ISSUE 43, SUMMER 2002 XCELL JOURNAL XILINX, INC.

Xcell journal

Issue 43
Summer 2002

THE AUTHORITATIVE JOURNAL FOR PROGRAMMABLE LOGIC USERS

Cover Story

Xilinx CEO Wim Roelandts Reveals How to Survive a Recession and Gain Market Leadership

ANALYSIS

Total Cost Management

TECHNOLOGY

How to Build Efficient FIR Filters

Xyron Inc. Makes RTOS Breakthrough

PERSPECTIVE

Using DSP in Real-Time Video Processing

NEWS

IBM ASICs Incorporate Xilinx Technology

Virtex-II Pro – New Products, Reduced Pricing

XILINX®

We Did It Again – Only Better This Time

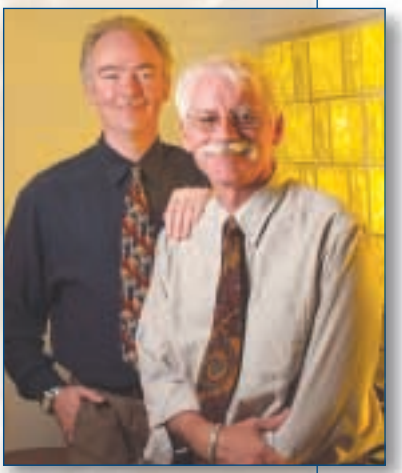

Last year, Xilinx debuted at #14 on the *FORTUNE* 100 Best Companies to Work For™ annual list. This year, we broke into the Top 10 with a ranking of #6 – the only Silicon Valley/San Francisco Bay area company to make the Top 10.

To help determine the new rankings, *FORTUNE* weighed how companies reacted to the economic downturn of 2001. Although revenue at Xilinx declined by 50 percent over the course of 2001, the company decided employee layoffs would come only as a last resort. As a result of significant cost-cutting measures – including tiered salary cuts for everyone except the lowest paid employees – Xilinx managed to avoid layoffs, continued to bring remarkable new products to market, and became larger than all other public programmable logic companies combined.

We believe the *FORTUNE* magazine ranking reflects the strong values we hold at Xilinx:

- Customer focus
- Respect
- Excellence
- Accountability
- Teamwork
- Integrity
- Very open communication
- Enjoying our work.

"Many Xilinx employees have told me how proud they are to receive this recognition. It speaks as much about the great culture we have created at Xilinx as is does about the quality of people who work here."

– Wim Roelandts, CEO, Xilinx Inc.

These CREATIVE values are a vital part of our culture, and they make a positive difference for our employees and our customers. Our values make Xilinx a great company to work for – and a great company to work with. Our technology, our service, and our customer satisfaction continue to lead the industry. We are always seeking new ways to make a real difference in our world. And, though technology drives our business, we've built our success on the strong partnerships we've created with our customers, our suppliers, and other technology leaders in our industry.

To learn more about the *FORTUNE* magazine's top 100 companies to work for, go to: *www.fortune.com/lists/bestcompanies/*.

Tom Durkin

Tom Durkin
Managing Editor

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

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

CONTRIBUTING EDITORS — Jane Bratun, Brendan Bridgford, LeRoy Miller, Ryan Wilson

DESIGN & ILLUSTRATION — Scott Blair, Dan Teie

Xcell journal

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

Cover Story Page 4

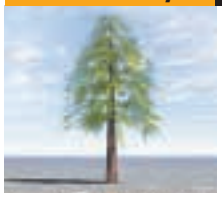

How to Survive a Recession and Gain Market Leadership

"I said from the beginning that we wanted to emerge from the recession a greater company – and I think we have achieved that."

Perspective Page 8

When Total Cost Management Counts, Xilinx PLDs Pay Off

CoolRunner-II CPLDs, Spartan-IIE FPGAs, and Virtex-II EasyPath Platform FPGAs can give you total system cost solutions superior to those of traditional ASICs.

Software Page 24

Exploring Hardware/Software Co-Design with Virtex-II Pro FPGAs

To explore the possibilities of hardware/software co-design, the Xilinx Design Services team created a reference design for a real-time operating system for telecommunications.

Technology Focus Page 42

Customize Virtual Private Networks with Spartan-IIE FPGA Solutions

Spartan-IIE FPGA solutions can help you to customize your virtual private network for highest performance, greatest security, and lowest cost.

Applications Page 55

How to Build an Efficient FIR Filter Using System Generator

Xilinx System Generator 2.2 software enables high-level modeling and implementation of DSP systems in Xilinx Platform FPGAs to outperform traditional DSP processors.

New Products Page 62

Develop MicroBlaze Applications with Three Flexible Hardware Evaluation Platforms

MicroBlaze Development Kits from Memec Design demonstrate the versatility of the MicroBlaze soft processor core in a variety of Xilinx FPGAs.

Reference Page 74

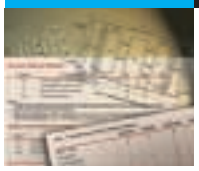

CoolRunner-II Data Sheet

RealDigital CPLDs consume only microwatts of power, yet provide very high performance – take a look at what's "under the hood."

Contents Summer 2002

Xcell journal

Subscribe to the Xilinx Xcell Journal at: *www.xilinx.com/forms/literature.htm.*

Visit Xcell Online for the latest news and information, as soon as it arrives at: *www.xilinx.com/publications/xcellonline/.*

How to Survive a Recession and Gain Market Leadership

"I said from the beginning that we wanted to emerge from the recession a greater company – and I think we have achieved that."

