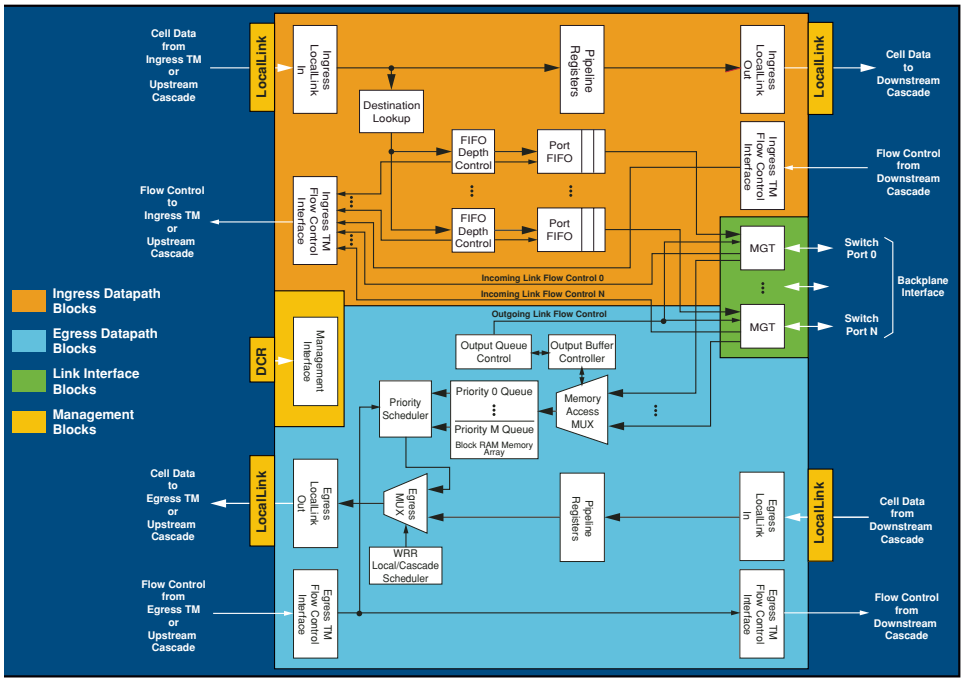

Figure 3 – MFRD block diagram

To achieve these goals, the MFRD implements a mesh switching architecture, as illustrated in Figure 2. The MFRD specifically implements a "mesh switch IP" element illustrated in each device in the figure. We will review the details of this IP, but for now let's focus on the bigger picture.

The MFRD implements a modular architecture that can be realized in one or more components. This enables configurations from four to 256 ports in any mix of Virtex-II Pro FPGAs and provides designers with exceptional flexibility in configuring their systems. For instance, you could implement a 16-port switch in a single 2VP50, in a combination of a 2VP20 and 2VP7, or in two 2VP7s. This flexibility is ideal for optimizing the price/performance of the solution to your specific needs.

Other aspects to note in Figure 2 are:

- The use of the standard LocalLink interface for switch ingress and egress
- The use of the device control register (DCR) bus for switch management by the Virtex-II Pro embedded PowerPC RISC processor
- The traffic management gasket: While a key element of any design, it is important to note that this interface will differ for every application and is therefore beyond the scope of the MFRD.

Internally, the MFRD is a cell-based switch architecture supporting 40 to 128 byte payloads. To understand its operation, let's look at a block diagram and follow the course of traffic from ingress through egress; in this way we can easily understand its features and capabilities. The basic structure of the MFRD is illustrated in Figure 3.

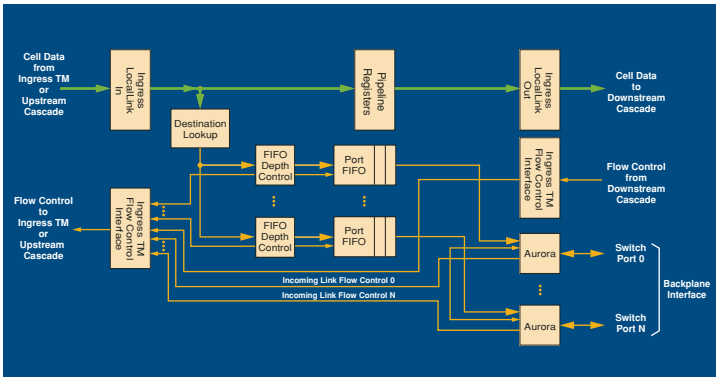

Figure 4A

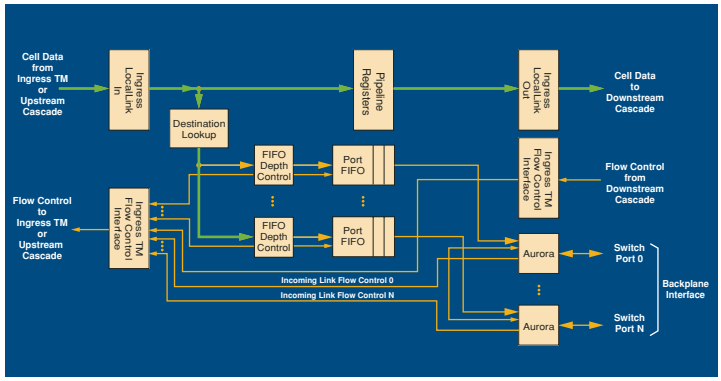

Figure 4B

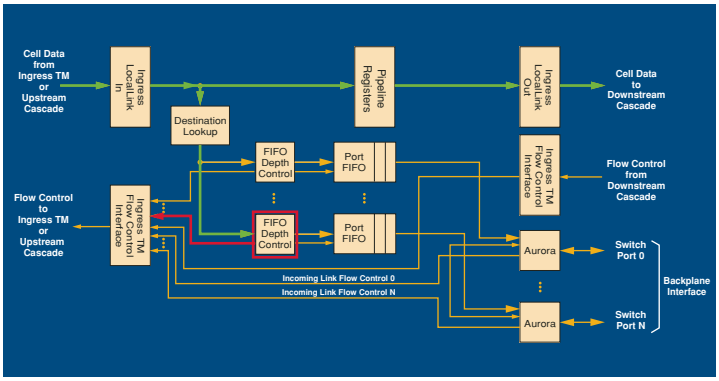

Figure 4C

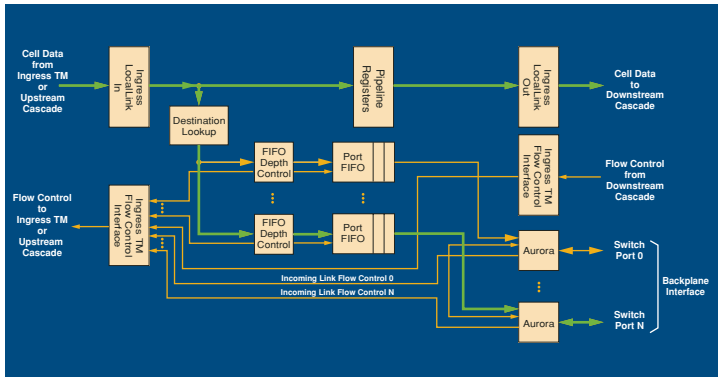

Figure 4D

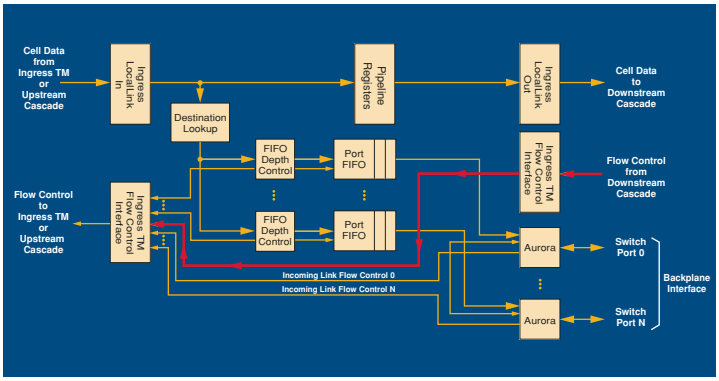

Figure 4E

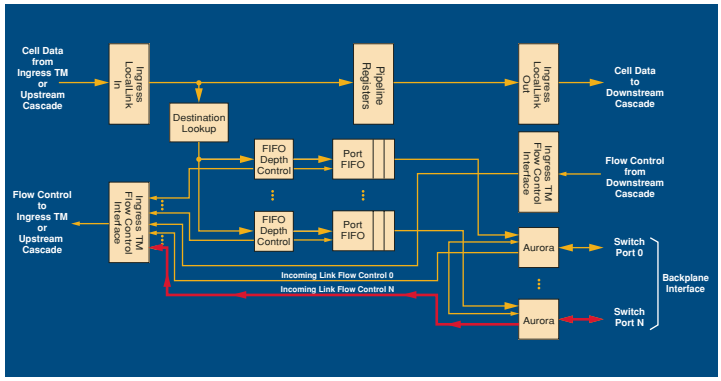

Figure 4F

Figure 4 – MFRD ingress datapath

The switch comprises four basic elements:

- The ingress datapath illustrated in the top half of the diagram
- The switch ports illustrated on the right side
- The egress datapath along the bottom
- The management interface on the left.

Also clearly visible is the use of LocalLink and the DCR bus as the interface standards in the architecture, as well as sideband signaling for flow control status on the cascade interfaces.

Figure 4 illustrates how data flows through the ingress datapath. Dataflow through the MFRD begins at the LocalLink ingress port at the top right side of Figure 4A. Incoming cells are simultaneously vectored to destination lookup and cascaded through the switch to any downstream devices in the configuration. This approach ensures efficient handling of broadcast and multicast traffic which traverse multiple devices.

In Figure 4B, destination lookup forwards the cell to the appropriate port (or multiple ports in the case of multicast or broadcast). On this path we first enter a FIFO depth control block, which is responsible for ingress flow control for this port. If this cell triggers a FIFO event entering the buffer immediately downstream, the logic generates port-specific backpressure to the ingress traffic manager over the cascade interface (Figure 4C). This logic does not exercise flow control. It merely signals the need for flow control as the packet is forwarded to the port, illustrated in Figure 4D.

Figure 4E shows how the cascade interface also aggregates port-specific backpres-

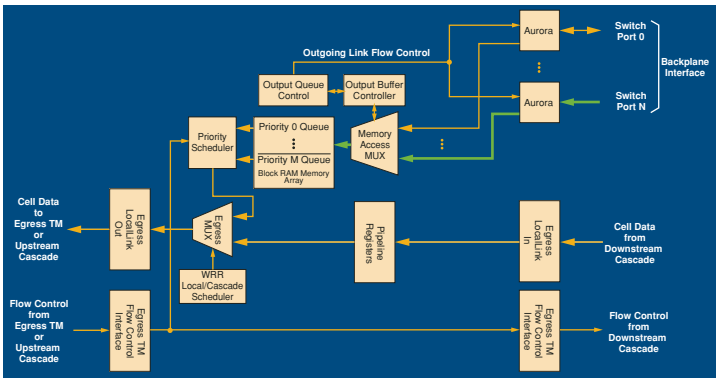

Figure 5A

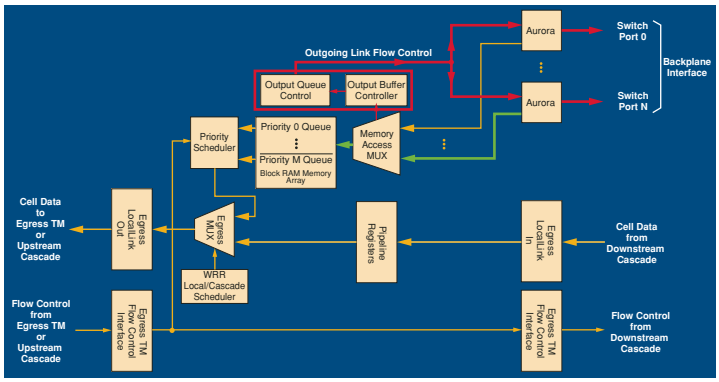

Figure 5B

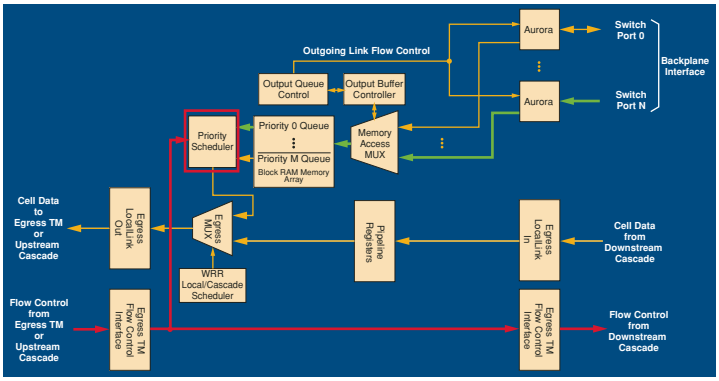

Figure 5C

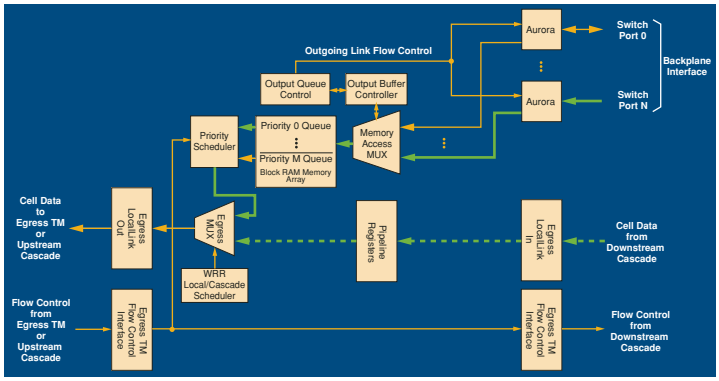

Figure 5D

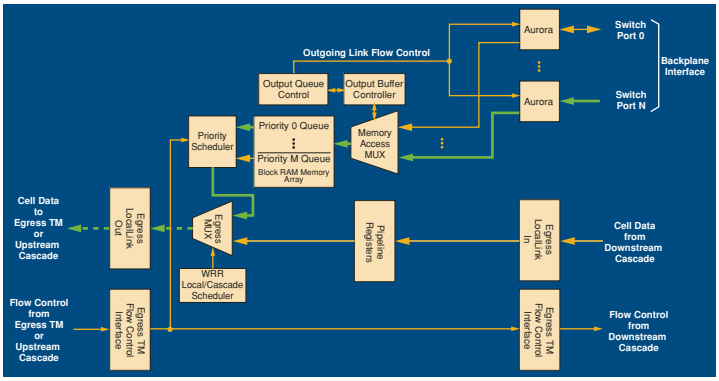

Figure 5E

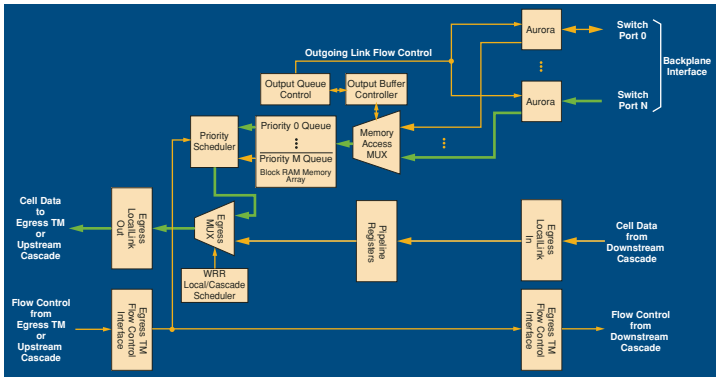

Figure 5F

Figure 5 – MFRD egress datapath

sure from downstream devices in the cascade chain, communicating flow control requirements for all ports to the ingress traffic manager. Figure 4F indicates that the architecture also supports the communication of flow control from the egress side of the switch across the serial links. This mechanism is able to refine backpressure to the ingress traffic manager with priority-specific information per port.

The egress datapath of MFRD is illustrated in Figure 5. Egress begins with the arrival of a cell at the switch port (Figure 5A). Immediately upon arrival, it is fed into a memory access multiplexer that places it into the appropriate priority queue. As shown in Figure 5B, this activity includes the generation of flow control messaging back to *all link partners* on the ingress side of the switch should this action trigger a buffer event in the target queue. This action communicates port- and priority-specific backpressure to all ingress traffic managers.

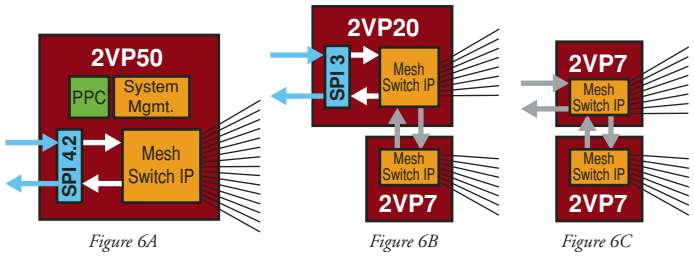

Figure 6A *Figure 6B* *Figure 6C*

Figure 6 – Design flexibility with the MFRD

Egress from the priority queues is controlled by the priority scheduler (Figure 5C). This block can be configured using either a strict priority or weighted round robin scheduling algorithm. The scheduler is tied into backpressure from the egress cascade interface, enabling the egress traffic manager to assert priority-based flow control on the scheduling algorithm. This ensures that the scheduler will not select a priority candidate that the egress traffic manager is not prepared to accept.

Once the scheduler selects a candidate cell for egress, it is forwarded to an egress multiplexer on the egress cascade interface (Figure 5D). This block is also responsible for forwarding traffic from downstream cascade devices and must therefore ensure fair access to egress bandwidth. This is achieved using weighted round robin scheduling through the egress multiplexer. Figures 5E and 5F illustrate how competing traffic is serialized through this mechanism.

Use Models

We have shown that the MFRD enables a great deal of flexibility to optimize the mesh switch implementation when designing your system. To illustrate this, consider the three configurations in Figure 6.

All three configurations support a 16-slot full mesh fabric. Figure 6A shows a fully integrated single-chip mesh fabric controller implementing a 10 Gb SPI4.2 interface to the application logic, a 15-port MFRD configuration, as well as processor IP suitable for implementing blade and even fully distributed shelf system management.

Figure 6B is a reduced-cost configuration of two devices that might be more suitable for supporting a 2.5 Gb SPI3-based application. Figure 6C illustrates a very low-cost solution for applications that would use the LocalLink cascade interface from another FPGA in the Virtex-II™ and Virtex-II Pro families – a very effective way to enhance an existing system architecture.

The Xilinx ATCA Development Platform

To facilitate mesh fabric development, Xilinx has also created a full mesh reference board for ATCA, a serial backplane standard developed by the PCI Industrial Computer Manufacturers Group (PICMG™). The ATCA Development Platform is an ideal prototyping ecosystem for mesh fabric systems (Figure 7).

The ATCA Development Platform features a Virtex-II Pro FPGA with 16 integrated MGTs, 4.2 Mb of block RAM, 53,000 cells of programmable logic, and embedded PowerPC 405 microprocessors. The card is routed as a 1X full mesh and includes IP for instantiating an MFRD demo configuration. IP for instantiating a PowerPC management complex and Linux board support package (BSP) is also available.

Programmable I/O suitable for SPI4.2, CSIX, or other interfaces is routed to personality module headers where you can integrate application-specific designs. The board also provides access to the ATCA update port and a rear transition module should your design require them.

Finally, the board features a Network Equipment Builders Specification (NEBS)-quality, dual feed, ATCA power subsystem delivering 30W to the base board and 170W to the personality module and rear transition module.

Figure 7 – The Xilinx ATCA Development Platform

Conclusion

Switch fabrics are the backbone of modern high-performance system architectures; MGT-based serial communications technology makes the benefits of mesh fabric configurations extremely accessible. With the introduction of the Virtex-II Pro Platform FPGA, Xilinx created a foundation for building such systems entirely with programmable logic. Now, with the availability of the Mesh Fabric Reference Design and ATCA Development Platform, Xilinx is making it even easier to exploit these developments and turbocharge your architectures.

For more information on these topics, please refer to the following resources:

- *www.xilinx.com/esp/networks_telecom/optical/xlnx_net/mfrd.htm*
- *www.xilinx.com/esp/networks_telecom/optical/xlnx_net/atca_dev.htm*
- *www.picmg.org/newinitiative.stm*

So Many Gates

So Few Dollars

The hard work is done! Now ASIC and FPGA designers can prototype logic designs for a fraction of the cost of existing solutions. Here are 6+ million gates (measured the ASIC way) on an easy to use, stand-alone, USB2.0-hosted board (a PCI/PCI-X interface is coming soon). The DN6000k10 supports up to 9, 2vp100 VirtexII-Pro FPGA's, with an incredible amount of FPGA to FPGA interconnect for easy logic partitioning. FPGA's are interconnected with rocket I/O's, enabling the movement of data between them at 100's of GB/s. In addition to 6M+ gates, the DN6000k10 also packs on-board:

- 2 PowerPC cores per FPGA (400MHz)
- Up to 8MB embedded RAM, 444, 18x18 multipliers — per FPGA
- 12 external 133MHz 32M x 16 DDR SDRAM's, 5 4Mx16 FLASH
- 480+ connections for daughter card and logic analyzer interfaces

Configuration is fast, easy, and robust using a SmartMedia-based FLASH card or, via the USB interface. Every tool, utility, driver, and support application that The Dini Group could imagine you might need is included. Please contact us for complete specifications, we are eager to show you how our hard work can make you job easier.

The DiNI Group

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

Programming Flash Memory from FPGAs and CPLDs Using the JTAG Port

A new, inexpensive tool from Ricreations makes it simple and easy to program small data files into Flash memory using Boundary Scan.

by Rick Folea
CTO
Ricreations, Inc.
rfolea@UniversalScan.com

The first prototype of a processor board with Flash memory on it always poses a bit of a problem: How do you get the first chunk of code/boot loader/RTOS into the PROM? You could pre-program the PROM before populating the board, but that assumes the code is ready in time and won't require any changes.

Most designers and lab technicians don't have access to or can't afford the high-end JTAG tools available today. Furthermore, they usually don't want to take the time to build the tests required to do the scan testing and Flash programming anyway. So, what do you do?

We have added a new tool to the popular Universal Scan™ JTAG test suite that makes Flash programming from your Xilinx FPGA or CPLD a snap. You just tell the Universal Scan tool which Xilinx pins are connected to the PROM, select the data file to put in the PROM, and then press PROGRAM.

That's it. What's more, the Universal Scan tool is compatible with your Xilinx parallel port download cable, so you don't even need special hardware to do it.

JTAG Background: How Does it Work?

The I/Os on all Xilinx FPGAs and CPLDs are connected to a giant shift register around the boundary of the device. From the JTAG port, you can shift test vectors into this boundary register using the TDI pin and then apply those vectors to the I/Os, independent of the logic inside the part. In fact, the part doesn't even have to be configured for this to work.

You can also unobtrusively monitor the I/O cells while your device is running by instructing the Boundary Scan chain to capture the state of the I/O cells, and then shift the result out on the TDO pin.

Simply stated, you would follow these steps to program your PROM:

1. Shift a vector into the JTAG chain to setup the address, data, and chip-enables (CEs) and apply it to the pins.
2. Shift the same vector into the JTAG chain to enable the write-enable (WE) signal to the PROM and apply it to the pins.
3. Shift the same vector into the JTAG chain with WE disabled and apply that to the pins.

Repeat these steps a few million times, throw in an occasional command or two to the PROM, and you're done. Sounds easy,

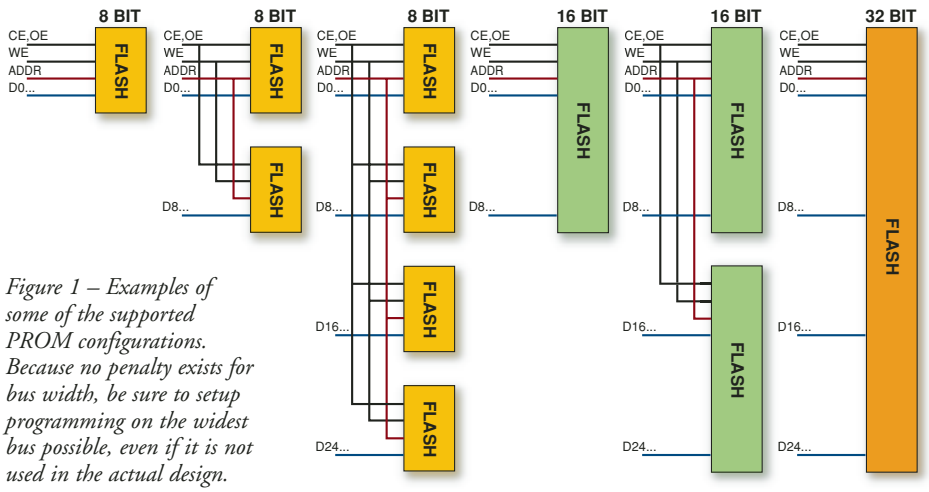

Figure 1 – Examples of some of the supported PROM configurations. Because no penalty exists for bus width, be sure to setup programming on the widest bus possible, even if it is not used in the actual design.

right? Unfortunately, dealing with the low-level details of the JTAG state machine is tedious and difficult.

Fortunately, Universal Scan takes care of these details for you, and knocks the PROM programming effort down to the absolute simplest model possible. You don't need any netlists, test executives, test vectors, special hardware, or anything else normally associated with JTAG test development.

A Simple Solution

Universal Scan supports any bus configuration: 8-, 16-, and 32-bit PROM data buses built from 8-, 16-, or 32-bit PROMs. For example, you can have a 32-bit bus that comprises four 8-bit PROMs in parallel, and Universal Scan will program all four in parallel.

Figure 1 shows example PROM configurations supported by Universal Scan. Universal Scan also supports direct connections between the Xilinx part and the PROM enables, or indirect enables through memory-mapped I/O. (Perhaps your PROM CEs are derived from address lines in a PLD that is not in the JTAG chain, as shown in Figure 2.)

As Figure 2 also illustrates, PROM signals don't have to come from a single device; they can be spread out among any of the devices in the chain.

Limitations and Design Considerations

All this shifting of data around the JTAG chain mentioned previously is very time consuming. To demonstrate this point, let's take a simple example of a Xilinx Spartan-IIE™ device in a 456-pin FBGA package (XC2S300E-FG456). Assume all of the PROM pins are connected directly to this part.

This device has roughly 1,200 Boundary Scan cells in the giant shift register around the boundary of the device. If we take the worst-case scenario of writing one byte at a time to the PROM (no buffered writes), then we need to:

1. Shift the 1,200 bits into the chain to setup data, address, and CEs.
2. Shift the 1,200 bits into the chain to enable the WE signal.
3. Shift the 1,200 bits into the chain to disable the WE signal.

If we use an Intel® algorithm for writing a single byte, the example must be preceded with a command, which doubles the overhead. Thus, a total of 7,200 bits must shift around the JTAG chain just to write one byte/word, as shown in Figure 3. And that doesn't include adding a command to check the results of the operation.

Now, if we assume we have a small 20 KB boot loader we want to put in the Flash memory, then we would need to repeat this 7,200 bit shift operation 20,000 times.

Because you typically get only a few hundred kilohertz bit rate out of a standard parallel port, that little 20 KB chunk of data takes about 10 minutes to program. If you happen to have a larger FPGA or a larger data file, it will take even longer. It all depends on the total length of the JTAG chain and the size of the data file.

Because all of the data is shifted in serially and then applied to a giant latch in parallel, there is no penalty for bus width. An 8-bit data bus programs at the same rate as a 16- or 32-bit bus. So if you are using a PROM that supports an 8- or 16-bit wide data bus as an 8-bit device, go ahead and connect the unused data lines and the BYTE control line to the FPGA. Even though they aren't used in the final design, you can use them to program the PROM through JTAG and cut the programming time in half. This works with 16- and 32-bit devices as well.

Other ways you can minimize programming time include:

- Connecting the PROM to the smallest JTAG device you can (the one with the shortest boundary register)

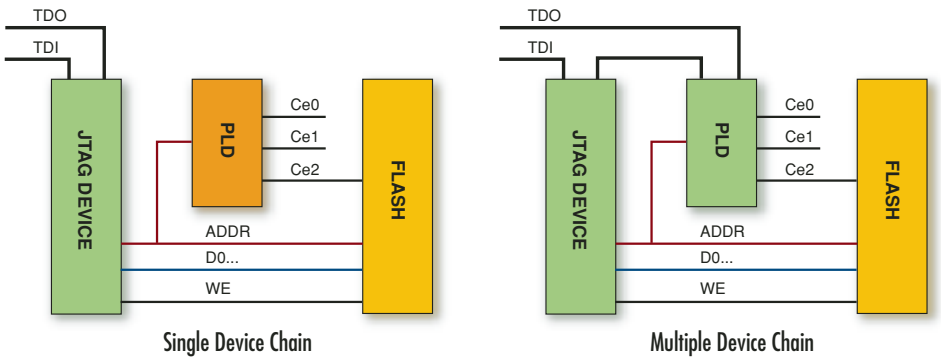

Figure 2 – Although it simplifies things to have all PROM pins connected to a single JTAG part, it is not a requirement. With Universal Scan you can program both memory-mapped PROMs and PROMs with signals from multiple JTAG devices.

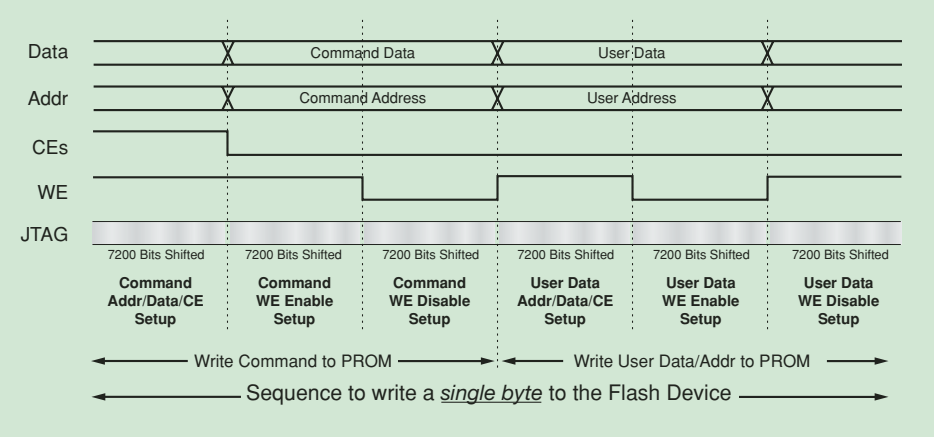

Figure 3 – Programming Flash memories is slow because each and every bus transition requires shifting the data through the entire JTAG chain. This shows the un-optimized example described in the text.

- Putting any JTAG parts not connected to the PROM into BYPASS mode to shorten the overall length of the JTAG chain
- Trying to connect all PROM signals to only one of the devices in the JTAG chain (maximizes the number of devices you can put into BYPASS)
- Choosing a Flash memory that supports buffered writes – these don't require that you write a command before every byte/word write-cycle and nearly doubles the throughput rate of the programming operation.

Also, be sure to connect all PROM signals to the JTAG chain so that you can control every aspect of the PROM's functionality through Boundary Scan – and don't forget to connect the VPEN signal to a pin under JTAG control.

Although this article focuses on Xilinx FPGAs and CPLDs, this method will work exactly the same way with any JTAG-enabled device: processors, DSPs, Ethernet switches, microcontrollers, and others.

If things don't go according to plan, it's easy to debug any issues you might be having with the JTAG chain or Flash programming, as the Flash programmer is part of the Universal Scan JTAG debugging tool. You can use Universal Scan to manually toggle signals between the PROM and your Xilinx device to isolate the issue quickly and efficiently.

Conclusion

The Universal Scan tool enables you to easily program small data files into Flash memory using Boundary Scan. Universal Scan does not replace high-end JTAG tools, which are great if you need to program large data files quickly or in large quantities. But if you want an inexpensive, simple, and flexible JTAG Flash memory programming tool for prototype and general lab development using small data files, then Universal Scan may be the perfect solution.

Universal Scan is available now as a free upgrade to Universal Scan 6.0 users with active registrations. A free, fully functional trial is also available on the Web at *www.UniversalScan.com*. Download it and start programming Flash from your Xilinx devices today.

You'll find more information on this tool on the Xilinx website at *www.xilinx.com*, under Products & Services > System Resources > Configuration Solutions > Automatic Test Equipment (ATE) and Boundary Scan Tools. You can also call your local Xilinx distributor for information or arrange a live demo at your business.

Check out these Xilinx approved suppliers if you are interested in full high-end or high-speed JTAG testing tools:

• JTAG Technologies	*www.jtag.com*
• Acculogic	*www.acculogic.com*
• Corelis	*www.corelis.com*
• Goepel	*www.goepel.com*
• Assett-Intertech	*www.assett-intertech.com*
• Intellitech	*www.intellitech.com*
• Flynn	*www.flynn.com*

For more information about Boundary Scan, please consult these resources:

Intel *(www.intel.com)*

Intel, "Designing for On-Board Programming Using the IEEE 1149.1 (JTAG) Access Port" Application Note #AP-630.

Intel, "Introduction to On-Board Programming with Intel Flash Memory" Application Note #AP-624.

Xilinx *(www.xilinx.com)*

Folea, Rick. "Got the BGA Blues?" *Xcell Journal* – Issue 46, Summer 2003.

Xilinx, "A Quick JTAG ISP Checklist" Application Note XAPP104.

Xilinx, "Using BSDL Files for Spartan-3 FPGAs" Application Note XAPP476.

Xilinx, "Using the XC9500/XL/XV JTAG Boundary Scan Interface" Application Note XAPP069.

Amazon *(www.amazon.com)*

Parker, Kenneth. 2003. *The Boundary-Scan Handbook.* Kluwer Academic Publishers: 3rd edition.

Developing the New Platform Flash PROM

Xilinx and ST Microelectronics have produced a "no compromise" PROM with extensive features at a reasonable cost.

by Anthony Le
Marketing Manager, Configuration Memory Products
Xilinx, Inc.
anthony.le@xilinx.com

Frank Toth
Marketing Manager, EasyPath Products
Xilinx, Inc.
frank.toth@xilinx.com

When Xilinx set out to design a high-performance and dense PROM that could configure a wide variety of FPGAs at a low cost, they needed to think "out of the box" to meet the monetary, device complexity, and schedule requirements. They also had several important requirements for a partner: world-class Flash Memory technology, system-level and silicon design expertise, and worldwide high-volume production.

After evaluating potential partners, Xilinx chose to collaborate with ST Microelectronics™ to design a new series of feature-rich, high-performance Platform Flash configuration PROMs (Table 1).

The Benefits of Partnership

The advantages of designing a complex device with a knowledgeable partner include:

- Reducing schedule timeframes and risk. Platform Flash PROMs use ST Microelectronics' state-of-the-art 0.15 micron flash technology. They were developed by a seasoned team of system-on-chip system and IC designers.
- Taking advantage of the core competencies and experience of each partner. Xilinx brought more than 20 years of FPGA configuration management expertise to the challenge of designing the Platform Flash configuration PROM. The new device includes multiple modes of configuration (serial, JTAG, and parallel) as well as the ability to easily manage multiple bitstreams and unique features like bitstream compression.

	XCF01S	XCF02S	XCF04S	XCF08P	XCF16P	XCF32P
Density	1 Mb	2 Mb	4 Mb	8 Mb	16 Mb	32 Mb
JTAG Prog	•	•	•	•	•	•
Serial Config			•	•	•	
SelectMap Config			•	•	•	
Compression			•	•	•	
VCC (V)	3.3	3.3	3.3	1.8	1.8	1.8
VCCO (V)	1.8 - 3.3	1.8 - 3.3	1.8 - 3.3	1.5 - 3.3	1.5 - 3.3	1.5 - 3.3
Clock (MHz)	33	33	33	40	40	40
Package	VO20	VO20	VO20	FS48 VO48	FS48 VO48	FS48 VO48

Table 1 – Platform Flash PROM family

- Designing and producing the right product for the market. With its extensive worldwide FAE force and design centers, Xilinx worked closely with FPGA users to define a PROM that meets the needs of sophisticated system designers.

  Platform Flash PROMs are cost-competitive products that have both the required baseline features as well as new capabilities to make FPGA systems more attractive and flexible.

- Taking advantage of sophisticated flexible worldwide manufacturing expertise. ST Microelectronics has the manufacturing capacity to produce PROMs in high volumes, along with extensive experience in very small form factor packages.

The Platform Flash PROM offers designers the smallest package board space area per megabit in the industry, such as the VO20 TSSOP (6.4 mm x 6.5 mm) for 1, 2, and 4 Mb density PROMs. Because board space (horizontal as well as vertical spacing) is always at a premium, larger Platform Flash devices with 8, 16, and 32 Mb densities come in small (8 mm x 9 mm) thin flat ball grid array packages.

During the product definition process for Platform Flash PROMs, several multiple chip packaging and stacked die approaches were carefully examined. None appeared to be competitive because the goal was to produce a configuration PROM with the highest reliability, lowest cost, and minimum board space. The result is among the world's most flexible, cost-effective configuration PROMs.

Compression

Higher density Platform Flash PROMs (Figure 1) employ an advanced bitstream compression technology from Xilinx. Compression allows you to store more information in the same memory space, thus reducing cost and board space.

