THE PLATFORM FOR POWERFUL SOLUTIONS.

Virtex-II Special Edition

The Future of Platform FPGAs

Virtex-II IP-Immersion™ Technology Enables Next-Generation Platform FPGAs

Discontinuity at the Gate – A New Era in FPGA Design

Extinct: Dinosaurs, Slide Rules, 8-Track Tapes, and now... External Termination Resistors

SystemIO Technology Promises High-Speed Connectivity Across Multiple I/O Standards

Virtex-II Platform FPGAs Support System Packet Interface Standards for Optical Networks

Video Demonstration Board

4G Wireless Systems in Virtex-II

XILINX®

INTRODUCING ISE 4.1i

All the speed you need.

ISE

The fastest programmable design software just got faster

Get ready to fly. Our ISE 4.1i software release completes your design twice as fast as our previous generation, and up to 6 times faster than competitive tools. With an array of industry-unique features, ISE 4.1i gives you unbeatable performance advantages, increasing your design's clock performance by up to 75%, and enabling I/O speeds of 3.125 Gbits/sec.

Performance and Speed with ProActive™ Timing Closure

ISE 4.1i features Xilinx ProActive Timing Closure - a set of new technologies including physical synthesis integration, HDL analysis, timing wizards with cross-probing, and the most advanced timing-driven place and route algorithms available in the industry today. This means ISE 4.1i will intelligently identify, analyze, and solve any potential timing bottlenecks. Also, our new Directed Routing technology delivers unprecedented improvements in performance predictability for IP.

VIRTEX-II

SPARTAN-II

XPLA3 CoolRunner

XC9500XV

A single design environment for *all* Xilinx devices

Combine the superior performance of ISE 4.1i with the speed and density leadership of the Virtex®-II Platform FPGAs, the Spartan®-II ASIC replacement FPGAs, the ultra-low power CoolRunner™ CPLDs, plus our high-performance XC9500™ CPLDs. With the first and only FPGA support for industry-leading formal verification tools, and clock speeds over 300MHz, Xilinx provides the totally integrated design solution.

Visit *www.xilinx.com/v2insert/swad* today for all the *speed* you need!

XILINX®
The Programmable Logic Company℠

www.xilinx.com/ise/xcell

FORTUNE® 2001
100 BEST COMPANIES TO WORK FOR

For more information visit xilinx.com/v2insert/swad

The Future of Platform FPGAs

by Wim Roelandts
CEO, Xilinx

A look at the Xilinx philosophy behind the Virtex family development.

The Virtex FPGA family is, without a doubt, the most advanced programmable logic solution ever conceived. Because of its system-level features, extreme density, and high performance, this family has given you design options that were never possible before, and has made it much easier for you to produce better designs in less time.

We designed the Virtex architecture, from the very beginning, as a technology platform on which we can build future generations of the Virtex family. This platform is optimized for use with both hard cores and soft cores, allowing us to offer you the flexibility of programmable logic along with the performance advantages of embedded hard logic, all tightly integrated with our high level development tools that significantly increase your productivity.

The Virtex family just keeps getting better, and the future looks very bright.

Performance Philosophy

Our first priority, with all of our devices, is to continuously improve the performance, density, and features. We work very closely with our manufacturing partners to refine our manufacturing processes, creating increasingly smaller geometry CMOS technologies that result in faster, denser devices, for less cost – we will introduce a new generation every year with increasingly advanced process technology. In addition, we are continuously developing new ways to improve our device architectures, to get better performance through enhanced routing and design features.

Today, you can purchase Virtex-II devices with up to 6 million system gates, a huge advancement in density over previous FPGAs. Yet, within three or four years we will offer 50 million gate devices – enough logic to build very complex, very high performance systems, on a single chip. In addition, Virtex-II devices now operate at internal clock speeds above 200 Mhz, the equal of many custom ASICs. Yet, every year, for the next four to five years, we expect a performance increase of 30% to 50%. (The inherent silicon performance increase is about 30% per year. Then, as we improve the routing infrastructure and so on, we expect up to another 20%.) As you can see, programmable logic technology is advancing very quickly, giving you more options, more capability, more flexibility, and more reasons to move away from ASICs and fixed logic designs.

The basic structure of an FPGA determines not only its capability and its ease of use, it also determines its ability to evolve as new technologies are developed and implemented. You want your FPGA family to grow with your needs without having to learn new tools, processes, and design techniques. That's why we developed a flexible, highly predictable, forward-thinking architecture that can easily integrate custom logic, soft cores, hard cores and mixed signal capability.

With the Virtex-II family you'll not only achieve high performance, you'll do it with a high degree of predictability and family stability which is key to your productivity. Plus, migrating to larger devices with higher performance, as they are developed, is easy.

It all starts with a strong FPGA platform, tightly integrated with fast, high-level development tools – and there is no end in sight to where this family can go.

Core Philosophy

As device densities keep increasing, it becomes even more important that we provide a wide range of intellectual property or cores, which help you quickly develop your design. Without cores, it would take many engineer-years to complete a 8-million gate design. With cores you can quickly create key parts of your system using proven, reliable designs.

Our "platform" philosophy is to provide both hard and soft cores that take full advantage of our Virtex architecture. Hard cores (such as the PowerPC) are actually fixed logic designs that we incorporate into the FPGA device architecture. Wherever possible, we'll offer soft cores to solve your design challenges because they are more flexible and are used on an as needed basis. We'll offer hard cores when there is a performance or density advantage. The Virtex architecture allows you to easily integrate both types of cores into your designs, giving you the maximum flexibility and performance

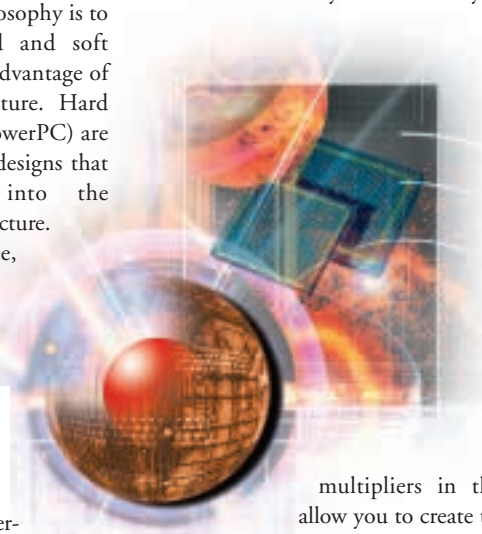

Hard Cores

As we move forward, we'll integrate more and more hard cores into our FPGA platform to increase the performance and ease of use. Examples include central processors, memory blocks, clock managers, multipliers, and high speed I/O systems.

Processor cores save you a lot of development time and they give you a known, reliable design. Our philosophy is to tightly integrate our processor cores into the FPGA fabric so you can achieve tremendous performance advantages that would not be available if you used a separate processor chip.

Memory is a critical part of most designs. The ratio of memory to logic gates in the Virtex family will continue to increase over time because our customers are demanding more and more memory. The amount of block RAM and distributed RAM will increase, as well as the ability to access off-chip memory – as memory standards evolve, so will our memory interface capability.

Clock management is another critical factor in large designs. The Virtex-II Delay Locked Loop Digital Clock Manager is already the most advanced, feature rich, clock manager in the industry. It eliminates clock skew, provides very flexible clock synthesis capabilities, and gives you the ability to drive and synchronize clocks both on and off chip, thus eliminating external components and simplifying your design. It will evolve as we develop even more advanced techniques.

The embedded multipliers in the Virtex family allow you to create the fastest possible DSP designs. Our customers are achieving unprecedented speeds of well over 0.5 TeraMACs (Multiply-Accumulate Cycles per second). Many customers are pushing the limits of performance and density in their networking designs, requiring very sophisticated DSP algorithms to extract the data from the noise. Because these hard core multipliers are so useful, in a wide range of applications, they will be added to every Virtex-family FPGA.

We will continue to develop hard and soft cores that make full use of the Virtex family architecture, to bring you all the ease-of-use and performance advantages that make Platform FPGAs so attractive.

I/O Philosophy

Over the last few years we have made tremendous progress in the I/O capabilities of our devices. In the past, I/O blocks were

very standard, and they could usually support just one voltage. Today with the Virtex-II family we support the vast majority of I/O standards in the industry. If you look at a combination of different standards and drive currents, we have 49 different ways that you can program every single I/O pin. And we will continue to add new I/O capabilities as standards evolve.

High Speed Serial I/O

The demands of high speed networking, and other high performance systems, requires the use of gigabit-per-second serial I/O capability for interconnecting devices, backplanes, and systems. In addition, some of the new communications standards and backplane standards are based on these high speed serial I/O capabilities, including POS-PHY4, FlexBus4, HyperTransport™, InfiniBand™, Fibre Channel, Gigabit Ethernet, and so on. With the Virtex-II I/O capability you can connect directly to a backplane without external components.

Through our Conexant Skyrail™ licensing agreement, we gained access to the highest speed I/O technology available – currently giving us a serial transceiver capability of up to 3.125 Gbps. With our RocketChips® acquisition, we expect our serial transceiver technology to reach 10 Gbps or more, as our process and design technologies continue to improve. Future Virtex families will allow you to make full use of this important capability.

Processor Philosophy

We intend to offer you a choice of processors using both hard and soft cores; all using the same peripherals so you can easily combine processors in your designs.

PowerPC Hard-Core Processor

Our PowerPC hard core is being developed in partnership with IBM. It gives you a well-known, very high performance architecture. We will embed this processor within our programmable logic fabric, so all of the processor I/O pins are available to the internal programmable logic for maximum flexibility. Plus, the processor I/O pins do not take up valuable FPGA I/O resources, unless you need them. This allows you to move data much faster than the competition. Our competition's embedded processor does not have the same performance or flexibility as our PowerPC core.

Our philosophy is to provide all of the peripherals and so on as soft cores so they require no resources if you don't need them; we chose to provide the PowerPC processor as a hard core because it gives you a performance advantage. As IBM continues to improve the performance of the PowerPC, Xilinx will continue to offer the latest PowerPC processor core technology, optimized for the Virtex architecture.

MicroBlaze Soft-Core Processor

Our MicroBlaze™ soft-core processor was developed by Xilinx. It uses only about 800 logic cells, requires about the same physical space as the PowerPC, and runs at 125 Mhz. By next year, it will be running at over 150 Mhz.