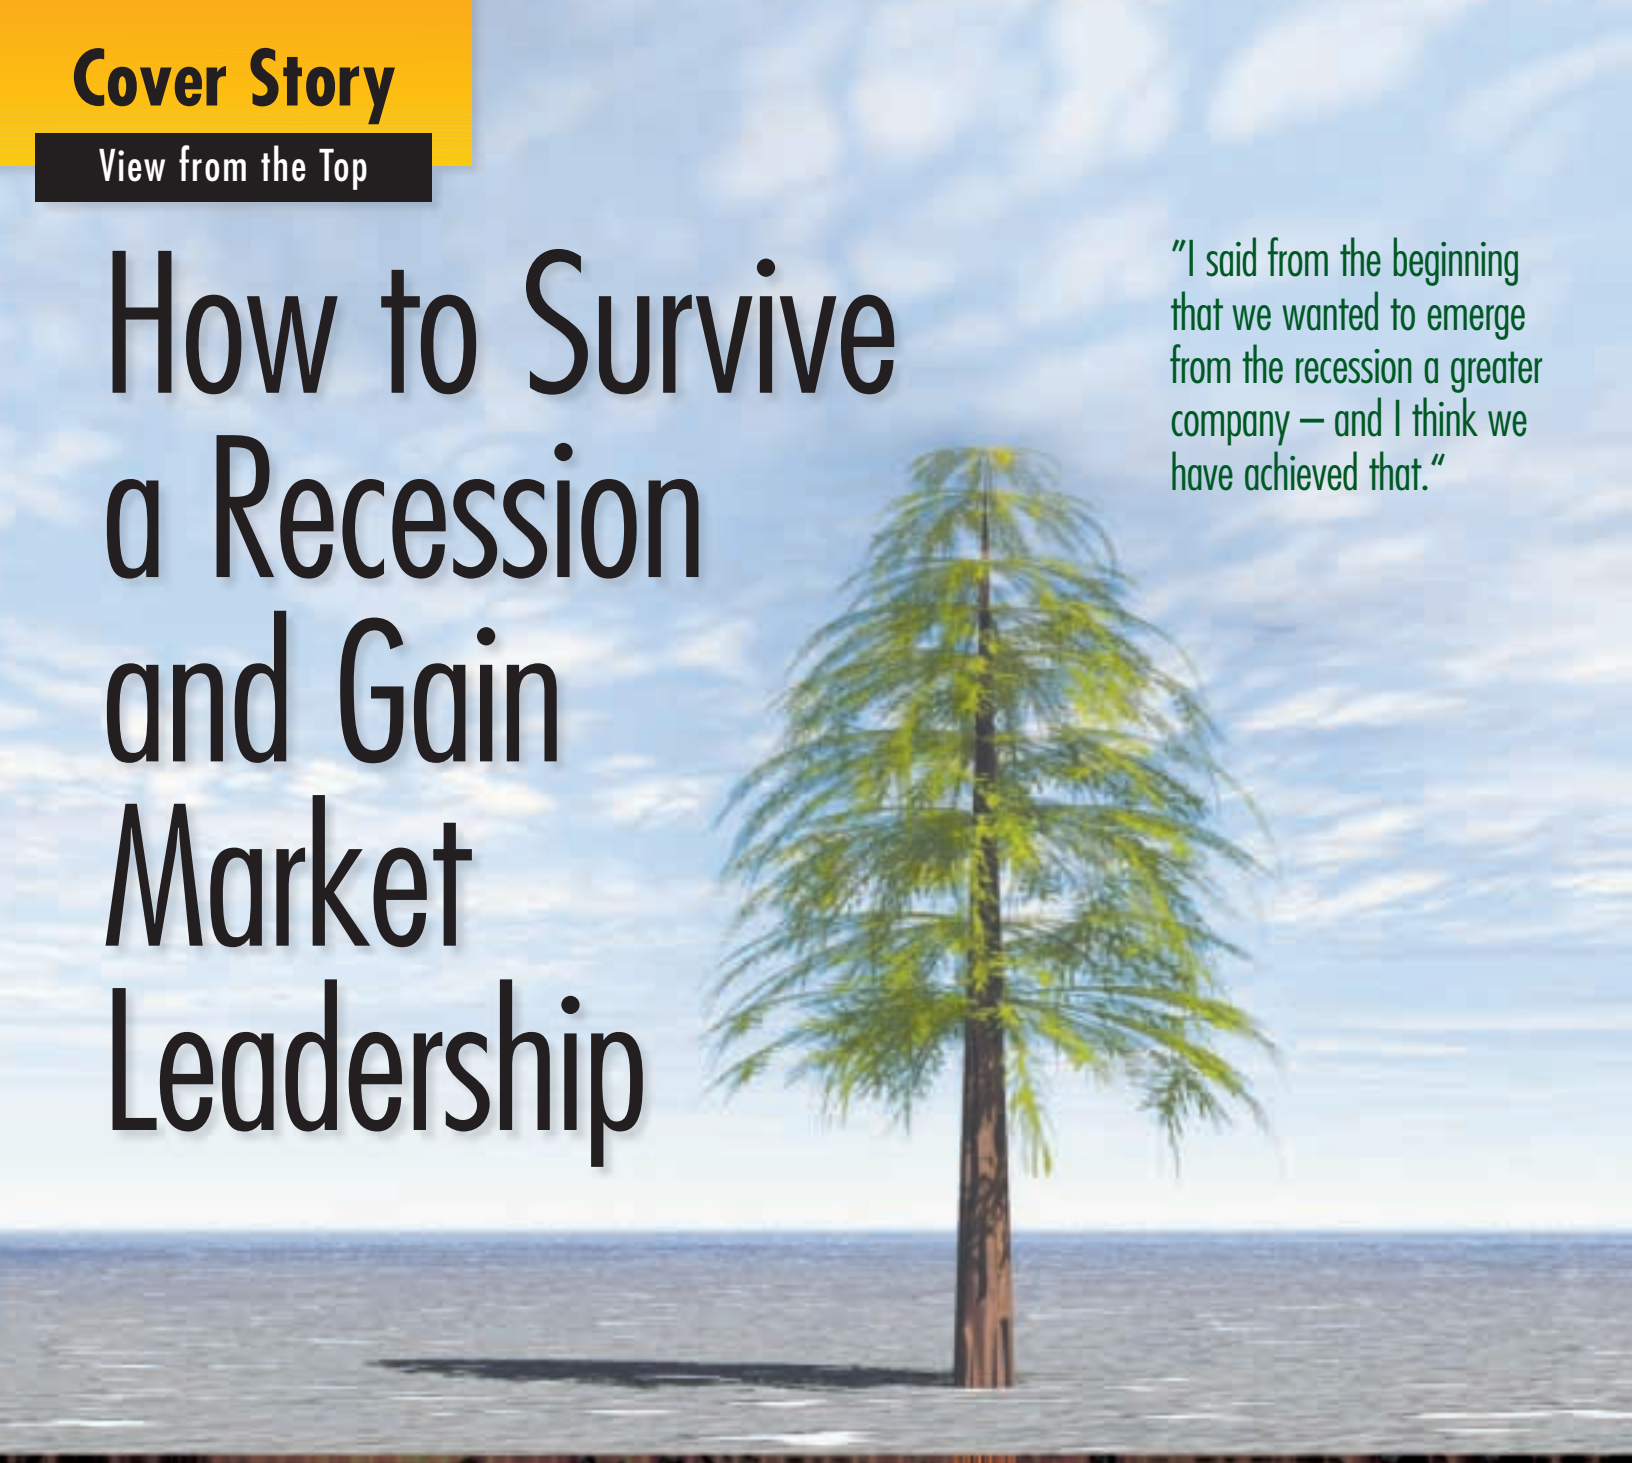

by Wim Roelandts
CEO, Xilinx, Inc.

Last year was a very difficult year for the semiconductor industry. In fact, it was the worst recession ever in the history of semiconductors. In 1985, the industry decreased in a year-over-year basis by 17% to 18%. In 2001, the decline was 32% to 33% – almost twice as bad as 1985.

Nevertheless, I believe that it's in the recessions that the good companies distance themselves from the not so good companies. Even though our business was very weak, we gained market share during this recession in quite a dramatic way. Last quarter, for the first time in our history, we became larger than all the other public programmable logic device companies combined.

From a strategic point of view, we want to win through product leadership. That means we have to have the best products – always from the customer's point of view – which means the highest densities, the fastest performance, and the software and service to support them.

When you are an industry that is new and that is still changing very rapidly, you have to continue to innovate. You have to continue to bring new products into the market.

Despite the recession, the last part of 2001 and the early part of 2002 was one of our most productive periods for new product introductions. In about five months' time, we introduced a new family in every one of our product categories – Spartan™, CoolRunner™, and Virtex™. Each, in its own way, is unique in its capabilities. In fact, two of these new products – the CoolRunner-II RealDigital CPLD and the Virtex-II Pro™ Platform FPGA – have no competition. Literally, there are no comparable products on the market that can compete directly with them.