Bitstream file(s) are compressed using Xilinx's ISE design software. The compressed file(s) are then programmed into the Platform Flash PROM, just like any other bitstream file.

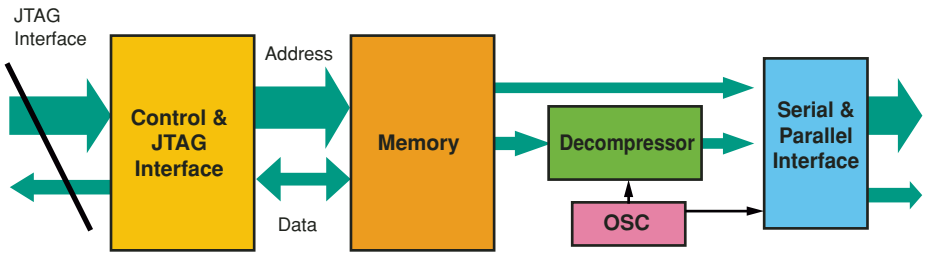

Figure 1 – Platform Flash block diagram

The Platform Flash PROM has a built-in decompressor that automatically senses when a compressed file is stored, and decompresses compressed bitstream information on the fly.

Typically, you can get 50% more bits using this advanced compression technology, fitting, for example, a 48 Mb FPGA design (such as a Virtex-II™ 2VP100 design) into a 32 Mb Platform Flash PROM.

Upgrade Management

You can accomplish in-system programming and upgrades via the JTAG port using the industry-standard four-wire Test Access Port interface (IEEE 1149.1). Platform Flash PROMs are IEEE 1532-compliant, adding to the flexibility and total system integration that allows them to be used with other Xilinx IEEE 1532-compliant devices such as CoolRunner™-II CPLDs and Virtex-II Pro™ FPGAs.

The Platform Flash PROM architecture integrates unique controls that give it the ability to store multiple bitstreams (Figure 2). A microprocessor or other configuration engine can then activate a bitstream at any time, allowing system administrators to access previous versions of system configuration in the event that a problem with a newly transmitted configuration arises.

Safe updates can be achieved with the ability to store multiple bitstreams. Updated bitstreams can be programmed into the free memory blocks of the Platform Flash PROM without losing the original bitstream.

Conclusion

Developed jointly with ST Microelectronics, the Platform Flash family of configuration PROMs offers users the ultimate in low-cost, feature-rich system options, including compression and bitstream upgrade management.

For more information about the new Platform Flash PROMs, visit *www.xilinx.com/xlnx/xil_prodcat_product.jsp?title=PFP.* Σ

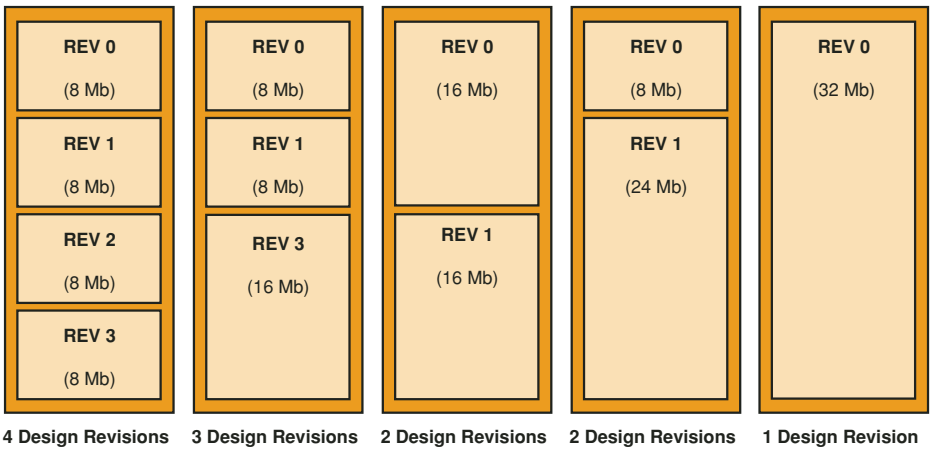

Figure 2 – Platform Flash block design revisions

Accelerate and Verify Algorithms with the XtremeDSP Development Kit-II

The XtremeDSP Kit-II is an ideal development platform for DSP design without VHDL programming.

by Daniel Denning
Research Engineer
Nallatech Inc.
d.denning@nallatech.com

The XtremeDSP™ Development Kit-II, developed by Nallatech in partnership with Xilinx, provides an ideal development platform for high-performance signal processing applications such as software defined radio, networking, HDTV, 3G wireless, and video imagery. The kit provides entry into scalable DIME-II systems from Nallatech.

The combination of the XtremeDSP kit, Simulink™ environment in MATLAB™, and Xilinx System Generator 6.1i software offers a complete design framework for FPGA and DSP designers to get partial or entire systems running on hardware quickly and efficiently. This approach is redefining time to market by rapidly producing high-performance systems with design flexibility, as well as the possibility of reconfiguring and upgrading the system without changing the physical hardware.

Features

The XtremeDSP kit includes an on-board user-programmable Xilinx XC2V3000 FPGA, two 14-bit ADC channels with as many as 65 mega samples per second (MSPS) per channel, and two 14-bit DACs with as many as 160 MSPS per channel. A Spartan-II™ FPGA is pre-configured with 32 bit/33 MHz PCI or USB 1.1 firmware. One bank of 1 Mb ZBT SRAM, configured as 512K x 16, is also available. An external power supply allows you to power the kit on its own; JTAG configuration headers and status LEDs provide feedback. Figure 1 shows the XtremeDSP kit. Figure 2 shows a block diagram of the kit.

The XtremeDSP kit comes with Nallatech's Field Upgradeable Systems Environment (FUSE) FPGA management software. This software provides the ability to control and configure the FPGA and transfer data between the kit and the host PC thorough a GUI or C-based API. Additional options allow the use of Java or MATLAB M-code script control.

Figure 1 – XtremeDSP (DIME-II) kit hardware

Designing with System Generator

The first step in creating a System Generator model is as follows:

1. Open a Simulink workspace and drop a System Generator token in at the top level of the model.
2. Open up the token. This allows you to select various options, such as FPGA device, package, system clock, location of the generated VHDL, and type of synthesis tool required. You can also drop the token into a subsection of the System Generator model.

This functionality allows you to construct a system quickly by placing and connecting traditional System Generator blocks. You do not have to write any VHDL, although a black box provides this capability if required.

Other user-friendly features include the ability to drop Xilinx efficient handcrafted IP blocks into the model, such as fast Fourier transforms (FFTs), multipliers, direct digital synthesizers (DDSs), linear feedback shift registers (LFSRs), and finite impulse response (FIR) filters.

Each block within the model has associated options that allow for customization. For example, Figure 3 depicts the different options available for block RAM placement, such as size, initial values, writing options, and distribution of memory.

Grouping various System Generator blocks together generates subsystems, creating a hierarchy within the model. Each subsystem has its own associated input and output ports. Placing low-level blocks into the Simulink environment creates a bottom-up design approach, which allows functional verification of each subsection before it is included in the system model.

During the design process, the Simulink environment allows you to test and verify the model, providing an array of test facilities for this purpose, including scopes, graphs, and displays. For further pre- and post-processing and verification, you can import to and extract the data from the MATLAB environment. When extracting data, MATLAB functions offer extensive post-processing options, including three-dimensional visualization graphs, plotting images, and a vast collection of computation algorithms.

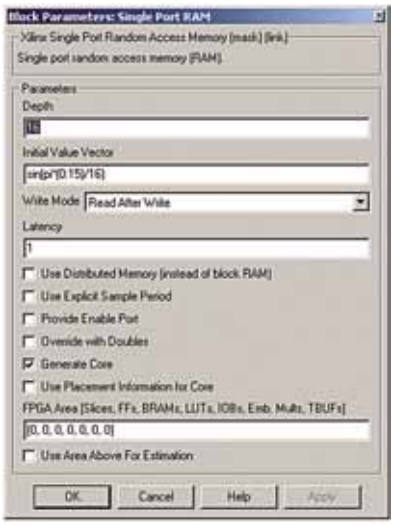

Figure 3 – Example of System Generator block options with block RAMs

Hardware Co-Simulation on the XtremeDSP Kit-II

Verification of the software model means that you can then test and verify the model in hardware by creating a hardware co-simulation block, executed by the System Generator token in the model, which controls the design flow. Select the compilation target for hardware co-simulation in this block and choose the XtremeDSP kit as the hardware target. This will generate an equivalent hardware Simulink co-simulation library block.

This block is effectively an FPGA bitstream, the result of a synthesis tool such as

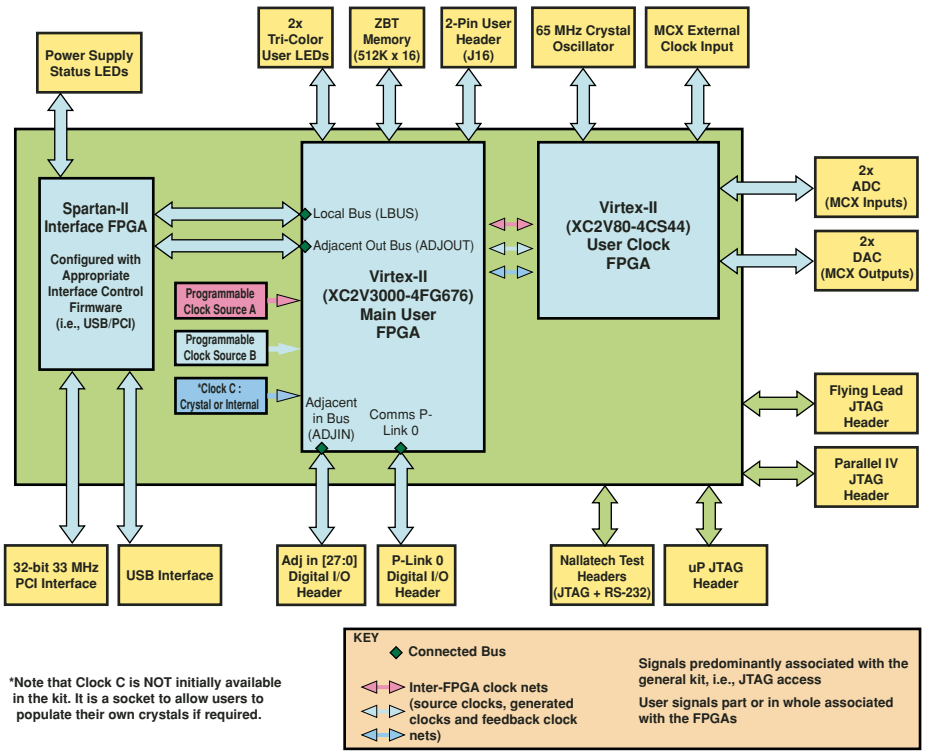

Figure 2 – Functional block diagram of the XtremeDSP kit

XST (Xilinx Synthesis Tool). The block takes care of board operations such as device configuration, data transfers, and clocking. Each port in the software model maps to the relevant hardware co-simulation library block.

The co-simulation library block now offers options for hardware co-simulation with the DIME-II XtremeDSP Kit-II. When Simulink simulates the model, it takes the results from the FPGA on the XtremeDSP kit. You can treat the library block exactly the same as any other in the Simulink library.

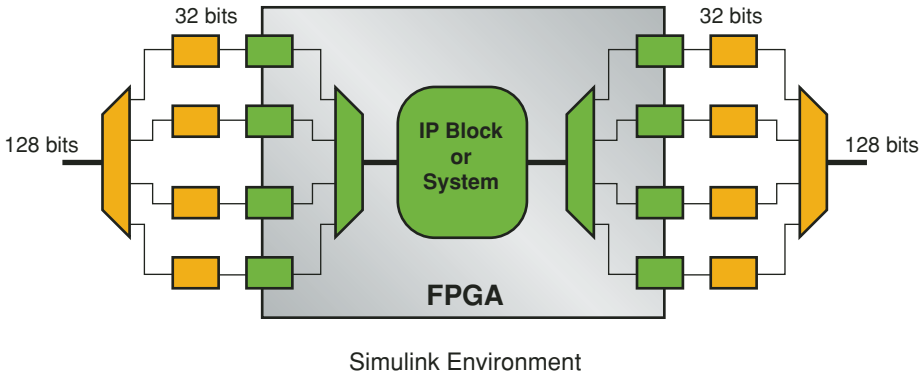

Figure 4 – Hardware co-simulating with designs greater than 32 bits

TCP/IP Hardware Co-Simulation

Having produced a hardware co-simulation library block, your next step is to select which bus configuration over which co-simulation will take place. The XtremeDSP kit offers the following standard co-simulation options: PCI and JTAG. An enhanced option is available from Nallatech to add support for TCP/IP, which allows you to share the XtremeDSP kit for hardware co-simulation through another workstation.

As the XtremeDSP kit is a derivative of Nallatech's DIME-II BenONE motherboard and BenADDA module, this combination of hardware makes the connection to an Ethernet module possible. This provides the capability to power the board on its own without the need for additional workstations and their associated licenses. With this module interface connection on the BenONE motherboard, the board can now be run on its own anywhere – providing that it has a TCP/IP location.

Therefore, if a design team wants to efficiently utilize the FPGA development board, the board no longer needs to be swapped from workstation to workstation, as each designer can change their location, providing that the appropriate software licenses are available.

For an extreme proof-of-concept example, take a transatlantic hardware co-simulation of the AES-128 (Advanced Encryption Standard). The board runs on its own with a loosely coupled connection to the Internet, located in a laboratory in San Jose, California, while the design of the AES-128 encryption core in System Generator is completed in the United Kingdom. Simulation speeds were increased by around a factor of two, but more importantly, the FPGA engineer did not need a board for functional hardware verification to occur.

This capability provides significant benefits when working with shared hardware in both the development and debug phases of a product's lifecycle. To have the ability to connect to a remote piece of hardware and carry out co-simulation aids the development of products within a team environment, where independent and geographically scattered team members need access to a single physical piece of hardware at a particular site. A principle example of this could be development and support work for base station components.

Hardware Co-Simulation Greater than 32 Bits

When simulating designs in System Generator without an FPGA co-simulation, you can simulate any number of binary widths by using the library blocks in System Generator. This is not currently possible when converting the model for co-simulation on the FPGA.

Once the I/O bit widths become greater than 32 bits, System Generator will indicate that an error has occurred. To avoid such an error, split the bit widths down to 32 bits so that each I/O port on the co-simulation library block has a width of 32 bits.

Incorporating a part-System Generator, part-FPGA wrapper keeps the same bit widths within the model and the FPGA. Slicing the data path on entry to the co-simulation block allows concatenation back into the FPGA for the IP core or system. This data path will split again before leaving the FPGA, finally concatenating back into the Simulink model.

Figure 4 illustrates this technique for co-simulating bit widths of 128 for the AES-128. The gray section in Figure 4 represents the co-simulated area on the FPGA achieved during the simulation process. The area outside the gray region shows the fixed order in which the bit widths appear in the System Generator model. This method applies to any number of bit widths: for example, 64, 128, 256, and 77.

Conclusion

The XtremeDSP Kit-II offers the ideal development platform for DSP designers. Combining the kit with System Generator provides the capability to design systems without the need for programming in VHDL, although the option exists to add VHDL and even MATLAB M-code for HDL auto generation.

By including a TCP/IP interface, unless there is a need to see or connect inputs to the board, this type of connection is a viable option in hardware co-simulation and offers cost-effective, efficient utilization and time-saving benefits.

To learn more about DIME-II, please visit *www.nallatech.com/solutions/products/embedded_systems/dime2/index.asp*, or e-mail *contact@nallatech.com*. For more information about the XtremeDSP Development Kit-II, please visit *www.xilinx.com/ipcenter/dsp/development_kit.htm*.

Increase Image Processing System Performance with FPGAs

Using FPGAs instead of DSPs to perform common image-processing functions can offer a wide range of benefits.

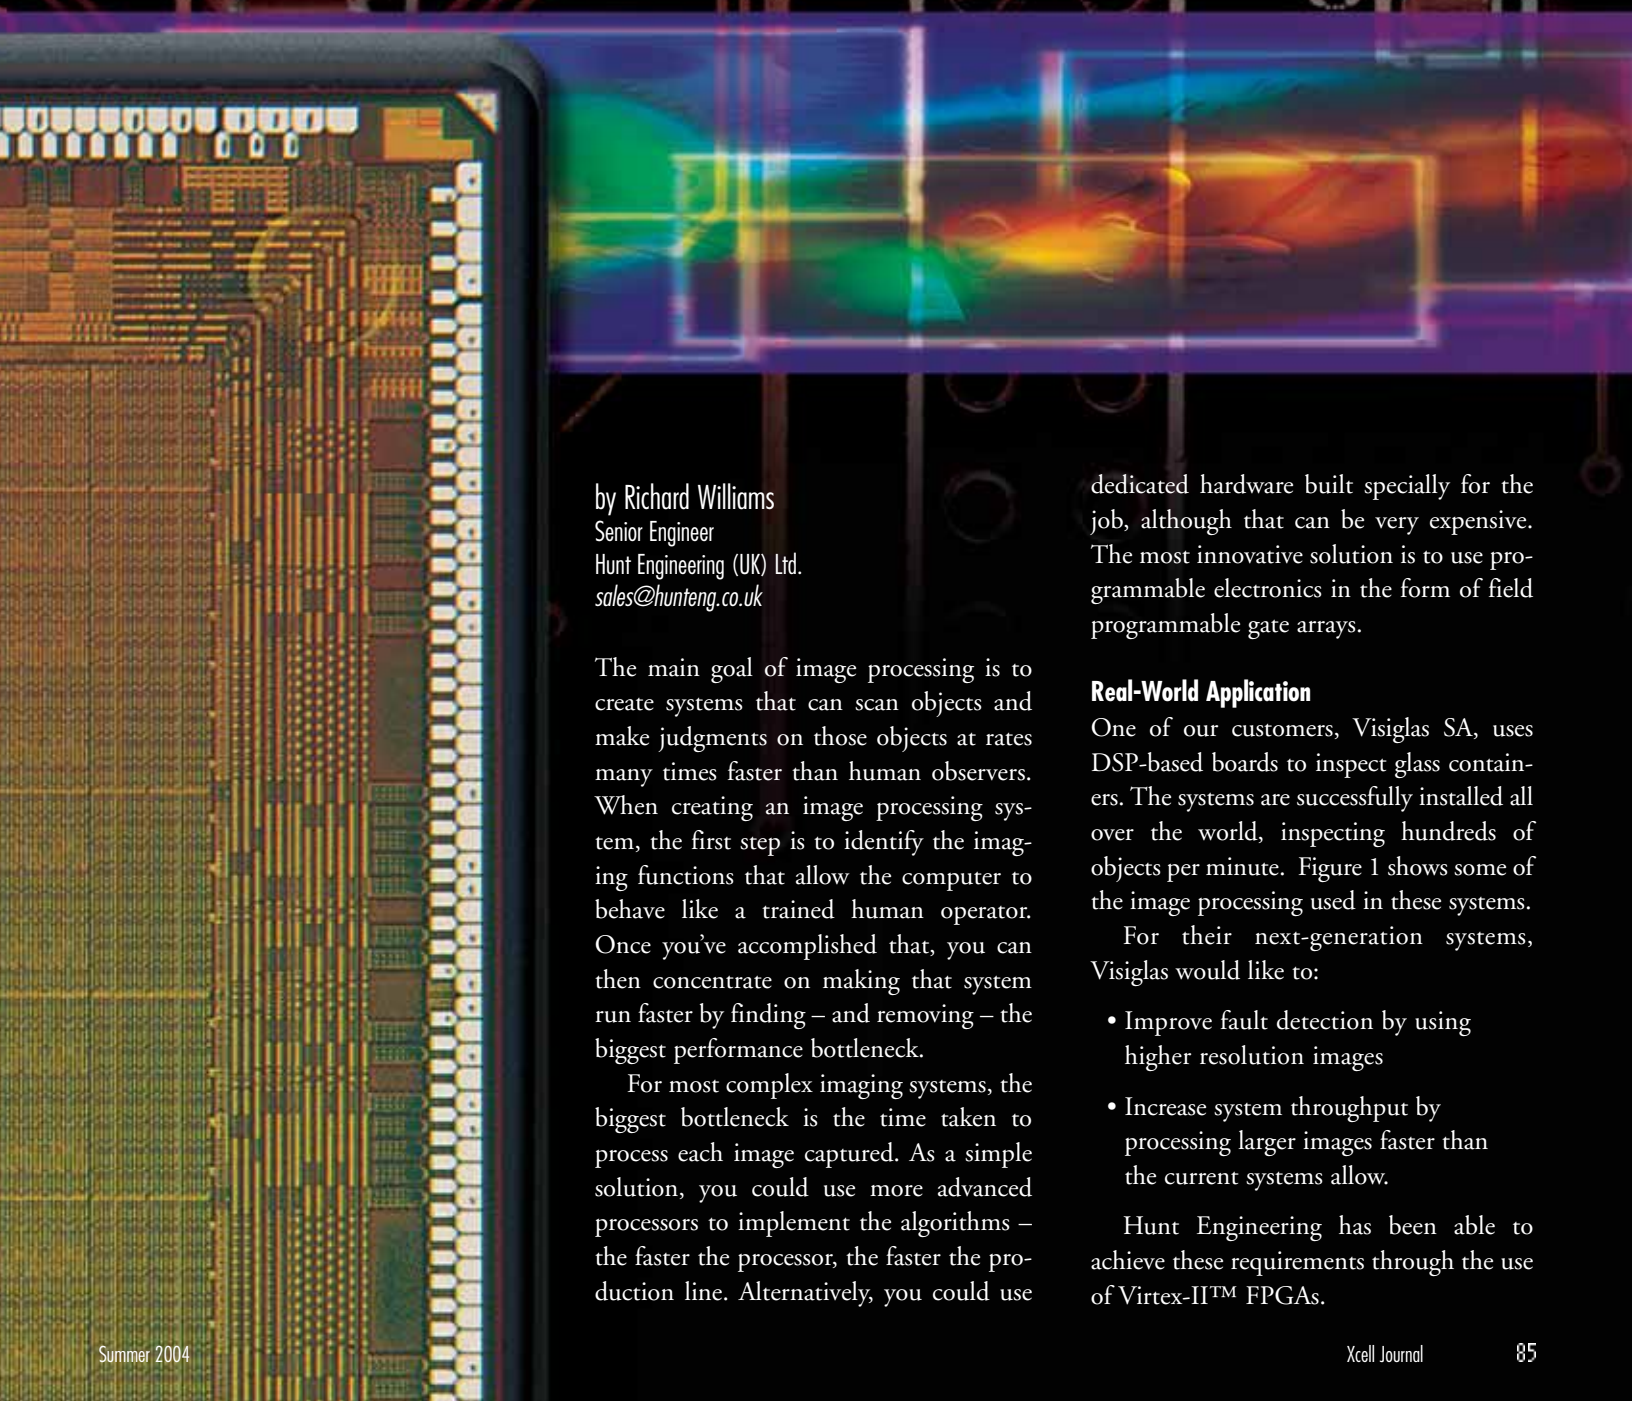

by Richard Williams
Senior Engineer
Hunt Engineering (UK) Ltd.
sales@hunteng.co.uk

The main goal of image processing is to create systems that can scan objects and make judgments on those objects at rates many times faster than human observers. When creating an image processing system, the first step is to identify the imaging functions that allow the computer to behave like a trained human operator. Once you've accomplished that, you can then concentrate on making that system run faster by finding – and removing – the biggest performance bottleneck.

For most complex imaging systems, the biggest bottleneck is the time taken to process each image captured. As a simple solution, you could use more advanced processors to implement the algorithms – the faster the processor, the faster the production line. Alternatively, you could use dedicated hardware built specially for the job, although that can be very expensive. The most innovative solution is to use programmable electronics in the form of field programmable gate arrays.

Real-World Application

One of our customers, Visiglas SA, uses DSP-based boards to inspect glass containers. The systems are successfully installed all over the world, inspecting hundreds of objects per minute. Figure 1 shows some of the image processing used in these systems.

For their next-generation systems, Visiglas would like to:

- Improve fault detection by using higher resolution images
- Increase system throughput by processing larger images faster than the current systems allow.

Hunt Engineering has been able to achieve these requirements through the use of Virtex-II™ FPGAs.

The Mathematics of Image Processing

Image processing typically involves applying the same repetitive function to each pixel in the image to create a new output image. We can categorize the techniques involved into three types:

1. Where one fixed-coefficient operation is performed identically on each pixel in the image.
2. Where there are two input images rather than one. In this type of operation, the mathematics performed may be the same as for the fixed coefficients, but now the operation is based on the position of the pixel in the image.
3. Neighborhood processing, or convolution. There is only one input frame, and the result created for each pixel location is related to a window of pixels centered at that location.

So although the exact mathematical operation may vary, all three techniques require repetitive functions to be performed across the entire image. Thus, this kind of processing is ideally suited to a hardware pipeline that can perform fixed mathematical operations over and over on a stream of data.

DSPs versus FPGAs

DSPs typically must execute several instructions to perform an image processing function. Because it is a sequential device, these instructions will probably take several processor clock cycles to complete. Add to that the cycles needed to fetch the image data, store the results, and handle interrupts, and you have a large number of clock cycles needed to process each pixel.

Because the majority of image processing can be broken down into highly repetitive tasks, FPGAs present a very interesting alternative to DSPs. Additionally, you can use FPGAs to perform lots of steps in parallel, using dedicated logic for each step.

Through the use of Virtex-II FPGAs, we can implement image-processing tasks at very high data rates, reaching hundreds of megahertz. These functions can be directly performed on a stream of camera data as it arrives without introducing extra processing delays – significantly reducing and sometimes removing performance bottlenecks.

In particular, you can map more complex functions such as convolution very successfully to FPGAs. When convolving an image, a window of pixels is treated with a mask, where individual locations in the window are "weighted" according to a set of previously defined coefficients. For each position of the window, all pixels are multiplied against their respective coefficients. The final result is then scaled to produce a single output pixel for the center location of the window.

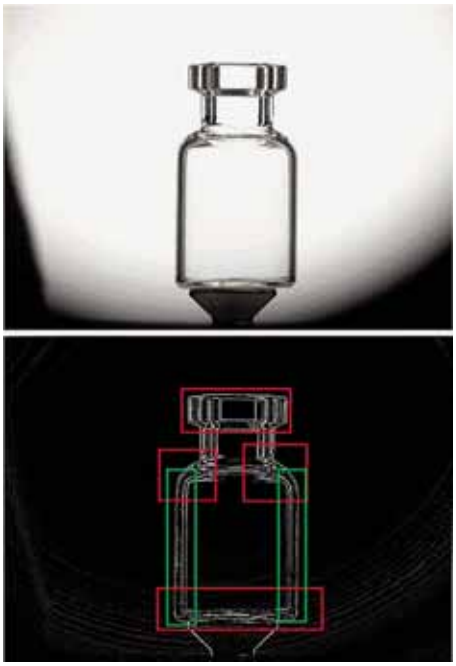

Figure 1 – A real-world application of an image processing system where FPGA processing can significantly increase performance

In essence, the whole convolution process is a matrix-multiplication, and as such requires several multiplications for each pixel. The exact number of multipliers required is dependent on the size of window used for convolution. For example, a 3x3 kernel (window) requires nine multipliers; a 5x5 kernel requires 25 multipliers.

Conventional DSPs have a fixed number of multiplication units inside the processor core – fewer multiplier units than what are needed to perform the matrix multiplication in one step. Thus, a DSP would introduce a performance drop by reusing multiplier units to complete the matrix multiplication.

FPGAs, however, can implement as many multipliers as necessary to calculate one pixel at the full input data rate, whether the convolution uses a 3x3 kernel or a larger 5x5. With the one-million-gate Virtex-II, 40 multipliers are available; in the eight-million-gate version, this number increases to 168. By mapping convolution to FPGAs that already provide dedicated multipliers among their sea of gates, it becomes easy to build a processing pipeline that can convolve at very high data rates.

A Role for the DSP?

Although a large proportion of image processing algorithms are simply highly repetitive processes, there is still a role for the DSP. In a system that can benefit from the performance advantages of FPGAs, there is a point in the data flow where a decision has to be made. This decision will often take the form of "if, then, else" logic rather than a pixel-by-pixel iteration.

For control loops and complex branches in operation, DSPs can still prove to be highly effective. Implementing equivalent logic in FPGAs can quickly eat up the available gates and reduce the overall data rate.

A simple solution is to use both types of resources in a single system: a high-data-rate FPGA as the data-reducing engine, feeding results downstream to a DSP as the accept/reject, pass/fail decision maker.

Image Acquisition and Processing with HERON

The HERON module range from Hunt Engineering provides a flexible, high-performance solution to image processing. HERON-FPGA modules, which include the Virtex-II series of FPGAs, present resource nodes that are suited to a wide range of tasks, particularly the repetitive tasks of image processing.

These FPGA modules can also be directly connected to cameras, accepting data in formats such as Camera Link and RS422. Combine that with HERON processor modules based around Texas Instruments'™ TMS320C6000 DSP series, and a complete imaging solution becomes possible.

In addition to the hardware resources required at the heart of the system, firmware and software are also necessary to implement the appropriate algorithms in the FPGA and DSP. Hunt Engineering offers imaging libraries for both DSPs (in C) and FPGAs (in VHDL), downloadable from *www.hunteng.co.uk*. These libraries enable you to quickly and easily assemble the key algorithm components into a working imaging system.

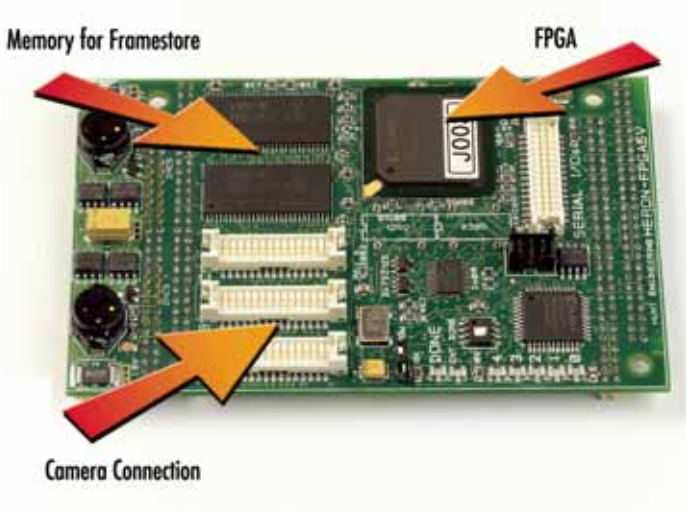

Figure 2 – Hunt Engineering's HERON-FPGA5 module with a Virtex-II FPGA, 256 Mb of SDRAM, and digital I/Os for connecting cameras

Memory Requirements

If you use the multi-frame operations provided in our VHDL imaging libraries (such as the addition of two images), you must have an area of available memory that can store an entire frame.

Unless the size of one frame is very small, the FPGA's internal RAM resources will be insufficient for this type of operation. In this situation, you could use a module like the HERON-FPGA5 (Figure 2). The reference image is stored in SDRAM external to the FPGA and read into the FPGA as required.

Because separate dedicated logic is used to receive the incoming image, access the stored image, and perform the processing, image processing can still be performed at pixel rates greater than 100 megapixels/sec. With a processor-based approach, the processor has to access both images from memory, and these operations will be slower than pixel-based operations when using a DSP.

Neighborhood processing, on the other hand, requires several lines of image data to be stored before processing can begin. The image size determines the amount of storage required per line, and the kernel size of the operation determines the number of lines. It's possible to use the FPGA's internal block RAM for this storage, but the amount available depends on the size of the FPGA and the design requirements.

For example, a one-million-gate Virtex-II FPGA has 90 Kb of block RAM. If nothing else in the design requires block RAM, then the convolution can use all 90 Kb. With 8-bit monochrome data, you can store 90 Kpixels. If the image is 2K pixels per line, then 45 lines of data is more than enough for a large convolution function.

If the FPGA design uses block RAM for other functions, using hardware like the HERON-FPGA5 enables you to store the image in off-chip SDRAM.

Conclusion

Many key imaging functions break down into highly repetitive tasks that are well suited to modern FPGAs, especially those with built-in hardware multipliers and on-chip RAM. The remaining tasks require hardware more suited to control flows and decision making, such as DSPs.

To gain the full benefit of both approaches, systems can effectively combine both FPGAs and DSPs. With the addition of standard imaging functions written in either VHDL or C, all of the key building blocks are available to create an image processing system. Hunt Engineering has developed a demonstration framework of such a system, shown in Figure 3.

For our customer Visiglas, a system such as the one shown in Figure 4 allows them to achieve their performance goals.

The next logical step is an addition to the HERON module range of devices for Virtex-II Pro™ FPGAs. With a PowerPC™ processor core, a sea of gates, built-in multipliers, and on-chip RAM, a self-contained high-performance imaging solution becomes possible in a single chip.

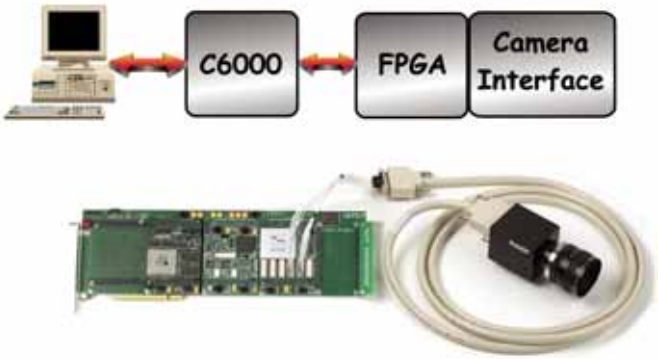

Figure 3 – Using HERON to combine a DSP and an FPGA for image processing

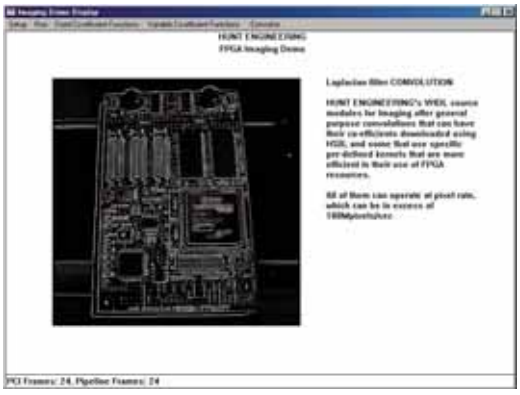

Figure 4 – The imaging framework provided by Hunt Engineering to demonstrate FPGA image processing at frame rate

Enabling Low-Cost DSP Co-Processing with Spartan-3 FPGAs

Embedding high-performance DSP functions within FPGA fabric is now a genuine low-cost option.

by Suhel Dhanani
Senior Solutions Marketing Manager
Xilinx, Inc.
suhel.dhanani@xilinx.com

Steve Zack
Signal Processing Engineer
Xilinx, Inc.
steve.zack@xilinx.com

FPGAs have been used in DSP applications for years as logic aggregators, bus bridges, and peripherals. More recently, FPGAs have gained considerable traction in high-performance DSP applications and have also emerged as ideal co-processors for standard DSP devices.

In these latter roles, FPGAs provide tremendous computational throughput by using highly parallel architectures. Because the hardware is re-configurable, you can develop customized architectures for ideal implementation of your algorithms.

The new generation of Spartan-3™ low-cost FPGAs, developed using 90 nm process technology, not only creates an effective way to implement high-performance DSP functions but provides an even more economical solution. Their low cost means that you can use them to implement high-performance DSP co-processing functions in conjunction with a conventional DSP device – typically integrating pre- and post-processing functions in a cost-effective manner.

Key Advantages

FPGA architectures are well suited for highly parallel implementations of DSP functions, allowing for very high performance. And user programmability allows you to trade off device area versus performance by selecting the appropriate level of parallelism to implement your functions.

FPGAs are essentially arrays of uncommitted logic and signal processing resources. These signal processing resources allow you to implement DSP functions using highly scalable, parallel processing techniques.

For example, whereas a traditional DSP solution would implement multiple multiply accumulate (MAC) functions in a serial manner, an FPGA allows you to implement these in parallel using dedicated multipliers and registers that are now available in the Spartan-3 family.

As another example, consider a 256-tap finite impulse response (FIR) filter. By using resources available in the FPGA fabric, you can design a highly parallel implementation and achieve higher performance (Figure 1).

Because FPGAs are completely hardware-configurable, you have the flexibility to only use the necessary resources that the algorithm demands.

Figure 2 shows the different ways of implementing four MAC functions. By using four embedded multipliers within the FPGA fabric, you can complete these implementations at maximum speed. Alternatively, you can opt to conserve area and implement the same function at a lower performance by using only one multiplier, one accumulator, and a register, or use the semi-parallel approach.

Although FPGAs bring significant benefits to DSP, it is important to analyze the effective cost of implementing DSP functions within the FPGA fabric. For the purpose of this analysis, the new Spartan-3 FPGA family is considered because of its low cost and system features for DSP.

Spartan-3 Devices: Optimized for DSP

Spartan-3 FPGAs use 90 nm manufacturing technology to achieve low silicon die costs. These devices are also the only low-cost FPGAs that have all of the features required for efficiently implementing DSP functions – features that were once the exclusive domain of high-end FPGAs (Table 1).