MicroBlaze is fully integrated with the Core Connect architecture of IBM which means it can use the peripheral modules we're developing for the PowerPC processor. In fact, you can use it in addition to the power PC processor. For example, you can have a combination of the PowerPC and one or more MicroBlaze cores spread around it, all using the same memory and peripherals. The possibilities are limitless.

No other company has this flexible multiprocessor capability. Plus, our MicroBlaze soft core runs almost as fast as our competition's hard core processor

Board Integration Philosophy

With each new generation of the Virtex family we will integrate more and more of the discrete components that are required to create a working system, making your PC boards simpler and less expensive. Our goal is to make our I/O structure so extensive that you will never have to use glue logic or understand the intricacies of each new standard. For example, by integrating a variety of different memory interfaces into our FPGAs, you can easily connect any known memory device without having to create your own custom interface designs.

All of these trends will continue in the future – as new I/O standards are introduced, we'll make them available on our FPGAs.

XCITE

The Xilinx Controlled Impedance TEchnology (XCITE) is another example of PC board simplification and improved signal integrity. XCITE places digitally controlled termination resistors on the FPGA, so you don't have to manually terminate your signals with huge numbers of discrete external resistors. This not only saves you a lot of board space and cost, it makes board layout much simpler. This built-in termination adjusts itself for temperature and voltage variations as well, so your boards are not only less expensive, they are also more reliable. XCITE solves the signal integrity issues that both circuit and PCB designers are now dealing with, allowing you to run your PC boards at full speed and get them to market quickly.

Conclusion

Our FPGAs add more than just logic; they are tremendously more valuable because, they make your design simpler, they eliminate other components on your board, and they continue to decrease your development time and costs.

If you want to build the systems of the future, and keep your costs down, you need a solid foundation on which your designs can grow and evolve as technology advances. You need a logic solution that will grow with you and help you solve the problems that have yet to be encountered. You need the devices, software tools, and company support that make a complete solution. That's what you get with the Xilinx Virtex-II Platform FPGA family. It's already the industry leader, by far, and it just keeps getting better.

I hope you enjoy this Special Virtex-II supplement.

Virtex-II IP-Immersion™ Technology Enables Next-Generation Platform FPGAs

Innovative technology allows the integration of discrete silicon components within Platform FPGAs.

by Erich Goetting
Vice President,
Product Development, Xilinx
erich@xilinx.com

The Field Programmable Gate Array Revolution began when Ross Freeman, a founder of Xilinx, conceived the FPGA architecture. Abandoning the restrictions of sum-of-products architecture, Ross utilized a host of 16-bit LUTs (Look-Up-Tables), each accompanied by a flip-flop circuit, and all interconnected with programmable routing pathways. This revolutionary formula, first deployed in 1984 in the Xilinx XC2000 family, is still the basis of all FPGA devices today, despite the unprecedented growth in the scale of programmable logic and continual advances in the complexity of the device architecture. It is a testament to the power of Ross's architectural vision that the FPGA has withstood the test of time, and many proposed alternatives in the marketplace.

Introducing the Revolutionary Platform FPGA

Today, in 2001, however, another revolution in programmable logic has begun with the introduction of Xilinx Virtex®-II Platform FPGAs. At the heart of this revolution is the ability to integrate the functions of other discrete silicon devices, such as microprocessors, within an FPGA platform. The integration provided by the Platform FPGA architecture delivers these advantages:

- Increased performance made possible by high-bandwidth, low-latency coupling of intellectual property (IP) blocks
- Enhanced architectural flexibility by virtue of immersion within a high-performance programmable fabric
- Reduced board space, power, and cost.

Although Platform FPGAs represent a revolutionary step forward, they still retain the fundamental advantages over ASICs (Application Specific Integrated Circuits) – namely, reprogrammability, off-the-shelf availability, and zero non-recurring engineering (NRE) costs. These advantages comprise the infrastructure that has enabled the development of the innovative IP-Immersion superstructure.

Enhanced Performance Through IP-Immersion

The Virtex-II IP-Immersion™ architecture embodies the concept that high-bandwidth, hard-IP blocks – implemented in full-custom or ASIC-style standard cell logic – can be immersed within the matrix of FPGA CLBs (Configurable Logic Blocks). The two-dimensional array of Virtex-II CLBs is ideally suited for this task, because the array possesses three all-important properties:

1. **Configurable layout** – Because hard-IP blocks (such as microprocessors) have a particular shape, the designer can "cut out" just the right amount of CLBs, creating an empty space for the IP-block. As a result, high performance hard-IP blocks, such as the IBM PowerPC 405 32-bit RISC CPU, can be implemented using advanced circuit design and layout techniques that maximize performance and minimize silicon area. The Virtex-II IP-Immersion architecture accommodates virtually any pre-defined rectilinear shape.

IP-Immersion Technology

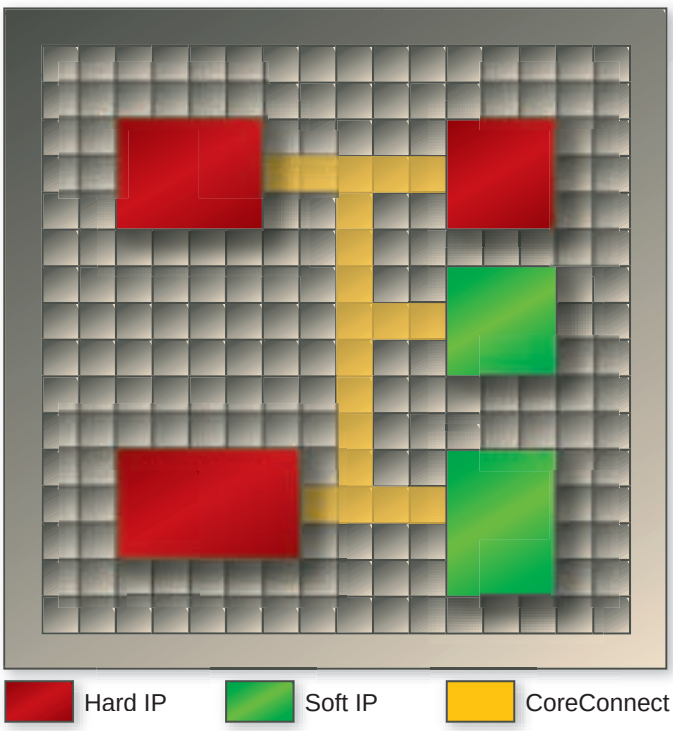

2. **Programmable routing** – Through its step-and-repeat of CLB tiles, the segmented routing of FPGA architecture allows the creation of specific "on-ramps" and "off-ramps" at every CLB border. In other words, because the Virtex-II routing architecture has some wiring segments that start within every CLB, these starting segments provide an ideal way for the hard-IP block to interconnect with the logic, memory, and I/Os of the FPGA platform. To provide the transition between the platform fabric and the hard-IP, Virtex-II devices introduce a new tile type: an "immersion tile." The immersion tile allows programmable interconnections between the IP-block and the fabric – much like the interconnection of discrete devices on a printed circuit board. For example, a designer can instantiate a large block and wire it to other parts of the system. In wiring this block, the designer can choose to connect an output pin to a net, leave an output unconnected, tie inputs to fixed one or zero levels, or connect an input to a particular net. In this way, the designer has full design flexibility in using IP blocks.

3. **High performance functionality** – The high wiring density and fully active nature of the Virtex-II routing architecture allows connections to occur in large quantities and at high speed, thus enabling the high-bandwidth interconnect necessary to fully exploit the potential of on-chip IP blocks.

Taken together, these three properties constitute the key ingredients of the new Virtex-II IP-Immersion architecture.

Conclusion

The Virtex-II series of Platform FPGAs are engineered to provide leading-edge functionality in logic, routing, clocking, DSP, memory, and I/O. Thanks to the innovative IP-Immersion architecture and development relationships with leading companies such as IBM, Xilinx Virtex-II Platform FPGAs are facilitating the next generation of advanced system designs.

Discontinuity at the Gate – A New Era in FPGA Design

The Xilinx vision in programmable logic will change how you do digital design. Mentor Graphics has recognized this and is committed to the FPGA market.

by Walden C. Rhines
Chairman and CEO, Mentor Graphics Corporation
walden_rhines@mentor.com

What happens when you give the design community a field-programmable hardware platform that contains 8 million system gates, 300+ MHz internal clock speeds, gigabit serial I/O performance, and IP immersion technology? You create an opportunity for an entirely new complex design methodology. By providing such robust capabilities in a programmable format, you create a discontinuity in the industry that removes previous cost and technology barriers from the product development process. The Xilinx Virtex-II Platform FPGA family embodies this emerging technology, and has provided the catalyst to change our current design methodologies.

Today it is estimated that there are 20,000 custom or semi-custom chip designers, and this number is growing very slowly, as shown in Figure 1. Under the current structure, issues including design styles, verification methodologies, NRE charges, risk, and software development inhibit how many circuits can be attempted each year. Creating a standard platform for development removes these limitations, and allows the number of potential designers to grow by an order of magnitude, to 200,000. This massive influx of potential designers creates a new brain pool for innovation.

FPGAs Coming of Age

The path to harmony between design tools and actual silicon is extremely challenging. For example, as shown in Figure 2, in the ASIC world it took 15 years to merge the silicon process with a solid design methodology based on reliable and functional EDA software. ASIC technology became the driving force in the industry. The process that began surrounding this technology created an effective solution for the electronics industry, which led to growth and innovation. But the ASIC process has matured to the point where it is applicable more for extremely high-end design, and it is slowly moving out of reach for the mass market.

In contrast, FPGA technology has taken only five years to get to the same level of functionality as ASICs. FPGA technology has uniquely solved the same problems the ASIC methodology addressed, but along the way it also minimized the NRE, risk, and manufacturing issues involved as illustrated in Figure 3. Today, the Xilinx Virtex-II technology is a legitimate ASIC replacement. When combined with leading-edge EDA software, Virtex-II FPGAs provide the electronics industry with a new, exciting path for growth and innovation. FPGAs are now the key vehicle driving the state-of-the-art for new electronic systems.

The Xilinx platform-based FPGA technology brings programmability to the System-on-Chip (SoC) methodology. Using a 0.15µ, 8-layer metal process with 0.12µ high-speed transistors, Virtex-II FPGAs provide designers with the performance and density they need to create an SOC design. With features like digital clock management, select I/O ultra technology, and active interconnect, designers can spend more time on functional verification, knowing that the main silicon issues and problems have already been solved. In addition, with IP immersion technology, FPGA designers can now work at a much higher level of abstraction and move the "gates per day" metric to a level where silicon utilization is maximized.