Total Solutions

I'm very excited about the programmable logic industry. It's an industry that's still in the early phases, but over time, it will become the dominant way to implement integrated circuit designs.

The inherent flexibility of programmable logic allows you to gain time to market, to introduce new products very quickly, and to upgrade your product, even after the sale, through field-upgradeable technology.

We're not just focusing on the chips, however. Our chips must also have the software, the intellectual property cores, the training, and the support to deliver total solutions to our customers.

And we have to continue to be innovative in all these areas of IC technology. My whole belief is that Xilinx must excel in a continuous push to innovation – throughout the company, be it in chip design, software development, or marketing. Innovation is absolutely critical – especially in times of recession.

No Layoffs – Not Just a Nice Thing to Do

This recession has put a tremendous strain on the industry. When I saw this decline coming, I concluded that it would be best for us if we could avoid layoffs – even though the economic consequences would be very difficult. I wanted to avoid layoffs, because once you start layoffs, people get worried – and if people are worried, they cannot innovate.

Innovation means that you have to take risks that some things are not going to work, and therefore, if you are worried about your job, the mantra becomes, "Don't rock the boat. Don't push yourself or make yourself too visible" – which is totally opposite from what we want to see happening.

If people are worried about their jobs, clearly, they're not going to be as productive. Layoffs affect not only the people who get laid off. Very often the people who are still there go through trauma, because they see their friends disappear – and they always have this fear: "Maybe next time it will be me."

Whenever you lay off employees, some of the accumulated knowledge of the company – its "intellectual capital" – is going to walk out the door. Especially in the high-tech industry, the value of the company is in the heads of its engineers and support staff. At Xilinx, 70% or 80 % of our employees are highly educated people working in engineering and marketing and supply-chain management. They are all experts and specialists in their own way – and that's why it's so important to keep these people happy at Xilinx. Once people like that leave, that intellectual capital goes out the door with them, and that's a lot of value for a cutting-edge company to lose.

Several compelling studies have proved that if you go back two years later, the companies that laid off employees are generally in worse shape than companies that didn't. Layoffs have a very profound impact on the morale and the trust between management and employees. Furthermore, a lot of know-how has left the company. And it takes a long time to get it back.

For these reasons, and more, I decided to try to avoid layoffs.

Pay Cuts – An Innovative Approach

The reality of economics is that when business drops, you have to bring your labor costs in line. With layoffs not being an option, the way we cut our labor costs was by a very innovative program where we had a tiered salary reduction that was proportional to the size of each worker's salary. The lowest paid people had no reductions, and the highest paid people had the biggest reductions (myself included).

In implementing this policy, we discovered that we should give people some choices and make sure they get something in return. People had the choice to either take the salary reduction or to take vacation time – or take stock options. Stock options aren't money, but it is compensation for the work people put in. We can't ask the employees to sacrifice unless we can compensate them for the extra flexibility we're asking from them.

The tiered salary reduction was extremely well received in the company. The vast majority – and I'm talking about close to 100% of the employees – felt this was the right thing to do.

People even volunteered for the plan, which was important in Europe, because in Europe, companies cannot mandate salary reductions. It must be done on a voluntary basis. If people hadn't volunteered, we could not have done it, but pretty much all of our European employees volunteered for the salary reduction, because they felt solidarity with the company.

With an acceptance of fair pay cuts and a promise of no layoffs, our employees felt safe to continue to risk and innovate. They have kept their attention full-time on their jobs. The net result has been that we have been able to maintain our aggressive schedule of new product rollouts – and this in a period when our business was off by 50% in three quarters.

Moreover, we were not only able to remain in the black from a profit point of view, but during all this period, we were also able to remain cash-flow positive for the company.

The Importance of Being "Fabless"

One of the reasons we could afford to stay in the black is because Xilinx is a "fabless" semiconductor company. We don't own fabs [fabrication plants]. We don't have this big problem of having a fab that we need to keep busy. For us, if business goes down, we just buy fewer wafers.

The same thing is true for our sales force. We don't have a Xilinx sales force. We work through distribution and sales reps, which again means that our sales costs are by and large variable. Sales people get paid a commission. If they don't sell, there's no commission.

These two factors helped us tremendously in remaining in the black and remaining cash-flow positive, because a lot of our expenses go up and down with the business climate.

The good news is that we reached our lowest quarter in September of 2001.

Since then, business has been starting to improve. We were able to eliminate salary reductions early in 2002, and as of this moment, all salary reductions have been eliminated. The company continues to do very well from both a new product generation point of view, as well as a business point of view.

Innovation – Not Just for Engineers

Innovation is not just the engineering department. Even in the way we manage the company, we are very innovative – and it really has paid off for us. Morale is very high, and the attrition rate is extremely low. People have kept their attention on their jobs of designing great new products and marketing them and supporting them.

In general, I think Xilinx is very much an admired company in our industry. I think we are perceived as a leader in our industry – not just from a product point of view or a financial point of new, but also from a management point of view.

I want to stress, however, that all of this is possible because we are a fabless company. If we had our own fabs, I don't know if we could have done this.

Nevertheless, as a fabless company, our business model lends itself very well to go though these economic cycles that are infamous in the semiconductor industry. This last recession was the worst, but it also served to prove that our business model is a very, very solid model that allows us to go through these ups and downs without too much hurt.

Conclusion

In high technology, things change very, very rapidly. And companies often get in trouble because they don't change quickly enough. One of the reasons they don't change quickly is that the people who see the change happen don't act on it, and the people who can act on it don't even know the change is happening.

Therefore, what we are creating at Xilinx is an environment where everybody says: "Hey, if it changes, let's change with it. Let's anticipate change. Let's make the change happen ourselves. Let's be in charge of the change – not the victim of the change."

It is this kind of thinking that got us through the recession, and now that we are emerging from the recession, there is a tremendous sense of pride in the company. People are saying: "Hey, this is unique. We have done something that very few other companies have done. Even companies that didn't lay people off didn't do it the way we did."

There is absolutely no doubt in my mind that there is a much, much higher spirit of belonging, of community within the Xilinx people than there was before. When you go through a tough time together, it creates solid bonds.

And it shows that, indeed, if you can manage a company during a recession in a correct way, you can really differentiate yourself from the competition. You can gain market share and emerge a stronger company with a creative, dedicated – and intact – workforce at the end of recession. Σ

IBM ASICs Will Incorporate Xilinx FPGAs

New ASICs from IBM will contain Xilinx programmable logic technology.

by Xilinx Staff

IBM and Xilinx recently signed an agreement that could help you shave hundreds of thousands of dollars off the cost of creating custom chips. Under the agreement, IBM has licensed FPGA technology from Xilinx for integration into IBM's recently announced Cu-08 application-specific integrated circuits (ASICs). Cu-08 will support circuits as small as 90 nanometers – less than 1/1,000th the width of a human hair.

This news underscores our joint commitment to a technology relationship aimed at bringing innovative and flexible new "hybrid" chips to market, combining the best attributes of standard ASIC and flexible FPGA technology for use in communications, storage, and consumer applications.