With the Spartan-3 family, you can implement high-performance, complex DSP functions in a small portion of the total device, leaving the rest of the device free to implement system logic or interfac-

Spartan-3 Silicon Features	Customer Benefits
Embedded 18 x 18 Multipliers	Area-efficient implementation of multiply function
Distributed RAM	Local storage for DSP coefficients, small FIFOs
Shift Register Logic	16-bit shift register ideal for capturing high-speed or burst mode data and to store data in DSP applications
Up to 104 18 Kb Block RAM	Video line buffers, cache tag memory, scratch-pad memory, packet buffers, large FIFOs

Table 1 – These Spartan-3 features enable DSP functions in an area-efficient manner.

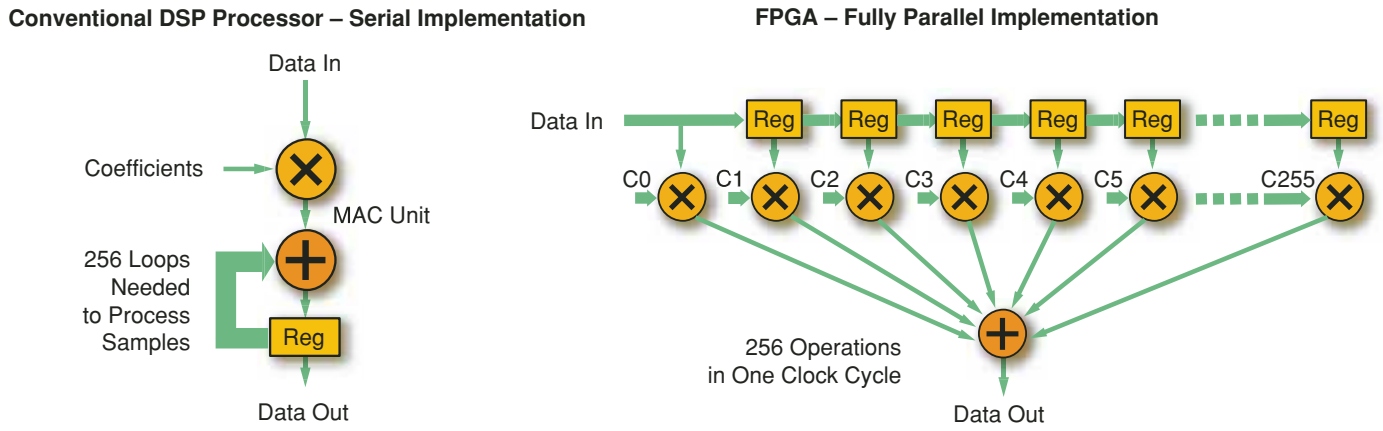

Figure 1 – An FPGA's parallel approach to DSP enables higher computational throughput.

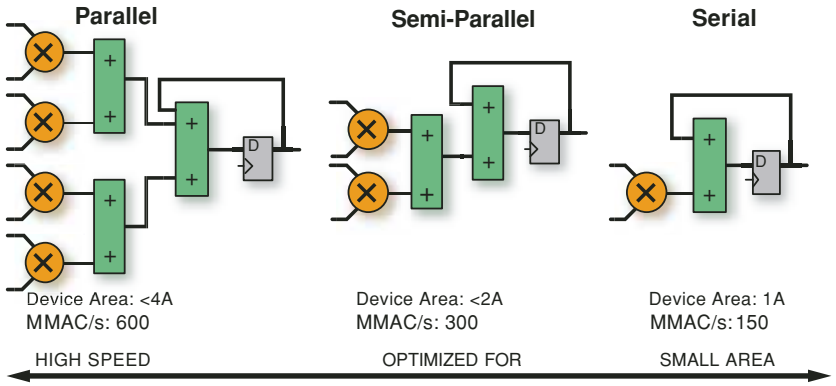

Figure 2 – You can customize an FPGA to suit your needs.

ing functions – providing both lower costs and higher system integration.

Table 2 demonstrates how the combination of advanced features and low cost work together to provide DSP capability at a low cost. The table shows a sampling of available Spartan-3 parts, the number of million multiply accumulate per second (MMAC/s), and the cost for MMAC/s in each device.

We calculated the MMAC/s column by multiplying the number of multipliers with their operating frequency, which for Spartan-3 FPGAs is 150 MHz in the slowest speed grade.

Then, looking at the published 50,000-unit price for the slowest speed grade of the appropriate device, we calculated the cost for MMAC/s. This is one of the quoted industry benchmarks, with the cost per MMAC/s reaching a quarter of a cent.

How to Achieve the Lowest DSP Function Cost

No standard currently exists to estimate the actual cost of implementing DSP functions onto FPGAs. For the purposes of this analysis, however, let's theorize that the effective cost is the cost based on percentage of silicon area utilized, multiplied by the unit device cost. This is a fair calculation, since the remainder of the FPGA is available for other system functions.

Device	Embedded Mults (18 x 18)	MMAC/second (Number of Mults x 150 MHz)	Cost for MMAC/s
XC3S50	4	600	$0.0055
XC3S200	12	1,800	$0.0024
XC3S400	16	2,400	$0.0030
XC3S1000	24	3,600	$0.0037
XC3S1500	32	4,800	$0.0044

Table 2 – Calculating the cost per MMAC/s

To calculate the effective cost of a DSP function when implemented in an FPGA, we considered the Spartan-3 XC3S1000 device, which is a mid-range member of the Spartan-3 family. In many cases, a given DSP function uses not only the FPGA logic but also embedded multipliers and block RAMs. In that case, we included the estimated amount of die space taken by these embedded functions and added that to the die area used by the logic.

Table 3 shows some of these functions and the cost of implementing these within the Spartan-3 silicon. (We have not included the cost for programming the PROM, because in many cases you can use the existing EPROM on-board to program the FPGA.)

Some of the most common functions used in DSP applications are fast Fourier transforms (FFTs) and FIR filters. A single channel 64-tap MAC FIR filter running at 8.1 mega samples per second (MSPS) can be implemented for an effective cost of $0.41. Note that this filter uses 200 logic slices and four embedded multipliers – approximately 3% of the die area.

You can also implement simple forward error correction DSP cores such as Viterbi and Reed Solomon functions at a low cost within the Spartan-3 device. A 32-channel, parallel mode Viterbi decoder running at 1.9 MSPS per channel has an effective cost of $5.06, or $0.16 per channel. A Reed Solomon G.709 decoder function running at 60 MHz

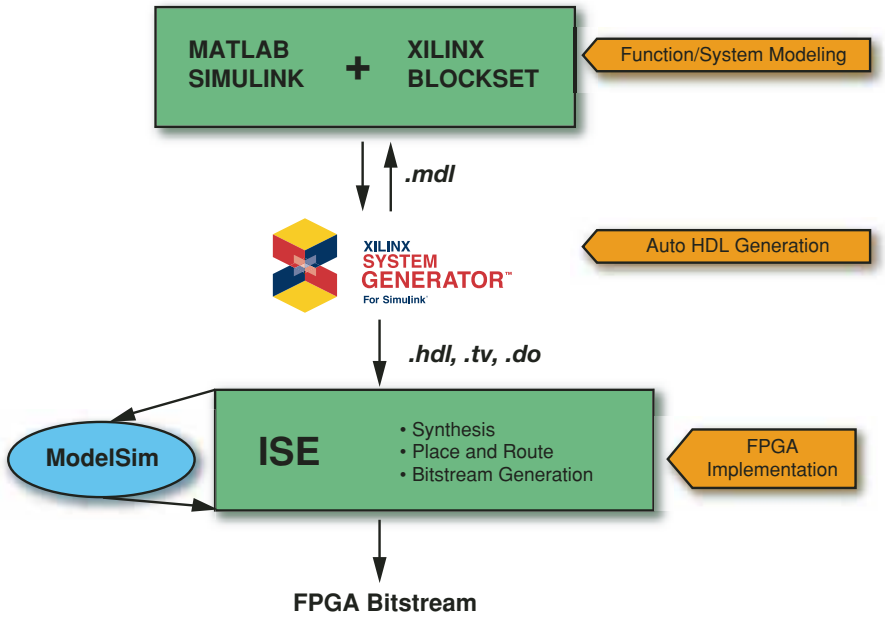

Figure 3 – A DSP design methodology that works within an existing tool flow

Functions	% of the XC3S1000 Device Utilized	Effective Cost (50K Units)	Key Specification	Other Specifications
1024-point complex FFT	24.1%	$3.23	20 µs transform	20 µs transform, burst I/O, 16-bit input and phase factor
Single channel 64-tap FIR filter	3.0%	$0.41	8.1 MSPS	16-bit data and co-efficient, MAC implementation, 8.1 MSPS
Digital down converter per channel	18.6%	$2.49	Sample rate 100 MSPS	
Digital up converter per channel	18.6%	$2.49	Sample rate 100 MSPS	
Viterbi decoder	37.8%	$5.06	1.9 MSPS per channel	Parallel mode, trace-back 42, constraint length = 7, 32-channel, 1.9 MSPS per channel
Reed Solomon G.709 encoder	1.3%	$0.17	120 MHz	
Reed Solomon G.709 decoder	6.9%	$0.92	60 MHz	

Table 3 – Effective costs of various DSP functions in a Spartan-3 device

takes only 6.9% of the same device (with an effective cost of $0.92).

Complex functions such as a digital down converter (DDC) or a digital up converter (DUC) – commonly used in wireless base stations – take less than 20% of the Spartan-3 XC3S1000 device (with an effective cost of $2.49).

Development Tool Flow

With Xilinx, you can use industry standard development tools for your DSP designs. Using MATLAB™ and Simulink™ from The MathWorks, coupled with Xilinx System Generator for DSP, you can now model, simulate, and verify your signal processing algorithms on your target hardware platform without leaving the Simulink environment.

The design flow typically involves the following steps:

1. A DSP designer develops and verifies the hardware model using industry-standard tools from The MathWorks in conjunction with Xilinx System Generator for DSP.

2. With a push of a button, Xilinx System Generator generates an HDL circuit representation that is bit- and cycle-true, meaning that the behavior is guaranteed to match the functionality seen in the Simulink/System Generator model.

3. The ISE design tools synthesize the design and produce a bitstream that can be used to program the FPGA.

The error-prone and time-consuming step of having an FPGA designer translate the system engineer's design into HDL is thus eliminated. Figure 3 shows a typical design flow using the Xilinx System Generator. With recent advances in this product, DSP designers can now generate an FPGA bitstream directly using Simulink/System Generator.

Conclusion

With its combination of low unit cost and architecture optimized for DSP functions, Spartan-3 FPGAs have the industry's lowest price points for high-performance DSP functions. Xilinx further enables embedded DSP functions by providing design tools that fit within your tool flow and enhance your productivity by automating the FPGA implementation process.

With the availability of Spartan-3 devices, associated design tools, and the increasing number of off-the-shelf DSP functions optimized for this fabric, you must evaluate embedding DSP functions within Spartan-3 FPGAs as a viable option.

For more information, visit *www.xilinx.com/spartan3/*, *www.xilinx.com/dsp/*, and *www.xilinx.com/ipcenter/*. Σ

Linear Technology–

The Power of Choice

for all your Xilinx FPGA designs

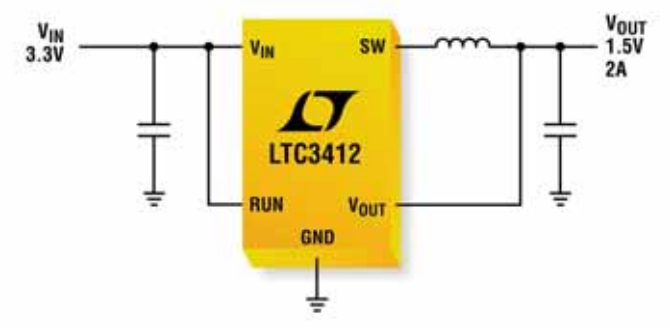

High-Performance Power Solutions.

Superior FPGA Performance.

Choose Linear Technology to power your next Xilinx FPGA design.

Now you can have the best of both worlds. The world's best FPGAs from Xilinx and the industry's best power solutions from Linear Technology. These solutions make it easy to optimize your FPGA power requirements at the beginning of your design cycle, eliminating last-minute surprises and delays.

Available today at Nu Horizons.

Nu Horizons offers everything you need to get started today, from Linear Technology samples to the FREE SwitcherCAD™ Spice simulator and more. To put the power of choice to work in your next design, contact us today.

www.nuhorizons.com/linecard/linear.html

1-888-747-NUHO

LINEAR TECHNOLOGY

NU HORIZONS ELECTRONICS CORP.

XILINX®
The Programmable Logic Company℠

Secure Your Consumer Design with CoolRunner-II CPLDs

CoolRunner-II CPLDs offer unique features to ensure a more secure design and reduce the risk of reverse engineering.

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

Product designs are a major investment. However, if a design is stolen (known as reverse engineering), that unique product can then be copied and sold for a lower price. The company that originally designed the product loses revenue and market share. Consumers may benefit in the short term, but in the long run companies will decide against major product designs and consumers will ultimately pay the price. Xilinx CoolRunner-II™ CPLDs offer a great way to protect consumer designs from reverse engineering. Of course, Xilinx CPLDs are non-volatile, so it's not necessary to configure the device at start up. But let's discuss more important security measures.

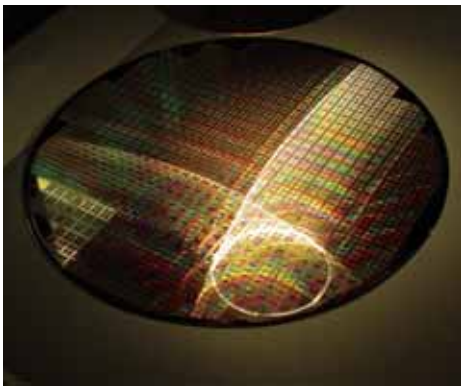

Figure 1 – CPLD wafer

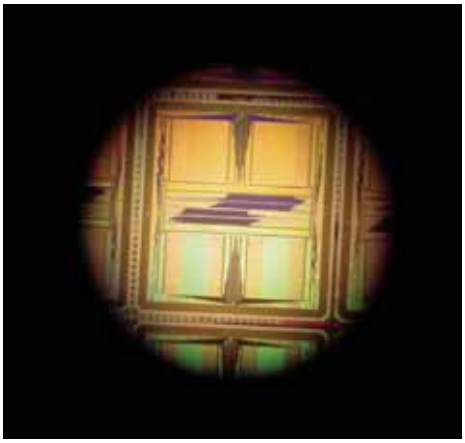

Figure 2 – Wafer below the top routing layers

Elusive Security Bits

Xilinx offers a unique capability in CoolRunner-II CPLDs: multiple security bits. These security bits are electrically erasable cells scattered throughout the device. You set the security bits by simply selecting the action within the Xilinx iMPACT programming software dialog box when the design is finished. Once these bits are programmed, the internal pattern remains fixed in the device and the program is protected from theft.

Somewhere above the substrate but beneath the metal are floating gates that hold the nonvolatile memory bit contents (Figure 1). If another company wanted to reverse engineer the design, they would have to de-program the security bits. To do that, they would first have to look through four or five metal layers to find them.

Even if they could see through four or five metal layers (Figure 2), they still couldn't "see" the bits because they are interspersed among the programming bits. Plus, they would have to figure out which ones are the security bits and which ones are the program bits. Figure 3 shows the underlying configuration cells beneath the architecture of CoolRunner-II devices.

So for someone to read-back the design, they would have to find the security bits (a very difficult task) and then erase them. If they attempted to erase them with a laser, they would have to know where to aim and how to erase each of them without erasing any other bits. They would also have to disconnect key signals for the chip operation and bit read-back. It would take many costly, time-consuming experiments to arrive at a solution.

Additional Protection Measures

After erasing the security bits, reverse engineers would still need to issue the correct demands and reverse the JEDEC file. The entire project would require a long, costly trial-and-error process, and would not be economically prudent.

Even after reading the JEDEC file, reverse engineers still need to understand the design. There are various tricks that you, as the original product designer, can use to make such an analysis prohibitively time-consuming. For example, double-data rate designs make analysis much more difficult to understand. You can also design using state machines, which are less predictable than processors. You can even build a unique CryptoBLAZE processor, based on the Xilinx PicoBlaze™ soft processor, with its own instruction set, non-volatility, and tricky timing. That would be a particularly difficult device to reverse engineer.

Additionally, the CoolRunner-II DataGATE feature can be used as a response to tampering. DataGATE is designed to dynamically and selectively block switching input signals that can draw power within CoolRunner-II devices. To increase design security, you can use the DataGATE feature to lock up the device when someone attempts to read the program.

For example, you can use a serial password from an external source, such as a keypad. If the password is correct, the device will run; if not, DataGATE will block all inputs and deny additional password attempts.

Conclusion

Considering how important maintaining design security is to your company, CoolRunner-II CPLDs offer an easy-to-implement solution to make reverse engineering CPLD designs nearly impossible. See for yourself how you can take advantage of this unique feature for your next project.

For more information about CoolRunner-II devices, visit *www.xilinx.com/cr2*. For a Quick Start presentation on security issues with Cool-Runner-II devices, visit *www.xilinx.com/products/cpldsolutions/module/cr2_security.pps*. Σ

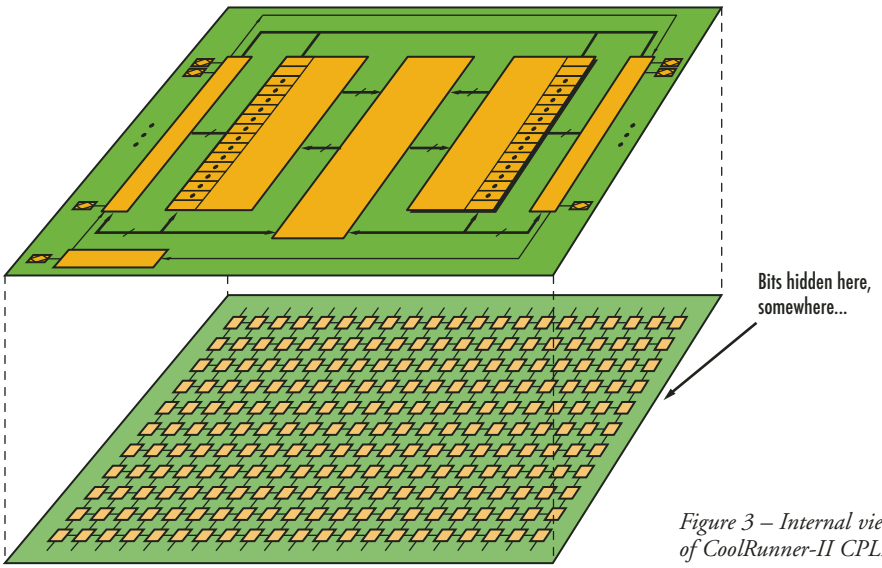

Figure 3 – Internal view of CoolRunner-II CPLD

Tips for Improving Synplify Pro Performance for FPGA Designs

Using simple setup and optimization techniques, the Synplify Pro synthesis tool helps you increase design speed and reduce chip area.

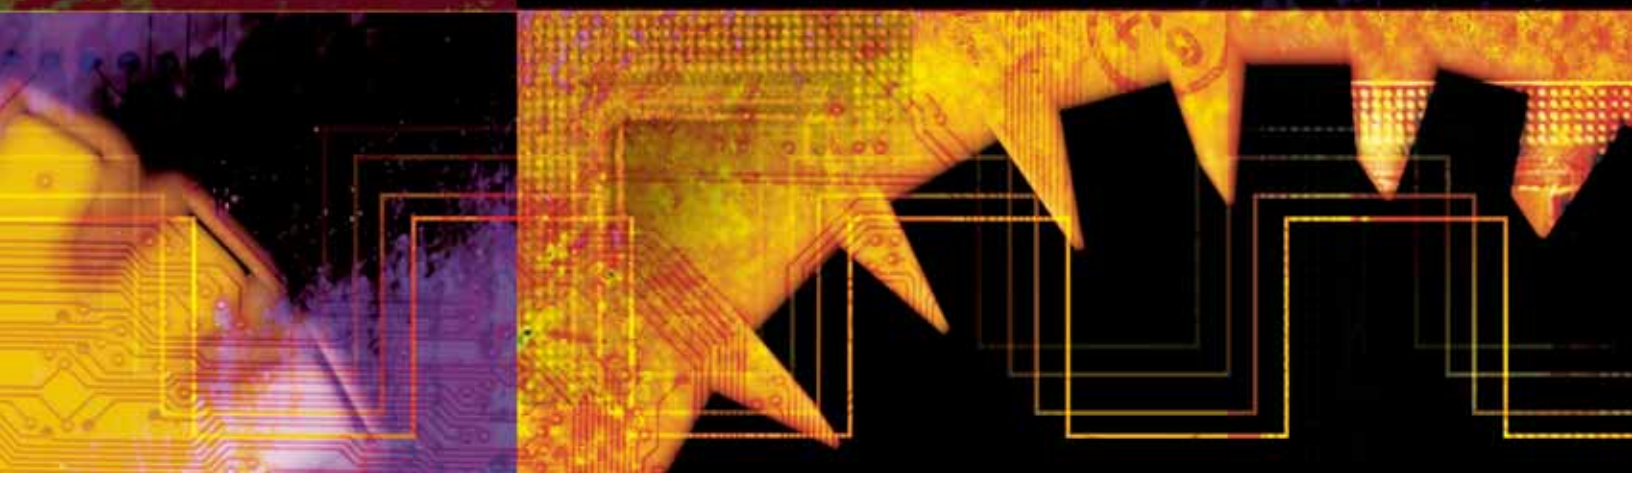

by Steve Pereira
Technical Marketing Manager
Synplicity, Inc.
stevep@synplicity.com

As system complexities advance, programmable logic follows suit. High-density FPGAs now contain millions of gates and operate at speeds in excess of 200 MHz. At this level, schedules, budgets, and FPGA design tools all begin to feel the burden.

You can incrementally increase performance or reduce the area of Xilinx devices using the Synplicity® Synplify Pro® tool in several ways. In this article, we'll describe four preferred ways to set up your design and four ways to fine-tune synthesis, all of which can be used together or independently.

Design Setup to Improve Timing or Area

Setting up your design correctly can result in huge performance increases or reductions in chip area. The following checklist describes the best practices you can use when setting up your design.

Include CORE Generator EDIFs or Timing Models for Black Boxes

If Xilinx CORE Generator™ EDIF files (*.edn) or black box timing models are provided during synthesis, the Synplify Pro tool knows the path timing and can alter the logic surrounding the boxes based on the timing constraints. If the design's critical path starts or ends in a black box, adding the EDN file usually results in better performance.

To demonstrate this point, we took an open-source design that included a black-box FIFO, with the critical path ending inside the FIFO. Without adding the CORE Generator EDN file to the Synplify Pro tool, the post PAR (place and route) results yielded an Fmax of 153 MHz. However, when we added the CORE Generator EDN file to the synthesis process, the clock frequency jumped to 171 MHz because of additional path optimization performed by Synplify Pro synthesis.

Provide Accurate Clock Constraints

Under- or over-constraining your design results in reduced performance. Do not over-constrain the clocks by more than 15%. For maximum performance, make sure that there is 10% negative slack on the critical clock. This ensures that critical paths are squeezed (see the Route Constraint section for more information).

The Fmax field on the front panel of the Synplify Pro software is fine for a quick run, but do not use it if you need maximum performance. Instead, you can put unrelated clocks in separate clock groups in the Synplify Pro synthesis design constraints file (*.sdc). If your clocks are in the same group, the Synplify Pro tool works out the worst-case setup time for the clock-to-clock paths.

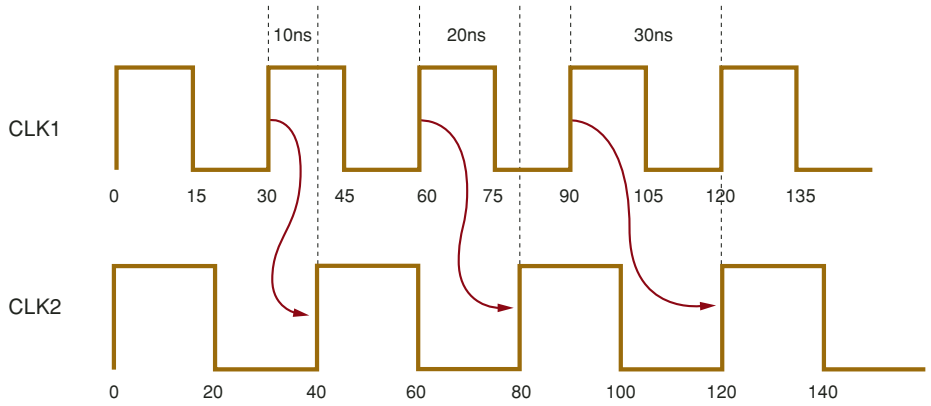

Figure 1 – Clock relationships when in the same group

Figure 1 shows a timing diagram for two clocks that are in the same clock group. Synplify Pro rolls the clocks forward until they match up again. The tool then calculates the minimum setup time between the clocks – in this case 10 ns.

If the clocks are very unrelated, several hundred clock periods may be required before the clocks match up again. This will probably result in the worst-case setup time being very small, such as 100 ps. You can check the setup time in the clock relationships table in the log file. If the setup time is too short, it is best to re-constrain the clocks so that they are more related.

Specify Timing Exceptions

You should provide all timing exceptions, such as false and multicycle paths, to the Synplify Pro tool. With this information, the tool can ignore these paths and concentrate on the actual critical paths.

As an example, in the Synplify Pro 7.3.3 tool we have enabled timing-driven, 3-state to MUX conversion. If a 3-state path is critical, the Synplify Pro tool automatically converts the logic to multiplexers, thus speeding up the path. Data on buses is usually not critical and can survive a few clock cycles because the bus master can wait for the data to become valid. In these situations, applying a multicycle constraint to the 3-state path causes the Synplify Pro tool to keep the TBUFs, thus saving area.

Constrain I/Os

If your design has I/O timing constraints, it is likely that the critical path is through the I/O buffer. The Synplify Pro tool recognizes these paths as the most critical and tries to optimize them. However, I/O paths cannot usually be physically optimized further. Therefore, the Synplify Pro tool prematurely stops optimizing the rest of the design.

A new switch has been added to the Synplify Pro 7.3 release called "Use clock period for unconstrained I/O." When enabled, the tool does not include any unconstrained I/O paths in timing optimizations, therefore allowing the optimization process to continue.

Fine-Tuning Designs to Improve Timing or Area

After setting up a design using the methods previously described, you can use additional options after synthesis to improve design performance or area utilization. Following these guidelines will usually save a device size or a speed grade, and in many cases, both.

The following optimization techniques are design-dependent. Not all designs benefit from enabling these features. The best method is to analyze the implementation of your design and see if the following optimizations improve performance.

Standard Optimization Techniques

Retiming and Pipelining

Enabling the retiming and pipelining options can improve your design performance by as much as 50%. Retiming attributes such as syn_allow_retiming let you refine your constraints by applying retiming to a single register.

Resource Sharing

With this option enabled, the software shares hardware resources, thus decreasing the area. If you disable this option, hardware resources are not shared, which will probably increase the area but yield higher performance.

FSM Compiler

This option extracts and optimizes finite state machines (FSMs) based on the number of states. As a rule, we find the following guidelines improve performance:

Number of States	Suggested Encoding Scheme
2-4	Sequential
5-40	Onehot
Over 40	Gray

FSM Explorer

If the previous methods of encoding do not produce the desired result, you can use timing-driven state encoding. The Synplify Pro tool automatically selects the best encoding for the specified timing constraints.

Resource Allocation

The use of dedicated macro blocks in Xilinx devices usually provides the best synthesis solution, but this is not always the case. A well-pipelined multiplier in logic can often provide a faster (but larger) solution. You can configure macro blocks within the Synplify Pro tool based on the design requirements. You can also force the tool to use a specific resource implementation by adding any of the following attributes:

Macro Block	Attribute (Options)
Multiplier	syn_multstyle {logic | block_mult}
RAM	syn_ramstyle {registers | select_ram | block_ram | no_rw_check}
ROM	syn_romstyle {logic | select_rom | block_rom}
Shift Register	syn_srlstyle {registers | select_srl | noextractff_srl}

These attributes are extremely design-dependent.

Aligning the routing delays almost always creates significantly better results.

Optimization Controls

The Synplify Pro tool provides user constraints to let you shape and control logic according to your design requirements. The following attributes and directives are the most commonly used.

- syn_keep (in the source code). Preserves an RTL net throughout synthesis and prevents LUT packing and replication. It is also useful for timing exceptions because you can apply a -thru constraint to it.
- syn_preserve. Disables sequential optimizations on registers, preventing removal, merging, inverter push-through, and FSM extraction.
- syn_replicate (in the constraint file). Prevents replication of registers.
- syn_maxfan (in the constraint file). Controls the maximum fanout limit, triggering register replication and buffering. This control is a hard limit on modules and instances but a soft limit when set globally.
- syn_direct_enable (in the constraint file). Forces a connection to the enable pin of the register; additional logic is moved to the D input path.

Route Constraint

The -route constraint is probably the most important but least known timing constraint. It can provide a 10% performance improvement with minimal effort, as well as drastically reduce area.

The -route constraint adds a specified delay to Synplify Pro's routing estimates. A positive value adds to the routing delay estimate and increases criticality. A negative value reduces the routing delay estimate and decreases criticality.

If the Synplify Pro timing estimate is different from the PAR value, the difference will prevent the Synplify Pro tool from optimizing the actual critical paths. The -route switch allows you to align synthesis estimates with the PAR delays. Aligning the routing delays almost always creates significantly better results.

Use the -route constraint to perform two functions:

- To make synthesis see the same critical path as PAR
- To make synthesis estimate the same slack as PAR.

If many clocks fail PAR timing, apply -route to the clock. If there are only a few paths failing PAR timing, apply -route to just those paths.

Conclusion

By setting up your design correctly and using features and constraints described in this article, you can meet and often surpass performance. We found that on 50% of the designs, using the following settings increased Fmax by more than 25%:

- Add the CORE Generator EDIF files to synthesis
- Apply the -route constraint to: paths or clocks
- Turn resource sharing off
- Turn pipelining/retiming on
- Turn use clock period for unconstrained I/O off.

For additional information about the Synplify Pro synthesis tool, contact Synplicity at (408) 215-6000, or visit *www.synplicity.com*.

Virtex-II and Spartan-3 Aid Ubiquitous Wireless Control Networking

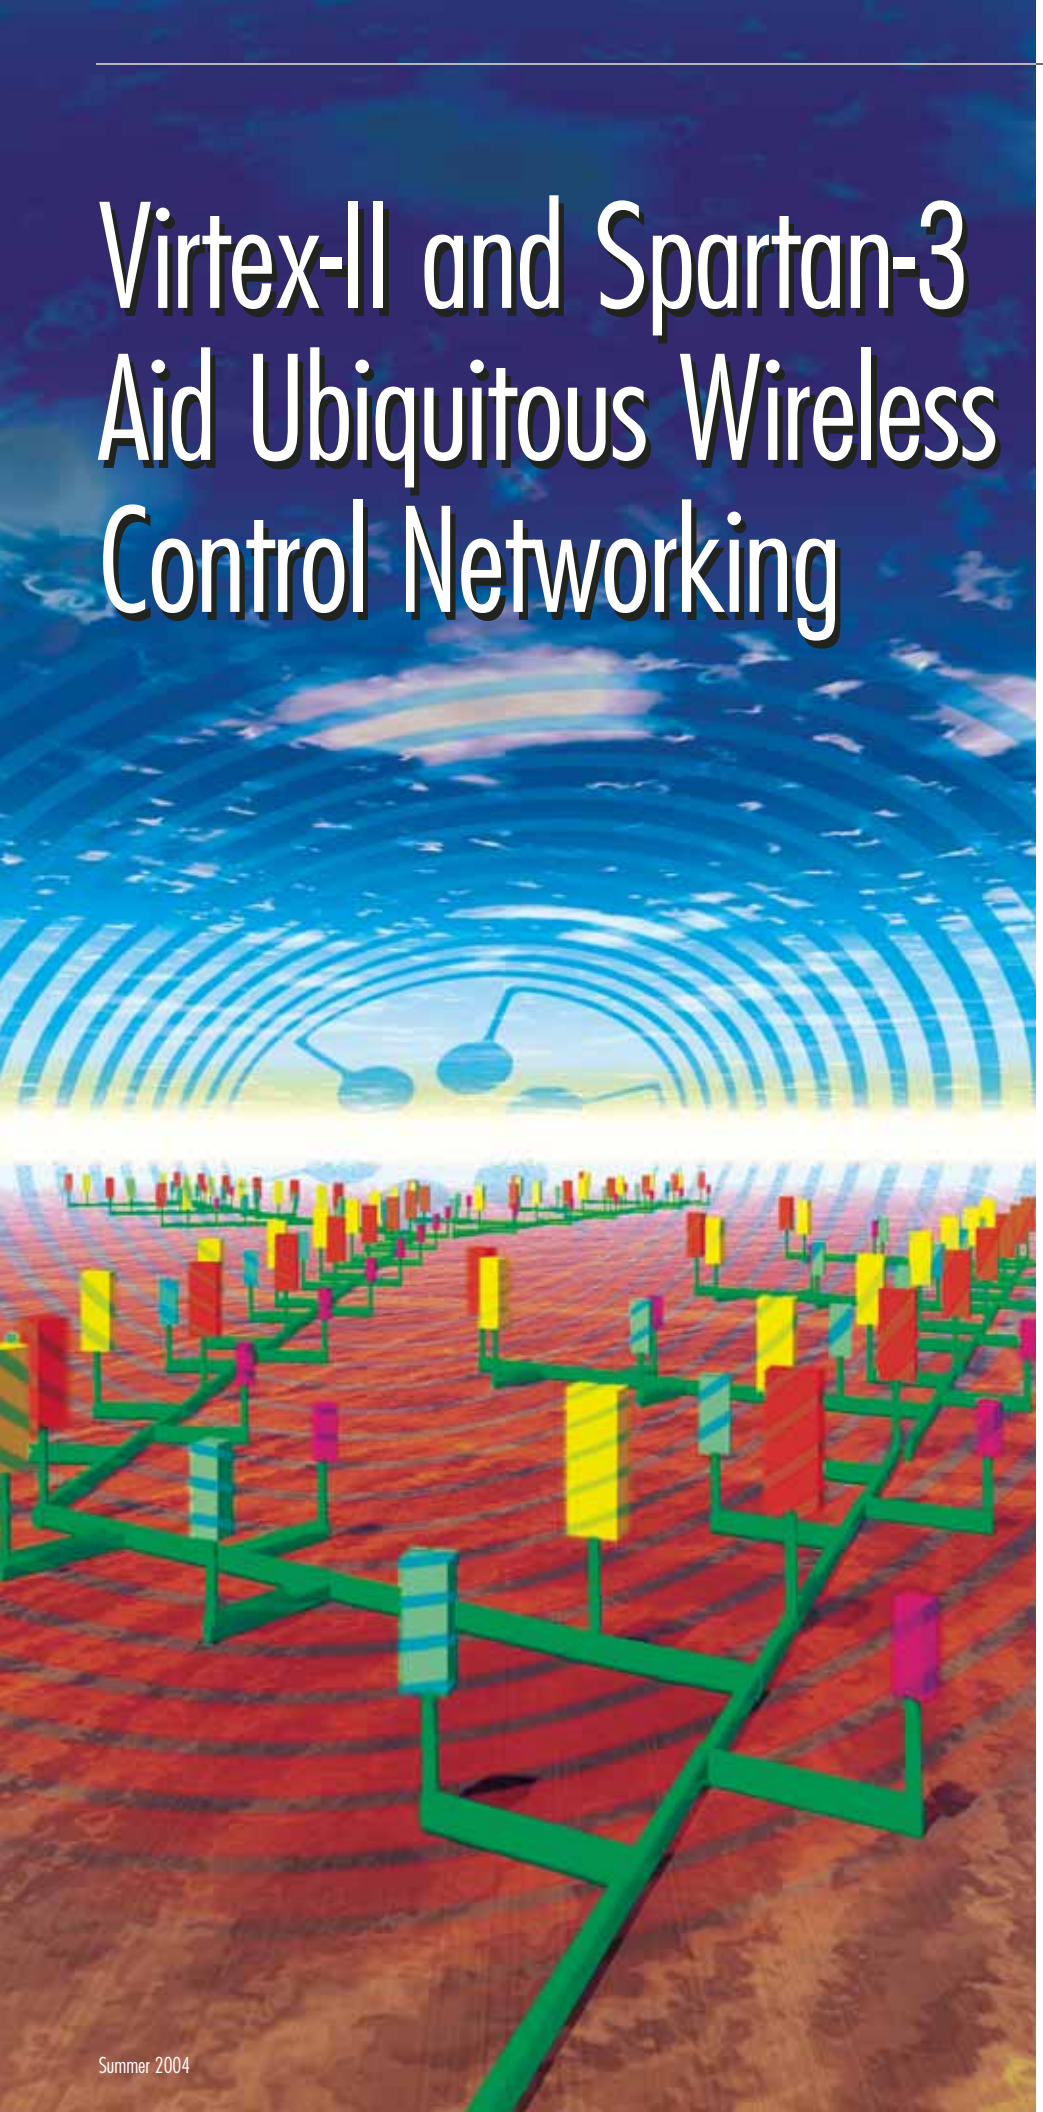

Development platforms and modules based on the new ZigBee and IEEE 802.15.4 wireless standards exploit FPGAs.

by Paul Marshall
Engineer
CompXs
p.marshall@compxs.com

A new wireless connectivity standard, IEEE 802.15.4, defines suitable media access control (MAC) and PHY layers to enable wireless control and sensing networks for low-data-rate applications. Opportunities include networked sensors in industrial, commercial, and health care applications, as well as low-cost toys and games.