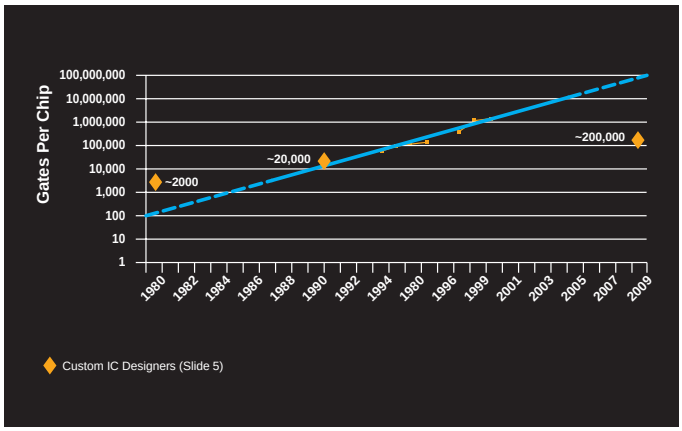

Figure 1 - Gate-to-designer ratio

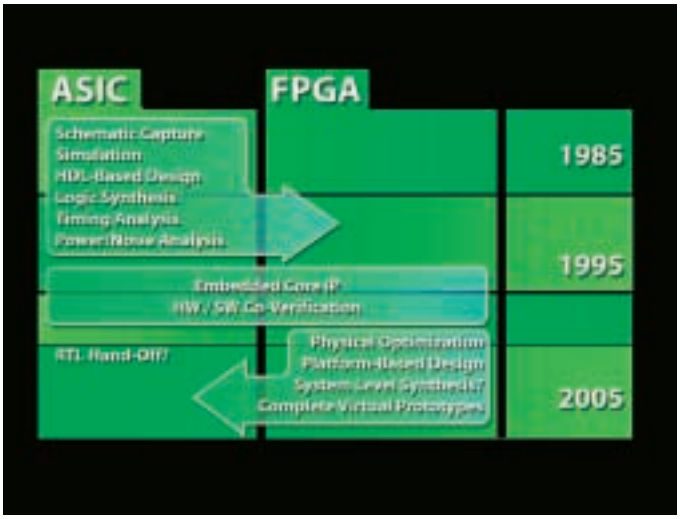

Figure 2. Evolution of ASIC and FPGA design methodologies

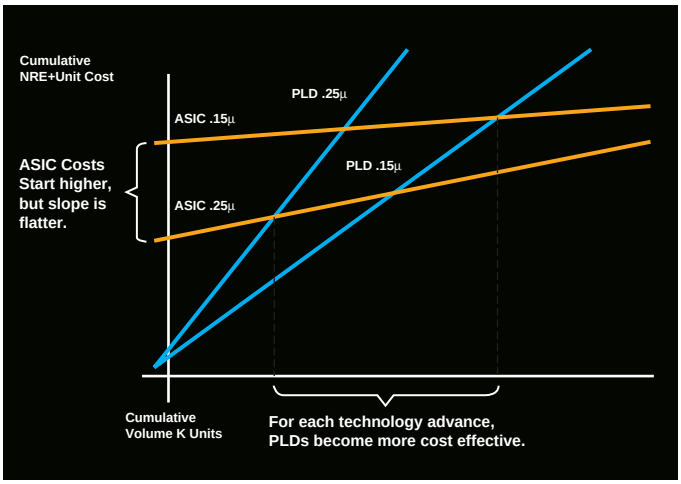

Figure 3. FPGA to ASIC crossover improves with process

Platform-Based FPGAs – The Value of History

Creating standards and putting boundary conditions on a design process can accelerate circuit development and shorten time to market. The ASIC SOC development process in place today is an open-ended approach to design with almost an infinite degree of freedom. This freedom provides flexibility and has enabled the creation of extremely complicated circuits, but it has also created a high risk methodology with a steep learning curve.

In contrast, platform-based FPGAs provide a structured approach to design. From the designer's point of view, the beauty of creating an FPGA is that the FPGA vendors worry about all the issues of the design process (silicon, methodology, and software).

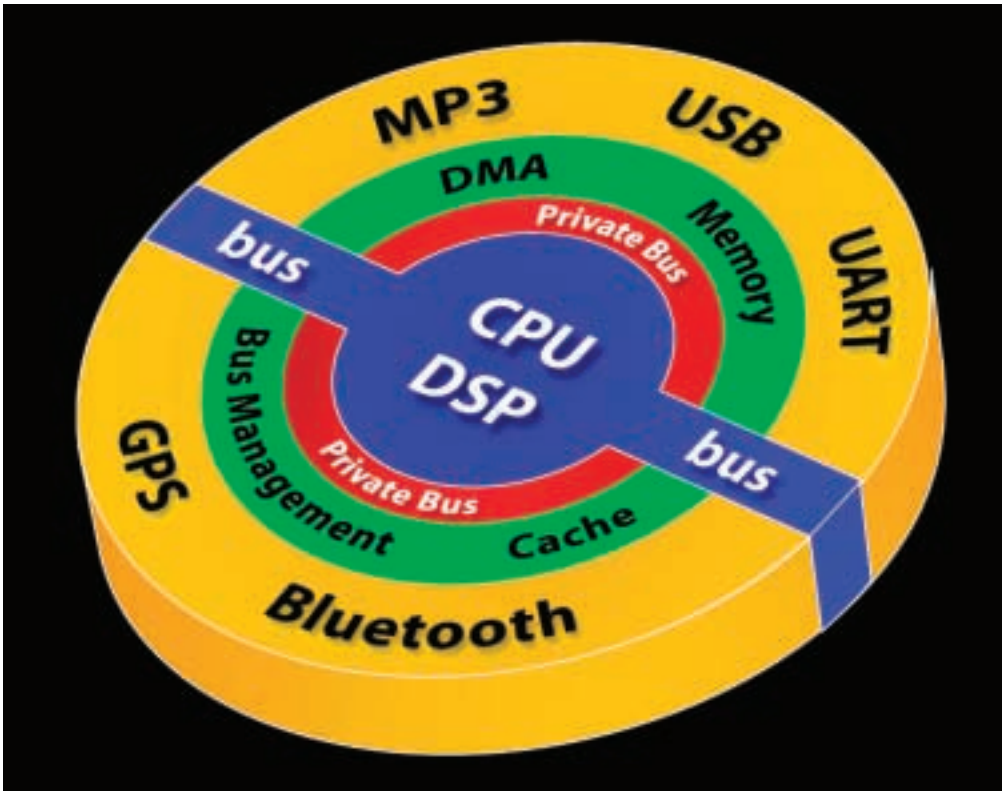

Figure 4 - Platform-based FPGAs integrate IP to user-defined logic

From the FPGA vendor's point of view, they need to provide their customers with a solution that enables high-quality, repeatable results.

Platform-based FPGAs have taken the problems found by previous ASIC SoC designers and minimized them by providing a proven, recommended path for success. In the history of electronic design, innovation and productivity are at their peak when proven methodologies permeate the industry.

IP Immersion

Designing an ASIC SoC requires a highly experienced team of engineers from multiple disciplines. This team must learn how to use a processor, develop software on it, and connect IP into the system. These learning curves are long and riddled with mistakes. And, because all design blocks come from different places, integration rarely works on the first try.

Platform-based FPGAs solve some of the big issues that limit the development of ASIC SoCs by addressing IP integration issues, as shown in Figure 4. The Xilinx IP immersion technology maximizes performance and density by providing a fixed and proven structure to integrate hard and soft IP into the silicon architecture.

Platform-based FPGAs using embedded processors will be the next key technology that will push FPGA design forward. By limiting support to only certain CPU architectures, all integration information is predefined for functions such as control signals, clocks, and data buses. Soft IP is pre-engineered to work with these predefined buses, so the designer just needs to connect these fixed IP objects to the buses. For the user-defined section of the chip, the design has been made easier because its boundary conditions are known.

To take advantage of these fixed-CPU platform-based FPGAs, designers will look for new applications and uses. Since the cost of implementation will be in everyone's reach, we should see a resurgence of the garage-shop mentality and new out-of-the-box thinking.

Enabling Innovation Together

At Mentor Graphics, we enjoy collaborating with a partner that looks at the big picture and asks, "How do I change the future?" The Xilinx vision in programmable logic will change how we do digital design. Mentor has recognized this and is committed to the FPGA market. We continue to develop point tools and solutions to solve the tough design problems. Mentor realizes that dedicated FPGA flows, tested and integrated tightly with the vendor software, will provide the technology that will "Enable Innovation" for the future of the electronic industry.

TODAY, THE XILINX VIRTEX-II TECHNOLOGY IS A LEGITIMATE ASIC REPLACEMENT. WHEN COMBINED WITH LEADING-EDGE EDA SOFTWARE, VIRTEX-II FPGAS PROVIDE THE ELECTRONICS INDUSTRY WITH A NEW, EXCITING PATH FOR GROWTH AND INNOVATION.