Engineers working on complex chip designs have been clamoring for ways to achieve high levels of integration, yet still have the ability to change "on the fly" late in the design cycle. By combining FPGAs (circuits that can be configured to perform a wide variety of digital electronic circuit functions) with standard ASICs, you get the flexibility of the FPGA, with the density, performance, and overall cost advantages of an ASIC – all on one chip.

"Savings here could be dramatic," said Michel Mayer, general manager, IBM Microelectronics Division. "When an ASIC takes on more function, you can reduce cost by eliminating one, two, or even more separate chips. With this technology, customers would be able to tweak designs and integrate new changes immediately, eliminating the need to restart a whole new design cycle, bringing tremendous time-to-market advantages."

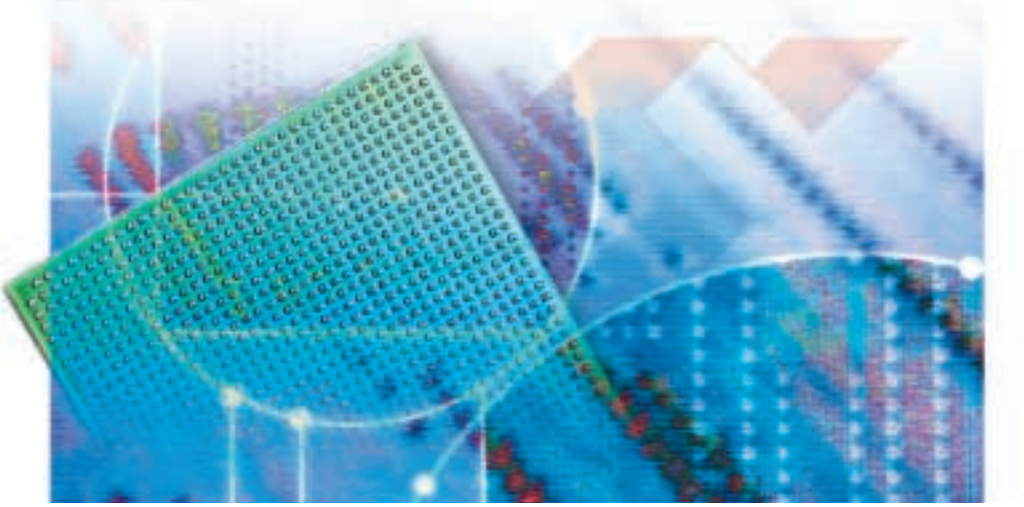

Mayer said changes that force an additional chip prototype can easily cost hundreds of thousands of extra dollars and can stretch design cycles out for several additional months. "This approach is expected to change the landscape entirely," he said.

Cu-08, with as many as eight layers of copper wiring separated by an advanced low-k insulation, will support up to 72 million gates for high-complexity IC solutions. As many as 400,000 ASIC gates may be dedicated to one or more of the FPGA cores on the ASIC.

"We've improved upon our delivery of programmable hardware by allowing reconfiguration using the same chip. Flexibility is the beauty of combining ASIC and FPGA technology," said Wim Roelandts, president and CEO of Xilinx. "Our latest agreement with IBM is a natural extension of and a significant milestone in our existing relationship that allows us to address new opportunities in the high-end ASIC market."

Even after this new custom chip is built into a product, you can add to its functionality, simply, easily, and effectively. A good example would be in the communication industry, where various protocol and interface specifications are constantly evolving. Cell phones, printers, set top boxes, and other consumer electronic products also are well suited for this new approach.

IBM is the number one ASIC manufacturer worldwide. Xilinx is the number one PLD manufacturer worldwide, with more than 15 years experience in FPGA design tools.

IBM and Xilinx already have a productive business relationship. IBM supplies embedded PowerPC™ processors for Xilinx Virtex-II Pro™ FPGAs. In March of this year, Xilinx signed a high-volume, multi-million dollar manufacturing agreement with IBM to manufacture these FPGAs using IBM's advanced 130 nanometer and 90 nanometer copper-based chip-making technology.

The new FPGA cores for the Cu-80 ASICS are now in development. They are expected to be available from IBM, embedded in an ASIC, in early 2004, following IBM's full release of its standard-setting Cu-08 technology.

When Total Cost Management Counts, Xilinx PLDs Pay Off

CoolRunner-II CPLDs, Spartan-IIE FPGAs, and Virtex-II EasyPath Platform FPGAs can give you total system cost solutions superior to those of traditional ASICs.