The IEEE 802.15.4 standard combines well with the new ZigBee™ network and application support layers. When these standards became available in 2003, designers demanded a suitable development platform almost immediately. Such a platform had to remain flexible and have a rapid design cycle.

CompXs introduced Blencathra, the first certified IEEE 802.15.4-compliant development system, in November 2003 (Figure 1). Blencathra leverages the capacity of Xilinx Virtex-II™ FPGAs to provide extensive stack monitoring and debug facilities within the FPGA. This helps development and compliance testing and allows our customers to observe the operation of the stack in real time.

CompXs also offers a MAC/PHY module built using the low-power, low-cost Spartan™-3 family. The module does not include the stack monitoring features of Blencathra, but is ideal for customers wishing to deploy their applications cost-effectively. It also includes an integrated 2.4 GHz radio.

Wireless Networking for Control Applications

ZigBee defines network and application support layers for wireless networks based on the MAC and PHY layers of IEEE 802.15.4. The IEEE standard uses the global 2.4 GHz ISM (industrial, scientific, and medical) band, as well as the American 915 MHz and equivalent European 868 MHz unlicensed bands. The maximum data rate in each band is 250 Kbps, 40 Kbps, and 20 Kbps, respectively. The range is typically 30 meters but can extend to 100 meters in optimal conditions.

The 802.15.4 physical layer uses direct sequence spread spectrum to spread the information over a range of frequencies. For devices that transmit infrequently, this allows for greater power conservation than Bluetooth's™ frequency-hopping scheme. At the MAC level, another advantage of 802.15.4 is that it has only two power modes: active or sleep. This greatly simplifies power management.

All devices have 64-bit IEEE addresses, allowing virtually unlimited devices in a network. This allows for massive sensor arrays and control networks, but the option also exists to allocate 16-bit addresses to reduce packet size.

IEEE 802.15.4 is well suited for periodic data (such as sensor outputs generated at a rate defined by the application), intermittent data generated externally by a switch, or repetitive low latency data allocated to a specific time slot (such as mouse data).

The ZigBee Alliance has defined the upper layers of the protocol stack to use the IEEE 802.15.4 MAC and PHY. ZigBee includes the parts of the protocol from the network layer to the application layer, including application profiles. The first profiles were published in mid-2003.

Blencathra

The Blencathra development platform allows developers to build and analyze ZigBee/802.15.4 designs quickly and at little design risk. Blencathra implements the entire 802.15.4 MAC and PHY in hardware using a Xilinx XC2V1500 Virtex-II FPGA.

Within the FPGA, CompXs' IP implements the MAC and PHY state machines, with shared MAC and PHY RAM. Timing, encryption, and modem functional blocks are also implemented in the device.

The 17,280 logic cells of the XC2V1500 FPGA provide vastly more capacity than needed to implement the 802.15.4 MAC and PHY, which are designed to have a very small footprint. The remainder of the device, more than 75% in fact, is used to implement compliance verification logic.

Figure 1 – The Blencathra 802.15.4 development platform with full stack tracing

By using the 864 KB of on-chip block RAM, pipelining the event log to implement a high-speed port on board the FPGA is easy. Through this port, you can inspect activity all the way up and down the 802.15.4 stacks in real time. This is an extremely valuable capability, because it shows very clearly how changes at the ZigBee layers affect behavior throughout the design.

Bowfell

CompXs has also created Bowfell, an 802.15.4 MAC/PHY module that combines easily with ZigBee software and includes an integrated 2.4 GHz radio. After proving the design using the Blencathra development system, you can use these cost-effective modules to quickly configure networks with many ZigBee nodes. As there is no need for on-board verification logic, the modules are built using the low-power, low-cost Xilinx Spartan-3 FPGA.

Spartan-3 devices support low power consumption, low cost, and fast time to market – the IEEE standards were published in October 2003, and by November the development system and turnkey modules were fully implemented using Xilinx devices.

An ASIC would likely have provided greater power savings, but the design cycle is far longer. Choosing the FPGA route enabled fully developed products to reach the market very soon after the standards were first published. Many of the details of this standard continue to change and evolve at this early stage. Therefore, the extra flexibility to reconfigure the hardware is valuable both to customer developers and to CompXs.

Application Development

You can use the Blencathra development platform on its own to develop a pure 802.15.4 wireless communication channel for links that require no network processing. A wireless keyboard or mouse, for example, requires no additional layer to handle network tasks such as routing. All you need is a radio block (CompXs has a suitable radio for development purposes), and you can set up a representative point-to-point link on the bench. The radio has been designed to ease development headaches by delivering strong performance.

For more complex applications requiring network processing capability, the ZigBee protocols add a network layer to the 802.15.4 system. A CompXs daughterboard plugs directly into Blencathra to facilitate this. The board, named Bannerdale, hosts network layer processing for a simple ZigBee-compliant network. On board is an eight-bit Flash microcontroller with EEPROM.

Note that the microcontroller has just one timer and serial peripheral interface (SPI), and only 8K of ROM for the network coordinator. This can easily support the network layer, demonstrating that you need only minimal microcontroller resources to implement ZigBee. The microcontroller

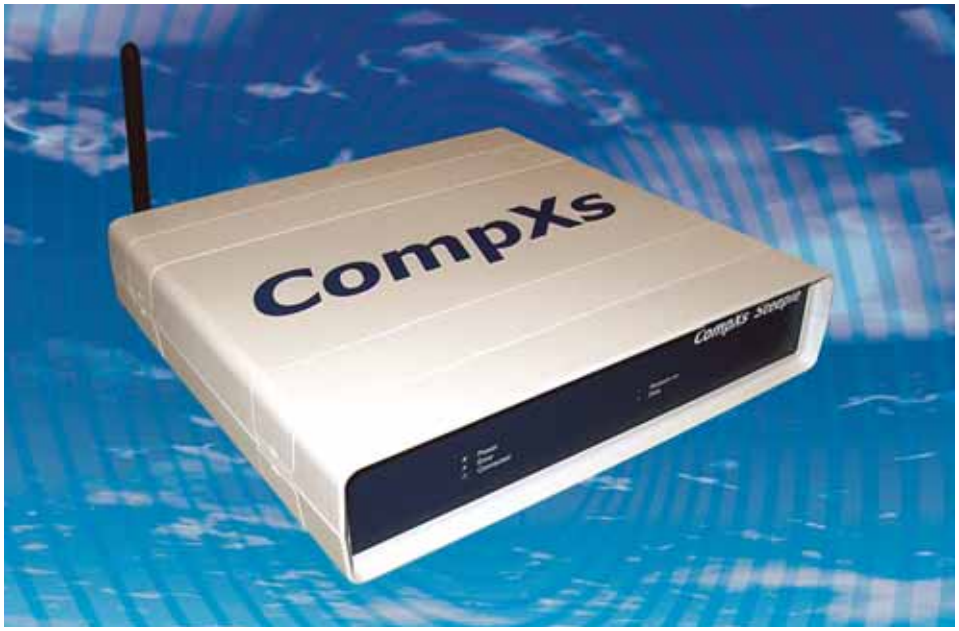

Figure 2 – The 802.15.4 packet sniffer from CompXs

may also be able to host the application if processing requirements allow.

Overall, ZigBee typically requires between 4 KB and 30 KB of RAM and ROM, depending on the complexity of the application. This compares with the 250 KB required by Bluetooth, for example. So ZigBee/802.15.4 not only simplifies the process of embedding wireless communications into products, but also makes for a considerably lower bill of materials in the final product.

Note that IEEE 802.15.4 is not dedicated to ZigBee as a network layer. If the network processing requirements are very simple and can be implemented quickly using very low memory resources, you can define your own network layer if you prefer.

Easing the Design Challenge

Off-the-shelf ZigBee software libraries will provide the fastest and easiest solution as they become more widely available. Current CompXs libraries include proprietary network layers as well as ZigBee version 0.7-compliant network and application support layers. These are ready to be integrated with IEEE 802.15.4 on Spartan-3 FPGA-based modules, or as part of a system-on-chip.

The network layer implemented on the Bannerdale daughter board for development purposes is also available as a linkable library to run on your target processor, or as source code. In fact, CompXs offers a complete set of development platforms, network modules, and tools. Available tools include an 802.15.4 platform stack analyzer hat that displays and logs activity to microsecond accuracy and a passive 802.15.4/ZigBee packet sniffer and analyzer (Figure 2).

The Steeple packet sniffer is based on the IP contained within the FPGA on Blencathra. Steeple will "sniff" all of the transmissions on a ZigBee/802.15.4 network and then display those transmissions in a convenient form on a PC (Figure 3). It recognizes valid and invalid transmissions and breaks down the packets of data, displaying them in an easily understood manner.

You can also quickly integrate proprietary application software with the ZigBee stacks via standard ZigBee APIs.

Conclusion

When a new networking standard emerges, developers first look for the easiest way to get a standard-compliant network up and running. A reconfigurable development platform is important, as well as large numbers of low-cost modules that can implement the standard-compliant elements with a good RF stage.

In the case of IEEE 802.15.4 and ZigBee, Xilinx FPGAs allowed for easy and rapid designs of suitable development tools. These tools will enable many new applications to benefit from low-cost wireless networking.

You can find more information on the products described here at the CompXs website, *www.compxs.com*. For details about the ZigBee organization and the standards it promotes, visit *www.zigbee.org*. And to learn how 802.15.4 and ZigBee can be used in your products, consider taking a training course. For information about introductory and in-depth/hands-on courses, visit *www.zigbeetraining.com*.

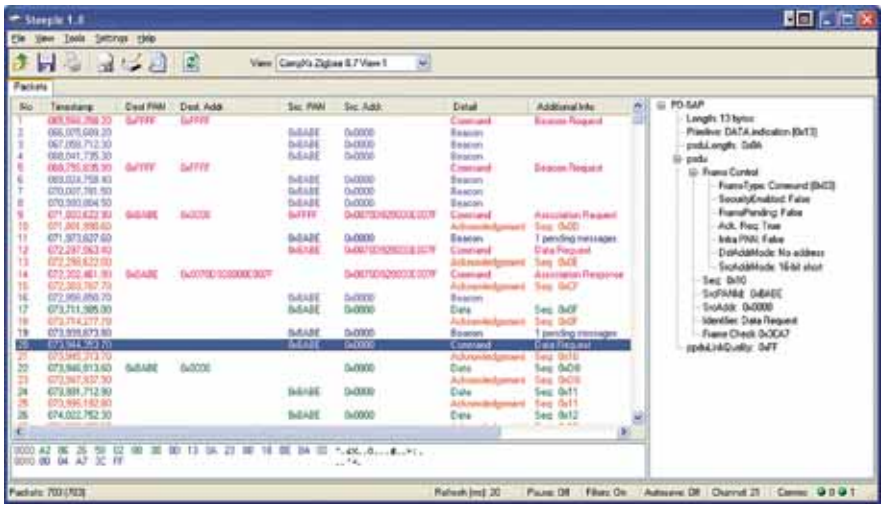

Figure 3 – A screenshot from Steeple showing a ZigBee network start up

Creating Pb-Free Packaging

Maintaining performance and reliability are key challenges as the industry complies with new Pb-free regulations.

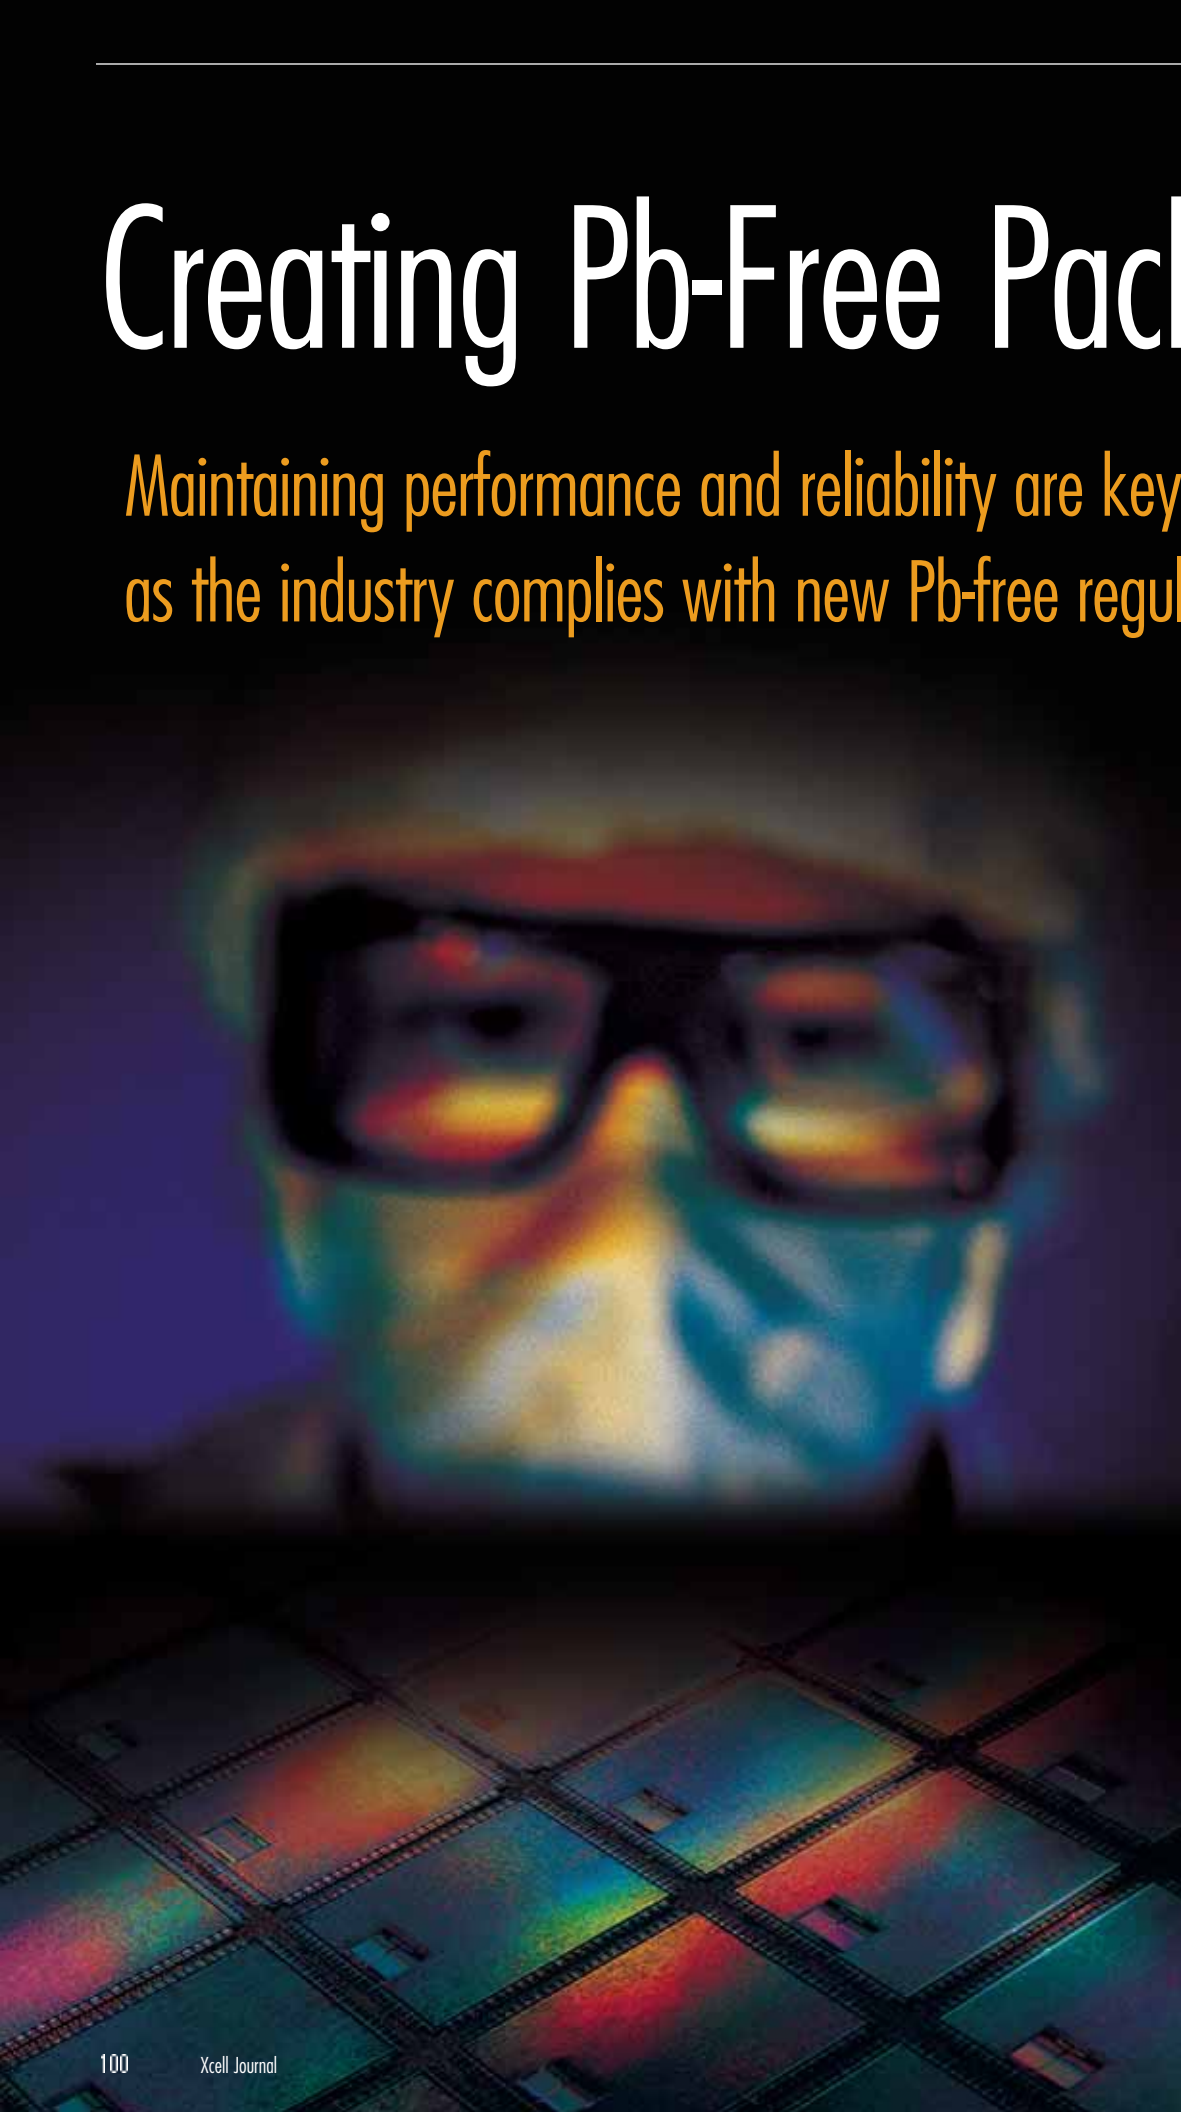

by Abhay Maheshwari
Director, Package Engineering
Xilinx, Inc.
abhay.maheshwari@xilinx.com

Lead (Pb)-free packaging is part of a concerted effort in the electronics industry to eliminate Pb in electronic assembly. Only ~0.2% of the Pb worldwide is used for electronic assembly; most of it is used in automotive applications. But when a hazardous substance ban is implemented, all products containing the material are equally affected.

The Pb-free movement in the electronics industry has accelerated since the European Union (EU) released a directive calling for the removal of Pb and other hazardous materials from electronics by 2006 as part of the RoHS (Restriction of Hazardous Substances).

The requirement for Pb-free is industry-wide and not exclusive to Xilinx. Although we wish to lead Pb-free implementation efforts to secure advantages in winning future designs, it is equally important to align with the rest of the industry in terms of Pb-free technical solutions, such as solder ball composition and lead finishes.

Xilinx has made a conscious effort to remain within industry-standard boundaries for lead finishes and therefore leverage the infrastructure to satisfy industry needs. However, there are some differences. As a PLD company, Xilinx is in a unique position because it has the industry's largest die sizes; the resulting stresses in packages are much higher. As a result, the most common package construction materials, such as die attach and mold compound offered for Pb-free packages, required a significant overhaul to ensure the same levels of performance and reliability for Xilinx applications.

Technical Challenges

Figure 1 shows the application of Pb-based alloys in electronic materials. Pb in tin (Sn)Pb solder is used for electronic assembly. The alloy is called eutectic solder, which consists of 63% Sn and 37% Pb. This material has been very well characterized and understood for the last 40 years, and a drop-in replacement of SnPb for soldering applications in electronics has not yet been found.

Typically, when the parts are placed on PCBs and sent through the board assembly process, the SnPb alloy melts to form a solder joint. This is what is referred to as the reflow process. The alloy melts at 183°C and the peak reflow temperature for this alloy is restricted to 220°C.

But Pb-free replacement alloys in the electronics industry have higher melting temperatures than the current SnPb eutectic alloy. As a result, the peak reflow temperatures have gone up by 25 to 40°C. This poses a significant restriction on the material capability of existing electronic packages today, as they must be able to withstand temperatures of 245 to 260°C. Figure 2 illustrates this issue.

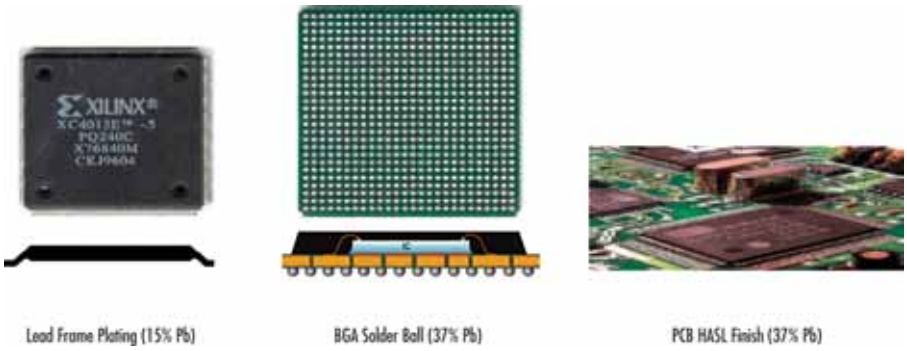

Figure 1 – Sources of Pb in the electronics industry

Today's Pb-free packages are capable of withstanding higher reflow temperatures during assembly. As a result, all package construction material such as die attach, substrate, and mold compound have improved significantly so that they too can withstand higher reflow temperature requirements.

Leader of the Pack(age)

The Pb-free program at Xilinx was established in 1999 as a proactive effort to ensure future leadership. Xilinx quickly formed partnerships with customers like Sony™ and Matsushita Electric Industrial for Pb-free beta-site implementations.

The initial focus for Pb replacement was unclear, and mostly viewed as an environmentally friendly product introduction offering a marketing advantage. Our implementation strategy today is far more serious than it was two years ago, given the EU directive with its "hard" compliance date of 2006.

The Xilinx plan is a global strategy of implementation and industry standardization. Xilinx has closely followed the activities of industry consortiums MEPTEC (Microelectronics Packaging and Test Engineering Council) and NEMI (National Electronic Manufacturing Initiatives Inc.) and standards organiza-

Figure 2 – Pb-free implications on electronic packages

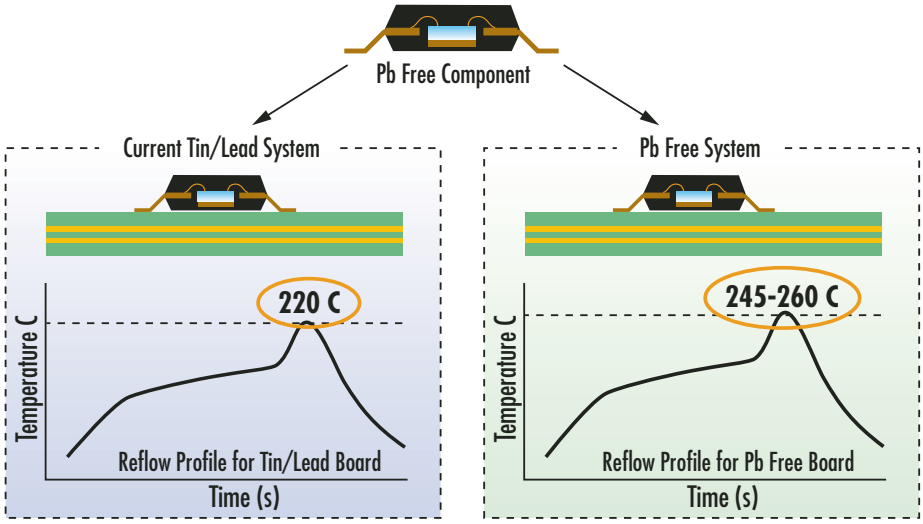

Figure 3 – Pb-free package backward compatibility

tions such as IPC and JEDEC (Joint Electron Device Engineering Council) to ensure that the solutions for Pb-free are accepted worldwide.

The primary assembly partner of choice was Amkor™, closely followed by Siliconware Precision Industries. The first Pb-free prototypes were shipped to beta customers for evaluation in 2001; since 2002, Xilinx has shipped Pb-free products in volume.

The extra letter "G" in current package designations easily identifies Xilinx Pb-free parts. For example, the Pb-free version of the VQ100 standard package is VQG100. This unique identification of the product simplifies inventory and supply chain management.

Backward Compatibility

In the electronics industry, the ideal conversion to Pb-free is a drop-in replacement of solder finishes and associated materials in electronic packages that mimic the current SnPb finishes relative to assembly and reliability. Unfortunately, no such product exists today that allows the industry to quickly switch over to Pb-free products.

There will always be a transition period for any major change requiring industry-wide implementation. Customers can expect today's standard and Pb-free products to coexist until the date of transition. From the perspective of component suppliers like Xilinx, this calls for carrying inventory of both packages for all product families targeted for the Pb-free market. This is a significant effort in terms of operations, with a supply chain that must be managed skillfully until all products are Pb-free.

This is a significant effort in terms of operations, with a supply chain that must be managed skillfully until all products are Pb-free.

From a user perspective, Xilinx would like to ensure that all system components are Pb-free compatible, including all components, passives, and connectors. This presents a significant challenge, however, as all suppliers are clearly not ready all at once to implement a Pb-free solution. Customers will have different timeframes for implementation dates using the Pb-free package solution, and companies supplying parts to the industry will be forced to carry dual inventories of packages for the same product.

A creative solution to this problem is backward compatibility – using Pb-free packages directly on existing PCBs with no changes in the assembly process or long-term reliability.

The schematic in Figure 3 illustrates the requirements and restrictions of backward compatibility. Lead frame Pb-free packages (TQ, PQ, VQ, SO, VO, and so on) are fully backward-compatible with the proposed Pb-free solutions.

The PBGA (plastic ball grid array) packages are a different story. To date, no obvious backward-compatible solution exists in the industry. This mandates a dual inventory of Pb-free PBGA parts until the industry is fully converted to Pb-free assembly solutions.

Lingering Legacy Issues

The Pb-free lead frame solution has a matte Sn finish on the leads. For the last 30 years, Sn plating has been used on terminal finishes for passive components. There are a few legacy issues relating to a whisker-like structure formation on Sn-plated leads.

When the electronics industry was in its infancy, bright Sn finishes were common. This bright Sn plating was shown to be susceptible to whisker growth (single crystals) in appropriate temperature/time conditions. The whiskers would eventually grow large enough to short out adjacent leads.

The new solution using appropriate matte Sn plating is implemented using an advanced, well-controlled Sn-plating bath with special additives resistant to whisker growth. Although no standard tests in the industry exist for whisker growth, Xilinx has worked very closely with industry consortiums and assembly partners to exhaustively test for whisker growth in the Sn-plating offered for Pb-free lead frame packages. So far, none of the known tests have shown significant whisker growth.

Although most of the industry, including large suppliers of microprocessors and controllers, have made clear commitments to move towards this preferred industry solution, the telecom, networking, storage, and aerospace industries are cautious about matt Sn as a single alternative for Pb-free. As a result, new tests are being proposed

Xilinx has successfully introduced Pb-free packaging solutions for wire-bonded parts, shipping in volume since 2002.

with several whisker growth mitigating solutions, although as yet no clear agreements exist among the industry players.

Pb-Free Flip-Chips

By definition, flip-chip packages are ball grid array packages with interconnect solutions based on area array solder bumps. The internal interconnection is through solder bumps, which are similar in composition to external solder balls.

Xilinx is evaluating its flip-chip packages with large die sizes, and our current focus is to introduce Pb-free flip-chip packages in two phases. The primary goal is to be compliant with RoHS.

The first phase for the introduction of Pb-free flip-chip packages will be based on eutectic Sn/Pb solder ball replacement with Pb-free solder balls only. The current schedule for this implementation is by the end of 2004. This will allow packages to be on customer boards until the 2006 deadline.

The second phase will ensure complete compliance with RoHS mandates by 2006. The primary focus during this stage will be compliance with the RoHS directive for the solder bumps inside the package at the silicon-to-substrate interconnect level. Xilinx plans to have a solution established at least a year before the 2006 deadline.

Industry Compliance

Europe

The European Union RoHS legal directive calls for the restriction of six primary materials from electrical and electronic products by 2006. These are Pb, mercury, hexavalent chromium, cadmium, and two types of flame retardant used in packages abbreviated as PBB and PDE.

There are some key exemptions in RoHS that have made conversion to Pb-free very complicated from an operations and business perspective; Figure 4 lists one example in which network and storage products are exempted. Similar exemptions have also been granted to high-Pb materials and ceramic packages.

Most companies in Europe are committed to implementing Pb-free solutions by the RoHS deadline. Evaluations are ongoing with samples of Pb-free parts on test boards to understand the manufacturing issues.

Japan

Japan has clearly been the leader for Pb-free implementations in most consumer products. Many companies such as Sony and NEC™ have issued mandates stating that all consumer products in 2004 will be Pb-free.

One key challenge has been to develop every component on the board (connectors, passives, boards, and associated materials) as Pb-free and capable of higher reflow temperatures. This elusive capability has forced many companies to push their implementation timelines back several times. Nevertheless, it has been clearly established that doing business in Japan requires a Pb-free solution for electronic packages.

North America

In North America, laws banning or restricting the use of Pb are already in place for many products, and there is an increasing demand for a total ban. However, the North American electronics industry has been slower than other global regions to adopt the move to Pb-free. This pattern is changing, however, with RoHS timelines for implementation set for July 2006.

Companies are scrambling to understand the implications of Pb-free solutions and have lobbied to extend the timelines for implementation to ensure that the robustness of the solution is proven to their satisfaction. Lobbying has yielded significant exemptions in RoHS, as described earlier. Many companies are working through industry consortiums such as NEMI and CALCE (Computer Aided Life Cycle Engineering) to understand the robustness of industry solutions, proposing new tests and evaluations to address specific concerns. Until these concerns are addressed beyond a doubt, Pb-free implementation timelines will inevitably be pushed back.

RoHS Key Exemptions

- Lead in high melting temperature type solders (i.e. tin-lead solder alloys containing more than 85% lead)
- Lead in solder for servers, storage, and storage array systems (exemption granted until 2110)
- Lead in solders for network infrastructure equipment for switching, signaling, transmission as well as network management for telecommunication
- Lead in electronic ceramic parts (e.g., piezoelectronic devices)

Figure 4 – RoHS key exemptions

Conclusion

Xilinx has successfully introduced Pb-free packaging solutions for wire-bonded parts, shipping in volume since 2002.

Yet current RoHS exemptions, coupled with non-backward-compatible PBGA solutions, pose a significant issue for the supply chain. These exemptions and restrictions may result in carrying dual inventory on specific products for an extended period of time and hence could have significant cost and logistical implications.

In the long term, it is expected that the industry will convert to a full Pb-free implementation on PCBs and convert all packages to Pb-free solutions. Σ

TechXclusives: A Valuable Source of Information

A frequent contributor to this section on *www.xilinx.com* gives a quick tour through recent articles.

by Peter Alfke
Director, Applications Engineering
Xilinx, Inc.
peter.alfke@xilinx.com

TechXclusives is an evolving series of technical articles and tutorials written by experts in Xilinx Applications. TechX articles are published on the Xilinx website under *www.xilinx.com/xlnx/xweb/xil_tx_home.jsp* and are also included in the quarterly DataSource CD.

These short papers cover a very broad range of subjects, from the evolution and best selection of Xilinx FPGAs and certain design styles to specific systems and circuit design tutorials, as well as advice on subtle issues such as metastability, single-event upsets, and PC-board signal integrity.

Below is a short description of all articles published before February 2004 organized by four topics: General, Tutorial, Applications, and Electrical.

General Topics

"Choices, Choices, and Opinions"
This summary of FPGA and CPLD devices helps you select the right technology and FPGA family.

"Evolution and Revolution: Recent Progress in Field-Programmable Logic"
FPGAs are now bigger, faster, and cheaper, with better software, faster compile times, and better technical support. This article provides a wide-reaching overview.

"Performance + Time = Memory"
You pay for silicon area. In this TechX, Ken Chapman suggests looking at time-sharing and sequential design as a third dimension to reduce cost.

Tutorials

"Moving Data Across Asynchronous Clock Boundaries"
This article explains how to design reliable and predictable asynchronous interfaces when multiple unrelated clocks access common data.

"Metastability Delay and Mean Time Between Failure in Virtex-II Pro Flip-Flops"
This article analyzes the metastable behavior of Virtex-II Pro™ flip-flops and provides quantitative data for calculating the metastable mean time between failure (MTBF).

"A Thousand Years Between Single-Event Upset (SEU) Failures"
SEUs can lead to data loss in configuration latches or user flip-flops. SEUs are caused by uncontrollable external forces, such as cosmic rays or neutrons. This article describes how Xilinx is running large-scale experiments that measure and document SEUs in a normal environment.

"IBIS Model Usage"
A TechX about I/O buffer information specification (IBIS) files, which extract SPICE parameters and present them to users, while protecting proprietary information.

"Magic Numbers"
Why do you need a 19.44 MHz clock signal? This article presents the derivations of these and other "magic" numbers in telecom and datacom.

"Digitally Removing a DC Offset (or 'DSP Without Math?')"
This article takes a gentle look at DSP and shows you how to optimize audio telecom functions using the SRL16E mode.

"Programmable Development and Test"
Learn some simple ways to use the programmable nature of Xilinx devices to help in product development and even accelerate testing on the production line.

"Expanding Virtex-II Multipliers"
Find out how to expand the natural bit-width capability of dedicated multipliers in a way that makes best use of Virtex-II™ resources.

"Asynchronous FIFO in Virtex-II FPGAs"
A FIFO is a popular memory structure to move data across clock boundaries. This article provides useful information to implement FIFOs in your design.

"Does Your Design Have Enough Slack?"
This article explores the factors that influence propagation delay, with ways to improve performance through simulation and better place and route.

"8 x 12 Does Not Equal 12 x 8"
How to optimize multipliers implemented in Virtex and Spartan™ CLBs.

"FPGAs Driving Voice-Data Convergence"
This article offers an overview of voice data convergence technologies, their benefits, and some of the significant challenges facing system designers.

"Color Space Conversion"
A discussion of color space conversion, used in broadcast-quality video systems. It may appear complicated in theory, but it can be reduced to a collection of basic functions that are well-suited for implementation in Xilinx FPGAs (adders, subtracters, multipliers, and delays).

Applications

"Six Easy Pieces (Non-Synchronous Circuit Tricks)"
With six proven designs, learn how to implement certain types of asynchronous functions, such as switch debouncer, RC-oscillator, Schmitt trigger, frequency doubler, and clock multiplexer.

"Creating Embedded Microcontrollers (Programmable State Machines)"
This TechX series describes the design of very small processor macros, bringing the advantages of a processor to a traditional design environment at minimum cost.

"Multiplexer Selection"
This article describes several ways to implement multiplexers in Xilinx FPGAs, from the straightforward method to alternative techniques using fewer device resources or achieving higher speed.

"Using Leftover Multipliers and Block RAM in Your Design"
How to use free multipliers as shifters and unused block RAMs as state machines, sine-cosine look-up tables, or 20-bit counters.

"Get Smart About Reset (Think Local, Not Global)"
Applying a global reset to your FPGA designs is not always a very good idea. This article discusses this highly controversial issue.

"Saving Costs with the SRL16E"
The SRL16E shift register is available in every look-up table in every CLB. This TechX explains how using this exciting mode can lead to significant cost savings.

"Timing Closure"
This article describes a proven methodology for timing closure in high-speed/high-density applications to meet given performance objectives.

"Relationally Placed Macros"
By placing logic blocks relative to each other, RLOC constraints allow you to increase speed and use die resources efficiently. This article explains the steps involved.

"Reconfiguring Block RAMs"
This TechX explores the use of the JTAG port to interrogate and update the contents of block RAMs and registers while a design is running.

Electrical Issues

"Signal Integrity Tips and Tricks"
Controlling crosstalk, ground bounce, ringing, noise margins, impedance matching, and decoupling is now critical to a successful design. This article covers the various techniques and design issues to ensure that signals are undistorted and do not cause problems.