— WALDEN C. RHINES, CHAIRMAN AND CEO, MENTOR GRAPHICS

Extinct: Dinosaurs, Slide Rules, 8-Track Tapes, and now... External Termination Resistors

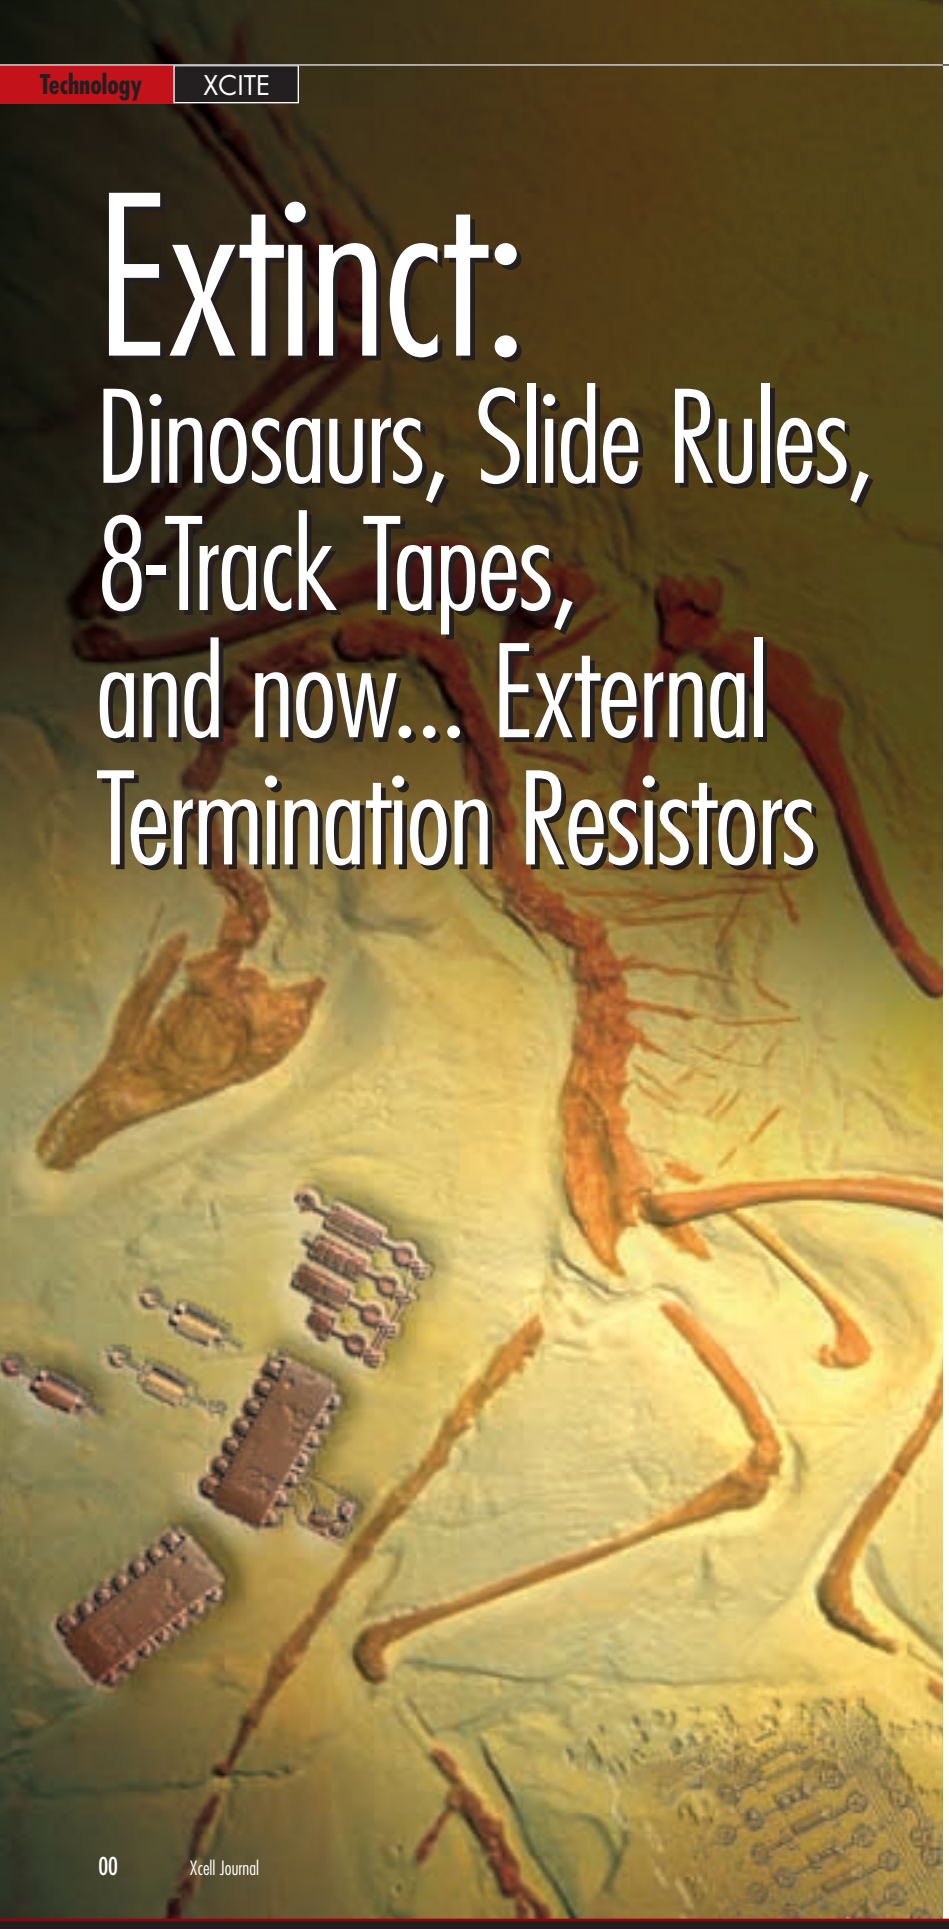

Xilinx Virtex-II Platform FPGAs now feature the world's first on-chip digitally controlled impedance-matching technology.

by Mark Alexander
Product Applications Engineer, Xilinx Inc.
mark.alexander@xilinx.com

Termination resistors as we know them are about to vanish. XCITE (Xilinx Controlled Impedance TEchnology) – the adaptive on-chip termination system now available in the Virtex®-II family of FPGAs – banishes external resistors, making PCBs (Printed Circuit Boards) less expensive, easier to design, and more reliable.

High performance systems require signal frequencies in the hundreds of megahertz, necessitating impedance matching between device I/Os and PCB traces. In the past, as clock speeds increased and I/O standards changed, system engineers and PCB designers learned signal termination techniques, using external resistors to match impedances. Over time, more and more resistors were required to account for the increasing widths of large data busses. With the advent of 1,500-pin packages, discrete resistor placement became a major challenge.

External termination resistors necessitate more board traces and increased PCB part counts, making the layout process more time consuming. These factors lead to higher manufacturing costs and longer times to market. Additionally, sky-high part counts do not help system reliability either – each additional part increases the risk of a system failure.

Here at Xilinx, our Virtex-II team wanted designers to have a worry-free solution to terminating high-speed signals, so we invented XCITE I/O. All Virtex-II I/O pins are equipped with XCITE, making impedance matching a pre-engineered solution.

XCITE can be used to terminate a variety of I/O standards. All variants of HSTL (High-Speed Transistor Logic), SSTL (Stub Series Terminated Logic), GTL (Gunning Transceiver Logic), and LVCMOS (Low Voltage Complementary Metal Oxide Semiconductor) are supported. Designers need only specify their signal I/O standard in the software. In the case of the bidirectional standard HSTL Class II, the external resistors usually required for the source and destination are implemented on-chip (see Figure 1). The PCB designer only has to route a 50Ω trace from the output of one Virtex-II device to the input of another (see Figure 2).

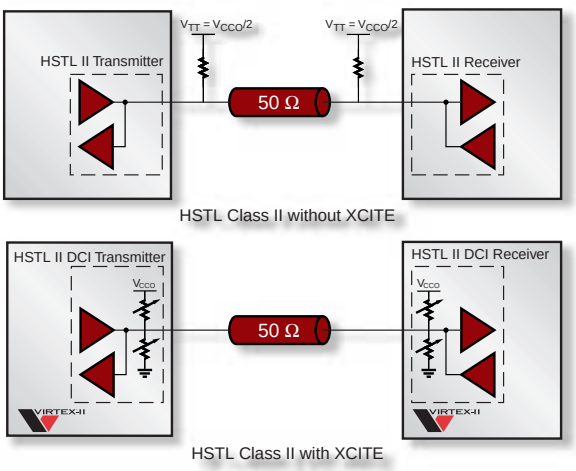

Figure 1 - Termination resistors on-chip

Not only does XCITE make PCBs less complicated, it actually improves signal quality. With termination residing inside the device instead of a few centimeters away, stub reflection is eliminated as a design factor.

XCITE is a better solution than discrete resistors because it continually adjusts the termination impedance to match the PCB trace impedance. In conventional systems, temperature and voltage variations can play havoc with the carefully engineered impedances of a signal path. With XCITE, the termination impedance of the driver or receiver is continually compared against a reference resistor. Over the full range of temperature, voltage, and process variations, XCITE maintains a tight impedance match.

I/O counts will continue to increase, and so will clock speeds. The Virtex-II architecture makes this situation livable by offering the latest high-speed I/O standards without the difficulties of external termination resistors (or punchcards).

For more information on XCITE, consult Chapter 2 of the Virtex-II Platform FPGA Handbook, which can be found online at ***xilinx.com/v2insert/xcite .***

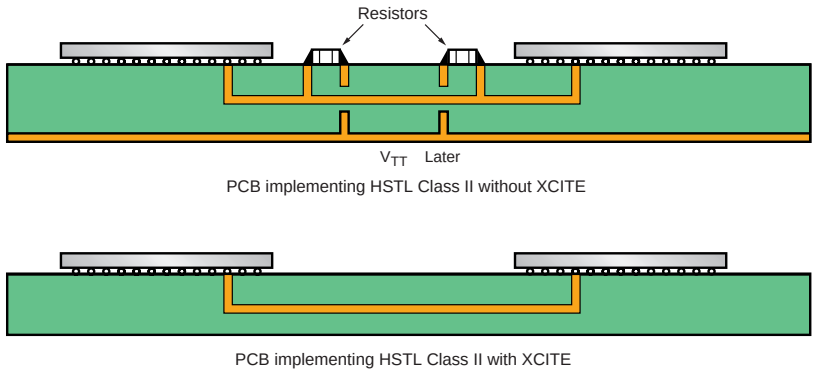

Figure 2

How XCITE Works

Both series and parallel termination schemes can be implemented with XCITE. While parallel termination is realized on-chip with pull-up and pull-down resistors, series termination is realized with a controlled impedance driver. The output impedance of this driver is set to equal the impedance of the PCB trace. XCITE transmitters can operate with parallel termination, a controlled-impedance driver, or a combination of the two; XCITE receivers can operate with parallel termination. Because of the way it is implemented, parallel XCITE termination may terminate to either the full VCCO voltage or to a VCCO/2 (as in the HSTL Class II standard).

XCITE matches its impedance to a pair of external reference resistors. These reference resistors are connected to dual-function pins on the Virtex-II device. Each of the eight I/O banks has two of these dual-function pins, VRN and VRP. If XCITE is not used in a bank, these pins are available for user I/O.