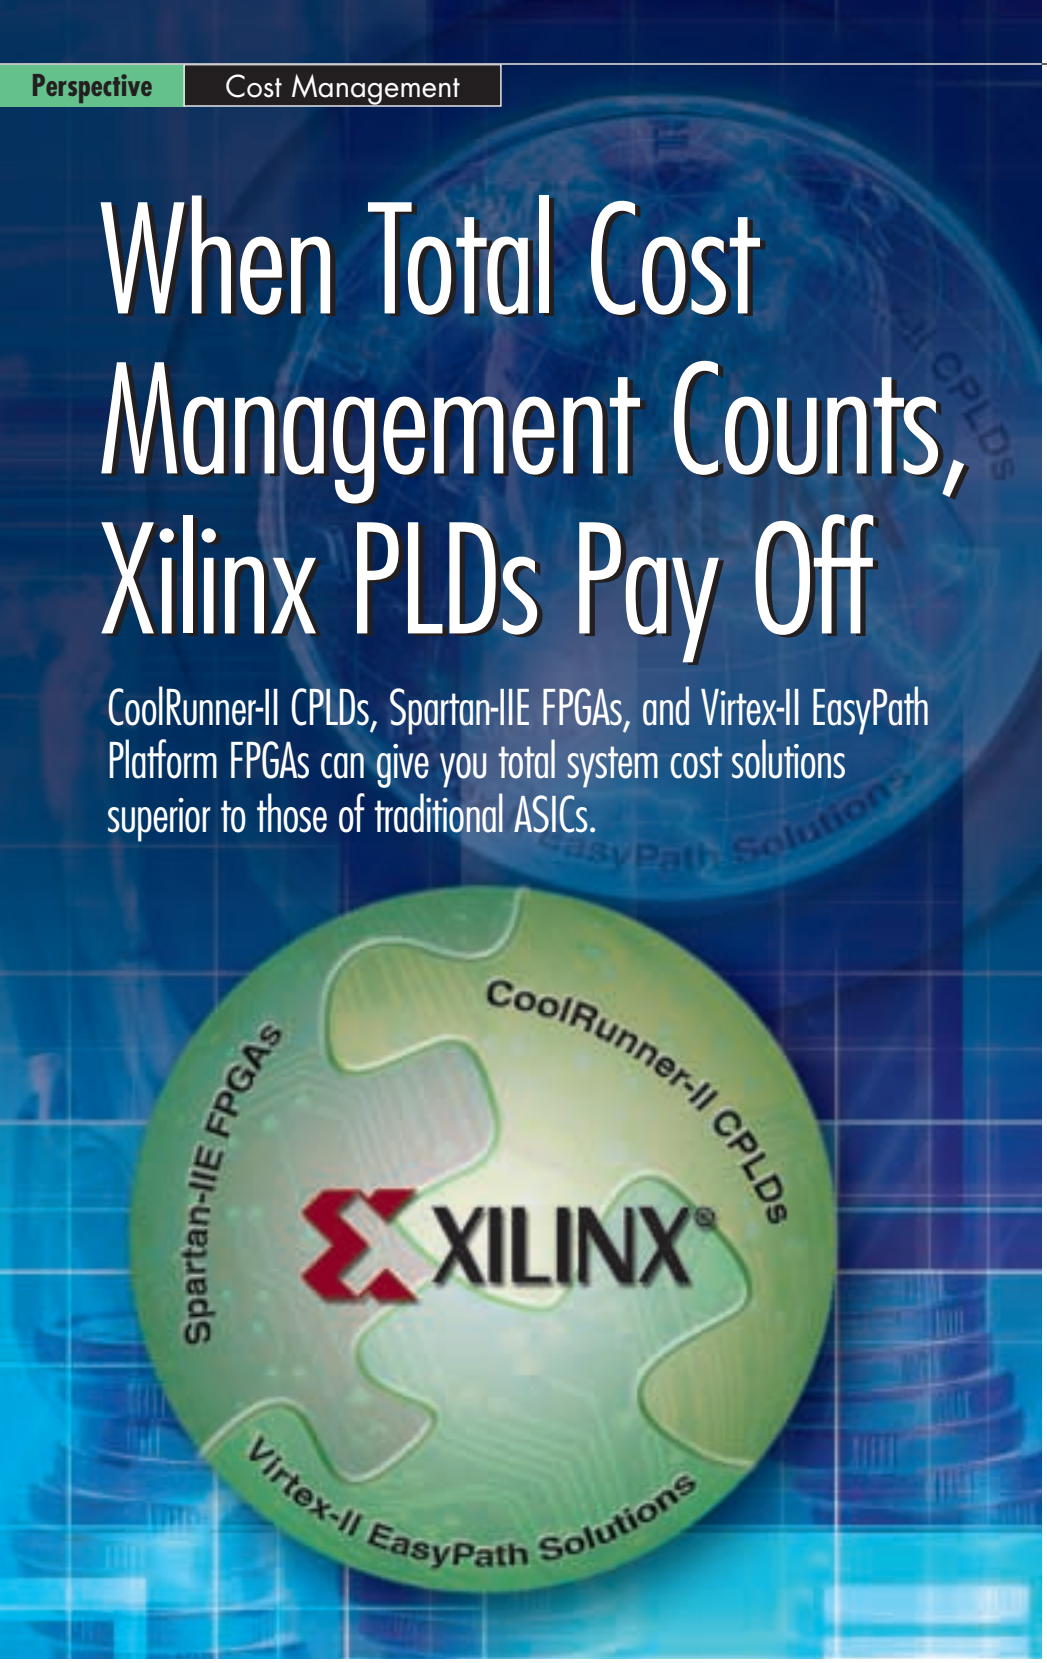

by Eric Thacker
Manager, Product Solutions Marketing
Xilinx, Inc.
eric.thacker@xilinx.com

Rapidly changing markets and new technologies are subverting traditional system design paradigms. Historically, in electronic systems development, the standard design procedure was to prototype a system using programmable logic devices (PLDs), and then redesign the system using ASICs as soon as possible for cost reduction. The idea was to leverage the flexibility and development time advantage of programmable logic until the design stabilized and then to convert the design to a lower-unit-cost ASIC to reduce overall system cost.

In the past, programmable logic was lower density, lower performance, and more limited in capabilities. When systems followed traditional product life cycle behavior, the PLD-to-ASIC design strategy worked quite well. In fact, in some market segments, characterized by relatively mature technologies, established standards, and a stable competitive environment, this strategy is still optimal today.

In dynamic, rapidly changing markets, however, this strategy becomes questionable. For many advanced electronic products – such as plasma televisions, home networking hardware, and digital audio/video equipment – the market is anything but stable. Baseline standards for advanced products are ever-changing, features demanded by customers shift constantly, new technological capabilities are continually being introduced, and competitors are always trying to gain a better market position – resulting in many short-lived products.

These volatile market pressures give a significant competitive advantage to flexible and first-to-market products. In these markets, ASICs – with their long lead times, high initial costs, high minimum order quantities, higher risks, and inflexibility – are not a compelling solution.

Therefore, many designers are tapping into programmable logic's benefits of instant reprogrammability, short lead times, greatly improved performance, expanded densities, and overall flexibility. Indeed, today's programmable logic devices offer many of the same features found in ASICs. Clock management, embedded processors, high-performance DSP, advanced interfacing, and a plethora of intellectual property (IP) cores all make today's PLDs strong candidates for a system's core logic.

With their emphasis on cost-optimization, the CoolRunner™-II CPLDs, Spartan™-IIE FPGAs, and Virtex-II EasyPath™ Solutions families from Xilinx bring these advanced benefits to your high volume applications.

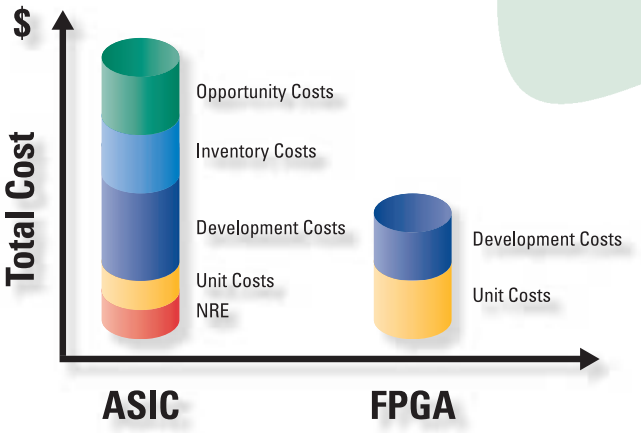

Figure 1 - Cost comparison of ASICs and FPGAs

Total Cost Solution

In the past, the primary motivation in converting a PLD design to one based on an ASIC was unit cost. ASICs, in general, are available at lower unit cost than PLDs. But looking at unit costs alone could actually lead to an overall higher cost solution. That is because unit costs are only one of the factors that contribute to total system cost. Total system costs can be broken down into three broad categories – design and development, production, and life cycle.

Design and Development

Design costs are primarily worker-hour and non-recurring engineering (NRE) costs. PLD designs require fewer worker hours, because design/debug changes can be made instantly, hardware and software can be developed concurrently, test vector generation is unnecessary, and product lead times are practically zero compared to ASICs. PLDs have no NRE costs, which can be a very important consideration as production volumes per design revision are much lower in rapidly evolving markets. In many cases, Xilinx offers free PLD design tools, thus increasing the PLD cost advantage.

There are also lost opportunity costs associated with ASICs. Systems using ASICS lose market penetration due to long development times. Furthermore, the inflexibility of ASIC-based systems prevents upgrading products to meet changing market needs. Overall, PLD design cost advantages alone can often outweigh any ASIC unit cost advantage.

Production

Unit cost is the most obvious cost in this category. ASICs usually do hold the unit cost advantage over PLDs, but this cost differential is shrinking.

Today's PLDs have aggressively driven down both die size and package costs to provide a persuasive ASIC alternative. The advanced system capabilities of modern PLDs allow designers to integrate discrete component functions into the PLD – reducing board and total component costs.