"Printed Circuit Board Considerations"
A discussion of PC board issues that applies to all modern systems with fast current and voltage transitions: VCC and ground planes, VCC decoupling, and transmission line reflections and terminations.

"Printed Circuit Board Modeling Issues"
Printed circuit board design has become an important part of a successful product. This article describes ways to avoid the pitfalls of the "build it and see if it works" method.

"It's Not Your Father's PCB Anymore"
A discussion of signal integrity design, now more necessary than ever if you want to save time and money and get things to work reliably.

"Those Tiny Little Vias Can Cause Bad Ground Bounce Problems"
This TechX shows how vias between PC board layers have become an insidious contributor to the problems caused by excessive ground bounce from simultaneously switching outputs (SSOs).

"What are Virtex and Spartan-II I/O Pins Doing?"
A detailed explanation about how I/O pins behave during power up, before and during configuration, and during normal operation.

"Jitter"
A discussion of the causes, measurement, and management of jitter.

"Power To The People – Not the FPGA!"
This article explains how excessive power-on current has finally been designed away, beginning with Virtex-II devices.

"The Old 35 pF Just Disappeared"
This TechX reflects on the traditional 35 pF lumped capacitive load, a meaningless and misleading model of today's outputs, which are now so fast that the load acts as a transmission line instead of a lumped capacitance.

Conclusion

You can find these and other TechXclusives at *www.xilinx.com/xlnx/xweb/xil_tx_home.jsp.*

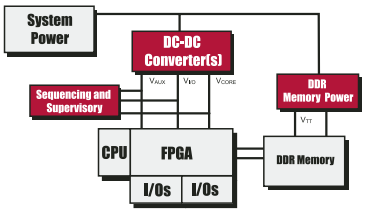

Power for Your Xilinx® FPGA and DDR Memory Designs

Intersil's switching regulators (PWMs) maintain efficiencies in excess of 90% in your Xilinx FPGA and DDR memory designs, even when the input and output voltages differ by a large amount and the current requirements range from a few milliamps to 100A. Intersil's regulators are available in several configurations including single-phase, multi-phase, integrated FETs up to 8A and up to 100A with external FETs.

www.intersil.com/xilinx

EL75XX Family of Integrated Switchers

Intersil's EL7536/EL7562/EL7554/EL7566 family of highly integrated switching regulators offer ease of use and very small footprint because the power FETs are integrated on the device. Input voltage of 3 to 6V are possible and continuous output currents of 1, 2, 4 or 6 amps respectively are available. These products make a very compact, and >90% efficient solution for either V_{DD} or OV_{DD}. Additionally, for +5V systems the 8A continuous output EL7558 is available.

www.intersil.com/EL7566

LinearLink™ Interface ICs

Space saving, high-performance data transceivers provide longer reach and higher data integrity than current competing products for data communications and industrial applications. LinearLink Interface ICs include:

RS-232 Interface
- 5v and 3.3v families
- High-speed, 1Mbps
- IEC61000 protected
- Low power shut down mode
- Ultra small package option

RS-485/RS-422 Interface
- 5v and 3.3v families
- 15kV HBM ESD protected
- High speed up to 40Mbps
- Up to 256 nodes on the bus
- Half and full duplex

www.intersil.com/interface

How Many Low Voltage Supplies Do You Need?

Multi-output DC-DC Converters from Intersil

Intersil Power Management Solutions

Each technology generation seems to create a new low voltage requirement: 2.5V, 1.8V, 1.5V, 1.25V, 1.2V, 0.9V and on it goes. Intersil offers a broad portfolio of power management ICs to easily generate the voltages you need.

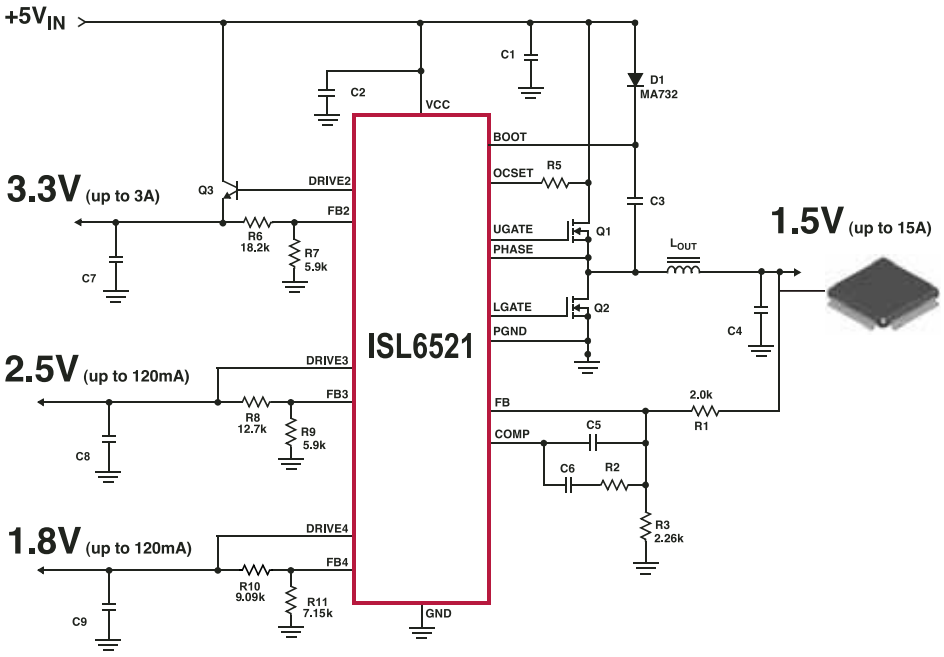

Device	Regulators PWMs	Regulators Linears	Vin	Package/Pin	# of Output Voltages
ISL6521	1	3	5V	SOIC-16	
HIP6021	1	3	5V, 12V	SOIC-28	
HIP6019B	2	2	5V, 12V	SOIC-28	4
ISL6537 (new)	2	2 + Ref	5V, 12V	QFN-28	
ISL6532A	1	2	5V, 12V	QFN-28	3
ISL6402/A (new)	2	1	4.5V to 24V	TSSOP-28, QFN-28	
ISL6539 (new)	2	0	5V to 15V	SSOP-28	
ISL6227 (new)	2	0	4.5V to 24V	SSOP-28	
ISL6444	2	Ref	5V to 24V	SSOP-28	2
ISL6530/1	2	Ref	5V	SOIC-24, QFN-32	
ISL6528	1	1	3.3V, 5V	SOIC-8	
ISL6529	1	1	3.3V to 5V, 12V	SOIC-14, QFN-16	

Learn more about this family and get free samples at www.intersil.com/multiDCDC

Get more technical info on Intersil's complete portfolio of High Performance Analog Solutions at
www.intersil.com/info

Amplifiers | Converters | Interface | **Power Management** | Switch/MUX | Voice Over IP | Military/Space

intersil® HIGH PERFORMANCE ANALOG

Intersil High Performance Analog Solutions

Virtex-II Pro™ LC UltraController™
Development Kit

Ready to Transform Your Embedded Processor System?

Learn to leverage the high performance embedded PowerPC™ inside the Xilinx Virtex-II Pro FPGA family, and master the design of the UltraController with this kit from Memec Design. Using the pre-configured, self-contained UltraController, you will be able to design your application, understand the UltraController hardware and software development flows, and see how easy it is to leverage the free processing power of UltraController. Embedded processing couldn't be easier, and the UltraController delivers the performance you need with minimal impact to your FPGA logic. What are you waiting for?

XILINX®

Call 888.488.4133 ext.268 or visit *www.memec.com/xilinxkits*

Memec Design

Support Across The Board.™

SpeedWay Design Workshops™

Hardware-based, laboratory-oriented workshops are now available from Avnet Cilicon for engineers interested in the newest Xilinx products and design tools such as Spartan™-3, Virtex-II Pro™, ISE 6.2i and EDK 6.2i.

Part Number	Description	Special Pricing for Workshop Attendees*
ADS-SPEEDWAY	SpeedWay Design Workshops™	$99.00 USD (free with kit purchase)
ADS-SPDWY-S3	Spartan-3 Evaluation Kit, populated with an XC3S400 device	$249.00 USD ($399.00 value)
ADS-SPDWY-S3DEV	Spartan-3 Developer Kit, Spartan-3 Evaluation Kit bundled with ISE BaseX Development System	$599.00 USD ($1,094.00 value)
ADS-SPDWY-V2P	Virtex-II Pro Evaluation Kit, with XC2VP7, -5 speed grade	$249.00 USD ($499.00 value)
ADS-SPDWY-V2P-DEV	Virtex-II Pro Developer Kit, Virtex-II Pro Evaluation Kit bundled with ISE Foundation Development System	$1,499.00 USD ($2,994.00 value)

*Pricing valid only within 60 days of attending a workshop.

SpeedWay Featured Kits

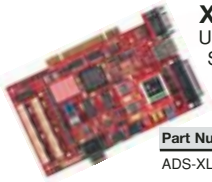

Xilinx Spartan™-3 Evaluation Kit

Use this kit to develop and test designs targeted to the Xilinx Spartan™-3 FPGA family. It has been optimized for the low-cost, consumer digital convergence market and can include a MicroBlaze Core License.

Part Number	Description	Resale Price
ADS-XLX-SP3-EVL400	Xilinx Spartan-3 Evaluation Kit, populated with an XC3S400 device	$399.00 USD ($249 for SpeedWay Attendees)
ADS-XLX-SP3-EVL1500	Xilinx Spartan-3 Evaluation Kit, populated with an XC3S1500 device	$499.00 USD
ADS-SP3-MB-EVL400	Xilinx Spartan-3 Evaluation Kit, populated with an XC3S400 device and bundled with MicroBlaze Core License	$799.00 USD
ADS-SP3-MB-EVL1500	Xilinx Spartan-3 Evaluation Kit, populated with an XC3S1500 device and bundled with MicroBlaze Core License	$899.00 USD

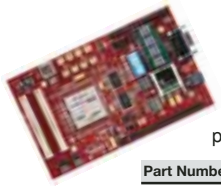

Xilinx Virtex-II Pro™ Evaluation Kit

This kit provides an environment for developing high-performance embedded processor and high-speed serial I/O intensive applications and is available with a Linux port and a Xilinx Aurora Protocol Engine design example.

Part Number	Description	Resale Price
ADS-XLX-V2PRO-EVLP7-5	Xilinx Virtex-II Pro Evaluation Kit, with XC2VP7, -5 speed grade	$499.00 USD ($249 for SpeedWay Attendees)
ADS-XLX-V2PRO-EVLP7-6	Xilinx Virtex-II Pro Evaluation Kit, with XC2VP7, -6 speed grade	$599.00 USD
ADS-XLX-V2PRO-LX-EVLP7-6	Xilinx Virtex-II Pro Evaluation Kit, with XC2VP7, -6 speed grade bundled with Linux, EDK, and Communications/Memory Module	$1,000.00 USD

SpeedWay Design Workshops™ Accelerate Time to Market

Running into roadblocks?

Avnet Cilicon's Speedway Design Workshops can help you break through.

Get hands-on training in new technologies.

Test-drive the latest hardware.

Simulate real-life development environments.

Solve problems before they cause the wheels to come off your new product design.

Bring your product idea to the finish line in record time.

To register for a SpeedWay Design Workshop™ in your area, visit www.em.avnet.com/speedway

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

Ready.
Set.
Go to market.™

AVNET® CILICON

Enabling success at the center of technology™

1 800 332 8638
www.em.avnet.com/semi

Supplier Authorized Distributor

XILINX®
The Programmable Logic Company℠

Xilinx Virtex™-II Series FPGAs

http://www.xilinx.com/devices/

Product Selection Matrix

				CLB Resources				Memory Resources			DSP	Clock Resources		I/O Features				Speed						
		Virtex-II Series EasyPath Solutions (see note 4)	System Gates (see note 1)	CLB Array (Row x Col)	Number of Slices	Logic Cells (see note 2)	CLB Flip-Flops	Max. Distributed RAM Bits (kbits)	# 18 kbits Block RAM	Total Block RAM (kbits)	# 18x18 Dedicated Multipliers	DCM Frequency (min/max)	# DCM Blocks (see note 3)	Digitally Controlled Impedance	Maximum Differential I/O Pairs	Maximum I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Configuration Memory (Bits)	RocketIO™ Transceiver Blocks (3.125 Gbps)	RocketIO™ X Transceiver Blocks (10.3125 Gbps)	PowerPC™ Processor Blocks	EasyPath
Virtex-II Pro Family – 1.5 Volt																								
Platform FPGAs	XC2VP2		*	16 x 22	1,408	3,168	2816	44	12	216	12	24/420	4	YES	100	204	LDT-25, LVDS-25, LVDSEXT-25, BLVDS-25, ULVDS-25, LVPECL-25, LVCMOS25, LVCMOS18, LVCMOS15, PCI33, LVTTL, LVCMOS33, PCI-X, PCI66, GTL, GTL+, HSTL I (1.5V,1.8V), HSTL II (1.5V,1.8V), HSTL III (1.5V,1.8V), HSTL IV (1.5V,1.8V), SSTL2I, SSTL2II, SSTL18 I, SSTL18 II	-5 -6 -7	-5 -6	1.3M	4		0	
	XC2VP4		*	40 x 22	3,008	6,768	6,016	94	28	504	28	24/420	4	YES	172	348		-5 -6 -7	-5 -6	3.0M	4		1	
	XC2VP7		*	40 x 34	4,928	11,088	9,856	154	44	792	44	24/420	4	YES	196	396		-5 -6 -7	-5 -6	4.4M	8		1	
	XC2VP20		*	56 x 46	9,280	20,880	18,560	290	88	1,584	88	24/420	8	YES	276	564		-5 -6 -7	-5 -6	8.2M	8		2	
	XC2VP30	XCE2VP30	*	80 x 46	13,696	30,816	27,392	428	136	2448	136	24/420	8	YES	372	644		-5 -6 -7	-5 -6	11.3M	8		2	
	XC2VP40	XCE2VP40	*	88 x 58	19,392	43,632	38,784	606	192	3,456	192	24/420	8	YES	396	804		-5 -6 -7	-5 -6	15.5M	0 or 12		2	
	XC2VP50	XCE2VP50	*	88 x 70	23,616	53,136	47,232	738	232	4,176	232	24/420	8	YES	420	852		-5 -6 -7	-5 -6	19.0M	0 or 16		2	
	XC2VP70	XCE2VP70	*	104 x 82	33,088	74,448	66,176	1,034	328	5,904	328	24/420	8	YES	492	996		-5 -6 -7	-5 -6	25.6M	16 or 20		2	
	XC2VP100	XCE2VP100	*	120 x 94	44,096	99,216	88,192	1,378	444	7,992	444	24/420	12	YES	572	1,164		-5 -6 -7	-5 -6	33.5M	0 or 20		2	
	XC2VP125	XCE2VP125	*	136 x 106	55,616	125,136	111,232	1,738	556	10,008	556	24/420	12	YES	644	1,200		-5 -6 -7	-5 -6	42.7M	0, 20 or 24		4	
Virtex-II Pro X Family – 1.5 Volt																								
	XC2VPX20		*	56 x 46	9,792	22,032	19,584	306	88	1584	88	24/420	8	YES	276	556	Same as above	-5 -6 -7	TBD	8.05M	0	8	1	
	XC2VPX70	XCE2VPX70	*	104 x 82	33,088	74,448	66,176	1034	308	5544	308	24/420	8	YES	492	996	Same as above	-5 -6 -7	TBD	25.60M	0	20	2	
Virtex-II Family – 1.5 Volt																								
	XC2V40		40K	8 x 8	256	576	512	8	4	72	4	24/420	4	YES	44	88	LDT-25, LVPECL-33, LVDS-33, LVDS-25, LVDSEXT-33, LVDSEXT-25, BLVDS-25, ULVDS-25, LVTTL, LVCMOS33, LVCMOS25, LVCMOS18, LVCMOS15, PCI33, PCI66, PCI-X, GTL, GTL+, HSTL I, HSTL II, HSTL III, HSTL IV, SSTL2I, SSTL2II, SSTL3 I, SSTL3 II, AGP, AGP-2X	-4 -5 -6	-4 -5	0.4M	–	–	–	–
	XC2V80		80K	16 x 8	512	1,152	1,024	16	8	144	8	24/420	4	YES	60	120		-4 -5 -6	-4 -5	0.6M	–	–	–	–
	XC2V250		250K	24 x 16	1,536	3,456	3,072	48	24	432	24	24/420	8	YES	100	200		-4 -5 -6	-4 -5	1.7M	–	–	–	–
	XC2V500		500K	32 x 24	3,072	6,912	6,144	96	32	576	32	24/420	8	YES	132	264		-4 -5 -6	-4 -5	2.8M	–	–	–	–
	XC2V1000		1M	40 x 32	5,120	11,520	10,240	160	40	720	40	24/420	8	YES	216	432		-4 -5 -6	-4 -5	4.1M	–	–	–	–
	XC2V1500		1.5M	48 x 40	7,680	17,280	15,360	240	48	864	48	24/420	8	YES	264	528		-4 -5 -6	-4 -5	5.7M	–	–	–	–
	XC2V2000		2M	56 x 48	10,752	24,192	21,504	336	56	1,008	56	24/420	8	YES	312	624		-4 -5 -6	-4 -5	7.5M	–	–	–	–
	XC2V3000	XCE2V3000	3M	64 x 56	14,336	32,256	28,672	448	96	1,728	96	24/420	12	YES	360	720		-4 -5 -6	-4 -5	10.5M	–	–	–	–
	XC2V4000	XCE2V4000	4M	80 x 72	23,040	51,840	46,080	720	120	2,160	120	24/420	12	YES	456	912		-4 -5 -6	-4 -5	15.7M	–	–	–	–
	XC2V6000	XCE2V6000	6M	96 x 88	33,792	76,032	67,584	1,056	144	2,592	144	24/420	12	YES	552	1,104		-4 -5 -6	-4 -5	21.9M	–	–	–	–
	XC2V8000	XCE2V8000	8M	112 x 104	46,592	104,832	93,184	1,456	168	3,024	168	24/420	12	YES	554	1,108		-4 -5	-4 -5	29.1M	–	–	–	–

* Logic cell counts are a more meaningful measurement of density for the Virtex-II Pro family since system gate count does not take into consideration the benefits of the immersed special blocks such as PowerPC processors and multi-gigabit transceivers.

Notes: 1. System Gates include 20-30% of CLBs used as RAM 2. Logic cell = One 4-Input Look Up Table (LUT) + Flip Flop + Carry Logic. 3. DCM = Digital Clock Management 4. Virtex-II Series EasyPath solution available to provide a no risk, no effort cost reduction path for volume production

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

XILINX®
The Programmable Logic Company℠

Xilinx Virtex™-II Series FPGAs

http://www.xilinx.com/devices/

Xilinx Virtex-II Series FPGAs and RocketPHY Physical Layer Transceivers

		Virtex-II Pro (1.5V)										Virtex-II Pro X (1.5V)		Virtex-II (1.5V)									
		XC2VP2	XC2VP4	XC2VP7	XC2VP20	XC2VP30	XC2VP40	XC2VP50	XC2VP70	XC2VP100	XC2VP125	XC2VPX20	XC2VPX70	XC2V40	XC2V80	XC2V250	XC2V500	XC2V1000	XC2V1500	XC2V2000	XC2V3000	XC2V4000	XC2V6000	XC2V8000
Pins[3]	Area	204	348	396	564	692	804	852	996	1164	1200	556	996	88	120	200	264	432	528	624	720	912	1104	1296
Chip Scale Packages (CS) – wire-bond chip-scale BGA (0.8 mm ball spacing)																								
144	12 x 12 mm													88	92	92								
BGA Packages (BG) – wire-bond standard BGA (1.27 mm ball spacing)																								
575	31 x 31 mm																	328	392	408				
728	35 x 35 mm																				516			
FGA Packages (FG) – wire-bond fine-pitch BGA (1.0 mm ball spacing)																								
256	17 x 17 mm	140	140											88	120	172	172	172						
456[4]	23 x 23 mm	156	248	248												200	264	324						
676[4]	27 x 27 mm				404	416	416												392	456	484			
FFA Packages (FF) – flip-chip fine-pitch BGA (1.0 mm ball spacing)																								
672	27 x 27 mm	204	348	396																				
896[5]	31 x 31 mm			396	556	556						556						432	528	624				
1152[5]	35 x 35 mm				564	564	692	692													720	824	824	824
1148[6]	35 x 35 mm						804	812																
1517	40 x 40 mm							852	964													912	1104	1108
1704	42.5 x 42.5 mm								996	1040	1040		996											
1696[6]	42.5 x 42.5 mm									1164	1200													
BFA Packages (BF) – flip-chip fine-pitch BGA (1.27 mm ball spacing)																								
957	40 x 40 mm																			624	684	684	684	

Notes: 1. Numbers in table indicate maximum number of user I/Os.
2. The number of I/Os for RocketIO MGTs are not included in this table.
3. Within the same family, all devices in a particular package are pin-out (footprint) compatible.
4. Virtex-II packages FG456 and FG676 are also footprint compatible.
5. Virtex-II packages FF896 and FF1152 are also footprint compatible.
6. RocketIO unavailable in this package.

Pb-free solutions are available for all packages. For more information contact your Xilinx sales representative or visit www.xilinx.com/pbfree.

	Transceiver Blocks											
	RocketIO (3.125Gbps)										RocketIO X (10Gbps)	
Package	XC2VP2	XC2VP4	XC2VP7	XC2VP20	XC2VP30	XC2VP40	XC2VP50	XC2VP70	XC2VP100	XC2VP125	XC2VPX20	XC2VPX70
FG256	4	4										
FG456	4	4	8									
FG676				8	8	8						
FG672	4	4	8									
FF896			8	8	8						8	
FF1152				8	8	12	16					
FF1148[6]						0	0					
FF1517							16	16				
FF1704								20	20	24		20
FF1696[6]									0	0		

	RocketPHY Family of 10Gbps Physical Layer Transceivers						
Device	SONET OC-192 SDH STM - 64	G . 709	10G ETHERNET	10G FIBRE CHANNEL	Applications	Parallel Interface	Package
XGC1120 - Ultra MSA	√	√	√	√	MSA Modules XFP Transceivers, SONET/SDH transmission systems, OTN system w/FEC, fiber optic test equipment	XSBI SFI-4	FT256
XGC1121 - 10G SONET/SDH	√				XFP Transceivers, SONET/SDH – based transmission systems, fiber optic test equipment	SFI-4	FT256
XGC1320 - 10GE/10GFC			√	√	XFP Transceivers, data transmission equipment, NICs, test equipment, edge routers, storage area networks	XSBI	FT256

For more information about the RocketPHY family, visit www.xilinx.com/rocketphy

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

XILINX®
The Programmable Logic Company℠

Xilinx Spartan™ Series FPGAs

http://www.xilinx.com/devices/

Product Selection Matrix

		CLB Resources				Memory Resources			DSP	CLK Resources				I/O Features				Speed		PROM
	System Gates (see note 1)	CLB Array (Row x Col)	Number of Slices	Logic Cells (see note 2)	CLB Flip-Flops	Max. Distributed RAM Bits	# Block RAM	Block RAM (bits)	Dedicated Multipliers	DCM Frequency (min/max)	# DCMs	Frequency Synthesis	Phase Shift	Digitally Controlled Impedance	Number of Differential I/O Pairs	Maximum I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Configuration Memory (Bits)
Spartan-3 Family – 1.2 Volt (see note 3)																				
XC3S50	50K	16 x 12	768	1,728	1,536	12K	4	72K	4	25/326	2	YES	YES	YES	56	124	Single-ended LVTTL, LVCMOS3.3/2.5/1.8/ 1.5/1.2, PCI 3.3V – 32/64-bit 33MHz, SSTL2 Class I & II, SSTL18 Class I, HSTL Class I, III, HSTL1.8 Class I, II & III, GTL, GTL+ Differential LVDS2.5, Bus LVDS2.5, Ultra LVDS2.5, LVDS_ext2.5, RSDS, LDT2.5, LVPECL	-4 -5	-4	.4M
XC3S200	200K	24 x 20	1,920	4,320	3,840	30K	12	216K	12	25/326	4	YES	YES	YES	76	173		-4 -5	-4	1.0M
XC3S400	400K	32 x 28	3,584	8,064	7,168	56K	16	288K	16	25/326	4	YES	YES	YES	116	264		-4 -5	-4	1.7M
XC3S1000	1000K	48 x 40	7,680	17,280	15,360	120K	24	432K	24	25/326	4	YES	YES	YES	175	391		-4 -5	-4	3.2M
XC3S1500	1500K	64 x 52	13,312	29,952	26,624	208K	32	576K	32	25/326	4	YES	YES	YES	221	487		-4 -5	-4	5.2M
XC3S2000	2000K	80 x 64	20,480	46,080	40,960	320K	40	720K	40	25/326	4	YES	YES	YES	270	565		-4 -5	-4	7.7M
XC3S4000	4000K	96 x 72	27,648	62,208	55,296	432K	96	1,728K	96	25/326	4	YES	YES	YES	312	712		-4 -5	-4	11.3M
XC3S5000	5000K	104 x 80	33,280	74,880	66,560	520K	104	1,872K	104	25/326	4	YES	YES	YES	344	784		-4 -5	-4	13.3M

Note: 1. System Gates include 20-30% of CLBs used as RAMs
2. For Spartan-3, a Logic Cell is defined as a 4-input LUT + flip-flop
3. Automotive Q-Grade Solutions for Spartan-3 will be available 2H2004.

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

XILINX®
The Programmable Logic Company℠

Xilinx Spartan™ Series FPGAs

http://www.xilinx.com/devices/

Product Selection Matrix

		CLB Resources				Memory Resources			CLK Resources				I/O Features			Speed				
		System Gates (see note 1)	CLB Array (Row x Col)	Number of Slices	Logic Cells (see notes 2 and 3)	CLB Flip-Flops	Max. Distributed RAM Bits	# Block RAM	Block RAM (bits)	DLL Frequency (min/max)	# DLLs	Frequency Synthesis	Phase Shift	Number of Differential I/O Pairs	Maximum I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grade (slowest to fastest)	Automotive Q-Grade Speed Grade	Configuration Memory (Bits)	Automotive Q-Grade Solutions (see note 4)
Spartan-IIE Family – 1.8 Volt																					
SPARTAN-IIE	XC2S50E	50K	16 x 24	768	1,728	1,536	24K	8	32K	25/320	4	YES	YES	83	182	LVTTL,LVCMOS2, LVCMOS18, PCI33, PCI66, GTL, GTL+, HSTL I, HSTL III, HSTL IV, SSTL3 I, SSTL3 II, SSTL2 I, SSTL2 II, AGP-2X, CTT, LVDS, BLVDS, LVPECL	-6 -7	-6	-6	0.6M	✔
	XC2S100E	100K	20 x 30	1,200	2,700	2,400	37K	10	40K	25/320	4	YES	YES	86	202		-6 -7	-6	-6	0.9M	✔
	XC2S150E	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	25/320	4	YES	YES	114	265		-6 -7	-6	-6	1.1M	✔
	XC2S200E	200K	28 x 42	2,352	5,292	4,704	73K	14	56K	25/320	4	YES	YES	120	289		-6 -7	-6	-6	1.4M	✔
	XC2S300E	300K	32 x 48	3,072	6,912	6,144	96K	16	64K	25/320	4	YES	YES	120	329		-6 -7	-6	-6	1.9M	✔
	XC2S400E	400K	40 x 60	4,800	10,800	9,600	150K	40	160K	25/320	4	YES	YES	172	410		-6 -7	-6	-6	2.7M	✔
	XC2S600E	600K	48 x 72	6,912	15,552	13,824	216K	72	288K	25/320	4	YES	YES	205	514		-6 -7	-6	-6	4.0M	✔
Spartan-II Family – 2.5 Volt																					
SPARTAN-II	XC2S15	15K	8 x 12	192	432	384	6K	4	16K	25/200	4	YES	YES	NA	86	LVTTL, LVCMOS2, PCI33 (3.3V & 5V), PCI66 (3.3V), GTL, GTL+, HSTL I, HSTL III, HSTL IV, SSTL3 I, SSTL3 II, SSTL2 I, SSTL2 II, AGP-2X, CTT	-5 -6	-5	-5	0.2M	✔
	XC2S30	30K	12 x 18	432	972	864	13.5K	6	24K	25/200	4	YES	YES	NA	132		-5 -6	-5	-5	0.4M	✔
	XC2S50	50K	16 x 24	768	1,728	1,536	24K	8	32K	25/200	4	YES	YES	NA	176		-5 -6	-5	-5	0.6M	✔
	XC2S100	100K	20 x 30	1,200	2,700	2,400	37.5K	10	40K	25/200	4	YES	YES	NA	196		-5 -6	-5	-5	0.8M	✔
	XC2S150	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	25/200	4	YES	YES	NA	260		-5 -6	-5	-5	1.1M	✔
	XC2S200	200K	28 x 42	2,352	5,292	4,704	73.5K	14	56K	25/200	4	YES	YES	NA	284		-5 -6	-5	-5	1.4M	✔
Spartan-XL Family – 3.3 Volt																					
SPARTAN-XL	XCS05XL	5K	10 x 10	100	238	200	3.1K	NA	NA	NA	NA	NA	NA	NA	77		-4 -5	-4	-4	0.05M	✔
	XCS10XL	10K	14 x 14	196	466	392	6.1K	NA	NA	NA	NA	NA	NA	NA	112		-4 -5	-4	-4	0.09M	✔
	XCS20XL	20K	20 x 20	400	950	800	12.5K	NA	NA	NA	NA	NA	NA	NA	160		-4 -5	-4	-4	0.18M	✔
	XCS30XL	30K	24 x 24	576	1,368	1,152	18.0K	NA	NA	NA	NA	NA	NA	NA	192		-4 -5	-4	-4	0.25M	✔
	XCS40XL	40K	28 x 28	784	1,862	1,568	24.5K	NA	NA	NA	NA	NA	NA	NA	224		-4 -5	-4	-4	0.33M	✔

Note: 1. System Gates include 20-30% of CLBs used as RAM
2. Logic cell = (1) 4 Input (LUT) and a register
3. For Spartan-IIE/II/XL, a Logic Cell is defined as a 4-input LUT + a register
4. Automotive Q-Grade Solutions are qualified to -40°C to +125°C junction temperature for FPGAs. Q-Grade products for Spartan-3 will be available 2H2004

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

XILINX®
The Programmable Logic Company℠

Xilinx Spartan™ Series FPGAs

http://www.xilinx.com/devices/

Package Options and User I/O[1]

		Spartan-3 (1.2V)									Spartan-IIE (1.8V)								Spartan-II (2.5V)						Spartan-XL (3.3V)				
			XC3S50	XC3S200	XC3S400	XC3S1000	XC3S1500	XC3S2000	XC3S4000	XC3S5000		XC2S50E	XC2S100E	XC2S150E	XC2S200E	XC2S300E	XC2S400E	XC2S600E	XC2S15	XC2S30	XC2S50	XC2S100	XC2S150	XC2S200	XCS05XL	XCS10XL	XCS20XL	XCS30XL	XCS40XL
Pins	Area[2]	I/Os	124	173	264	391	487	565	712	784	I/Os	182	202	265	289	329	410	514	86	132	176	196	260	284	77	112	160	192	224
PLCC Packages (PC) – wire-bond plastic chip carrier (1.27mm lead spacing)																													
84	30.2 x 30.2 mm																								61	61			
PQFP Packages (PQ) – wire-bond plastic QFP (0.5mm lead spacing)																													
208	30.6 x 30.6 mm		124	141	141							146	146	146	146	146				132	140	140	140	140			160	169	169
240	34.6 x 34.6 mm																											192	192
VQFP Packages (VQ) – very thin TQFP (0.5mm lead spacing)																													
100	16.0 x 16.0 mm		63	63															60	60					77	77	77	77	
TQFP Packages (TQ) – thin QFP (0.5mm lead spacing)																													
144	22.0 x 22.0 mm		97	97	97							102	102						86	92	92	92				112	113	113	
Chip Scale Packages (CS) – wire-bond chip-scale BGA (0.8 mm ball spacing)																													
144	12 x 12 mm																		86	92						112	113		
280	16 x 16 mm																											192	224
FGA Packages (FT) – wire-bond fine-pitch thin BGA (1.0 mm ball spacing)																													
256	17 x 17 mm			173	173	173						182	182	182	182	182	182												
FGA Packages (FG) – wire-bond fine-pitch BGA (1.0 mm ball spacing)																													
256	17 x 17 mm																				176	176	176	176					
320	23 x 23 mm				221	221	221																						
456	23 x 23 mm				264	333	333						202	265	289	329	329	329				196	260	284					
676	27 x 27 mm					391	487	489									410	514											
900	31 x 31 mm							565	633	633																			
1156	35 x 35 mm								712	784																			
BGA Packages (BG) – wire-bond standard BGA (1.27 mm ball spacing)																													
256	27 x 27 mm																											192	205

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

Automotive products are highlighted: -40C to +125C junction temperature for FPGAs.
Automotive Q-Grade products will be available in Spartan-3 family 2H2004.
For more information on IQ Solutions please visit **http://www.xilinx.com/automotive**

Xilinx IQ Solutions for Automotive Intelligence

Pb-free packaging available for all devices. For more information visit www.xilinx.com/pbfree

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

Product Selection Matrix – CoolRunner™ Series

		System Gates	Macrocells	Product Terms per Macrocell	Input Voltage Compatible	Output Voltage Compatible	I/O Features: Maximum I/O	I/O Features: I/O Banking	Speed: Min. Pin-to-pin Logic Delay (ns)	Speed: Commercial Speed Grades (fastest to slowest)	Speed: Industrial Speed Grades (fastest to slowest)	Speed: IQ Speed Grade	Clocking: Global Clocks	Clocking: Product Term Clocks per Function Block
	CoolRunner-II Family – 1.8 Volt													
CoolRunner-II	XC2C32	750	32	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	33	1	3	-3 -4 -6	-6	-6	3	17
	XC2C64	1500	64	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	64	1	4	-4 -5 -7	-7	-7	3	17
	XC2C128	3000	128	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	100	2	4.5	-4 -6 -7	-7	-7	3	17
	XC2C256	6000	256	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	184	2	5	-5 -6 -7	-7	-7	3	17
	XC2C384	9000	384	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	240	4	6	-6 -7 -10	-10	-10	3	17
	XC2C512	12000	512	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	270	4	6	-6 -7 -10	-10	-10	3	17
	CoolRunner XPLA3 Family – 3.3 Volt													
CoolRunner XPLA3	XCR3032XL	750	32	48	3.3/5	3.3	36		5	-5 -7 -10	-7 -10	-10	4	16
	XCR3064XL	1500	64	48	3.3/5	3.3	68		6	-6 -7 -10	-7 -10	-10	4	16
	XCR3128XL	3000	128	48	3.3/5	3.3	108		6	-6 -7 -10	-7 -10	-10	4	16
	XCR3256XL	6000	256	48	3.3/5	3.3	164		7.5	-7 -10 -12	-10 -12	-12	4	16
	XCR3384XL	9000	384	48	3.3/5	3.3	220		7.5	-7 -10 -12	-10 -12	-12	4	16
	XCR3512XL	12000	512	48	3.3/5	3.3	260		7.5	-7 -10 -12	-10 -12	-12	4	16

Package Options and User I/O

		CoolRunner-II						CoolRunner XPLA3					
Pins	Area[1]	XC2C32	XC2C64	XC2C128	XC2C256	XC2C384	XC2C512	XCR3032XL	XCR3064XL	XCR3128XL	XCR3256XL	XCR3384XL	XCR3512XL
PLCC Packages (PC) – wire-bond plastic chip carrier (1.27mm lead spacing)													
44	17.5 x 17.5 mm	33	33					36	36				
PQFP Packages (PQ) – wire-bond plastic QFP (0.5mm lead spacing)													
208	30.6 x 30.6 mm				173	173	173				164	172	180
VQFP Packages (VQ) – very thin TQFP (0.5mm lead spacing)													
44	12.0 x 12.0 mm	33	33					36	36				
100	16.0 x 16.0 mm		64	80	80				68	84			
TQFP Packages (TQ) – thin QFP (0.5mm lead spacing)													
144	22.0 x 22.0 mm			100	118	118				108	120	118*	
Chip Scale Packages (CP) – wire-bond chip-scale BGA (0.5 mm ball spacing)													
56	6 x 6 mm	33	45						48				
132	8 x 8 mm			100	106								
Chip Scale Packages (CS) – wire-bond chip-scale BGA (0.8 mm ball spacing)													
48	7 x 7 mm							36	40				
144	12 x 12 mm									108			
280	16 x 16 mm										164		
FGA Packages (FT) – wire-bond fine-pitch thin BGA (1.0 mm ball spacing)													
256	17 x 17 mm				184	212	212				164	212	212
FBGA Packages (FG) – wire-bond Fine-line BGA (1.0 mm ball spacing)													
324	23 x 23 mm					240	270					220	260

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

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

Automotive products are highlighted: -40C to +125C ambient temperature for CPLDs

Xilinx IQ Solutions for Automotive Intelligence

Product Selection Matrix – 9500 Series

	System Gates	Macrocells	Product Terms per Macrocell	Input Voltage Compatible	Output Voltage Compatible	I/O Features: Maximum I/O	I/O Features: I/O Banking	Speed: Min. Pin-to-pin Logic Delay (ns)	Speed: Commercial Speed Grades (fastest to slowest)	Speed: Industrial Speed Grades (fastest to slowest)	Speed: IQ Speed Grade	Clocking: Global Clocks	Clocking: Product Term Clocks per Function Block
XC9500XV Family – 2.5 Volt													
XC9536XV	800	36	90	2.5/3.3	1.8/2.5/3.3	36	1	5	-5 -7	-7	NA	3	18
XC9572XV	1600	72	90	2.5/3.3	1.8/2.5/3.3	72	1	5	-5 -7	-7	NA	3	18
XC95144XV	3200	144	90	2.5/3.3	1.8/2.5/3.3	117	2	5	-5 -7	-7	NA	3	18
XC95288XV	6400	288	90	2.5/3.3	1.8/2.5/3.3	192	4	6	-6 -7 -10	-7 -10	NA	3	18
XC9500XL Family – 3.3 Volt													
XC9536XL	800	36	90	2.5/3.3/5	2.5/3.3	36		5	-5 -7 -10	-7 -10	-10	3	18
XC9572XL	1600	72	90	2.5/3.3/5	2.5/3.3	72		5	-5 -7 -10	-7 -10	-10	3	18
XC95144XL	3200	144	90	2.5/3.3/5	2.5/3.3	117		5	-5 -7 -10	-7 -10	NA	3	18
XC95288XL	6400	288	90	2.5/3.3/5	2.5/3.3	192		6	-6 -7 -10	-7 -10	NA	3	18
XC9500 Family – 5 Volt													
XC9536	800	36	90	5	5	36		10	-5 -6 -10 -15	-7 -10 -15	-15	3	18
XC9572	1600	72	90	5	5	72		10	-7 -10 -15	-10 -15	-15	3	18
XC95108	2400	108	90	5	5	108		10	-7 -10 -15 -20	-7 -10 -15 -20	NA	3	18
XC95144	3200	144	90	5	5	133		10	-7 -10 -15	-10 -15	NA	3	18
XC95216	4800	216	90	5	5	166		10	-10 -15 -20	-10 -15 -20	NA	3	18
XC95288	6400	288	90	5	5	192		10	-10 -15 -20	-15 -20	NA	3	18

Pb-free solutions available for all packages. For more information contact your Xilinx sales representative or visit www.xilinx.com/pbfree.