Reference resistors are chosen to have the same impedance value as the PCB trace. XCITE can match any impedance from 25Ω to 100Ω. The XCITE I/O matches the reference value by selectively enabling or disabling parallel transistors. A coarse impedance adjustment is made during the device startup sequence, accounting for process variation. During device operation, fine adjustments are continually made to ensure that I/O impedance stays matched as temperature and voltage drift.

XCITE is implemented in software either by direct HDL instantiation or through the IOSTANDARD attribute in the constraints file. The following are examples of each:

VHDL

```
HSTL_DCI_buffer: OBUF_HSTL_I_DCI
port map (I => data_out, O => data_out_DCI);
```

UCF or NCF

```
NET <net name> IOSTANDARD =
OBUF_LVDCI_25;
```

(Where <net name> is the name of the net between the IPAD and IBUF, or OPAD and OBUF. For HDL designs, this name is the same as the port name.)

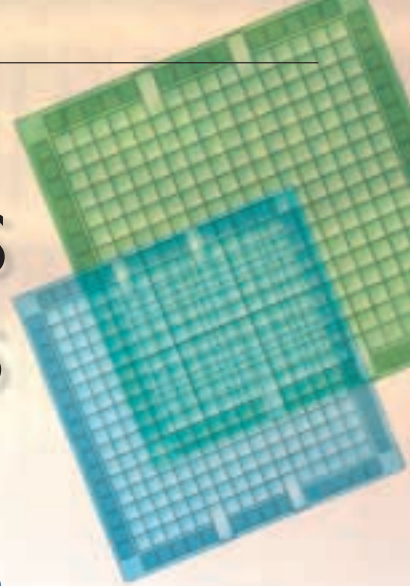

SystemIO Technology Promises High-Speed Connectivity Across Multiple I/O Standards

As implemented via Virtex-II Platform FPGAs, SystemIO technology addresses both the physical interfaces and networking protocols for high bandwidth connectivity and throughput.

by Rina Raman
APG Director of Applications Engineering, Xilinx
rina.raman@xilinx.com

Moore's Law, which predicts computer processor speeds will double approximately every 18 months, has proved to be remarkably accurate for more than 35 years. The exponential increases in processor speeds have enabled ever higher bandwidth networking. The problem now, however, is that I/O busses have only doubled their frequencies every three years. Thus, the level of performance of a system at the board level is now dictated not by the speed of processor but by the limits of the I/O bus. While delivering greater bandwidth, numerous emerging I/O standards – RapidIO™, POS-PHY™ Level 4, 10 Gb/s Ethernet, XAUI, HyperTransport™, Fibre Channel, among others – provide various choices and architectural options. Xilinx Virtex®-II Platform FPGAs address this proliferation of standards with the SystemIO™ solution, which provides system interconnectivity at both the physical I/O interface and the networking protocol levels.

I/O Signaling Standards

Virtex-II devices are designed to support many signaling standards, including essential interfaces to high performance systems. For instance, memory devices operating above 180 MHz can only use SSTL (Stub Series Terminated Logic) or HSTL (High Speed Transceiver Logic) I/O standards.

Many I/O standards – including the traditional switching standards of LVCMOS (Low Voltage CMOS), memory interfaces of HSTL and SSTL, and even LVDS (Low Voltage Differential Signaling) – are not keeping pace with the increased demand for more bandwidth. Designers have tried to overcome these bandwidth limits by using more pins and/or larger bus widths. This has made the traditional I/O standards more "pin-intensive." Significant problems arise as pin counts grow into the hundreds and thousands, creating routing congestion on printed circuit boards. Although Virtex-II Platform FPGAs can easily route inside the device, it is extremely difficult to interface with all the pins on a PCB. More layers of interconnect routing in the PCB cause dramatic increases in overall board costs. The number of pins running these standards also produces problematic electromagnetic interference.

Additionally, the most popular bus architecture is a "shared" medium where multiple entities use the same bus, each waiting in turn for its opportunity to complete its transaction. As the size of audio and video data streams increase, the waiting period gets longer. Ultimately, overall performance decreases along with performance predictability. Therefore, the industry trend is to move from a shared bus to a point-to-point link, typically configured as a switched fabric.

Beyond the Shared PCI Bus

The most popular shared bus is the PCI (Peripheral Component Interconnect) bus. It has become a general-purpose bus for personal computers and embedded systems supporting many different applications. A 33 MHz, 32-bit PCI bus can support one source and five destinations with an overall bandwidth of 1 Gb/s. At 66 MHz and 64 bits, the bandwidth rises to 4 Gb/s, but then the bus only supports one source and two destinations. The PCI shared bus structure has diminishing performance returns, is less predictable than a point-to-point solution, and is limited to internal systems.

Ideally, we would like to establish a link and then "burst" the data over a really wide data bus. This would maximize bus efficiency. Bus efficiency is highest with long bursts.

Decreased bus efficiency has many causes. Although we may have economies of scale for most applications, the shared PCI bus compromises graphics performance. For example, when using a PCI graphics card, the graphics card needs to refresh every few milliseconds. In order to do this, it needs to have immediate and frequent access to the PCI bus. As a result, other PCI cards cannot send huge data bursts. This is because the PCI arbiter is designed to ensure that no one component can send large bursts of data.

In graphics intensive applications, an AGP (Advanced Graphics Port) is a much better choice than a PCI bus. Although AGP is based on PCI technology, it is designed especially for the high throughput requirements of 3D graphics. Rather than using the shared PCI bus for graphics data transmission, the AGP introduces a dedicated point-to-point channel so that the graphics controller can directly access main memory. Virtex-II Platform FPGAs support both PCI and AGP I/O standards, as well as many others.

Another drawback of the PCI bus is that it has no termination – or rather, it is series terminated. This means that it relies on reflective wave switching. Although reflective wave switching is inexpensive and has relatively low power consumption, series termination requires the system to wait for the reflection, so we lose valuable time.

Switched Fabric Advantages

A switched fabric is more scalable for high performance and ultimately, offers a lower cost solution for high bandwidth applications. In the past, data rates weren't high enough to warrant using a switched fabric within a system, because this would require many point-to-point connections and would increase system complexity. With the dramatic increases in performance enabled by the new interface standards, however, switched fabric solutions are becoming cost-effective.

With the huge existing PCI infrastructure, bridges are needed to interconnect the various high-performance I/O standards. Virtex-II Platform FPGAs are key building blocks in the transition from shared buses to switched fabrics. The Virtex-II architecture supports universal switching capabilities to these new standards, making them the logical choice for system designers.

As we mentioned earlier, the key advantage of SystemIO technology is that it provides both physical interfaces as well as various cores that support the network protocols. With the transition from shared busses to switched fabrics comes a variety of different source-synchronous (parallel) protocols. The following are some of the leading source-synchronous standards that are emerging:

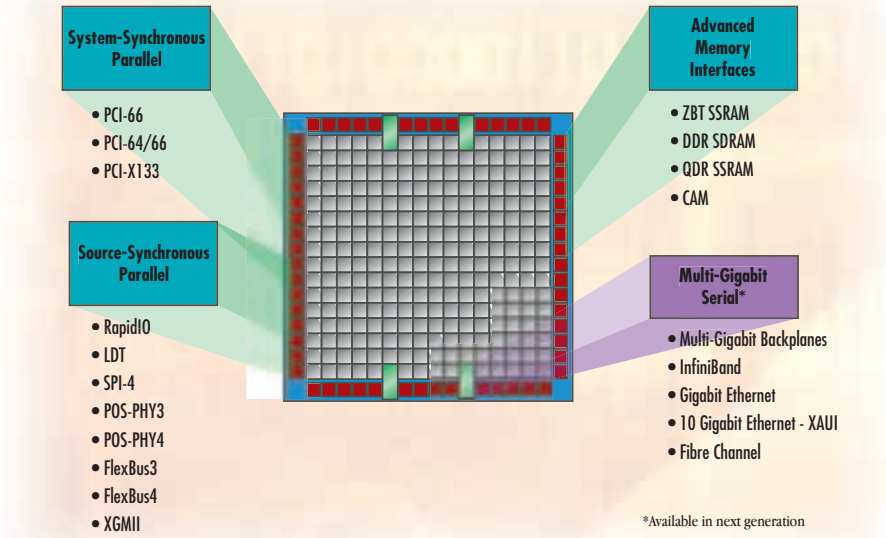

- **RapidIO** – Originally organized to support the processor and local bus markets, the RapidIO interconnect architecture has been embraced by the networking and storage markets.
- **InfiniBand™** – Founded by an industry consortium, InfiniBand targets remote storage, servers, and networking devices.
- **HyperTransport** – Formerly known as LDT (Lightning Data Transport), HyperTransport was jointly developed by AMD and API to replace PCI in high-speed computing applications. It has gained some acceptance in the networking space.
- **OIF - SPI4** (Optical Internetworking Forum - System Packet Interface Level 4) – This standard is defined for 10 Gb/s Ethernet applications and optical networking applications demanding OC-192 performance.
- **10 Gb/s Ethernet** – Using an XGMII interface, this IEEE standard is positioned to drive the convergence of LAN and WAN technologies.

In addition, there are emerging serial channel standards as well. Two of the most popular standards in this arena are:

- **Fibre Channel** – A favorite for the storage area networking (SAN) market, Fibre Channel uses optical fiber, coaxial cable, and/or twisted-pair telephone wire.
- **XAUI** (pronounced ZOW-ee) – This new standard targets 10 Gb/s serial channels by bonding four 3.125 Gb/s transceivers. XAUI targets the OC-192 and 10 Gb/s Ethernet markets for WAN and LAN routers.

Conclusion

Having a solution that supports all these various system interfaces is crucial to success in the marketplace. The Virtex-II Platform FPGAs' SystemIO solution offers exactly this – support for physical interfaces as well as cores that support the network protocols for all the common and emerging system I/O interfaces. Virtex-II FPGAs enable high performance interfaces to memories from Cypress, IDT, Micron, SiberCore, GSI Technology and others, as well as interfaces to networking ASSPs (Application Specific Standard Parts) from vendors such as AMCC, PMC Sierra, and Vitesse. Now, with Virtex-II SystemIO solution, you can pick any standard, and any vendor offering that standard in their ASSPs, and rest assured that Xilinx Virtex-II Platform FPGAs will support that standard.