Finally, unlike an inflexible ASIC, one line of PLDs can be inventoried to supply multiple applications. You can reprogram and redeploy PLDs as needed. Combined with much lower minimum order quantities, PLDs reduce inventory costs, as well as the risk of obsolescence, further offsetting PLDs' unit cost disadvantage. See Figure 1.

Life Cycle

Life cycle costs are probably the least considered costs when designing an electronic system, but they can have a significant impact on the total return from a given design. Using PLDs instead of ASICs means practically no risk of obsolete inventory, because standard-product PLDs can be redeployed to different applications as needed. Unlike ASICs, PLDs do not have to be scrapped if a specific product is cancelled.

In addition, because PLDs can be reprogrammed, it is possible to upgrade products in the field simply by downloading a new hardware configuration file to the PLD. Bug fixes and feature upgrades can be performed remotely without changing any hardware.

Finally, long-lived legacy products using ASICs can face the need for board redesign should the ASIC product or process be discontinued. PLD reprogrammability avoids this problem. Thus, legacy products with PLDs do not siphon design resources from next-generation products, and this presents a significant market cost savings.

As illustrated in the discussion of these three cost categories, we must look beyond unit cost when deciding on logic solutions based on cost. In many applications, the unit cost savings of ASICs are more than offset by the lower risks, design, production, and life cycle costs of PLDs.

Xilinx Solutions for Total Cost Management

Xilinx has taken the lead among PLD suppliers in providing low risk, low cost, high performance logic solutions for today's rapidly changing and highly competitive markets.

Low Power, Low Density: CoolRunner-II CPLDs

CoolRunner-II RealDigital CPLDs combine high performance, ultra low power, low cost, and integration of specialized capabilities to bring you a new class of programmable logic devices. The versatility and low cost of CoolRunner-II CPLDs allow them to serve a multitude of logic, interface, and control functions. See Figure 2.

Figure 2 - CoolRunner-II CPLDs deliver high performance at low cost.

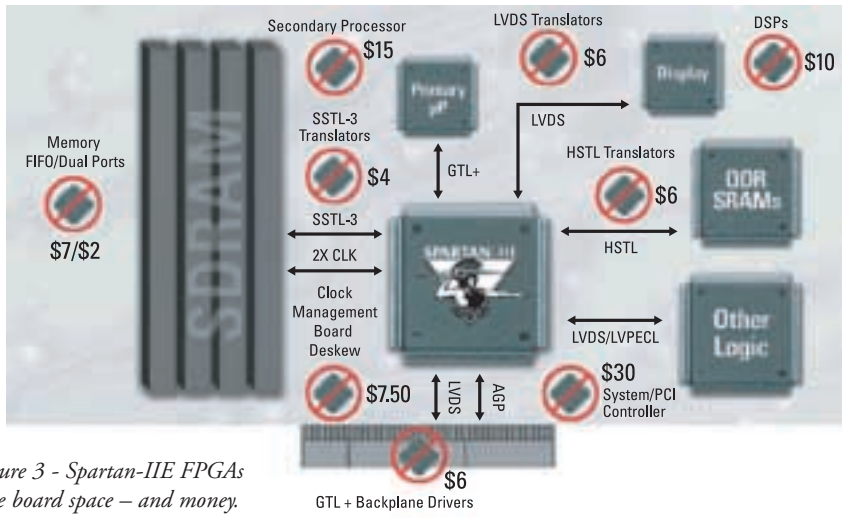

Figure 3 - Spartan-IIE FPGAs save board space – and money.

The 0.18-micron CoolRunner-II family utilizes second-generation Fast Zero Power™ technology for low power, high performance operation. These RealDigital CPLDs feature an all-digital core that eliminates power-hungry analog sense amplifiers. The CPLDs also offer increased capabilities, higher performance, and lower cost – with no performance penalty.

Available in densities from 32 to 512 macrocells, the CoolRunner-II family provides performance over 330 MHz, with standby power consumption of less than 100 µA. CoolRunner-II devices also deliver advanced system features that enable the integration of costly discrete system functions into a single reprogrammable device. This integration results in lower costs, further power reduction, increased reliability, faster time to market, and smaller designs. These system features include superior I/Os, advanced clocking features, and four levels of design security.

The in-system reprogrammable nature of CoolRunner-II CPLDs allows designers to change designs on the fly, correct designs, and test out alternate designs without changing any hardware. This saves significant engineering change orders (ECOs) and NRE charges compared to ASIC solutions.

The CoolRunner-II CPLDs' advanced I/O interface capabilities include DataGATE, a programmable on/off switch for power management; advanced interfaces such as HSTL, and SSTL; and Schmitt trigger inputs for input signal conditioning. All of these features reduce the need for external components, reducing part costs and board space.

To further reduce system costs and the need for additional components, the CoolRunner-II family provides a unique feature called CoolCLOCK. CoolCLOCK is a combination of a clock doubler and clock divider. The incoming clock is divided by two and then doubled at the output to maintain the same performance, while reducing the internal power consumption.

To increase your cost management advantage, Xilinx provides a free Internet-based ISE WebPACK™ design tool suite and a free WebFITTER™ design fitting tool to give you a complete, pushbutton design solution.

The combination of advanced interfacing, integrated system features, and free design tools makes CoolRunner-II CPLDs a compelling choice for logic functions in cost-sensitive applications. For more information, visit *www.xilinx.com/coolrunner2*.

Moderate Density, System Integration: Spartan-IIE FPGAs

Cutting-edge consumer electronics markets demand products that provide performance and flexibility in a cost-optimized format for home networking equipment, home theater systems, and other high-end personal electronic devices. These consumer products must accommodate rapidly evolving standards, shifting demand patterns, highly fluid competitive positioning, and the continuous introduction of new capabilities. Logic requirements for these designs include moderate density, advanced system features, predesigned and verified IP blocks, and high-quality design tools. Spartan-IIE FPGAs meet all of these requirements and more.

Xilinx designed its next-generation Spartan-IIE consumer FPGA with total cost management in mind. Spartan-IIE FPGAs range in density from 50,000 system gates to 300,000 system gates. To reduce FPGA component costs, Spartan-IIE FPGAs are produced by the industry's most advanced foundry process on 300 mm wafers. This process lowers die cost and enables advanced, low-cost packaging.

To reduce board cost and improve system performance, Spartan-IIE FPGAs have a number of built-in advanced system features that allow you to eliminate costly discrete translators and commodity chips from your boards, as shown in Figure 3.

Advanced interfacing – including LVDS, LVPECL, HSTL, and SSTL – remove the need for specialized interface components on the board. Furthermore, Spartan-IIE FPGAs offer four delay-locked loops (DLLs) for advanced clock management. Spartan-IIE FPGAs also contain both parameterizable block RAM (single or dual port) and distributed RAM (scalable and adjacent to logic). All these features can speed the design process, provide design flexibility, reduce design worker hours, deliver very high levels of performance, and accelerate time to market.

As consumer products increase their digital capabilities, digital signal processing (DSP) capability has become a critical performance factor. The Xilinx XtremeDSP™ initiative enables nearly one billion DSP multiply-and-accumulates (MACs) per second per dollar of performance in Spartan-IIE devices – delivering advanced DSP without the need for specialized board components.