Package Options and User I/O

		XC9500XV				XC9500XL				XC9500					
Pins	Area[1]	XC9536XV	XC9572XV	XC95144XV	XC95288XV	XC9536XL	XC9572XL	XC95144XL	XC95288XL	XC9536	XC9572	XC95108	XC95144	XC95216	XC95288
PLCC Packages (PC) – wire-bond plastic chip carrier (1.27mm lead spacing)															
44	17.5 x 17.5 mm	34	34			34	34			34	34				
84	30.2 x 30.2 mm										69	69			
PQFP Packages (PQ) – wire-bond plastic QFP (0.5mm lead spacing)															
100	23.3 x 17.2 mm										72	81	81		
160	31.2 x 31.2 mm											108	133	133	168
208	30.6 x 30.6 mm				168				168					166	168
VQFP Packages (VQ) – very thin TQFP (0.5mm lead spacing)															
44	12.0 x 12.0 mm	34	34			34	34			34					
64	12.0 x 12.0 mm					36	52								
TQFP Packages (TQ) – thin QFP (0.5mm lead spacing)															
100	16.0 x 16.0 mm		72	81			72	81			72	81	81		
144	22.0 x 22.0 mm			117	117			117	117						
Chip Scale Packages (CP) – wire-bond chip-scale BGA (0.5 mm ball spacing)															
56	6 x 6 mm														
132	8 x 8 mm														
Chip Scale Packages (CS) – wire-bond chip-scale BGA (0.8 mm ball spacing)															
48	7 x 7 mm	36	38			36	38			34					
144	12 x 12 mm			117				117							
280	16 x 16 mm				192				192						
BGA Packages (BG) – wire-bond standard BGA (1.27 mm ball spacing)															
256	27 x 27 mm								192						
352	35.0 x 35.0 mm								192					166	192
FGA Packages (FT) – wire-bond fine-pitch thin BGA (1.0 mm ball spacing)															
256	17 x 17 mm														
FBGA Packages (FG) – wire-bond Fine-line BGA (1.0 mm ball spacing)															
256	17 x 17 mm				192				192						
324	23 x 23 mm														

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

Automotive products are highlighted: -40C to +125C ambient temperature for CPLDs

Xilinx IQ Solutions for Automotive Intelligence

XILINX®
The Programmable Logic Company℠

Xilinx Defense and Aerospace

http://www.xilinx.com/devices

QPRO

Device	SMD	Voltage	System Gates	Logic Cells	RAMbus	CFG Bits	MULTs	DLLs	Flip-Flops	Max I/Os	Packages
XC3020*	5962-89948	5	1500	256	–	14779	–	–	256	64	PG84, CB100
XC3030*	None	5	2000	360	–	22176	–	–	360	80	PG84
XC3042*	5962-89713	5	3000	480	–	30784	–	–	480	96	PG84, PG132, CB100
XC3064*	None	5	4500	688	–	46064	–	–	688	120	PG132
XC3090*	5962-89823	5	6000	928	–	64160	–	–	928	144	PG175, CB164
XC4005*	5962-92252	5	9K	466	6272	151960	–	–	616	112	PG156, CB164
XC4010*	5962-92305	5	20K	950	12800	283424	–	–	1120	160	PG191, CB196
XC4013*	5962-94730	5	30K	1368	18438	393632	–	–	1536	192	PG223, CB228
XQ4005E	5962-97522	5	9K	466	6272	95008	–	–	616	112	PG156, CB164
XQ4010E	5962-97523	5	20K	950	12800	178144	–	–	1120	160	HQ208, PG191, CB196
XQ4013E	5962-97524	5	30K	1368	18438	247968	–	–	1536	192	HQ240, PG223, CB228
XQ4025E	5962-97525	5	45K	2432	32768	422176	–	–	2560	256	PG299, CB228
XQ4028EX	5962-98509	5	18K-50K	2432	32768	668184	–	–	2560	256	HQ240, BG352, PG299, CB228
XQ4013XL	5962-98513	3.3	10-30K	1368	18K	393632	–	–	1536	192	PQ240, BG256, PG223, CB228
XQ4036XL	5962-98510	3.3	22-65K	3078	42K	832528	–	–	3168	288	HQ240, BG352, PG411, CB228
XQ4062XL	5962-98511	3.3	40-130K	5472	74K	1433864	–	–	5376	384	HQ240, BG432, PG475, CB228
XQ4085XL	None	3.3	55-180K	7448	100K	1924992	–	–	7185	448	HQ240, BG432, CB228
XQV100	None	2.5	108904	2700	40K	781248	–	4	2400	180	PQ240, BG256, CB228
XQV300	5962-99572	2.5	322970	6912	64K	1751840	–	4	6144	316	PQ240, BG352, BG432, CB228
XQV600	5962-99573	2.5	661111	15552	96K	3608000	–	4	13824	512	HQ240, BG432, CB228
XQV1000	5962-99574	2.5	1124022	27648	128K	6127776	–	4	24576	512	BG560, CG560
XQV600E	None	1.8	986K	15552	288K	3961632	–	8	13824	512	BG432, CB228
XQV1000E	None	1.8	1569K	27648	384K	6587520	–	8	24576	660	BG560, CG560
XQV2000E	None	1.8	2542K	43200	640K	10159648	–	8	38400	804	BG560, FG1156
XQ2V1000	TBD	1.5	1000K	11520	720	4082592	40	8	10240	432	FG456, BG575
XQ2V3000	TBD	1.5	3000K	32256	1728	10494368	96	12	28672	720	BG728, CG717
XQ2V6000	TBD	1.5	6000K	76032	2592	21849504	144	12	67584	1104	CF1144

* Not for new designs

QPRO Radiation Tolerant

Device	Voltage	System Gates	Logic Cells	RAMbus	CFG Bits	MULTs	DLLs	Flip-Flops	Max I/Os	Packages	Total Ionizing Dose (TID) in KRADs
XQR4013XL	3.3	10-30K	1368	18K	393632	–	–	1536	192	CB228	60
XQR4036XL	3.3	22-65K	3078	42K	832528	–	–	3168	288	CB228	60
XQR4062XL	3.3	40-130K	5472	74K	1433864	–	–	5376	384	CB228	60
XQVR300	2.5	322970	6912	64K	1751840	–	4	6144	316	CB228	100
XQVR600	2.5	661111	15552	96K	3608000	–	4	13824	512	CB228	100
XQVR1000	2.5	1124022	27648	128K	6127776	–	4	24576	512	CG560	100
XQR2V1000	1.5	1000K	11520	720	4082592	40	8	10240	432	BG575	200
XQR2V3000	1.5	3000K	32256	1728	10494368	96	12	28672	720	CG717	200
XQR2V6000	1.5	6000K	76032	2592	21849504	144	12	67584	1104	CF1144	200

QPRO Configuration PROMs

Device	SMD	Voltage	System Gates	Logic Cells	RAMbus	CFG Bits	MULTs	DLLs	Flip-Flops	Max I/Os	Packages	Total Ionizing Dose (TID) in KRADs
XC1736D	None	5	–	–	–	32768	–	–	–	–	DD8	–
XC1765D	5962-94717	5	–	–	–	65536	–	–	–	–	DD8	–
XC17128D	None	5	–	–	–	131072	–	–	–	–	DD8	–
XC17256D	5962-95617	5	–	–	–	262144	–	–	–	–	DD8	–
XQ1701L	5962-99514	5	–	–	–	1048576	–	–	–	–	SO20, CC44	–
XQ17V16	TBD	3.3	–	–	–	16777216	–	–	–	–	VQ44, CC44	–
XQR1701L	–	3.3	–	–	–	1048576	–	–	–	–	SO20, CC44	60
XQR18V04	–	3.3	–	–	–	4194304	–	–	–	–	VQ44, CC44	40
XQR17V16**	–	3.3	–	–	–	16777216	–	–	–	–	VQ44, CC44	TBD

** Under development

Device Nomenclatures
XC = Qualified Prior to QML Certification

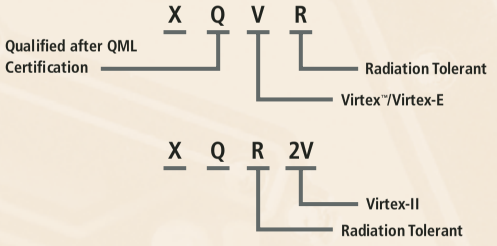

XILINX®
The Programmable Logic Company℠

Xilinx Configuration Storage Solutions

http://www.xilinx.com/configsolns

Product Selection and Package Option Matrix

Platform Flash Device Cross Reference

Platform Flash	PROM Solution
Spartan-3	
XC3S50	XCF01S
XC3S200	XCF01S
XC3S400	XCF02S
XC3S1000	XCF04S
XC3S1500	XCF08P
XC3S2000	XCF08P
XC3S4000	XCF16P
XC3S5000	XCF16P
Spartan-IIE	
XC2S50E	XCF01S
XC2S100E	XCF01S
XC2S150E	XCF02S
XC2S200E	XCF02S
XC2S300E	XCF02S
XC2S400E	XCF04S
XC2S600E	XCF04S
Spartan-II	
XC2S15	XCF01S
XC2S30	XCF01S
XC2S50	XCF01S
XC2S100	XCF01S
XC2S150	XCF01S
XC2S200	XCF02S
Spartan-XL	
XCS05XL	XCF01S
XCS10XL	XCF01S
XCS20XL	XCF01S
XCS30XL	XCF01S
XCS40XL	XCF01S

Platform Flash	PROM Solution
Virtex-II Pro	
XC2VP2	XCF02S
XC2VP4	XCF04S
XC2VP7	XCF08P
XC2VP20	XCF08P
XC2VP30	XCF16P
XC2VP40	XCF16P
XC2VP50	XCF32P
XC2VP70	XCF32P
XC2VP100	XCF32P*
XC2VP125	XCF32P**
Virtex-II	
XC2V40	XCF01S
XC2V80	XCF01S
XC2V250	XCF02S
XC2V500	XCF04S
XC2V1000	XCF04S
XC2V1500	XCF08P
XC2V2000	XCF08P
XC2V3000	XCF16P
XC2V4000	XCF16P
XC2V6000	XCF32P
XC2V8000	XCF32P

* assumes bitstream compression is used (XCF32P + XCF01S without compression).
** assumes bitstream compression is used (XCF32P + XCF16P without compression).

Note: For information regarding legacy PROMs, visit **http://www.xilinx.com/legacyproms**

Platform Flash Family Features and Packages

	XCF01S	XCF02S	XCF04S	XCF08P	XCF16P	XCF32P
Density	1Mb	2Mb	4Mb	8Mb	16Mb	32Mb
JTAG Prog	Y	Y	Y	Y	Y	Y
Serial Configuration	Y	Y	Y	Y	Y	Y
SelectMap Configuration				Y	Y	Y
Compression				Y	Y	Y
Design Rev				Y	Y	Y
VCC (V)	3.3	3.3	3.3	1.8	1.8	1.8
VCCO (V)	1.8 – 3.3	1.8 – 3.3	1.8 – 3.3	1.5 – 3.3	1.5 – 3.3	1.5 – 3.3
VCCJ (V)	1.8 – 3.3	1.8 – 3.3	1.8 – 3.3	1.5 – 3.3	1.5 – 3.3	1.5 – 3.3
Clock (MHz)	33	33	33	40	40	40
Package	VO20	VO20	VO20	FS48	FS48	FS48
				VO48	VO48	VO48

For multiple FPGA Configuration and for designs utilizing system level features, use SystemACE™.

SystemACE	Memory Density	Number of Components	Min Board Space	Compression	FPGA Config. Mode	Multiple Designs	Software Storage	Removable	IRL Hooks	Max Config. Speed	Non-Volatile Media
SystemACE CF	Up to 8 Gbit	2	25 cm²	No	JTAG	Unlimited	Yes	Yes	Yes	30 Mbit/sec	CompactFlash
SystemACE SC	16 Mbit 32 Mbit 64 Mbit	3	Custom	Yes	SelectMAP (up to 4 FPGA) Slave-Serial (up to 8 FPGA chains)	Up to 8	No	No	Yes	152 Mbit/sec	AMD Flash Memory

Pb-free packaging available for all Platform Flash devices. For more information visit www.xilinx.com/pbfree.

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

XILINX®
The Programmable Logic Company℠

Xilinx Packaging

http://www.xilinx.com/packaging

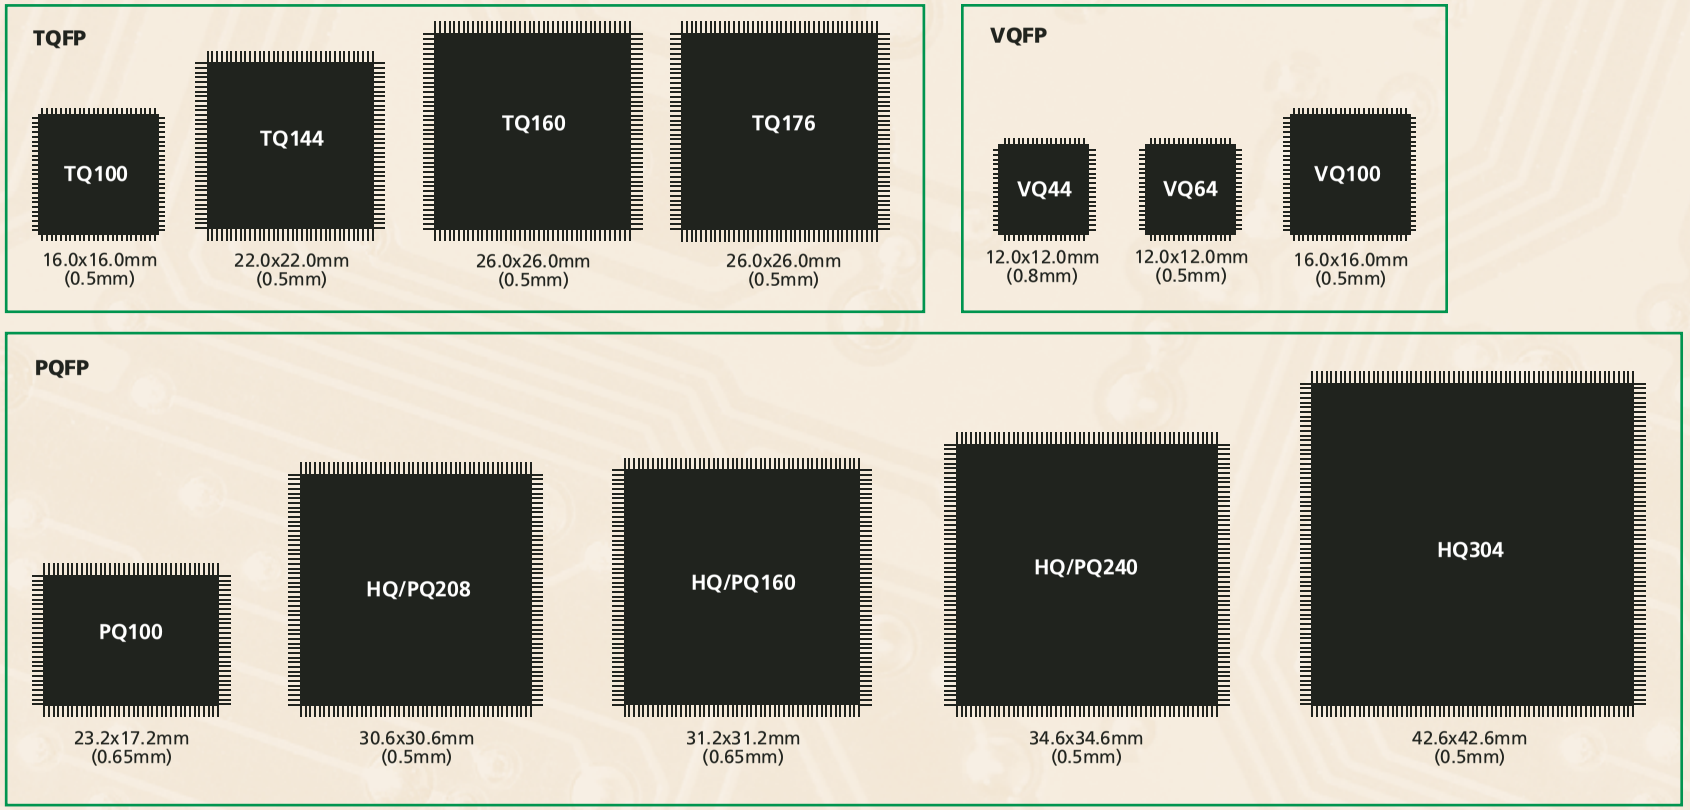

Note: 1. Package outlines shown are actual size.
2. For lead-frame packages, dimensions (D & E) shown are inclusive of leads. Dimensions in parenthesis represent package pitch.

XILINX®
The Programmable Logic Company℠

Xilinx Packaging

http://www.xilinx.com/packaging

PLASTIC OVERMOLDED BGA (CAVITY UP)

17.0x17.0mm (1.0mm) | 17.0x17.0mm (1.0mm) | 23.0x23.0mm (1.0mm) | 23.0x23.0mm (1.0mm) | 27.0x27.0mm (1.0mm) | 27.0x27.0mm (1.27mm) | 31.0x31.0mm (1.27mm) | 31.0x31.0mm (1.0mm)

PLASTIC OVERMOLDED BGA (CAVITY UP) CONTINUED

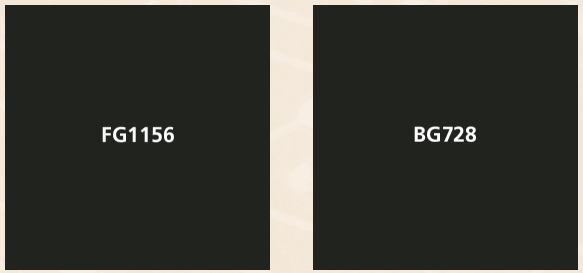

35.0x35.0mm (1.0mm) | 35.0x35.0mm (1.27mm)

METAL BGA (CAVITY DOWN)

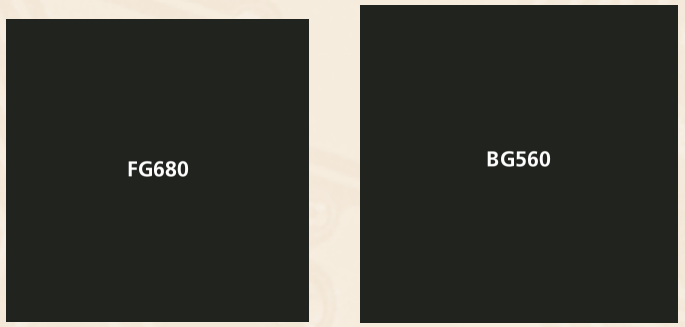

40.0x40.0mm (1.0mm) | 42.5x42.5mm (1.27mm)

Note: 1. Package outlines shown are actual size.
2. Dimesions referenced in parenthesis represent package pitch.

FLIP-CHIP BGA

27.0x27.0mm (1.0mm) | 31.0x31.0mm (1.0mm) | 35.0x35.0mm (1.0mm) | 40.0x40.0mm (1.0mm) | 40.0x40.0mm (1.27mm) | 42.5x42.5mm (1.0mm)

XILINX®
The Programmable Logic Company℠

Xilinx Software – ISE 6.2i

http://www.xilinx.com/ise

	Feature	*ISE WebPACK™*	*ISE BaseX*	*ISE Foundation*	*ISE Alliance*
Devices	Virtex™ Series	**Virtex-E:** V50E -V300E **Virtex-II:** 2V40 - 2V250 **Virtex-II Pro:** 2VP2	**Virtex:** V50 - V600 **Virtex-E:** V50E - V600E **Virtex-II:** 2V40 - 2V500 **Virtex-II Pro:** 2VP2, 2VP4, 2VP7	ALL	ALL
	Spartan™ Series	**SpartanII/IIE:** ALL (except XC2S400E and XC2S600E) **Spartan-3:** 3S50, 3S200, 3S400	**Spartan-II/IIE:** ALL **Spartan-3:** 3S50, 3S200, 3S400	**Spartan-II/IIE:** ALL **Spartan-3:** ALL	**Spartan-II/IIE:** ALL **Spartan-3:** ALL
	CoolRunner™ XPLA3 CoolRunner-II	ALL	ALL	ALL	ALL
	XC9500 Series	ALL	ALL	ALL	ALL
Design Entry	Schematic Editor	Yes	MS Windows and Linux only	MS Windows and Linux only	No
	HDL Editor	Yes	Yes	Yes	Yes
	State Diagram Editor	Yes	MS Windows only	MS Windows only	MS Windows only
	CORE Generator System	No	Yes	Yes	Yes
	PACE (Pinout and Area Constraint Editor)	Yes	Yes	Yes	Yes
	Architecture Wizards DCM - Digital Clock Management MGT - Multi-Gigabit Transcievers	Yes	Yes	Yes	Yes
	3rd Party RTL Checker Support	Yes	Yes	Yes	Yes
	Xilinx System Generator for DSP	No	Sold as an Option	Sold as an Option	Sold as an Option
Embedded System Design	GNU Embedded Tools GCC - GNU Compiler GDB - GNU Software Debugger	No	Yes (Available with optional EDK)	Yes (Available with optional EDK)	Yes (Available with optional EDK)
	WindRiver Xilinx Edition Development Tools Diab C/C++ Compiler SingleStep Debugger visionPROBE II target connection	No	Sold as an Option	Sold as an Option	Sold as an Option
Synthesis	Xilinx Synthesis Technology (XST)	Yes	Yes	Yes	No
	Synplicity Synplify/Pro	Integrated Interface	Integrated Interface (MS Windows only)	Integrated Interface (MS Windows only)	Integrated Interface (MS Windows only)
	Synplicity Amplify Physical Synthesis Support	Yes	Yes	Yes	Yes
	Mentor Graphics Leonardo Spectrum	Integrated Interface	Integrated Interface	Integrated Interface	Integrated Interface
	Mentor Graphics Precision RTL	EDIF only	EDIF only	EDIF only	EDIF only
	Synopsys FPGA Compiler II	EDIF Interface	EDIF Interface	EDIF Interface	EDIF Interface
	ABEL	CPLD	CPLD (MS Windows only)	CPLD (MS Windows only)	CPLD (MS Windows only)

XILINX®
The Programmable Logic Company℠

Xilinx Software – ISE 6.2i

http://www.xilinx.com/ise

	Feature	*ISE WebPACK™*	*ISE BaseX*	*ISE Foundation*	*ISE Alliance*
Implementation	FloorPlanner	Yes	Yes	Yes	Yes
	Constraints Editor	Yes	Yes	Yes	Yes
	Timing Driven Place & Route	Yes	Yes	Yes	Yes
	Modular Design	No	Yes	Yes	Yes
	Timing Improvement Wizard	Yes	Yes	Yes	Yes
Programming	iMPACT	Yes	Yes	Yes	Yes
	System ACE Configuration Manager	Yes	Yes	Yes	Yes
Board Level Integration	IBIS Models	Yes	Yes	Yes	Yes
	STAMP Models	Yes	Yes	Yes	Yes
	HSPICE Models*	Yes	Yes	Yes	Yes
Verification	HDL Bencher™	Yes	MS Windows only	MS Windows only	MS Windows only
	ModelSim® Xilinx Edition (MXE II)	ModelSim XE II Starter**	ModelSim XE II Starter**	ModelSim XE II Starter**	ModelSim XE II Starter**
	Static Timing Analyzer	Yes	Yes	Yes	Yes
	ChipScope™ Pro	Sold as an Option	Sold as an Option	Sold as an Option	Sold as an Option
	FPGA Editor with Probe	No	Yes	Yes	Yes
	ChipViewer	Yes	Yes	Yes	Yes
	XPower (Power Analysis)	Yes	Yes	Yes	Yes
	3rd Party Equivalency Checking Support	Yes	Yes	Yes	Yes
	SMARTModels for PPC and RocketIO	No	Yes	Yes	Yes
	3rd Party Simulator Support	Yes	Yes	Yes	Yes
Platforms		PC (MS Windows 2000/MS Windows XP)	PC (MS Windows 2000/MS Windows XP), Linux	PC (MS Windows 2000/MS Windows XP), Sun Solaris, Linux	PC(MS Windows 2000/MS Windows XP), Sun Solaris, Linux
IP/CORE	For more information on the complete list of Xilinx IP products, visit the Xilinx IP Center at http://www.xilinx.com/ipcenter				

*HSPICE Models available at the Xilinx Design Tools Center at www.xilinx.com/ise.
**MXE II supports the simulation of designs up to 1 million system gates and is sold as an option.

For more information, visit the Xilinx Design Tools Center at www.xilinx.com/ise

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
Spartan-3				
ADS-S3-MB-EVL400	Spartan-3 Evaluation Kit w/MicroBlaze and Communications/ Memory Module	Avnet Design Services	XC3S400	Networking, telecommunication, data communication, embedded and consumer markets
ADS-XLX-SP3-EVL1500	Spartan-3 Evaluation Kit with 3S1500 device	Avnet Design Services	XC3S1500-4FG676	Very cost effective Spartan-3 evaluation platform; allows experimentation with the advanced features of the Spartan-3 1500 device
ADS-XLX-SP3-EVL400	Spartan-3 Evaluation Kit with 3S400 device	Avnet Design Services	XC3S400	Very cost effective Spartan-3 evaluation platform; allows experimentation with the advanced features of the Spartan-3 400 device
DS-KIT-3SLC400*	Spartan-3 LC Development Kit	Insight (Memec)	XC3S400-4PQ208	General purpose Spartan-3 development platform
DS-KIT-3SLC400-BAS	Spartan-3 LC Development Kit w/ISE Foundation and JTAG Cable	Insight (Memec)	XC3S400-4PQ208	General purpose Spartan-3 development platform
DS-KIT-3SLC400-PAK*	Spartan-3 LC Development Kit w/ WebPACK CD and JTAG Cable	Insight (Memec)	XC3S400-4PQ208	General purpose Spartan-3 development platform
DS-KIT-3SMB1500*	Spartan-3 MB 3S1500 Development Kit	Insight (Memec)	XC3S1500-4FG676	General purpose Spartan-3 development platform
DS-KIT-3SMB1500-ISE	Spartan-3 MB 3S1500 Development Kit w/ISE Foundation and JTAG Cable	Insight (Memec)	XC3S1500-4FG676	General purpose Spartan-3 development platform
DS-KIT-MB-3SLC400	Spartan-3 MicroBlaze Development Kit	Insight (Memec)	XC3S400-4FG676	Embedded microprocessor
DS-KIT-MB-3SMB1500	Spartan-3 MicroBlaze Development Kit	Insight (Memec)	XC3S1500-4FG676	Embedded microprocessor
HW-AFX-SP3-1500-DB	NuHo 3S1500 Development Board	NuHorizons	XC3S1500	Prototyping, MicroBlaze Soft Processor Development, DSP, Industrial Systems, Data Communications / Telecommunications
HW-AFX-SP3-400-DB	NuHo 3S400 Development Board	NuHorizons	XC3S400-4PQ208C	Prototyping, MicroBlaze Soft Processor Development, DSP, Industrial Systems, Data Communications / Telecommunications
Spartan-IIE				
ADS-S2E-US2-SOL	Xilinx & Cypress USB 2.0 to SCSI System Solution Kit	Avnet Design Services	XC2S300E-PQ208	Complete solution for interfacing a SCSI drive to a host computer using Universal Serial Bus Specification Revision 2.0 - provides the designer with a Windows ready USB 2.0 to SCSI demonstration design
ADS-SP2E-MB-EVL	Spartan-IIE Evaluation Kit w/MicroBlaze & Communications/ Memory Module	Avnet Design Services	XC2S200E	Prototyping, Digital Video, Multimedia, Telecom/Datacom
ADS-XLX-SP2E-EVL	Spartan-IIE Evaluation Kit	Avnet Design Services	XC2S200E-6FT256C	Prototyping, Telecom/Datacom
DS-KIT-MB-S2E3LC	Spartan-IIE MicroBlaze Kit	Insight (Memec)	Spartan-IIE	Embedded microprocessor
DS-KIT-MB-S2E6LC	Spartan-IIE MicroBlaze Kit	Insight (Memec)	Spartan-IIE	Embedded microprocessor
DS-KIT-S2E3LC*	Spartan-IIE LC 2S300E Development Kit	Insight (Memec)	XC2S300E-6FG456C	General purpose Spartan-3 development platform
DS-KIT-S2E3LC-ISE-BAS	Spartan-IIE LC Development Kit w/ISE BaseX	Insight (Memec)	XC2S300E-6FG456C	
DS-KIT-S2E6LC*	Spartan-IIE LC 2S600E Development Kit	Insight (Memec)	XC2S600E-6FG456C	
DS-KIT-S2E6LC-ISE-BAS	Spartan-IIE LC 2S600E Development Kit w/ISE BaseX	Insight (Memec)	XC2S600E-6FG456C	
iDEV2	CAN, LIN and TTCAN Development Platform and Starter Kit	Intelliga Integrated Design, Ltd.	XC2S300E	Automotive, Industrial
HW-AFX-DIGI-2S200E	Spartan-IIE Development Board	NuHorizons	XC2S200E	Prototyping, DSP, Multimedia, Reconfigurable Computing

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits

XILINX ON BOARD

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
Spartan-II				
ADS-SP2-MB-EVL	Spartan-II Evaluation Kit w/MicroBlaze & Communications/ Memory Module	Avnet Design Services	XC2S150-5PQ208	Networking, telecommunication, data communication, embedded and consumer markets
ADS-XLX-SP2-EVL	Spartan-II Evaluation Kit	Avnet Design Services	XC2S150-5PQ208	Very cost effective evaluation and prototyping platform to develop and test designs that are targeted to the Xilinx Spartan-II FPGA family - helps shorten and simplify the design cycle
2D Fabric Board	2D Fabric Evaluation and Demo Board	Crossbow Technologies, Inc.	XC2S150	Multi-processing system development; Networking, wireless basestations, VoP gateways
DS-KIT-2S100*	Spartan-II 100 Development Kit	Insight (Memec)	XC2S100	DSP, Digital Video, IP-Based Systems, Image Processing
DS-KIT-2S200*	Spartan-II PCI 2S200 Development Kit	Insight (Memec)	XC2S200-6FG456C	DSP, Digital Video, IP-Based Systems, Image Processing, PCI, Reconfigurable Computing
DS-KIT-2S200-ISE-BAS	Spartan-II PCI 2S200 Development Kit w/ISE BaseX and JTAG Cable	Insight (Memec)	XC2S200-6FG456C	DSP, Digital Video, IP-Based Systems, Image Processing, PCI, Reconfigurable Computing
DS-KIT-2S200-PAK*	Spartan-II PCI 2S200 Development Kit w/WebPACK CD and JTAG Cable	Insight (Memec)	XC2S200-6FG456C	DSP, Digital Video, IP-Based Systems, Image Processing, PCI, Reconfigurable Computing
DS-KIT-2SLC100*	Spartan-II LC 2S100 Development Kit	Insight (Memec)	XC2S100-5PQ208C	General purpose Spartan-II development platform
DS-KIT-2SLC100-ISE-BAS	Spartan-II LC 2S100 Development Kit w/ISE BaseX and JTAG Cable	Insight (Memec)	XC2S100-5PQ208C	General purpose Spartan-II development platform
DS-KIT-2SLC100-PAK*	Spartan-II LC 2S100 Development Kit w/WebPACK CD and JTAG Cable	Insight (Memec)	XC2S100-5PQ208C	General purpose Spartan-II development platform
DS-KIT-MBLAZE-S2	Spartan-II MicroBlaze Kit	Insight (Memec)	Spartan-II	Embedded microprocessor
DS-KIT-PCI32S-200	Spartan-II 200 PCI Development Kit w/Xilinx Single Project PCI325 License	Insight (Memec)	XC2S200	DSP, Digital Video, Image Processing, PCI, Reconfigurable Computing
RT-DAC4/PCI-D-OMNI	Digital OMNI Board	INTECO	XC2S100, XC2S150, XC2S50	ASIC Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Low power designs, Multimedia, Navigation, PCI, Reconfigurable Computing, Serial/ Deserialization, Telecom / Datacom, Telematics, VoIP
RT-DAC4/PCI-OMNI	OMNI Board	INTECO	XC2S100, XC2S150, XC2S50	ASIC Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Low power designs, Multimedia, Navigation, PCI, Reconfigurable Computing, Serial/ Deserialization, Telecom / Datacom, Telematics, VoIP
MicroEngine V	CPU + FPGA (Virtex/Spartan-II) MicroEngine Cards	Intrinsyc, Inc.	Spartan-II	Embedded Systems
NPE565-M	Spartan-II Power-PC Engine	North Pole Engineering	XC2S200	Prototyping, DSP, Embedded System, Industrial Automotive, Reconfigurable Computing
NV-FPSD-001	FPGA SDRAM Controller Evaluation Board	Novtech	XC2S50	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Low power designs, Multimedia, Navigation, PCI, Reconfigurable Computing, Serial/Deserialization, Telecom / Datacom, Telematics, VoIP
SPCMUSB-EVL	USB 2.0 Mass Storage Class Reference Design	Specsoft Consulting, Inc.	XC2S200	Prototyping, Data Storage, Embedded System, IP-Based Systems
SPCVUSB-EVL	USB 2.0 Video Class Reference Design	Specsoft Consulting, Inc.	XC2S200	Prototyping, Data Transmission & Manipulation, Embedded System, IP-Based Systems

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits

XILINX ON BOARD

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
Virtex-II Pro				
ADM-XPL	Reconfigurable Computer Based on Virtex-II Pro	Alpha Data	XC2VP7, XC2VP20	DSP, Reconfigurable Computing , Image Processing , ASIC Prototyping
PCI Platform	Virtex-II Pro PCI Platform FPGA Development Board	Amirix Systems, Inc.	XC2VP7,XC2VP20, XC2VP30	Networking (SAN, VoIP, Bridges), Communications, DSP, Image Processing, Industrial Controls, Instrumentation, test and measurement
ADS-XLX-V2PRO-DEVP20-6	Virtex-II Pro Development Kit w/XC2VP20, -6 speed grade	Avnet Design Services	XC2VP20	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Navigation, PCI, Reconfigurable Computing, Serial / Deserialization
ADS-XLX-V2PRO-DEVP30-6	Virtex-II Pro Development Kit w/XC2VP30, -6 speed grade	Avnet Design Services	XC2VP30	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Navigation, PCI, Reconfigurable Computing, Serial / Deserialization
ADS-XLX-V2PRO-DEVP7-5	Virtex-II Pro Development Kit w/XC2VP7, -5 speed grade	Avnet Design Services	XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Navigation, PCI, Reconfigurable Computing, Serial / Deserialization
ADS-XLX-V2PRO-DEVP7-6	Virtex-II Pro Development Kit w/XC2VP7, -6 speed grade	Avnet Design Services	XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Navigation, PCI, Reconfigurable Computing, Serial / Deserialization
ADS-XLX-V2PRO-EVLP7-5	Virtex-II Pro Evaluation Kit w/XC2VP7, -5 speed grade	Avnet Design Services	XC2VP7	Very cost effective evaluation and prototyping platform to develop and test designs targeted to the Virtex-II Pro device
ADS-XLX-V2PRO-EVLP7-6	Virtex-II Pro Evaluation Kit w/XC2VP7, -6 speed grade	Avnet Design Services	XC2VP7	Very cost effective evaluation and prototyping platform to develop and test designs targeted to the Virtex-II Pro device
V2PRO Kit	Virtex-II Pro™ Development Kit	Avnet Design Services	XC2VP7, XC2VP20, SC2VP30	Packet switching; network security; SAN, servers and super computers; video computing/ transmission; High-speed serial interfaces
D6PC	Danube 6-PaC	BittWare, Inc.	XC2VP20	DSP, Data Transmission & Manipulation, Embedded System, Image Processing, Multimedia, Reconfigurable Computing, Telecom/Datacom
DN6000K10S	DN6000K10S ASIC Prototyping Engine	DINI Group	XC2VP125, XC2VP70	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Low power designs, Multimedia, Navigation, PCI, Reconfigurable Computing, Serial/Deserialization, Datacom / Telecom, Telematics, VoIP
DS-KIT-2VP20FF1152*	Virtex-II Pro FF1152 P20 Development Kit	Insight (Memec)	XC2VP20 in FF1152	
DS-KIT-2VP30FF1152*	Virtex-II Pro FF1152 P30 Development Kit	Insight (Memec)	XC2VP30 in FF1152	
DS-KIT-2VP4FF672*	Virtex-II Pro FF672 P4 Development Kit	Insight (Memec)	XC2VP4 in FF672	
DS-KIT-2VP4FG456*	Virtex-II Pro P4 FG456 Development Kit	Insight (Memec)	XC2VP4 in FG456	DSP, Digital Video, Embedded System, IP-Based Systems, Image Processing
DS-KIT-2VP7FF672*	Virtex-II Pro FF672 P7 Development Kit - EURO	Insight (Memec)	XC2VP7 in FF672	
DS-KIT-2VP7FG456*	Virtex-II Pro P7 FG456 Development Kit	Insight (Memec)	XC2VP7 in FG456	DSP, Digital Video, Embedded System, IP-Based Systems, Image Processing
BenPRO-2VP7-6	BenPRO	Nallatech	XC2VP7-XC2VP20	Data Networks & Digital Signal processing, Embedded System, Telecom
SMT387	Disk Storage Module	Sundance Multiprocessor Co.	XC2VP20	DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Encryption Devices, Image Processing, Reconfigurable Computing
TB-V2P-20-MGT	Virtex-II Pro RocketIO Evaluation Board	Tokyo Electron Device Limited	XC2VP20-6FF896	Prototyping, DSP, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Reconfigurable Computing, Serial / Deserialization
TB-V2P-30-MGT	Virtex-II Pro RocketIO Evaluation Board	Tokyo Electron Device Limited	XC2VP30-6FF896	Prototyping, DSP, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Reconfigurable Computing, Serial / Deserialization

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits. For Virtex-II Pro, 2VP40 and 2VP50 boards are available upon customer request.