Virtex-II Platform FPGAs Support System Packet Interface Standards for Optical Networks

The production release of SPI-4 Phase 2 cores to Xilinx communication customers worldwide, is a critical technology boost for multi-service, packet, and cell-based networking equipment.

by Ron DiGiuseppe
System Interfaces Product Marketing Manager, Xilinx
ron.digiuseppe@xilinx.com

Chris Ebeling
System Logic and Networking Staff Enginee, Xilinxr
chris.ebeling@Xilinx.com

Xilinx has developed a suite of LogiCORE™ intellectual property blocks to perform the System Packet Interface (SPI) function between the Physical (PHY) and Data Link layer devices for POS/SDH (Packet Over SONET/Synchronous Digital Hierarchy) fiber optic applications. The cores address the exploding demands of network IP (Internet Protocol) traffic by ensuring Xilinx devices are compatible with the Optical Internetworking Forum's (OIF) SPI-4 Phase 2 standard as well as the SATURN® Development Group POS-PHY™ Level 4 (PL4) interface. The cores assure compliance with the OC-192 data transfer standard by moving IP packets at a data rate in excess of 10Gb/s.

Just in Time to Market

The interface cores, referred to as PL4 cores, make use of unique features available only in the Xilinx Virtex™-II Platform FPGA architecture – including DCM (Digital Clock Manager), enhanced Block RAM, and high-speed LVDS I/O buffers. Combined with Platform FPGA DDR (Double Data Rate) registers, the PL4 cores can support data rates up to 832Mb/s per pin pair.

Xilinx is working with the OIF and the ATM (Asynchronous Transfer Mode) Forum to promote the SPI-4 Phase 2 standard. Along with other industry-leading networking developers Xilinx is facilitating the design and deployment of data switching and routing products using interoperable optical networking technologies. The improved efficiencies and lower cost per Mbit of POS/SDH packet transfer makes it an enabling technology for gigabit routers, terabit and optical cross-connect switches, and a wide range of multi-service DWDM (Dense Wave Division Multiplexing) and SONET/SDH-based transmission systems. The PL4 cores implemented in Virtex-II FPGAs allow next-generation network developers to reduce their system time to market.

In addition to providing fully standard-compliant cores, Xilinx is collaborating with leading network device developers – including PMC-Sierra, AMCC, and Conexant – to ensure interoperability between the Xilinx networking cores and the latest industry products. By combining the leading-edge performance of Virtex-II devices, Xilinx PL4 cores, and PMC-Sierra's or Conexant's OC-192 PHY devices, a complementary solution is available to our mutual customers. With IP traffic on network backbones doubling every six to nine months, it is critical to provide a high performance, scalable, system solution.

Interfacing the PHY and Data Link Layers

The POS/SDH Physical Layer Level 4 (POS-PHY L4) interface allows the interconnection of Physical Layer devices to Data Link Layer devices in 10Gb/s POS, ATM, and Ethernet applications. While the Xilinx PL4 core can perform the interface functions on both sides of the PL4 bus as shown in Figure 1, the FPGA implementation is generally intended to operate on the Data Link Layer side.

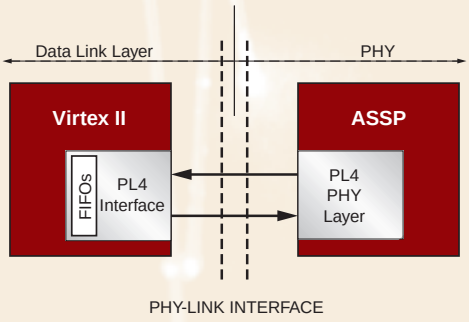

Figure 1 - OIF SPI-4 Phase 2 System Reference Model

The SPI-4 (PL4) interface has the following general characteristics:

- Point-to-point connection (such as between a single PHY Layer and a single Data Link Layer device)

- Support for 256 ports, suitable for STS-1 granularity in SONET/SDH applications (168 ports) and Fast Ethernet granularity in Ethernet applications (100 ports)
- Transmit/receive data path: 16-bits wide
- Source synchronous clocking where the source of the data provides a data clock
- In-band port address, start/end-of-packet indication, error-control code
- LVDS I/O (IEEE 1596.3 – 1996 [1], ANSI/TIA/EIA-644-1995 [2])
- 622 Mb/s minimum data rate per line using double data rate I/O
- Packet address, delineation information, and error-control coding sent in-band with data in both transmit and receive modes.

In addition to supporting the listed PL4 interface features, the Xilinx PL4 cores were developed with configurable FIFO buffers using Virtex-II Block RAM. The Virtex-II Block RAM provides high-performance data read/write access times and a four times increase in density over previous architectures. The Block RAM allows appropriate buffering to match the required channel support for the external PHY ASSP.

The cores use LVDS I/O buffers paired with dedicated DDR registers in the data path and LVTTL I/O buffers in the FIFO status path. The internal data rate is reduced by expanding the 16-bit words in DDR format on the PL4 interface to a 64-bit (four-word) single-edge clocked format running at half the PL4 clock rate. The core utilizes the Virtex-II DCM as shown in Figure 2. The DCM generates internal and external clocks to meet the aggressive system jitter requirements.

The cores implement "static alignment" of the received data to the clock by using the DPS (Digital Phase Shift) function of the DCM. The DPS module permits very fine-grained adjustments (under 50 ps) of the RDCLK (Received Clock) relative to the RDAT (Received Data). The fine resolution allows the RDCLK to be adjusted to the optimal sampling point relative to the RDATs eye pattern.

Not shown in Figure 2 are the PL4 FIFO interface blocks for implementing the single or multiple links or ports. The PL4 sink block stores data received for a particular link in a single FIFO buffer along with the link address information decoded from the control word preceding the data burst. When data is received, the address information is extracted from the received control word. The address and data are written into the sink FIFO. Two-bit dual-port Block RAMs are used to pass the per-channel FIFO status between the PL4 interface and the user's application. The PL4 sink block transmits the FIFO status information according to the contents of the FIFO status memories.

The PL4 source section decodes the FIFO status channel (flow control) information and writes it to the dual-port block RAM for each channel. The flow control information is used by the PL4 core to determine the channel status of the bus. Once the core receives the flow control information, it determines the link (address) and amount of data to send. The core monitors the fill level of the source FIFO to determine whether to send data or idle control words on the PL4 interface.

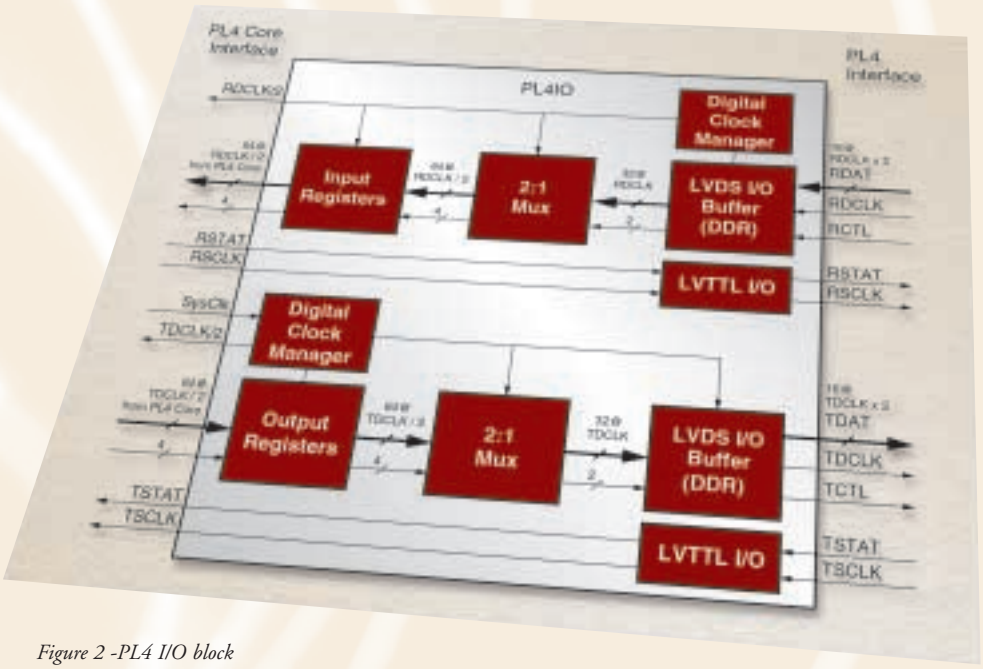

Figure 2 -PL4 I/O block

Conclusion

The Xilinx POS-PHY Level 4 cores are available as fixed netlists designed to interoperate with industry leading POS/ATM framers and mappers to achieve carrier-class performance. The cores have been configured to interface to a single-channel OC-192 device, a 10-channel by 1Gb/s device, and a 4-channel by 2.5Gb/s device.

Working in collaboration with engineering teams from PMC-Sierra and others, Xilinx is verifying the cores by using reference designs provided by each standard product developer. The Xilinx team solves system developers' 10 Gbps performance requirements by offering interoperable, standards-compliant, Packet-Over-SONET cores.

Video Demonstration Board

A glimpse at broadcast video router/mixer functions inside a Virtex-II Platform FPGA

by Gregg C. Hawkes
Senior Staff Applications Engineer, Xilinx, Inc.
gregg.hawkes@xilinx.com

Virtex-II FPGAs are the ideal platform for developing video applications. No other FPGA can provide the combination of 18x18 2's complement signed block multipliers, Digital Clock Managers (DCM), glitch-free global clock multiplexers, Bus LVDS I/O, and DDR I/O, which are all essential for the pixel-rate math and the high bandwidth needed for managing and manipulating video data streams.

To demonstrate the many video-friendly features and the video IP (application module software) of the Virtex-II FPGA family, we've developed a demonstration board that provides a Virtex-II FPGA interfaced to the essential video support functions such as:

- Large, fast, frame buffer memories
- Video inputs (4 NTSC or PAL, CCIR 601/656 4:2:2 format)
- Video outputs (1 NTSC or PAL, CCIR 601/656 4:2:2 format)
- Video outputs (1 RGB, 24 bit format)
- Network connection
- System configuration devices.

There are many routing/mixing functions currently available as intellectual property for the demonstration board, with room for future exploration of more advanced video algorithms. This board can also be used as a convenient video IP development and experimentation system.

Brief Description

The Virtex-II video demonstration board is a simple version of a video router/mixer. The board allows conversion of several, high bandwidth video streams from various video sources, into a common format and color space. Then, using high-speed, pixel-rate, pipelined math, you can manipulate and merge the video streams. The Virtex-II FPGA provides high bandwidth access to devices and large memories, high data rate arithmetic, and the necessary control logic.