Xilinx offers more than 200 Spartan-family IP cores to speed logic design and reduce engineering resources needed to complete the system. These IP cores are also opti-

mized for die area, which may permit the use of smaller density Spartan-IIE parts, resulting in further cost savings.

Additionally, the free ISE WebPACK software provides the industry's most advanced design tool platform to support the entire Spartan series.

Die and packaging, advanced system features, extensive IP offerings, enhanced flexibility – all these Spartan-IIE cost control features give you the performance, features, and time to market advantage necessary to compete in today's rapidly evolving electronics markets. For more information, visit *www.xilinx.com/spartan*.

High Performance, High Density: Virtex-II EasyPath FPGAs

The Virtex-II™ FPGA is the first embodiment of the Platform FPGA concept. Virtex-II devices deliver SystemIO interfaces to bridge emerging standards, XtremeDSP performance up to 100X conventional DSP solutions, multi-gigabit transceivers, and Empower! embedded processor technology. Together with the industry's most advanced design tools, Virtex-II Platform FPGAs offer a feature set that is unparalleled in the industry.

Virtex-II EasyPath FPGAs bring the unprecedented capabilities of the Virtex-II family to cost-sensitive applications by reducing the cost of high-density FPGA designs. Virtex-II EasyPath Solutions use the exact same FPGA silicon as the Virtex-II FPGA family, but Virtex-II EasyPath Solutions make use of a specialized testing program to verify the FPGAs for a specific customer application, resulting in higher yields and significantly lower costs.

Virtex-II EasyPath solutions provide an FPGA volume conversion strategy with no risk, no investment of customer engineering resources, and the fastest conversion time of any competing high-volume strategy for high-density FPGA designs. The main benefits of Virtex-II EasyPath solutions are:

- Lower unit costs: Customers can expect a 30-90% price reduction compared to standard FPGAs.
- Lower up-front costs: These costs can be as low as 30% of the NRE costs of an ASIC or gate array.
- No added engineering resources: Unlike other FPGA conversion strategies, Virtex-II EasyPath FPGAs require no additional engineering resources or tools. Customers are not required to contribute any resources for conversion, verification, qualification, or board redesign.
- No risk: Virtex-II EasyPath Solutions have none of the performance, timing, functionality, or architecture match risks of alternative FPGA conversion strategies, because the conversion FPGA fabric is the same as the standard Virtex-II FPGA.
- Minimal lead times: With Virtex-II EasyPath FPGAs, initial orders are filled within six weeks and restocking orders are filled within days of being placed. Compare that to lead times of multiple months for both initial and follow-on orders for alternative conversion strategies.

Because Virtex-II EasyPath FPGAs are tested for a specific design configuration, testing throughput and yields are significantly higher, thus driving down the unit costs of the FPGA, as shown in Figure 4. You can initially design your system using standard Virtex-II FPGAs. Once the design has stabilized, you can use Virtex-II EasyPath Solutions to convert to a lower-cost platform with none of the risk, lower NRE costs, and none of the additional engineering resources required by conversion to ASICs. For more information, visit *www.xilinx.com/virtex2*.

Conclusion

From low density, ultra low power CoolRunner-II CPLDs to mid-range, low cost Spartan-IIE FPGAs to high density Virtex-II EasyPath Platform FPGA Solutions, Xilinx presents a portfolio of products that can play a critical role in your total cost management of system design. These advantages include:

- Cost-optimized silicon and packaging to minimize unit costs
- Advanced system features to reduce costs through system integration
- Reprogrammability to streamline design and debug
- No-cost system upgrade capabilities
- Reduced risk of obsolete inventory
- An extensive IP library to improve design efficiency and reduce system components
- Free design tools that maximize design efficiency.

Together, Xilinx CPLDs and FPGAs give you a cost-optimized solution for managing system costs while maximizing system performance, functionality, and flexibility.

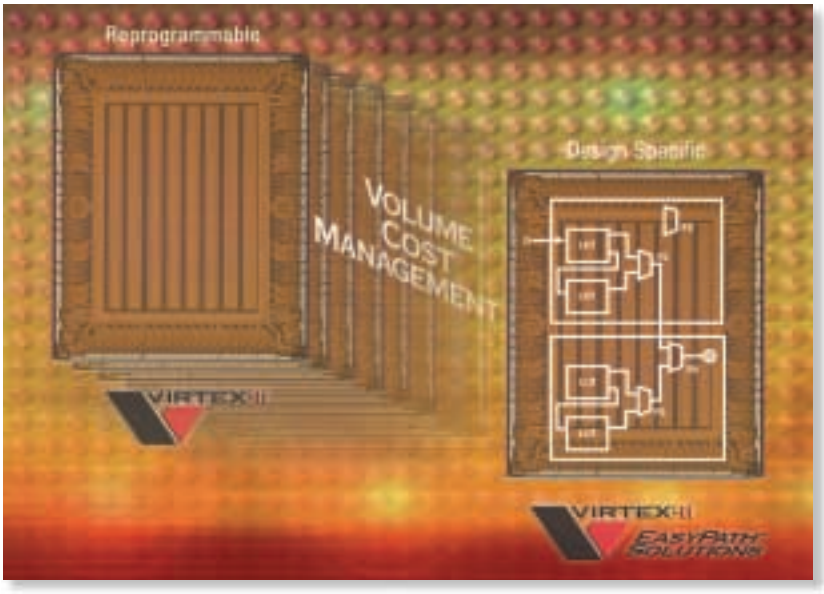

Figure 4 - Virtex-II EasyPath Solutions are design-specific and cost-effective.

New Virtex-II Pro Devices — Increased Capability, Lower Cost

Now you get programmable system features at programmable logic prices.

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

In a significant move, Xilinx recently expanded the Virtex-II Pro™ Platform FPGA family two ways:

- Five new devices offering increased capability
- Lower cost versions for the entire family in a new reduced-speed series.

With this, the new Virtex-II Pro family expands to ten devices, and sets industry records in every conceivable category, providing the advanced system-level features and high performance you need, at remarkably low prices.

The new Virtex-II Pro FPGA solution, illustrated in Figure 1, provides everything you need, on a single, reprogrammable device, including:

- The world's highest logic and memory density
- The world's highest I/O capacity
- The world's fastest DSP capability
- The world's fastest multi-gigabit serial I/O connectivity
- The world's most flexible parallel I/O connectivity
- The world's only embedded processing FPGA solution using IBM® PowerPC™ processor cores.