XILINX ON BOARD

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
Virtex-II Pro				
TB-V2P-7-MGT	Virtex-II Pro RocketIO Evaluation Board	Tokyo Electron Device Limited	XC2VP7-6FF896	Prototyping, DSP, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Reconfigurable Computing, Serial / Deserialization
XSP-016	Virtex-II Pro PowerPC&MicroBlaze Evaluation Board	Tokyo Electron Device Limited	XC2VP7-5FG456	Prototyping, DSP, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Reconfigurable Computing
DO-V2P-ML300-EC	Virtex-II Pro ML300 Evaluation Platform- EC version	Xilinx Online Store	XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Serialization / Deserialization, Telecom
DO-V2P-ML300-UK	Virtex-II Pro ML300 Evaluation Platform- UK version	Xilinx Online Store	XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Serialization / Deserialization, Telecom
DO-V2P-ML300-USA	Virtex-II Pro ML300 Evaluation Platform- US version	Xilinx Online Store	XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Serialization / Deserialization, Telecom / Datacom, Telematics, VoIP
DO-V2P-ML300-WRS-EC	Virtex-II Pro ML300 Evaluation Platform with WindRiver tools - EC version	Xilinx Online Store	XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Serialization / Deserialization, Telecom
DO-V2P-ML300-WRS-UK	Virtex-II Pro ML300 Evaluation Platform with WindRiver tools - UK version	Xilinx Online Store	XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Serialization / Deserialization, Telecom
DO-V2P-ML300-WRS-USA	Virtex-II Pro ML300 Evaluation Platform with WindRiver tools - US version	Xilinx Online Store	XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Serialization / Deserialization, Telecom / Datacom, Telematics, VoIP
HW-AFX-FF1152-300	Virtex-II Pro Proto Board	Xilinx Online Store	XC2VP20, XC2VP30, XC2VP40, XC2VP50	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Navigation, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, Telematics, VoIP
HW-AFX-FF672-300	Virtex-II Pro Proto Board	Xilinx Online Store	XC2VP2, XC2VP4, XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Navigation, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, Telematics, VoIP
HW-AFX-FG456-300	Virtex-II Pro Proto Board	Xilinx Online Store	XC2VP2, XC2VP4, XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Navigation, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, Telematics, VoIP
HW-AFX-SMA-HSSDC2	SMA To HSSDC2 Conversion Module	Xilinx Online Store	Supports boards with SMA connectors	
HW-AFX-SMA-RJ45	SMA To RJ45 Conversion Module	Xilinx Online Store	Supports boards with SMA connectors	
HW-AFX-SMA-SATA	SMA To SATA Conversion Module	Xilinx Online Store	Supports boards with SMA connectors	
HW-AFX-SMA-SFP	SMA To SFP Conversion Module	Xilinx Online Store	Supports boards with SMA connectors	
DO-V2P-ML300-USA	Virtex-II Pro ML300 Evaluation Platform- US version	Xilinx Sales Offices	XC2VP7	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Procesing, Multimedia, Navigation, PCI, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, Telematics, VoIP
HW-V2PRO-XLVDS	Virtex-II Pro XLVDS	Xilinx Sales Offices	XC2VP20	Data Transmission & Manipulation, Embedded System, Encryption Devices, IP-Based Systems, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, Test Equipment

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits

XILINX ON BOARD

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
Virtex-II				
ADM-XPL/2VP7-5	ADM-XPL	Alpha Data	XC2V1000, XC2VP20, XC2VP7	Prototyping, DSP, Image Processing, Reconfigurable Computing
ADM-XRC-II	ADM-XRC-II	Alpha Data	XC2V3000, XC2V4000, XC2V6000, XC2V8000	Prototyping, Compression, DSP, Encryption, Reconfigurable Computing, Software Radio, Video / Image Processing, XML processing
ADS-V2-MB-DEV4000XP	Virtex-II XC2V4000XP Development Kit with MicroBlaze	Avnet Design Services	XC2V4000	Prototyping. Data Transmission & Manipulation, Highend DSP, IP-Based Systems, PCI
ADS-V2-MB-DEV6000XP	Virtex-II XC2V6000 Development Kit with MicroBlaze	Avnet Design Services	XC2V6000	Prototyping. Data Transmission & Manipulation, Highend DSP, IP-Based Systems, PCI
ADS-XLX-MB-DEV1500	Virtex-II XC2V1500 Development Kit with MicroBlaze	Avnet Design Services	XC2V1500	Prototyping, Data Transmission & Manipulation, High End DSP, IP-based Systems, PCI
ADS-XLX-MB-DEV4000	Virtex-II XC2V4000 Development Kit with MicroBlaze	Avnet Design Services	XC2V4000	Prototyping, Data Transmission & Manipulation, High End DSP, IP-based Systems, PCI
ADS-XLX-PMC-IRL PMC IRL	Reference Design Kit	Avnet Design Services	XC2V1000-4FG456C/FF896C	Consumer, Industrial, IRL, PAVE, Telecom/Datacom, Telecommunications
ADS-XLX-V2-DEV1500	Virtex-II Development Kit w/XC2V1500 device	Avnet Design Services	XC2V1500	Complete hardware environment to develop, prototype, and test designs targeted to the Virtex-II FPGA family
ADS-XLX-V2-DEV4000	Virtex-II XC2V4000 Development Kit w/XC2V4000 device	Avnet Design Services	XC2V4000	Prototyping, Data Transmission & Manipulation, High End DSP, IP-based Systems, PCI
ADS-XLX-V2-DEV4000XP	Virtex-II XC2V4000XP Development Kit w/high-current power supply	Avnet Design Services	XC2V4000	Prototyping. Data Transmission & Manipulation, Highend DSP, IP-Based Systems, PCI
ADS-XLX-V2-DEV6000XP	Virtex-II XC2V6000 Development Kit w/high-current power supply	Avnet Design Services	XC2V6000	Prototyping. Data Transmission & Manipulation, Highend DSP, IP-Based Systems, PCI
ADS-XLX-V2-EVL1000	Virtex-II XC2V1000 Evaluation Kit	Avnet Design Services	XC2V1000-4FG256	Prototyping, Data Transmission & Manipulation, High End DSP, IP-based Systems
ADS-XLX-XV2-EVL	Virtex-II High-Speed Evaluation Kit	Avnet Design Services	XC2V40	Data Transmission & Manipulation, DSP
BCPM	Barracuda-PMC+	BittWare, Inc.	XC2V1000	DSP, Data Transmission & Manipulation, Embedded System, Image Processing, Multimedia, Reconfigurable Computing, Telecom/Datacom
RFPM	Reef-PMC+	BittWare, Inc.	XC2V1000	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Low power designs, Multimedia, Navigation, PCI, Reconfigurable Computing, Serial/Deserialization, Datacom / Telecom, Telematics, VoIP
RMPM	Remora-PMC+	BittWare, Inc.	XC2V1000	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Low power designs, Multimedia, Navigation, PCI, Reconfigurable Computing, Serial/Deserialization, Datacom / Telecom, Telematics, VoIP
CYL2T0201-DVK	MetroLink2T-2 Link Layer Evaluation Board	Cypress Semiconductor Corporation	XC2V2000	Data Transmission & Manipulation, IP-Based Systems, Telecom / Datacom
PF3100	PC/104-Plus Reconfigurable Module Board (PF3100)	Derivation Systems, Inc.	XC2V1000, FG256	Internet appliance, industrial control
DN3000K10	DN3000K10 ASIC Prototyping Engine	DINI Group	XC2V4000, XC2V6000, XC2V8000	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Low power designs, Multimedia, Navigation, PCI, Reconfigurable Computing, Serial/Deserialization, Telecom/ Datacom, Telematics, VoIP
DN3000K10	DN3000K10	DINI Group	XC2V4000, XC2V6000, XC2V8000	Prototyping, Algorithmic Acceleration, Logic Emulation, PCI / PCI-X, Reconfigurable Computing
CHM2-VME-604-SZ	Chameleon II VME Reconfigurable Computing Board	DRS Tactical Systems West, Inc.	XC2V6000	Prototyping, DSP, IP-Based Systems, Image Processing, Reconfigurable Computing, Telecom / Datacom, Telematics
APB-2V1000	Virtex-II Prototyping Board for 2V1000	ErSt Electronics	XC2V1000	Prototyping, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, VoIP

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits

XILINX ON BOARD

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
Virtex-II				
APB-2V1500	Virtex-II Prototyping Board for 2V1500	ErSt Electronics	XC2V1500	Prototyping, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, VoIP
APB-2V2000	Virtex-II Prototyping Board for 2V2000	ErSt Electronics	XC2V2000	Prototyping, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, VoIP
APB-2V3000	Virtex-II Prototyping Board for 2V3000	ErSt Electronics	XC2V3000	Prototyping, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, VoIP
APB-2V4000	Virtex-II Prototyping Board for 2V4000	ErSt Electronics	XC2V4000	Prototyping, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, VoIP
APB-2V6000	Virtex-II Prototyping Board for 2V6000	ErSt Electronics	XC2V6000	Prototyping, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, VoIP
APB-2V8000	Virtex-II Prototyping Board for 2V8000	ErSt Electronics	XC2V8000	Prototyping, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Reconfigurable Computing, Serial / Deserialization, Telecom / Datacom, VoIP
GVA-300	DSP Hardware Accelerator, Virtex-II (GVA-300)	GV & Associates, Inc.	XC2V1500-4, 2000-4 or 3000-4	DSP
GVA-325	DSP Hardware Accelerator, Virtex-II (GVA-325)	GV & Associates, Inc.	XC2V1500-4, 2000-4 or 3000-4	DSP
GVA-350	DSP Hardware Accelerator, Virtex-II (GVA-350)	GV & Associates, Inc.	XC2V4000-4, 6000-4 or 8000-4	DSP
DS-KIT-MBLAZE-V2	Virtex-II MicroBlaze Kit	Insight (Memec)	Virtex-II	Embedded microprocessor
DS-KIT-V2LC1000*	Virtex-II LC1000	Insight (Memec)	XC2V1000	Digital Video, Image Processing, Telecom / Datacom
DS-KIT-V2LC1000-ISE	Virtex-II LC1000 w/ISE Foundation and JTAG Cable	Insight (Memec)	XC2V1000	Digital Video, Image Processing, Telecom / Datacom
DS-KIT-V2MB1000*	Virtex-II MB 2V1000 Development Kit	Insight (Memec)	XC2V1000-4FG456C	DSP, Digital Video, Embedded System, Image Processing, Reconfigurable Computing, Telecom / Datacom,
DS-KIT-V2MB1000-ISE	Virtex-II MB 2V1000 Development Kit w/ISE Foundation and JTAG Cable	Insight (Memec)	XC2V1000-4FG456C	DSP, Digital Video, Embedded System, Image Processing, Reconfigurable Computing, Telecom / Datacom,
MicroEngine V-II	CPU + FPGA (Virtex-II) MicroEngine Cards	Intrinsyc, Inc.	Virtex-II	Embedded Systems
ASPE-1000	Advanced Signal Processing Engine (ASPE)	Multiple Access Communications	XC2V1000	DSP, Reconfigurable Computing, Telecom / Datacom
BenADDA-2V250-4-xx	BenADDA	Nallatech	XC2V250 - XC2V6000	Infrared Processing, Mobile Communications Systems, Multi-channel, Multi-mode receivers, Wideband Cable Systems
BenBLUE-II-2V4000-4-01	BenBLUE-II	Nallatech	XC2V4000 - XC2V8000	ASIC Prototyping, Image Processing, Reconfigurable Computing, Software Defined Radio Data Processing
BenDATA-DD-2V3000	BenDATA-DD	Nallatech	XC2V3000-XC2V8000	Data Transmission & Manipulation, Image Processing
BenDATA-WS-2V4000-4-01	BenDATA-WS	Nallatech	XC2V4000 - XC2V8000	Image Processing
42-055-01	SONET / SDH ATM-POS Board	NetQuest Corporation	XC2V4000	Data Transmission & Manipulation, Embedded System, IP-Based Systems, Multimedia, Telecom / Datacom

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits

XILINX ON BOARD

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
Virtex-II				
42-059-01	Gigabit IP Content Processor Board	NetQuest Corporation	XC2V4000	Data Transmission & Manipulation, Embedded System, IP-Based Systems, Multimedia, Telecom / Datacom
42-060-01	SONET / SDH / PDH Groomer Board	NetQuest Corporation	XC2V3000	Data Transmission & Manipulation, Embedded System, IP-Based Systems, Multimedia, Telecom / Datacom
XL9000	XL 9000 Multi-Port Camera Link PCI Frame Grabber	Novtech	XC2V2000 (can be substituted by XC2V2000-XC2V8000)	Data Storage, Data Transmission & Manipulation, Digital Video, Image Processing
GR-PCI-XC2V	LEON PCI Virtex-II Development Board	Pender Electronic Design	XC2V3000	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Low power designs, Multimedia, Navigation, PCI, Reconfigurable Computing, Serial/Deserialization, Telecom / Datacom, Telematics, VoIP
JockoBoard	JockoBoard SOC Virtex-II Prototyping Platform	RealFast Operating Systems AB	XC2V1000	Embedded Systems
600-00422	PRO-3100 Virtex-II FPGA Processing Engine	Spectrum Signal Processing	XC2V3000, XC2V6000	DSP, Data Transmission & Manipulation, Reconfigurable Computing
650-00075	ePMC-8120	Spectrum Signal Processing	XC2V6000	DSP, Data transmission & Manipulation, Reconfigurable Computing
SMT351-G	Memory Module 351-G	Sundance Multiprocessor Co.	XC2V1000	Base Band Radar Sampling, Cellular / PCS Base Stations, Communications, Digital Radio Receivers, General Data Logging, IF Radar Sampling, Multi-Channel Receivers, Spectrum Analysers
SMT351-M	Memory Module 351-M	Sundance Multiprocessor Co.	XC2V1000	Base Band Radar Sampling, Cellular / PCS Base Stations, Communications, Digital Radio Receivers, General Data Logging, IF Radar Sampling, Multi-Channel Receivers, Spectrum Analysers
SMT365E	DSP Module	Sundance Multiprocessor Co.	XC2V6000	Image Processing, Industrial, Medical, Telecomunications
SMT398-1000-4-Z1	FPGA Module 1000-4-Z1	Sundance Multiprocessor Co.	XC2V1000	Base Band Radar Sampling, Cellular / PCS Base Stations, Communications, Digital Radio Receivers, General Data Logging, IF Radar Sampling, Imaging, Multi-Channel Receivers, Spectrum Analysers
SMT398-2000-4-Z1	FPGA Module 2000-4-Z1	Sundance Multiprocessor Co.	XC2V2000	Base Band Radar Sampling, Cellular / PCS Base Stations, Communications, Digital Radio Receivers, General Data Logging, IF Radar Sampling, Imaging, Multi-Channel Receivers, Spectrum Analysers
SMT398-3000-4-Z4-Q2	FPGA Module 3000-4-Z4-Q2	Sundance Multiprocessor Co.	XC2V3000	Base Band Radar Sampling, Cellular / PCS Base Stations, Communications, Digital Radio Receivers, General Data Logging, IF Radar Sampling, Imaging, Multi-Channel Receivers, Spectrum Analysers
SMT398-8000-4-Z4-Q2	FPGA Module 8000-4-Z4-Q2	Sundance Multiprocessor Co.	XC2V8000	Base Band Radar Sampling, Cellular / PCS Base Stations, Communications, Digital Radio Receivers, General Data Logging, IF Radar Sampling, Imaging, Multi-Channel Receivers, Spectrum Analysers
SMT370	Dual Channel A/D and D/A Module	Sundance Multiprocessor Technology Ltd.	XC2V1000	Base Band Radar Sampling, Cellular / PCS Base Stations, Digital Radio Receivers, General Data Logging and I/O Control, IF Radar Sampling, Instrumentation, Multi-Channel Receivers, Sonar, Spectrum Analysers
XT2000-X	XTENSA Microprocessor Emulation Kit (XT2000-X)	Tensilica, Inc.	XC2V6000-4	Prototyping, Embedded Systems
DO-V2000-MLTA	Virtex-II Multimedia Board	Xilinx Online Store	XC2V2000	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded Microprocessor, Embedded System, IP-Based Systems, Image Processing, Multimedia, Reconfigurable Computing, Telecom / Datacom, VoIP

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits

XILINX ON BOARD

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
Virtex-II				
HW-AFX-FF1152-200	FF1152-200 Proto Board	Xilinx Online Store	XC2V3000, XC2V4000, XC2V6000, XC2V8000	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded System, Multimedia, PCI, Reconfigurable Computing, Telecom / Datacom, Telematics, VoIP
HW-AFX-FG256-200	FG256-200 Proto Board	Xilinx Online Store	XC2V1000, XC2V250, XC2V40, XC2V500, XC2V80	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded System, Multimedia, PCI, Reconfigurable Computing, Telecom / Datacom, Telematics, VoIP
HW-AFX-FG456-200	FG456 Virtex-II Proto Board	Xilinx Online Store	XC2V1000, XC2V250, XC2V500	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded System, IP-Based Systems, Multimedia, PCI, Reconfigurable Computing, Telecom / Datacom, Telematics, VoIP
HW-AFX-FG676-200	FG676 Virtex-II Proto Board	Xilinx Online Store	XC2V1500, XC2V2000, XC2V3000	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Embedded System, IP-Based Systems, Multimedia, PCI, Reconfigurable Computing, Telecom / Datacom, Telematics, VoIP
DO-DI-DSP-DK2	XtremeDSP Development Kit	Xilinx Online Store	2V3000	High Performance DSP
DO-DI-DSP-DK2-SG	XtremeDSP Development Kit + System Generator for DSP	Xilinx Online Store	2V3000	Hardware in the loop from Simulink using System Generator software and high performance DSP board
Virtex and Virtex-E				
ADS-VE-MB-DEV	Virtex-E Development Kit w/MicroBlaze	Avnet Design Services	XCV1000E	Networking, telecommunication, data communication, embedded and consumer markets
ADS-VE-MB-EVL	Virtex-E Evaluation Kit w/MicroBlaze	Avnet Design Services	XCV100E	Networking, telecommunication, data communication, embedded and consumer markets
ADS-XLX-VE-DEV	Virtex-E Development Kit	Avnet Design Services	XCV1000E	Communications, Embedded Control, LAN Routers, LAN Switch, Networking, WAN Access, xDSL Equipment
ADS-XLX-VE-EVL	Virtex-E Evaluation Kit	Avnet Design Services	XCV100E	Very cost effective evaluation and prototyping platform to develop and test designs targeted to the Virtex-E device
DN2000K10	DN2000K10	DINI Group	XCV1000, XCV1000E, XCV1600E, XCV2000E	Prototyping, Algorithmic Acceleration, Logic Emulation, PCI / PCI-X
GVA-290	DSP Hardware Accelerator, Virtex-E (GVA-290)	GV & Associates, Inc.	2-XCV1000E, 1600E or 2000Es	DSP
BenADIC-2000E	BenADIC	Nallatech	XCV2000E, XCV600E	Multichannel Data Acquisition & Software Radio, Phased Array Radar, Smart Antenna Arrays
BenERA-1000E-6-A	BenERA	Nallatech	XCV1000E-XCV2000E	Communications & Real-Time systems, DSP, Image Processing, Reconfigurable Computing
BenFAD-2000E	BenFAD	Nallatech	XCV2000E	Broadband wireless/satellite communications, Data Acquisition/signal analysis systems
42-047-01	SONET / SDH POS PCI NIC Card	NetQuest Corporation	XCV600	Data Transmission & Manipulation, IP-Based Systems, Multimedia, Telecom / Datacom
42-053-01	SONET / SDH ATM PCI NIC Card	NetQuest Corporation	XCV600	Data Transmission & Manipulation, IP-Based Systems, Multimedia, Telecom / Datacom
HW-AFX-BG352-100	BG352-100 Proto Board	Xilinx Online Store	XCV150, XCV200, XCV300	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Telecom / Datacom, Telematics, VoIP
HW-AFX-BG432-100	BG432-100 Proto Board	Xilinx Online Store	XCV300/E, XCV400/E, XCV600/E, XCV800	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Telecom / Datacom, Telematics, VoIP

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits

XILINX ON BOARD

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
Virtex-II				
HW-AFX-BG560-100	BG560-100 Proto Board	Xilinx Online Store	XCV1000/E, XCV1600E, XCV2000E, XCV400/E, XC405E, XCV600/E, XCV800, XCV812E	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Telecom / Datacom, Telematics, VoIP
HW-AFX-PQ240-100	PQ240-100 Proto Board	Xilinx Online Store	XCV100, XCV150, XCV200, XCV300, XCV50	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Telecom / Datacom, Telematics, VoIP
HW-AFX-PQ240-110	PQ240-110 Proto Board	Xilinx Online Store	XCV100, XCV150, XCV200, XCV300, XCV50	Prototyping, DSP, Data Storage, Data Transmission & Manipulation, Digital Video, Embedded System, IP-Based Systems, Image Processing, Multimedia, Navigation, PCI, Reconfigurable Computing, Telecom / Datacom, Telematics, VoIP
RocketPHY				
HWK-RPHY2XFP-1	RocketPHY XFP Kit	Xilinx Sales Offices	XGC1120 10G Ultra MSA	Data Storage, Data Storage (10 Gigabit Fibre Channel), Data Transmission & Manipulation, Datacom (10 Gigabit Ethernet), Serial / Deserialization, Telecom (Sonet OC-192 / FEC), Telecom / Datacom
HWK-RPHY2XFP-M	RocketPHY XFP Kit with XFP Module	Xilinx Sales Offices	XGC1120 10G Ultra MSA	Data Storage, Data Storage (10 Gigabit Fibre Channel), Data Transmission & Manipulation, Datacom (10 Gigabit Ethernet), Serial / Deserialization, Telecom (Sonet OC-192 / FEC), Telecom / Datacom
HWK-RPHY-DVLP	RocketPHY Development Kit	Xilinx Sales Offices	XC2VP20, XGC1120 10G Ultra MSA	Data Storage, Data Storage (10 Gigabit Fibre Channel), Data Transmission & Manipulation, Datacom (10 Gigabit Ethernet), Encryption Devices, Routers, Serial / Deserialization, Telecom (Sonet OC-192 / FEC), Telecom / Datacom, Test Equipment
CoolRunner-II				
ADS-XLX-CR2-EVL	CoolRunner-II Evaluation Kit	Avnet Design Services	XC2C256-VQ100	Low power designs
Digilab-XC2	Digilent CoolRunner-II Development Board	Digilent	XC2C256	Low power designs, Telecom / Datacom
DS-KIT-2C256*	CoolRunner-II Development Kit	Insight (Memec)	XC2C256	Low power designs
DS-KIT-2C256-PAK*	CoolRunner-II Development Kit w/WebPACK CD and JTAG Cable	Insight (Memec)	XC2C256	Low power designs
MXCK-100-003	Mechatronics CoolRunner-II Prototyping Board	Mechatronics Test Equipment	XC2C64	Embedded System, Low power designs
HW-AFX-COOL2-256MC	CoolRunner-II Evaluation Board	NuHorizons	XC2C256	Low power designs

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits

XILINX ON BOARD

XILINX®
The Programmable Logic Company℠

Xilinx On Board – Development Reference Boards

http://www.xilinx.com/board_search

Board Part Number	Board Description	Supplier	Xilinx Device Support	Application
CoolRunner XPLA3				
ADS-XLX-X3-EVL	XPLA3 Evaluation Kit	Avnet Design Services	XCR3256XL	Low power designs
XCR	Digilab CoolRunner Development Board	Digilent	XCR3064XL	Low power designs
ET-XCR	Emulation Technology XCR Development Board	Emulation Technology, Inc.	XCR3064XL	Low power designs
DS-KIT-XPLA3	CoolRunner XPLA3 Development Kit	Insight (Memec)	XCR3256XL	Low power designs
DS-KIT-XPLA3-PAK	CoolRunner XPLA3 Development Kit w/WebPACK CD and JTAG Cable	Insight (Memec)	XCR3256XL	Low power designs
MXCK-100-002	Mechatronics CoolRunner Development Board	Mechatronics Test Equipment	XCR3064XL, XCR3128XL	Low power designs
HW-AFX-DIGI-XCR3064XL	Digilab XCR board	NuHorizons	XCR3064XL	Low power designs
XC9500 Series				
PBX-84	PBX84 Xilinx Prototyping Board	AL Williams	XC95108, XC9572	Embedded System
DB-CPLD-PQ	ASICentrum XC9500 Development Kit	ASICentrum	XC9572	Embedded System
XC95	Digilab XC9500 Development Board	Digilent	XC95108	Embedded System
XCR-DEV-BRD	XCR Development Board	Digilent	XC95108, XCR3064XL	Automotive, Navigation, Telematics
DS-KIT-95XL*	XC9572XL Development Kit	Insight (Memec)	XC9572XL-10VQ64	Low power designs
DS-KIT-95XL-PAK*	XC9572XL Development Kit w/WebPACK CD and JTAG Cable	Insight (Memec)	XC9572XL-10VQ64	General purpose Spartan-II development platform
DS-KIT-95XL-PAK*	XC95144XV Development Kit w/WebPACK CD and JTAG Cable	Insight (Memec)	XC95144XV-10TQ144C	Low power designs
DS-KIT-95XV*	XC95144XV Development Kit	Insight (Memec)	XC95144XV-10TQ144C	Low power designs
84-0050	LogicFlex	JK Microsystems, Inc	XC9572XL	DSP, Data Storage, Data Transmission & Manipulation, Embedded Microprocessor, Embedded System, IP-Based Systems, Low power designs, Serial / Deserialization, Telecom / Datacom
MXCK-044-001	CPLD Junior Board	Mechatronics	XC9536, XC9572	Everything, Datacom
MXCK-084-004	CPLD Prototyping Board	Mechatronics	XC9572, XC95144	Telecom, industrial controls, instrumentations, etc.
MP1000TX	Gigabit Ethernet Phy Prototyping Board	Metanetworks	XC95144XL	Data Transmission & Manipulation, Telecom / Datacom
SBX2	SBX2 Systronix	Systronix	XC9572	Embedded System
LogiCRAFT	LogiCRAFT Evaluation System	Xylon d.o.o.	XCV95144-VQ100, XC9572-VQ64; XC2S150-BG256 on ICU daughtercard	Automotive, Industrial, Human-Machine Interfaces

* Insight Products: Append -EURO to part number for international kits (ex. DS-KIT-2SLC100-EURO); Power Supply not included in -EURO kits

XILINX ON BOARD

Function	Vendor Name	Virtex-II Pro	Virtex-II	Virtex-E	Spartan-3	Spartan-IIE	Spartan-II
Audio, Video & Image Processing							
Burst Locked PLL (BURST_PLL)	Pinpoint Solutions, Inc.		V-II	V-E	S3	S-IIE	
Color Space Converter, RGB2YCrCb (CSC)	CAST, Inc.		V-II	V-E		S-IIE	S-II
Huffman Decoder (HUFFD)	CAST, Inc.		V-II	V-E		S-IIE	S-II
JPEG Fast Codec (JPEG_FAST_C)	CAST, Inc.	V-IIP	V-II	V-E	S3		
JPEG, 2000 Encoder (JPEG2K_E)	CAST, Inc.	V-IIP	V-II	V-E			
JPEG, Fast color image decoder (FASTJPEG_C DECODER)	Barco-Silex		V-II	V-E			
JPEG, Fast Decoder (JPEG_FAST_D)	CAST, Inc.	V-IIP	V-II	V-E		S-IIE	
JPEG, Fast Encoder (JPEG_FAST_E)	CAST, Inc.	V-IIP	V-II	V-E			
JPEG, Fast gray scale image decoder (FASTJPEG_BW DECODER)	Barco-Silex		V-II	V-E			
JPEG, Motion Codec V1.0 (CS6190)	Amphion Semiconductor, Ltd.		V-II	V-E			
JPEG, Motion Decoder (CS6150)	Amphion Semiconductor, Ltd.		V-II	V-E			
JPEG, Motion Encoder (CS6100)	Amphion Semiconductor, Ltd.		V-II	V-E			
Motion JPEG Decoder (JPEG Decoder)	4i2i Communications Ltd.	V-IIP	V-II	V-E	S3	S-IIE	
Motion JPEG Encoder (JPEG Encoder)	4i2i Communications Ltd.	V-IIP	V-II	V-E	S3	S-IIE	
MPEG-2 HDTV I & P Encoder (DV1 HDTV)	Duma Video, Inc.		V-II				
MPEG-2 SDTV I & P Encoder (DV1 SDTV)	Duma Video, Inc.		V-II				
NTSC Color Separator (NTSC-COSEP)	Pinpoint Solutions, Inc.		V-II	V-E	S3	S-IIE	
Communication & Networking							
ADPCM, 1024 Channel Simplex (CS4190)	Amphion Semiconductor, Ltd.		V-II				
ADPCM, 256 Channel Simplex (CS4130)	Amphion Semiconductor, Ltd.		V-II	V-E			
ADPCM, 512 Channel Duplex (CS4180)	Amphion Semiconductor, Ltd.		V-II				
ADPCM, 768 Channel	Amphion Semiconductor, Ltd.		V-II				
AES Decryption Family (CS5200)	Amphion Semiconductor, Ltd.		V-II	V-E			S-II
AES Encryption	CAST, Inc.	V-IIP	V-II	V-E	S3	S-IIE	
AES Encryption Family (CS5200)	Amphion Semiconductor, Ltd.		V-II	V-E			S-II
AES Standard Encryptor/Decryptor	Helion Technology Limited	V-IIP	V-II	V-E	S3	S-IIE	
AES Tiny Encryptor/Decryptor	Helion Technology Limited	V-IIP	V-II	V-E	S3	S-IIE	
ATM Adaption Layer 1 (AAL1)	ModelWare, Inc.		V-II	V-E			
ATM Cell Processor (CC200)	Paxonet Communications, Inc.	V-IIP	V-II				
Bit Stream Analyzer and Data Extractor (Parser)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
Bluetooth Baseband Processor (BOOST Lite)	NewLogic GmbH		V-II				
CAM for Internet Protocol (IPlogiCAM)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
Convolutional Encoder (CONV_ENC)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
CRC-32 for 10 Gbps OC192 systems (CORE-CRC-128)	Calyptech Design Services				S3		
CRC-32 for 40 Gbps OC-768 systems (CORE-CRC-256)	Calyptech Design Services	V-IIP	V-II				
DES and DES3 Encryption Engine (MC-XIL-DES)	Memec Design		V-II	V-E		S-IIE	S-II
DES Encryption	CAST, Inc.		V-II	V-E			S-II
DES3 Encryption	CAST, Inc.	V-IIP	V-II	V-E	S3	S-IIE	
Distributed Sample Descrambler (DSD)	Telecom Italia Lab S.p.A.						S-II
Distributed Sample Scrambler (DSS)	Telecom Italia Lab S.p.A.						S-II
DVB Satellite Modulator (MC-XIL-DVBMOD)	Memec Design		V-II	V-E			S-II
Email Trigger	Amirix Systems, Inc.	V-IIP					
Ethernet MAC, 1 Gigabit Full Duplex (PE-GMAC0)	Mentor Graphics Corporation		V-II				
Ethernet MAC, 10/100	Zuken, Inc.		V-II	V-E			
Ethernet MAC, 10/100 (MAC)	CAST, Inc.	V-IIP	V-II		S3	S-IIE	
Ethernet MAC, 10/100 (PE-MACMII)	Mentor Graphics Corporation		V-II				S-II
Ethernet MAC, 10G (CC410)	Paxonet Communications, Inc.		V-II				
Ethernet PCS, 10G (CC411)	Paxonet Communications, Inc.	V-IIP	V-II				
Ethernet PCS, 10G (MC-XIL-10GEPCS)	Memec Design		V-II				
Framer, 1.25 Gb/s GFP (CC224)	Paxonet Communications, Inc.	V-IIP	V-II		S3	S-IIE	
Framer, 2.5 Gb/s GFP (CC226)	Paxonet Communications, Inc.	V-IIP	V-II				
Framer, 8-Bit Multichannel GFP (CC225)	Paxonet Communications, Inc.	V-IIP	V-II				
Framer, 8-Bit Transparent GFP (CC124)	Paxonet Communications, Inc.	V-IIP	V-II				
Framer, E1 (CC303)	Paxonet Communications, Inc.	V-IIP	V-II			S-IIE	
Framer, OC12 (CC351)	Paxonet Communications, Inc.	V-IIP	V-II				
Framer, OC192/10 GB/s GFP (CC327)	Paxonet Communications, Inc.	V-IIP	V-II				
Framer, OTU2 (CC481)	Paxonet Communications, Inc.	V-IIP	V-II				