The Virtex-II architectural features highlighted by the demonstration board include:

- 18x18 2's complement, signed block multipliers which provide the high speed math capability
- Digital Clock Managers (DCM) which provide clock de-skew, frequency synthesis, clock phase shifting, and EMI reduction logic
- Global clock multiplexer buffers which provide clock multiplexing, clock buffering, and distribution
- A Network connection to a Local Area Network (LAN) and the Internet so you can remotely update the Virtex-II internal algorithms
- Configuration from compact flash memory providing a way to update and store future changes to the programming bit stream of the Virtex-II device.

Supported Effects

Using the Video Demonstration board, you can mix video streams, from many different sources, in interesting ways. For example, you can easily perform the all-too-familiar video fade or alpha blend from one scene to another, where the current video stream (such as a basketball game live feed from a satellite), slowly disappears and a new scene appears (such as a commercial).To accomplish this, the pixels in one scene are multiplied by a fraction (alpha) while the pixels in the other scenes are multiplied by "one minus the fraction" (1-alpha). Varying the fraction from zero to one produces the blending effect.

The master controller or technician viewing the different video input streams and the resulting video output executes video commands to manipulate the input streams. Thus, just as in a typical video production, the master controller queues up an effect, such as "going to a commercial" and the FPGA executes the mathematics behind the command.

The types of broadcast video effects currently supported are:

- Fade to/from black
- Fade through black
- Dissolve
- Horizontal wipe
- Vertical wipe.

Over time, more effects will be available. These expanded effects will appear on the Xilinx website as video application notes.

Board-Level Block Diagrams

The video demonstration board has a number of input sources of live video and a number of separate frame buffers to support the increased amount of storage and bandwidth needed by the extra live streams. An audio codec is supported for embedded audio as well as supporting potential audio algorithm updates from a network connection. The board can also drive a TV monitor. A block diagram is shown in Figure 1.

Board Features

- Four sources of live video input (either NTSC or PAL)
- Composite and S-Video inputs (from a camcorder)
- Separate fixed graphic image loaded from Compact Flash memory
- Real-time video output (XVGA touch-screen and/or NTSC/PAL output)
- Composite and S-Video NTSC/PAL video and RGB output
- Video effects (fades, dissolves, wipes, and so on)
- Compact Flash FPGA configuration
- Touch-screen-selected video source and effects. The touch screen is enabled by RS232 port; video source selection and effects are also pushbutton enabled
- Audio switching
- 10 Base-T and 100 Base-TX Ethernet support
- XC2V6000FF1517 Platform FPGA support
- Universal power supply module.

Verilog Modules

The following list of functions, written initially in Verilog, are available for use with the Virtex-II Video Demonstration Board:

- User Interface – Push button scan affects what is seen on the output screen
- ZBT memory interface controllers – Drives data to and from the FPGA and ZBT RAM
- XVGA controller – Outputs data to a regular computer monitor (self adapts for NTSC or PAL). The module will work with any resolution given the right amount of memory; the 4-channel version will support 1024 x 768, the 1-channel version runs at 800 x600
- Clock generation – Generates four different clock rates, supporting various video functions, the audio codec at 25.576MHz (if the DCM will work with the required ratios), and 25MHz for Ethernet
- Line-field decoder – Assists in identifying frame and field parameters
- I^2C serial interface standard – Loads initialization parameters from the FPGA to the video peripheral chips
- On chip line buffers – Allowing algorithm pixels to be processed vertically
- Interlace fields to non-interlace frames conversion
- Color space conversion
- 4:2:2 to 4:4:4 format conversion
- Up/down scaling of thumbnails
- Blend and fade between frames (video processing).

Check the Video Applications website for the latest new functions. Future algorithms may include compression of video data, detection and enhancement of video imagery via DSP functions, and additional flexibility to support the constantly emerging video standards.

Conclusion

The advanced system-level features and the growing list of video-related Intellectual Property make the Virtex-II FPGA family an ideal choice for video applications. And now, the Virtex-II Video Demonstration Board gives you everything you need to quickly and easily explore video applications. For more information on the demonstration board, visit the Xilinx website at: www.xilinx.com/v2insert/boards.

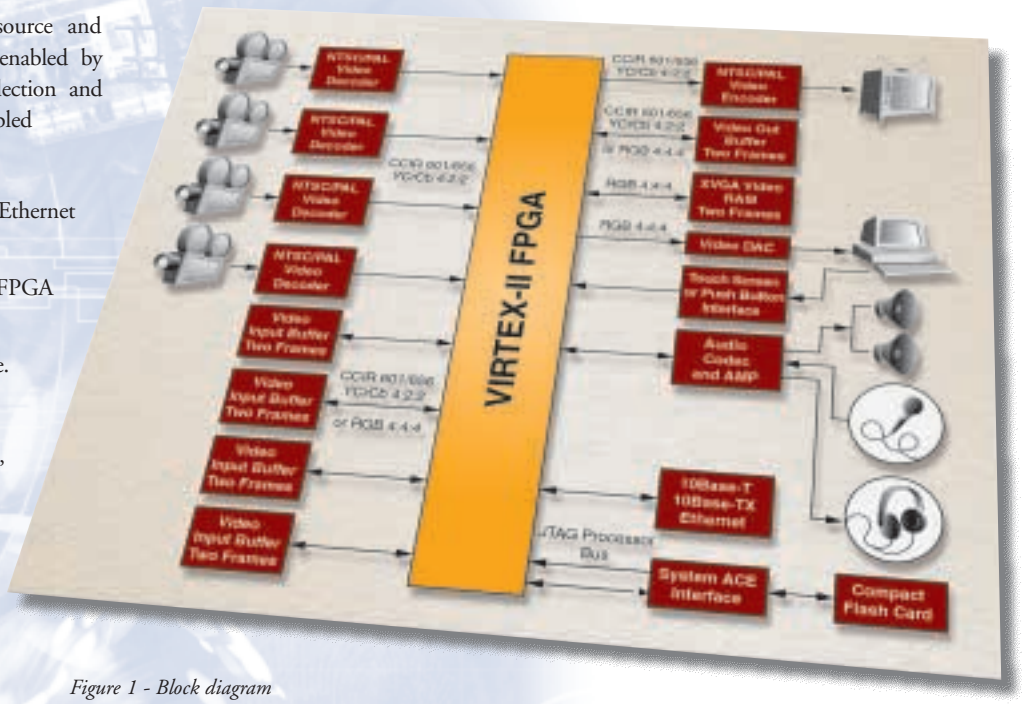

Figure 1 - Block diagram

4G Wireless Systems in Virtex-II

Designers of the 4G wireless systems infrastructure are confronted with challenging product development issues, including the uncertainty about fundamental system architectural standards such as the air interface, encryption protocols, planetary interoperability, and so on. Because there are unresolved uncertainties, you must pay close attention to risk management – making sure your designs can evolve with the changing standards.

by James A. Watson
Manager, Applications Engineering, Xilinx, Inc.
jim.watson@xilinx.com

Virtex®-II FPGAs are an ideal platform for designing with ambiguous or evolving standards. Because of the inherent flexibility, reprogrammability, and extremely high performance (approximately 0.5 TeraMACs) of Virtex-II devices, you can easily test different air-interface schemes and variants in-system, and you can quickly assess the system performance. In particular, Virtex-II FPGAs make it easy to develop hybrid systems such as multi-carrier CDMA or QAM-modulated OFDM.

Orthogonal Frequency Division Multiplexing (OFDM)

Currently, there are two principal 4G development technologies contending for attention: CDMA and OFDM. Code Division Multiple Access is a well-known standard and has been used for several years. However, OFDM is relatively new. OFDM, with many technical variants, is endorsed by Nokia, Cisco, Lucent, and Philips Semiconductor, and is represented as the successor to frequency hopping and direct sequence CDMA. It is also positioned as the technique of choice for next generation wireless LANs and metropolitan networks. The capability of OFDM to cancel multipath distortion in a spectrally efficient manner without requiring multiple local oscillators has won adherents in the IEEE 802.11a and 802.16 working groups. However, despite the support of many key industry players, OFDM is not actually deployed in mainstream wireless systems.

Todd Carothers, vice president of marketing for Adaptive Broadband, recently stated "We've developed a commercial OFDM system for one application, and we think OFDM has real advantages in the mobile arena, but we don't see it for fixed point. We think that adaptive time division multiple access is still the best solution for fixed point-to-multipoint, and I'll state that we still have the fastest system out there and the most extensively deployed."

Philip Gee of WiLAN said recently, "There is no question that OFDM and CDMA are in contention for some of the same wireless markets. We believe that OFDM enjoys a number of significant advantages, however."

How it Works

OFDM is fundamentally different from other modulation schemes. In fact, it should probably not be considered a modulation scheme at all, because it may be transmitted via AM, FM, QAM (Quadrature Amplitude Modulation), and so on. OFDM is properly defined as a mathematically elegant technique for the generation and demodulation of radio waves. Although its origins date back to the second World War, its application to wireless communications is new.

In OFDM the subcarrier pulse shape is a square wave. The task of pulse forming and modulation can be performed by a simple Inverse Discrete Fourier Transform (IDFT) which can be implemented very efficiently in Virtex-II FPGAs as an Inverse Fast Fourier Transform (IFFT). To decode the transmission, a receiver need only implement an FFT.

As you can see in Figure 1, the spectra of the subcarriers overlap. By using an IFFT, the spacing of the subcarriers is varied in such a way that, at the target frequency of the received signal (indicated as arrows), all other signals are zero. This is known as "frequency orthogonality." This contrasts with Direct Sequence CDMA, which uses a Walsh code to achieve code orthogonality.

OFDM and the Virtex-II Architecture

Virtex-II FPGAs offer several architectural advances that allow you to create extremely efficient implementations of OFDM systems.

Multipliers

Virtex-II FPGAs contain a number of 18x18 2's complement signed multipliers associated with the block SelectRAM™ memory.

This association allows high-speed access to complex multiplicand coefficients, thus supporting extremely high-performance arithmetic. To see why the multipliers are so valuable, consider the nature of the FFT algorithm itself. It essentially decomposes into a series of multiply-accumulate functions.