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

Function	Vendor Name	Virtex-II Pro	Virtex-II	Virtex-E	Spartan-3	Spartan-IIE	Spartan-II
Framer, STS192/STM64 (CC314)	Paxonet Communications, Inc.	V-IIP	V-II				
Framer, T1 (CC302)	Paxonet Communications, Inc.	V-IIP	V-II			S-IIE	
Framer/Digital Wrapper, STS48 OTN (CC381)	Paxonet Communications, Inc.	V-IIP	V-II				
G.709 Compliant FEC Core (CC345)	Paxonet Communications, Inc.	V-IIP	V-II				
HDLC, Single-Channel (MC-XIL-HDLC)	Memec Design	V-IIP	V-II		S3		
HyperTransport Cave 16-Bit	GDA Technologies, Inc.	V-IIP	V-II				
Interleaver Deinterleaver (INT_DEINT)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
Inverse Multiplexer for ATM (IMA)	ModelWare, Inc.		V-II	V-E			
Mapper, E1 (CC333)	Paxonet Communications, Inc.	V-IIP	V-II			S-IIE	
MD5 Message Digest Algorithm	CAST, Inc.	V-IIP	V-II		S3	S-IIE	
Noisy Transmission Channel Model (CHANNEL)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
Path Processor, OC12c (CC321)	Paxonet Communications, Inc.	V-IIP	V-II				
Path Processor, STS192/STM64 (CC324)	Paxonet Communications, Inc.	V-IIP	V-II				
Reed Solomon Decoder (MC-XIL-RSDEC)	Memec Design		V-II			S-IIE	S-II
Reed Solomon Encoder (MC-XIL-RSENC)	Memec Design		V-II			S-IIE	S-II
Reed-Solomon Decoder (RS_DEC)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
SDLC Controller (SDLC)	CAST, Inc.	V-IIP	V-II	V-E	S3	S-IIE	
SHA-1 Encryption Processor	CAST, Inc.		V-II	V-E		S-IIE	S-II
SPI 4.2 Interface (CC401)	Paxonet Communications, Inc.	V-IIP	V-II				
Turbo Decoder (TURBO_DEC)	Telecom Italia Lab S.p.A.		V-II				
Turbo Decoder, 3GPP	SysOnChip, Inc.		V-II	V-E			
Turbo Decoder, 3GPP (S3000)	iCoding Technology, Inc.		V-II	V-E		S-IIE	
Turbo Decoder, DVB-RCS (S2000)	iCoding Technology, Inc.		V-II	V-E			
Turbo Decoder, DVB-RCS (TC1000)	TurboConcept		V-II	V-E			
Turbo Decoder, UMTS Hardwired Interleaver	Telecom Italia Lab S.p.A.	V-IIP	V-II	V-E			
Turbo Decoder, UMTS Mother Interleaver (UMTS_ADDRESS_GEN)	Telecom Italia Lab S.p.A.	V-IIP	V-II	V-E		S-IIE	
Turbo Encoder (TURBO_ENC)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
Turbo Encoder, DVB-RCS (S2001)	iCoding Technology, Inc.		V-II	V-E			
Turbo Product Code Decoder, 160 Mbps (TC3404)	TurboConcept		V-II	V-E			
Turbo Product Code Decoder, 25 Mbps (TC3000)	TurboConcept		V-II	V-E			
Turbo Product Code Decoder, 30 Mbps (TC3401)	TurboConcept			V-E		S-IIE	
ATM Utopia Level 2 Master and Slave w/OPB Interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
UTOPIA Level-2 PHY Side RX Interface (UTOPIA L2 PHY Rx)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
UTOPIA Level-2 PHY Side TX Interface (UTOPIA L2 PHY Tx)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
UTOPIA RX Level 2 Master Interface (UTOPIA2M_RX)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
UTOPIA TX Level 2 Master Interface (UTOPIA2M_TX)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
Viterbi Decoder (VITERBI_DEC)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
3G FEC Package	Xilinx		V-II	V-E			
8b/10b Decoder	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
8b/10b Encoder	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
AWGN - Additive White Gaussian Noise	Xilinx	V-IIP	V-II				
Convolutional Encoder	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Ethernet 1000BASE-X PCS/PMA	Xilinx	V-IIP					
Ethernet MAC, 1 Gigabit Half/Full duplex with GMII or 1000BASE-X PCS/PMA	Xilinx	V-IIP	V-II	V-E		S-IIE	
Ethernet MAC, 10 Gigabit Full Duplex with XGMII or XAUI	Xilinx	V-IIP	V-II				
Interleaver/De-interleaver	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Reed Solomon Decoder	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Reed Solomon Encoder	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
SPI-3 (POS-PHY L3) Link Layer Interface, 1-Ch	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
SPI-3 (POS-PHY L3) Link Layer Interface, 2-Ch	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
SPI-3 (POS-PHY L3) Link Layer Interface, 4-Ch	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
SPI-3 (POS-PHY L3) Link Layer Interface, Multi-Channel	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
SPI-3 (POS-PHY L3) Physical Layer Interface	Xilinx			V-E			
SPI-4.1 (Flexbus 4) Interface Core, 1-Channel	Xilinx		V-II				
SPI-4.1 (Flexbus 4) Interface Core, 4-Channel	Xilinx		V-II				
SPI-4.2 (POS-PHY L4) Multi-Channel Interface	Xilinx	V-IIP	V-II				
SPI-4.2 (POS-PHY L4) to SPI-4.1 (Flexbus 4) Bridge	Xilinx		V-II				
SPI-4.2 (POS-PHY L4) to XGMII (10GE MAC) Bridge	Xilinx		V-II				

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

Function	Vendor Name	Virtex-II Pro	Virtex-II	Virtex-E	Spartan-3	Spartan-IIE	Spartan-II
SPI-4.2 Lite (POS_PHY L4)	Xilinx		V-II		S-3		
Turbo Decoder, Convolutional, 3GPP Compliant	Xilinx		V-II	V-E			
Turbo Decoder, Convolutional, 3GPP2/CDMA2000	Xilinx	V-IIP	V-II		S-3		
Turbo Decoder, Product Code	Xilinx	V-IIP	V-II		S-3		
Turbo Encoder, Convolutional, 3GPP Compliant	Xilinx		V-II	V-E			
Turbo Encoder, Convolutional, 3GPP2/CDMA2000	Xilinx	V-IIP	V-II		S-3		
Turbo Encoder, Product Code	Xilinx	V-IIP	V-II		S-3		
Viterbi Decoder, General Purpose	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Viterbi Decoder, IEEE 802-compatible	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
XAPP 289: Common Switch Interface CSIX-L1							
Reference Design	Xilinx	V-IIP	V-II				
XAUI	Xilinx	V-IIP					
Digital Signal Processing							
Discrete Cosine Transform (eDCT)	eInfochips Pvt. Ltd.		V-II	V-E		S-IIE	S-II
Discrete Cosine Transform, 2D Inverse (IDCT)	CAST, Inc.		V-II	V-E			S-II
Discrete Cosine Transform, Combined 2D Forward/Inverse (DCT_FI)	CAST, Inc.		V-II	V-E			S-II
Discrete Cosine Transform, Forward 2D (DCT)	CAST, Inc.		V-II	V-E			S-II
Discrete Cosine Transform, Forward/Inverse (FIDCT)	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
Discrete Cosine Transform, forward/inverse 2D (DCT/IDCT 2D)	Barco-Silex		V-II	V-E			S-II
Discrete Wavelet Transform, Combined 2D Forward/Inverse (RC_2DDWT)	CAST, Inc.		V-II	V-E			S-II
Discrete Wavelet Transform, Line-based programmable forward (LB_2DFDWT)	CAST, Inc.		V-II	V-E			S-II
FIR Filter using DPRAM	eInfochips Pvt. Ltd.		V-II	V-E		S-IIE	S-II
FIR Filter, Parallel Distributed Arithmetic	eInfochips Pvt. Ltd.		V-II	V-E		S-IIE	S-II
TMS32025 DSP Processor (C32025)	CAST, Inc.		V-II	V-E			S-II
Bit Correlator	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
Cascaded Integrator Comb (CIC) Filter	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
CORDIC	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Digital Down Converter (DDC)	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
Digital Up Converter (DUC)	Xilinx	V-IIP	V-II		S-3		
Direct Digital Synthesizer (DDS)	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
DOCSIS ITU-T J.83 Modulator	Xilinx	V-IIP	V-II		S-3		
Fast Fourier Transform	Xilinx	V-IIP	V-II		S-3		
FFT/IFFT for Virtex-II, 1024-Point Complex	Xilinx		V-II				
FFT/IFFT for Virtex-II, 16-Point Complex	Xilinx		V-II				
FFT/IFFT for Virtex-II, 256-Point Complex	Xilinx		V-II				
FFT/IFFT for Virtex-II, 64-Point Complex	Xilinx		V-II				
FFT/IFFT, 1024-Point Complex	Xilinx			V-E			
FFT/IFFT, 16-Point Complex	Xilinx		V-II	V-E			
FFT/IFFT, 256-Point Complex	Xilinx		V-II	V-E			
FFT/IFFT, 32-Point Complex	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
FFT/IFFT, 64-, 256-, 1024-Point Complex	Xilinx	V-IIP	V-II		S-3		
FFT/IFFT, 64-Point Complex	Xilinx			V-E			
FIR Filter, Distributed Arithmetic (DA)	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
FIR Filter, MAC	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
LFSR, Linear Feedback Shift Register	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Math Function							
Floating Point Adder (DFPADD)	Digital Core Design		V-II				S-II
Floating Point Comparator (DFPCOMP)	Digital Core Design		V-II				S-II
Floating Point Divider (DFPDIV)	Digital Core Design		V-II				S-II
Floating Point Multiplier (DFPMUL)	Digital Core Design		V-II	V-E			S-II
Floating Point Square Root Operator (DFPSQRT)	Digital Core Design		V-II	V-E			S-II
Floating Point to Integer Converter (DFP2INT)	Digital Core Design		V-II				S-II
Integer to Floating Point Converter (DINT2FP)	Digital Core Design		V-II				S-II
Accumulator	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Adder Subtracter	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Divider, Pipelined	Xilinx		V-II	V-E		S-IIE	S-II
Multiply Accumulator (MAC)	Xilinx	V-IIP	V-II		S-3	S-IIE	S-II
Multiply Generator	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

Function	Vendor Name	Virtex-II Pro	Virtex-II	Virtex-E	Spartan-3	Spartan-IIE	Spartan-II
Sine Cosine Look Up Table	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Twos Complementer	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Memories & Storage Element							
RLDRAM Memory Controller	Avnet Design Services		V-II				
SDRAM Controller, DDR (EP525)	Eureka Technology	V-IIP	V-II		S3	S-IIE	
SDRAM Controller, DDR (MC-XIL-SDRAMDDR)	Memec Design		V-II	V-E			S-II
Block Memory, Dual-Port	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Block Memory, Single-Port	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Content Addressable Memory (CAM)	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Distributed Memory	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
FIFO, Asynchronous	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
FIFO, Synchronous	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Microprocessor, Controller & Peripheral							
16450 UART (H16450)	CAST, Inc.		V-II	V-E		S-IIE	S-II
16450 UART w/Synchronous Interface (H16450S)	CAST, Inc.		V-II	V-E		S-IIE	S-II
16550 UART w/FIFOs & synch interface (H16550S)	CAST, Inc.		V-II	V-E		S-IIE	S-II
16550 UART w/FIFOs (H16550)	CAST, Inc.		V-II	V-E		S-IIE	S-II
2910A Microprogram Controller (C2910A)	CAST, Inc.						S-II
2D Multiprocessing Interface Fabric (2D-fabric402C)	Crossbow Technologies, Inc.	V-IIP	V-II			S-IIE	S-II
68000 Compatible Microprocessor (C68000)	CAST, Inc.	V-IIP	V-II	V-E	S3	S-IIE	
80186 Compatible Microprocessor (e80186)	eInfochips Pvt. Ltd.		V-II				
8051 Base Compatible Microcontroller (DR8051BASE)	Digital Core Design		V-II				S-II
8051 Compatible Microcontroller (C8051)	CAST, Inc.	V-IIP	V-II	V-E	S3	S-IIE	S-II
8051 Compatible Microcontroller (FLIP8051 Thunder)	Dolphin Integration	V-IIP	V-II			S-IIE	
8051 Microcontroller, PicoBlaze Emulated (PB8051-MX/TF)	Roman-Jones, Inc.	V-IIP	V-II		S3	S-IIE	
8051 RISC Microcontroller (DR8051)	Digital Core Design		V-II				S-II
80515 High-speed 8-bit RISC Microcontroller (R80515)	CAST, Inc.						
8052 Compatible Microcontroller (DR8052EX)	Digital Core Design		V-II				S-II
80C51 Compatible RISC Microcontroller (R8051)	CAST, Inc.			V-E			S-II
8237 Programmable DMA Controller (C8237)	CAST, Inc.		V-II	V-E			S-II
8250 UART (H8250)	CAST, Inc.		V-II	V-E		S-IIE	S-II
8254 Programmable Interval Timer/Counter (C8254)	CAST, Inc.		V-II	V-E			S-II
8254 Programmable Interval Timer/Counter (e8254)	eInfochips Pvt. Ltd.		V-II				S-II
8255 Programmable I/O Controller (e8255)	eInfochips Pvt. Ltd.		V-II				S-II
8259A Programmable Interrupt Controller (C8259A)	CAST, Inc.		V-II	V-E			S-II
Compact Video Controller (logiCVC)	Xylon d.o.o.		V-II				S-II
CRT Controller (C6845)	CAST, Inc.		V-II	V-E		S-IIE	S-II
FPU for Microblaze (Quixilica)	QinetiQ Limited	V-IIP	V-II				
Internet Appliance (socPiP-1A_Platform)	SoC Solutions, LLC		V-II	V-E			
Java Processor, 32-bit (Lightfoot)	Digital Communications Technologies, Ltd.		V-II				S-II
Java Processor, Configurable (LavaCORE)	Derivation Systems, Inc.		V-II				
MIPS system controller (ES500)	Eureka Technology		V-II	V-E			
Motor Controller - 3 phase (MLCA_4)	MEET Ltd.					S-IIE	S-II
Operating System Accelerator (Sierra S16)	RealFast Operating Systems AB		V-II			S-IIE	S-II
PIC125x Fast RISC Microcontroller (DFPIC125X)	Digital Core Design		V-II				S-II
PIC1655x Fast RISC Microcontroller (DFPIC1655X)	Digital Core Design		V-II				S-II
PIC165X Compatible Microcontroller (C165X)	CAST, Inc.	V-IIP	V-II	V-E	S3	S-IIE	
PIC165x Fast RISC Microcontroller (DFPIC165X)	Digital Core Design		V-II	V-E			S-II
PIC16C55X Compatible RISC Microcontroller (C1655x)	CAST, Inc.		V-II	V-E		S-IIE	S-II
PowerPC Bus Master (EP201)	Eureka Technology	V-IIP	V-II		S3	S-IIE	
PowerPC Bus Slave (EP100)	Eureka Technology	V-IIP	V-II		S3	S-IIE	
RISC Processor, 16-bit Proprietary (AX1610)	Loarant Corporation	V-IIP	V-II		S3	S-IIE	
UART, Generic Compact (MC-XIL-UART)	Memec Design	V-IIP	V-II		S3		
XTENSA-V Configurable 32-bit Microprocessor	Tensilica, Inc.		V-II				
Z80 Compatible Microprocessor (CZ80CPU)	CAST, Inc.		V-II	V-E	S3		S-II
Z80 Compatible Programmable Counter/Timer (CZ80CTC)	CAST, Inc.	V-IIP	V-II	V-E	S3	S-IIE	
Z80 Peripheral I/O Controller (CZ80PIO)	CAST, Inc.		V-II	V-E		S-IIE	S-II
1 Gigabit Ethernet MAC w/PLB interface	Xilinx	V-IIP					

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

Function	Vendor Name	Virtex-II Pro	Virtex-II	Virtex-E	Spartan-3	Spartan-IIE	Spartan-II
10/100 Ethernet MAC Lite w/OPB interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
10/100 Ethernet MAC w/OPB interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
16450 UART w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
16550 UART w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Arbiter and Bus Structure w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Arbiter and Bus Structure w/PLB interface	Xilinx	V-IIP					
ATM Utopia Level 2 Master and Slave w/OPB Interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
ATM Utopia Level 2 Master and Slave w/PLB Interface	Xilinx	V-IIP					
BRAM Controller w/LMB interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
BRAM Controller w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
BRAM Controller w/PLB interface	Xilinx	V-IIP					
BSP Generator (SW only)	Xilinx	V-IIP					
DCR Bus Structure	Xilinx	V-IIP					
External Memory Controller (EMC) w/OPB interface (Includes support for Flash, SRAM, ZBT, System ACE)	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
External Memory Controller (EMC) w/PLB interface (Includes support for Flash, SRAM, ZBT, System ACE)	Xilinx	V-IIP					
GPIO w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
HDLC Controller (Single Channel) w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
HDLC Controller (Multi (256) Channel) w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
IIC w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Interrupt Controller (IntC) w/DCR interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
Interrupt Controller (IntC) w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
IPIF Address Decode w/OPB interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
IPIF DMA w/OPB interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
IPIF Interrupt Controller w/OPB interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
IPIF Master/Slave Attachment w/OPB interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
IPIF Read/Write Packet FIFO w/OPB interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
IPIF Scatter/Gather w/OPB interface	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
IPIF Slave Attachment w/PLB interface	Xilinx	V-IIP					
JTAG UART w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Memory Test Utility (SW only)	Xilinx	V-IIP					
MicroBlaze Soft RISC Processor	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
MicroBlaze Source Code	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
ML300 VxWorks BSP (SW only)	Xilinx	V-IIP					
OPB2DCR Bridge	Xilinx	V-IIP					
OPB2OPB Bridge (Lite)	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
OPB2PCI Full Bridge (32/33)	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
OPB2PLB Bridge	Xilinx	V-IIP					
PicoBlaze (XAPP 213: PicoBlaze 8-bit Microcontroller for Virtex-E and Spartan-II/E Devices) - REFERENCE DESIGN	Xilinx			V-E		S-IIE	S-II
PicoBlaze (XAPP 627: PicoBlaze 8-bit Microcontroller for Virtex-II and Virtex-II Pro Devices) - REFERENCE DESIGN	Xilinx	V-IIP	V-II				
PLB2OPB Bridge	Xilinx	V-IIP					
SDRAM Controller w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
SDRAM Controller w/PLB interface	Xilinx	V-IIP					
SPI Master and Slave w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Systems Reset Module	Xilinx	V-IIP					
Timebase/Watch Dog Timer (WDT) w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
Timer/Counter w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
UART Lite w/OPB interface	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
UltraContoller Solution: A lightweight PowerPC Microcontroller (XAPP672) - REFERENCE DESIGN	Xilinx	V-IIP					
VxWorks Board Support Package (BSP)	Xilinx	V-IIP					
Standard Bus Interface							
LIN - Local Interconnect Network Bus Controller (iLIN)	Intelliga Integrated Design, Ltd.	V-IIP	V-II	V-E	S-3	S-IIE	
PCI 64-bit/66-MHz master/target interface (EC240)	Eureka Technology		V-II	V-E			
PCI Host Bridge (EP430)	Eureka Technology		V-II	V-E			

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

Function	Vendor Name	Virtex-II Pro	Virtex-II	Virtex-E	Spartan-3	Spartan-IIE	Spartan-II
PCI, 64-bit Target Interface (PCI-T64)	CAST, Inc.	V-IIP	V-II			S-IIE	
PCI-PCI Bridge (EP440)	Eureka Technology	V-IIP	V-II		S-3	S-IIE	
Serial Protocol Interface Slave (SPI Slave)	CAST, Inc.		V-II	V-E		S-IIE	
Two-Wire Serial Interface - I2C (MC-XIL-TWSI)	Memec Design	V-IIP	V-II	V-E	S-3	S-IIE	S-II
USB 1.1 Function Controller (CUSB)	CAST, Inc.	V-IIP	V-II		S-3	S-IIE	
USB 2.0 Function Controller (CUSB2)	CAST, Inc.	V-IIP	V-II	V-E	S-3	S-IIE	
Arbiter	Telecom Italia Lab S.p.A.	V-IIP	V-II			S-IIE	
CAN 2.0 B Compatible Network Controller (LogiCAN)	Xylon d.o.o.	V-IIP	V-II		S-3	S-IIE	
CAN Bus Controller (MC-XIL-OPB-XCAN)	Memec Design	V-IIP	V-II	V-E	S-3	S-IIE	S-II
CAN Bus Controller 2.0B	CAST, Inc.	V-IIP	V-II		S-3	S-IIE	
CAN Bus Controller with 32 Mail Boxes	Robert Bosch GmbH	V-IIP	V-II		S-3	S-IIE	
HyperTransport Cave, 8-bit	GDA Technologies, Inc.	V-IIP	V-II				
I2C Bus Controller (I2C)	CAST, Inc.		V-II	V-E	S-3		S-II
I2C Bus Controller Master (DI2CM)	Digital Core Design		V-II	V-E			S-II
I2C Bus Controller Slave (DI2CS)	Digital Core Design		V-II	V-E			S-II
I2C Bus Controller Slave Base (DI2CSB)	Digital Core Design		V-II	V-E			S-II
I2C Two-Wire Serial Interface Master-Only (MC-XIL-TWSIMO)	Memec Design		V-II	V-E		S-IIE	S-II
I2C Two-Wire Serial Interface Master-Slave (MC-XIL-TWSIMS)	Memec Design		V-II	V-E		S-IIE	S-II
HyperTransport Single-Ended Slave Core	Xilinx	V-IIP	V-II				
Advanced Switching Endpoint Core	Xilinx	V-IIP					
PCI Express Endpoint Core	Xilinx	V-IIP					
PCI32 Interface Design Kit (DO-DI-PCI32-DKT)	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
PCI32 Interface, IP Only (DO-DI-PCI32-IP)	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
PCI32 Single-Use License for Spartan (DO-DI-PCI32-SP)	Xilinx				S-3	S-IIE	S-II
PCI64 & PCI32, IP Only (DO-DI-PCI-AL)	Xilinx	V-IIP	V-II	V-E	S-3	S-IIE	S-II
PCI64 Interface Design Kit (DO-DI-PCI64-DKT)	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
PCI64 Interface, IP Only (DO-DI-PCI64-IP)	Xilinx	V-IIP	V-II	V-E		S-IIE	S-II
PCI-X 64/133 Interface for Virtex-II (DO-DI-PCIX64-VE). Includes PCI 64 bit interface at 33 MHz	Xilinx	V-IIP	V-II				
PCI-X 64/66 Interface for Virtex-E (DO-DI-PCIX64-VE). Includes PCI 64 bit interface at 33 MHz	Xilinx			VE			
RapidIO 8-bit port LP-LVDS Phy Layer (DO-DI-RIO8-PHY)	Xilinx	V-IIP	V-II				
RapidIO Logical (I/O) and Transport Layer (DO-DI-RIO8-LOG)	Xilinx	V-IIP	V-II				
RapidIO Phy Layer to PLB Bridge reference design - REFERENCE DESIGN	Xilinx	V-IIP					
XAPP653: Virtex-II Pro/Spartan-3 3.3V PCI Reference Design - REFERENCE DESIGN	Xilinx	V-IIP			S-3		
Backplanes and Gigabit Serial I/O							
Aurora 201, 401and 804 Designs - REFERENCE DESIGN	Xilinx	V-IIP					
WP160: Emulating External SERDES Devices with Embedded RocketIO Transceivers - WHITE PAPER	Xilinx	V-IIP					
XAPP 649: SONET Rate Conversion in Virtex-II Pro Devices - REFERENCE DESIGN	Xilinx	V-IIP					
XAPP 651: SONET and OTN Scramblers/Descramblers - REFERENCE DESIGN	Xilinx	V-IIP					
XAPP 652: Word Alignment and SONET/SDH Framing - REFERENCE DESIGN	Xilinx	V-IIP					
XAPP660: Partial Reconfiguration of RocketIO Attributes using PPC405 core (DCR Bus) - REFERENCE DESIGN	Xilinx	V-IIP					
XAPP661: RocketIO Transceiver Bit-Error Rate Tester (BERT) - REFERENCE DESIGN	Xilinx	V-IIP					
XAPP662: Partial Reconfig. of RocketIO Attributes using PPC405 core (PLB or OPB bus) + RocketIO Transceiver Bit-Error Rate Tester (BERT) - REFERENCE DESIGN	Xilinx	V-IIP					

Visit the Xilinx IP Center for more details at *www.xilinx.com/ipcenter*

XILINX®
The Programmable Logic Company℠

Xilinx Global Services

http://www.xilinx.com/support/gsd/index.htm

Part Number	Product Description	Duration
Education Services		hours
FPGA13000-6-ILT	Fundamentals of FPGA Design	8
FPGA23000-6-ILT	Designing for Performance	16
FPGA33000-6-ILT	Advanced FPGA Implementation	16
LANG11000-5-ILT	Introduction to VHDL	24
LANG21000-5-ILT	Advanced VHDL	16
LANG12000-5-ILT	Introduction to Verilog	24
PCI8000-4-ILT	PCI CORE Basics	8
PCI28000-4-ILT	Designing a PCI System	16
PCI2900-5-ILT	Designing for PCI-X	16
DSP2000-3-ILT	DSP Implementation Techniques for Xilinx FPGAs	24
DSP-100001-5-ILT	DSP Design Flow	24
RIO22000-6-ILT	Designing with Multi-Gigabit Serial I/O	16
EMBD-21000-5-ILT	Embedded Systems Development	16
Platinum Technical Service		hours
SC-PLAT-SVC-10	Seat Platinum Technical Service w/10 education credits	N/A
SC-PLAT-SITE-50	Platinum Technical Service site license up to 50 customers	N/A
SC-PLAT-SITE-100	Platinum Technical Service site license for 51-100 customers	N/A
SC-PLAT-SITE-150	Platinum Technical Service site license for 101-150 customers	N/A
Titanium Technical Service		hours
PS-TEC-SERV	Titanium Technical Service (minimum 40 hours)	N/A
DC-DES-SERV	Design Services Contract	N/A
Xilinx Productivity Advantage		hours
DS-XPA-10K	Custom XPA for $0 – $10,000	N/A
DS-XPA-25K	Custom XPA for $10,001 – $25,000	N/A
DS-XPA	Custom XPA for $25,001 – $100,000	N/A
DS-XPA-10K-INT	Custom XPA (International) for $0 – $10,000	N/A
DS-XPA-25K-INT	Custom XPA (International) for $10,001 – $25,000	N/A
DS-XPA-INT	Custom XPA (International) for $25,001 – $100,000	N/A
DS-ISE-ALI-XPA	XPA Seat, ISE Alliance	N/A
DS-ISE-FND-XPA	XPA Seat, ISE Foundation	N/A
Promotion Packages		hours
PROMO-5003-6-ILT	Designing for Performance, Live Online	9
PROMO-5004-6-ILT	FPGA Essentials: Includes both Fundamentals of FPGA Design and Designing for Performance classes	24
DO-EDK-T	EDK Bundle (EDK + 10 TCs towards Embedded Systems Dev. Course)	NA
DS-SYSGEN-4SL-T	DSP Bundle (Sysgen + 15 TCs towards DSP Design Flow Course)	NA

MySupport

- Sign up for personalized email alerts at mysupport.xilinx.com
- Search our knowledge database
- Troubleshoot with Problem Solvers
- Consult with engineers in Forums

Xilinx Productivity Advantage (XPA) Program

The Xilinx Productivity Advantage (XPA) delivers everything your customer needs to create their best design using Xilinx Programmable Solutions, including Software, Support, Services and IP cores. A single PO allows designers to get what they need, when they need it, at a great price.

Two types of XPAs are available. The custom XPA, tailored to your customers specific requirements, and the XPA Seat, an off-the-shelf pre-packaged solution for the individual designer.

http://support.xilinx.com/support/gsd/xpa_program.htm

Xilinx Design Services (XDS)

- XDS provides extensive FPGA hardware and embedded software design experience backed by industry recognized experts and resources to solve even the most complex design challenge.
- **System Architecture Consulting** — Provide engineering services to define system architecture and partitioning for design specification.
- **Custom Design Solutions** — Project designed, verified, and delivered to mutually agreed upon design specifications.
- **IP Core Development, Optimization, Integration, Modification, and Verification** — Modify, integrate, and optimize customer intellectual property or third party cores to work with Xilinx technology. Develop customer-required special features to Xilinx IP cores or third party cores. Perform integration, optimization, and verification of IP cores in Xilinx technology.
- **Embedded Software** — Develop complex embedded software with real-time constraints, using hardware/software co-design techniques.
- **Conversions** — Convert ASIC designs and other FPGAs to Xilinx technology and devices.

Titanium Technical Service

- Dedicated Application Engineer
- Service at customer site or at Xilinx
- Design flow methodology coaching
- Engineer/expertise specific to customer need
- Service guaranteed via a contract

Platinum Technical Service

- Access to a dedicated team of Senior Applications Engineers
- Dedicated toll free number*
- Priority case resolution
- Ten Education Credits
- Electronic newsletter
- Formal escalation process
- Service Packs and Software updates
- Application Engineer to customer ratio is 2x the Gold Level

**Toll free number available in US only, dedicated local numbers available across Europe*

Education Services

- Reduce time to knowledge
- Improve design efficiency
- Reduce overall development costs

Public and private classes available worldwide

http://www.xilinx.com/education

Xilinx Global Services
Finish Faster

Education Services Contacts

North America: 877-XLX-CLAS (877-959-2527)

Europe: +44-870-7350-548
eurotraining@xilinx.com

Japan: +81-3-5321-7750

Asia Pacific: +852-2424-5200
education_ap@xilinx.com

www.xilinx.com/education

Design Services Contacts

North America:
Richard Fodor: 408-626-4256
Mike Barone: 512-238-1473

Europe:
Alex Hillier: +44-870-7350-516
Martina Finnerty: +353-1-403-2469

Asia Pacific:
David Keefe: +852-2401-5171

designservices@xilinx.com

XPA Contacts

North America: 800-888-FPGA (3742)
fpga.xilinx.com

Europe:
Stuart Elston: +44-870-7350-532

Asia Pacific:
David Keefe: +852-2401-5171

Titanium Technical Service Contacts

North America:
Telesales: 800-888-3742

Europe:
Stuart Elston: +44-870-7350-632

Japan:
+81-3-5321-7730
or japantitanium@xilinx.com

Asia Pacific:
David Keefe: +852-2401-5171

Avoid messy timing mistakes. Use Mentor Graphics® FPGA design tools.

FPGA Design

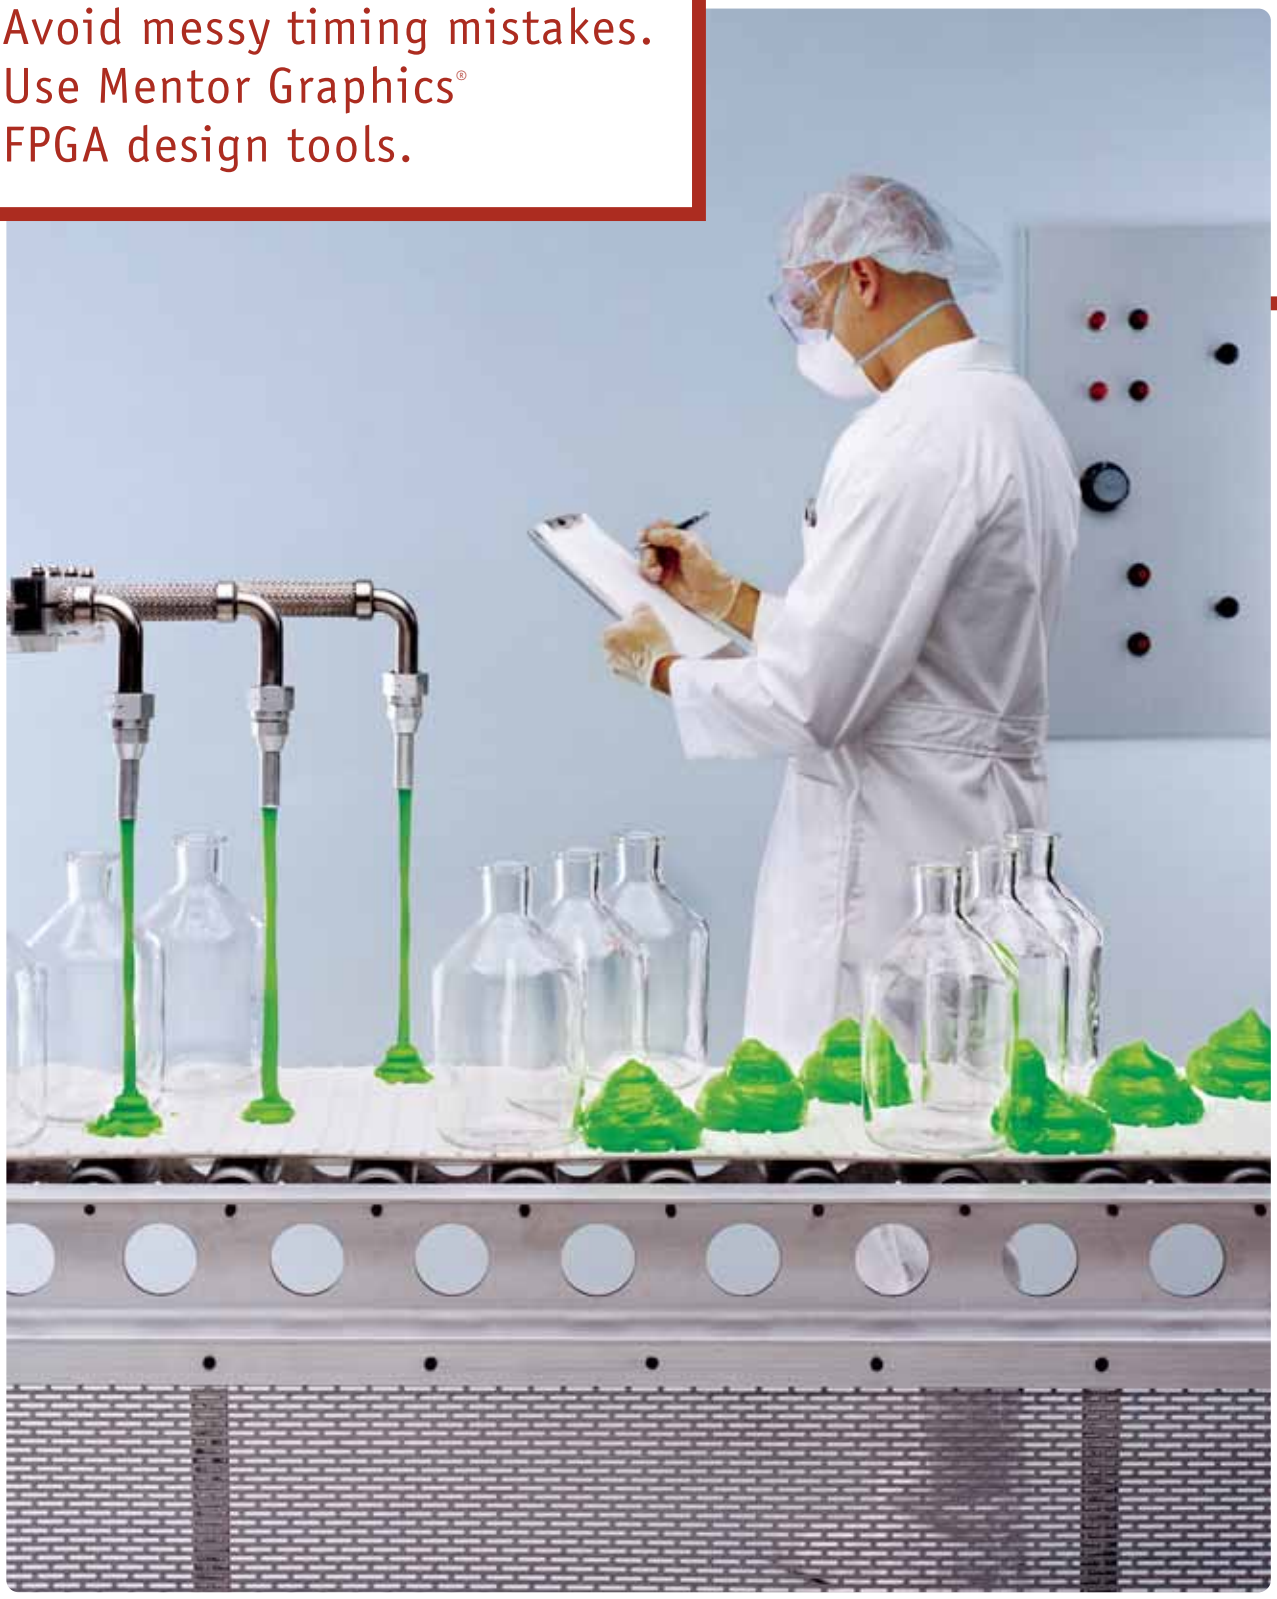

PRECISION™ SYNTHESIS | As any FPGA designer can tell you, achieving precise timing is their single most critical challenge. That's why Mentor Graphics® created Precision™ Synthesis, a powerful new tool suite that lets designers close on timing faster than ever. Precision™ Synthesis provides the most comprehensive analysis for complex FPGAs with the only complete, built-in incremental timing analysis. Quickly find and optimize the most critical paths and avoid frustrating trial-and-error design iterations. Visit **www.mentor.com/fpga** today or call **877.387.5873** for more information on how Precision™ Synthesis can help you find the fastest path to a completed design.

Mentor Graphics®

LOWEST SYSTEM COST. HIGHEST SYSTEM PERFORMANCE.

Virtex-II Pro™ FPGAs deliver more capabilities and performance than any other FPGA.

In a single device, you get the highest density, most memory, best performance, and at no additional charge 400MHz PowerPC™ processors and 3.125 Gbps RocketIO™ or 10.3125 Gbps RocketIO X serial transceivers. This gives you the fastest DSP, connectivity and processing solutions in the industry.

The Power to Develop Your Design

The Xilinx ISE software tools are the easiest to use solution for high-density logic. With over 200 IP cores, ChipScope Pro debug environment, and compile times up to 6x faster than our nearest competitor, you can quickly take your Virtex-II Pro FPGA design from concept to reality.

Guaranteed 25% And Up To 80% Savings With Virtex-II Pro EasyPath Solutions

For high-volume system production, Xilinx offers a lower cost path with Virtex-II Pro EasyPath. You can immediately take advantage of dramatic cost reduction at no risk and no effort. Only Xilinx can offer you this flexibility for design and production.

VIRTEX-II PRO™ EASYPATH SOLUTIONS

Visit ***www.xilinx.com/virtex2pro*** today to get the price and performance you're looking for.

XILINX®
The Programmable Logic Company℠

www.xilinx.com/virtex2pro

Earth Friendly
Pb-free devices available now

FORTUNE® 2004 100 BEST COMPANIES TO WORK FOR

20 YEARS XILINX

In Production Now

PN 0010782