Digital Clock Management

To successfully implement OFDM, the receiver and the transmitter must be in perfect synchronization. Synchronizing to the transmitter's data clock is always necessary, whereas, carrier recovery is only necessary in coherent detection receivers. The data clock must be recovered so that the receiver will sample the transmitted data symbols at the appropriate time.

An algorithmic approach such as times-two, early-late, or zero-crossing clock recovery can be implemented in a Virtex-II device; all of these functions are performed in the digital domain. These algorithmic approaches are perfect applications for the Virtex-II Digital Clock Manager (DCM). For example, the DCMs in the Virtex-II devices, along with a DDS (Direct Digital Synthesis) core, can provide the complex sinusoids necessary for demodulating the incoming data. The timing/phase of these complex sinusoids is directed by the data recovery clock and easily adjusted by the DCM's timing controls. The DCM can also perform other functions vital to synchrony of the transmitter and receiver including clock deskew and frequency synthesis.

The DCM can also de-skew the received signal relative to the local receiver frequency by adding digital delay. This results in a signal that is delayed but has perfect phase alignment to the local receiver frequency.

Virtex-II DCMs can drive global clock resources, general logic interconnect, and I/O pads simultaneously. This provides maximal flexibility when placing logic.

High Performance

The most valuable feature of the Virtex-II family, for implementation of advanced wireless systems, is the extremely high-performance. This gives you a great degree of freedom that is not available with alternative implementations, such as ASICs. To understand the value of this advantage, consider the following scenario.

OFDM Field Deployment Example

In this example, an OFDM system is deployed on an experimental basis by a wireless service provider. It is located in an urban market, but there are many business factors that must be addressed if the venture is to succeed, including:

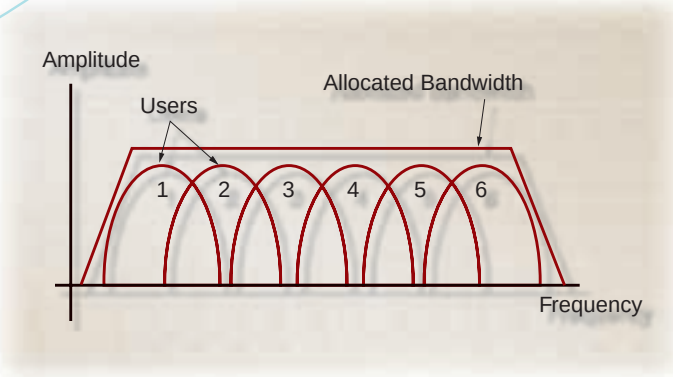

Figure 1 - OFDM allows for greater spectral efficiency

- Which of the emerging broadband wireless services will generate high demand?
- What is the peak bandwidth per subscriber?
- What is the average bandwidth per subscriber per service?
- What Quality of Service factors can be used to differentiate the new service?

There is no substitute for field trials to answer these questions, and there is no better platform for field trials than Virtex-II FPGAs.

Consider a case where a standards body adopts a new variant of OFDM, a very likely scenario. With conventional ASICs, the upgrade path is painful. Utility workers must climb telephone poles, climb to the tops of buildings, and so on, and manually upgrade circuit boards in the base stations. Keep in mind that many of these base stations are deployed in regions of the world that suffer climactic extremes. This expensive and potentially dangerous upgrade path is part of the cost of ownership of a base station constructed around conventional ASICs.

Now consider a base station constructed with a Virtex-II Platform FPGA; the configuration of the cellular base station may be completely altered simply by transmitting a new Virtex-II configuration file from the comfort of a central office. This technique is extensible, so that an entire network can be reconfigured, automatically, without replacing any hardware. This capability allows you to more rapidly introduce the product to the market and helps to protect the base station architecture against obsolescence.

There is another degree of flexibility which the Virtex-II-based OFDM base station provides, the ability to trade off silicon area for performance. Consider the market success factors described above. If it became evident through field trials and customer testing that the average customer was not a heavy consumer of bandwidth, the OFDM algorithm could be re-targeted to use more general purpose logic. Properly done, this would result in the ability to support more channels in the same device. Essentially, the Virtex-II Platform FPGA allows you to dynamically trade off silicon area for performance.

Conclusion

The Virtex-II product family is uniquely suited to the demanding digital signal processing that will be required to roll out next-generation broadband wireless services. Its powerful suite of dedicated high-performance logic functions such as the high-speed multipliers and DCM, along with extremely versatile high-performance general logic, define an optimal solution for wireless designs.

XILINX®
The Programmable Logic Company℠

VIRTEX-II

PRODUCT SELECTION MATRIX

	CLB Resources						BLK RAM		CLK Resources					I/O Features				Speed			
	System Gates	CLB Array (Row X Col)	Number of Slices	Logic Cells	CLB Flip-Flops	Max. Distributed RAM Bits	# Block RAM Blocks	Block RAM Bits	# Dedicated Multipliers	Digital Clock Manager (Frequency min/max)	# DCM's	Frequency Synthesis	Phase Shift	Digitally Controlled Impedance	Number of Differential I/O Pairs	Max. I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Serial PROM Family	Config. Memory (Bits)
Virtex-II Family — 1.5 Volt																					**.15um Eight Layer Metal Process**
XC2V40	40K	8 x 8	256	576	512	8K	4	72K	4	1/420	4	YES	YES	YES	44	88	LDT-25, LVPECL-33,	-4 -5 -6	-4 -5	ISP	0.4M
XC2V80	80K	16 x8	512	1,152	1,024	16K	8	144K	8	1/420	4	YES	YES	YES	60	120	LVDS-33, LVDS-25,	-4 -5 -6	-4 -5		0.6M
XC2V250	250K	24 x16	1,536	3,456	3,072	48K	24	432K	24	1/420	8	YES	YES	YES	100	200	LVDSEXT-33, LVDSEXT-25,	-4 -5 -6	-4 -5		1.7M
XC2V500	500K	32 x 24	3,072	6,912	6,144	96K	32	576K	32	1/420	8	YES	YES	YES	132	264	BLVDS-25, ULVDS-25,	-4 -5 -6	-4 -5		2.8M
XC2V1000	1M	40 x 32	5,120	11,520	10,240	160K	40	720K	40	1/420	8	YES	YES	YES	216	432	LVTTL, LVCMOS33,	-4 -5 -6	-4 -5		4.1M
XC2V1500	1.5M	48 x 40	7,680	17,280	15,360	240K	48	864K	48	1/420	8	YES	YES	YES	264	528	LVCMOS25, LVCMOS18,	-4 -5 -6	-4 -5		5.7M
XC2V2000	2M	56 x 48	10,752	24,192	21,504	336K	56	1008K	56	1/420	8	YES	YES	YES	312	624	LVCMOS15, PCI33, PCI66,	-4 -5 -6	-4 -5		7.5M
XC2V3000	3M	64 x 56	14,336	32,256	28,672	448K	96	1728K	96	1/420	12	YES	YES	YES	360	720	PCI-X, GTL, GTL+, HSTL I,	-4 -5 -6	-4 -5		10.5M
XC2V4000	4M	80 x 72	23,040	51,840	46,080	720K	120	2160K	120	1/420	12	YES	YES	YES	456	912	HSTL II, HSTL III, HSTL IV,	-4 -5 -6	-4 -5		15.7M
XC2V6000	6M	96 x 88	33,792	76,032	67,584	1056K	144	2592K	144	1/420	12	YES	YES	YES	552	1104	SSTL21, SSTL211,	-4 -5 -6	-4 -5		21.9M
XC2V8000	8M	112 x 104	46,592	104,832	93,184	1456K	168	3024K	168	1/420	12	YES	YES	YES	554	1108	SSTL3 I, SSTL3 II	-4 -5	-4		29.1M

PACKAGE OPTIONS AND USER I/O

			Virtex-II (1.5V)										
			XC2V40	XC2V80	XC2V250	XC2V500	XC2V1000	XC2V1500	XC2V2000	XC2V3000	XC2V4000	XC2V6000	XC2V8000
Pins	**Body Size**	**I/O's**	88	120	200	264	432	528	624	720	912	1104	1296
Chip Scale Packages — wire-bond chip-scale BGA (0.8 mm ball spacing)													
144	12 x 12 mm		88	92	92								
BGA Packages (BG) — wire-bond standard BGA (1.27 mm ball spacing)													
575	31 x 31 mm						328	392	408				
728	35 x 35 mm								456	516			
FGA Packages (FG) — wire-bond fine-pitch BGA (1.0 mm ball spacing)													
256	17 x 17 mm		88	120	172	172	172						
456	23 x 23 mm				200	264	324						
676	27 x 27 mm							392	456	484			
FFA Packages (FF) — flip-chip fine-pitch BGA (1.0 mm ball spacing)													
896	31 x 31 mm						432	528	624				
1152	35 x 35 mm									720	824	824	824
1517	40 x 40 mm										912	1104	1108
BFA Packages (BF) — flip-chip fine-pitch BGA (1.27 mm ball spacing)													
957	40 x 40 mm								624	684	684	684	684

Virtex-II packages FG456 and FG676 are footprint compatible.
Virtex-II packages FF896 and FF1152 are footprint compatible.
Important: Verify all Data with Device Data Sheet

Numbers indicated in the matrix are the maximum number of user I/O's for that package and device combination.

VIRTEX™-II

THE PLATFORM FOR POWERFUL SOLUTIONS.

THE FIRST PLATFORM FPGA

It's here. The next generation of the most remarkable programmable logic device ever introduced: the Virtex-II FPGA.

The first embodiment of the Platform FPGA, the Virtex-II solution is the ultimate system design platform, delivering SystemIO™ interfaces to bridge emerging standards and address all aspects of system connectivity. For the first time in the industry, designers facing the challenges of signal integrity, system timing, EMI issues and design security have a programmable platform that heralds a new era in high-performance designs.

The world's first Digitally Controlled Impedance Technology (XCITE™) dramatically simplifies board design and maximizes system performance. IP Immersion™ architecture enables easy integration of hard and soft IP. XtremeDSP delivers over 600 billion MACs/s of processing power. Unique Digital Clock Managers allow unparalleled levels of flexibility in high-speed customized clock design.

Watch for more on the Platform FPGA . . . it's where the system design world is going.

XILINX®

The Programmable Logic Company℠

www.xilinx.com

FORTUNE® 2001
100 BEST COMPANIES TO WORK FOR

PN 0010617

For more information visit xilinx.com/v2insert/solutionsad