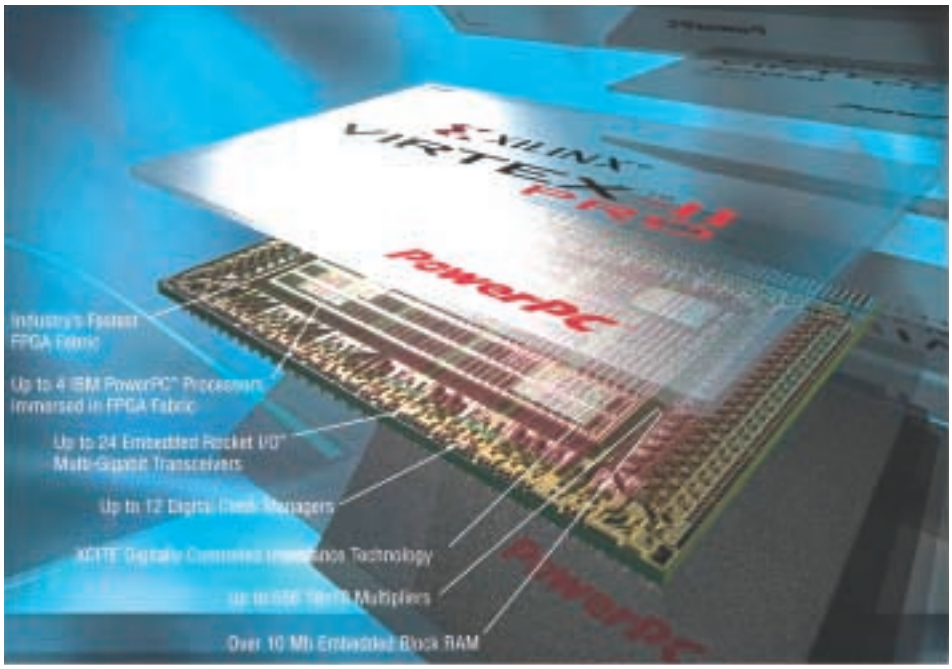

Figure 1 - Virtex-II Pro Platform FPGA

Your demands, both for increased capability as well as lower cost versions of our Virtex-II Pro devices, were made clear at our Programmable World 2002 event. And now, we can deliver the performance and the pricing you need due to the success of our initial devices and development tools, as well as advances in our process technology which includes 300-mm wafers.

These new, lower cost devices not only deliver tremendous cost savings over comparable competing programmable device offerings, but they also bring Virtex-II Pro capabilities within the reach of a much wider range of applications. With the Virtex-II Pro family you now get multiple IBM PowerPC processors and multi-gigabit serial transceivers at about the same cost as our Virtex-II™ FPGAs (which don't offer these features), effectively providing programmable system features at programmable logic prices.

The Highest Performance and the Lowest Cost

The new Virtex-II Pro XC2VP125 FPGA, shown in Table 1, is the world's highest capacity, highest performance, and highest capability programmable logic device, delivering unique system performance such as one trillion MACs/sec DSP performance and 120 Gb/s aggregate payload bandwidth. There is no programmable device in the world that competes with it in features, performance, or cost. With such unparalleled leadership advantage, the Virtex-II Pro Platform FPGAs put you on a platform of success, closer to your design goals.

Expanding the Market and Target Applications

The new -5 speed grade delivers cost-effective 2Gb/s transceiver performance, with over 266 MHz embedded PowerPC processor performance. In comparison, our fastest devices (-7 speed grade) provide up to 3.125 Gb/s transceiver performance, with over 300 MHz embedded PowerPC processor performance. The -5 speed grade is therefore ideal for medium to high-speed applications including storage, digital broadcast, networking, and wireless networking. For example, in networking, as you move from the core of the network to the edge, the line rate and processing speed of the equipment decreases while the need for intelligent processing increases. In addition, price pressure increases, because the edge equipment is more price sensitive than the core equipment. This analogy can be applied to almost any application and market segment. Xilinx is providing -5 devices specifically to address such price-sensitive applications that benefit from system features that include embedded processors and multi-gigabit transceivers.

Virtex-II Pro FPGA Architecture Review

The Virtex-II Pro family, shipping since February of this year, is the world's first 130 nm, 9-layer copper, FPGA family and the first and only one to offer integrated IBM PowerPC microprocessor technology and 3.125 Gb/s serial transceivers. Its architecture is built upon the award-winning Virtex-II FPGA fabric, which already has thousands of satisfied users.

"The Virtex brand has been the number one choice of designers worldwide, based on its industry-leading capacity, performance, capabilities, and cost-effectiveness. In each generation of families, our mission is to deliver significantly enhanced functionality while driving down prices. We continue this strategy in the Virtex-II Pro family by including state-of-the-art multi-gigabit serial transceivers and high-performance PowerPC RISC CPUs, in the standard feature set, at no charge. Thus, we enable the creation of next-generation systems designs at previously unattainable price points."

Erich Goetting — Vice President and General Manager of the Xilinx Advanced Products Division

With 10 devices, multiple packages, and multiple speed grades to choose from, the Virtex-II Pro FPGAs give you the right features and performance to meet your specific needs. The family, as shown in Table 1, includes:

- Immersed IBM PowerPC processors (0 to 4)
- Integrated 3.125 Gb/s Rocket I/O™ transceivers (0 to 24)
- The award-winning Virtex-II FPGA fabric.

IBM PowerPC Processors

Each immersed IBM PowerPC processor runs at up to 300+ MHz and 420+ Dhrystone MIPS providing high-performance embedded processing. The PowerPC processor is supported by the IBM CoreConnect™ bus technology. A complete set of embedded software tools are available to help you quickly develop and debug your PowerPC designs – through an OEM agreement, Xilinx provides software tools from Wind River Systems, customized for Virtex-II Pro FPGAs.

"With the immersion of PowerPC cores, Xilinx has developed the optimal platform to implement our storage network designs. Virtex-II Pro devices provide the critical elements of system-level design, including high-speed serial I/O, making it an ideal combination for our next generation products. Moreover, the unique architectural synthesis design environment allows us to make hardware and software tradeoffs throughout the design cycle."

Sandy Helton – Executive VP and CTO, San Valley

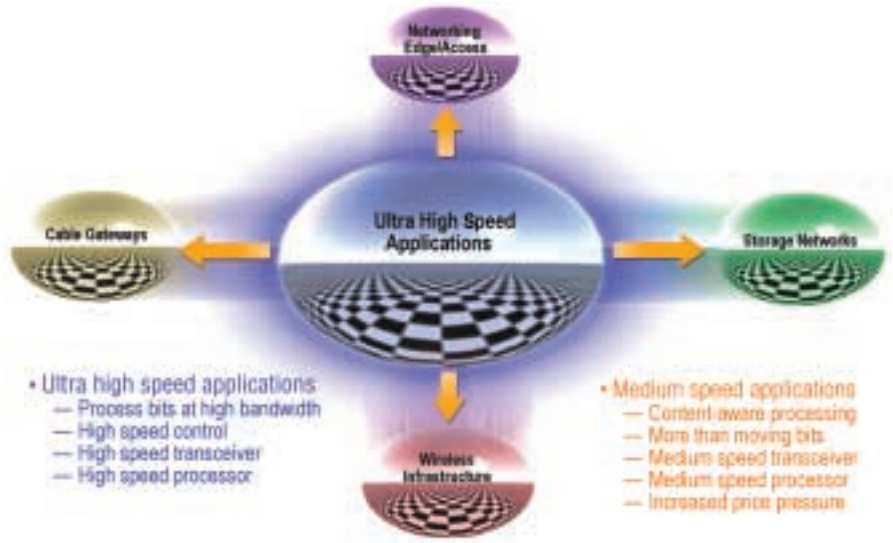

Figure 2 - Expanding the market from ultra-high speed to medium-high speed applications

Rocket I/O Multi-gigabit Transceivers

The Rocket I/O™ multi-gigabit transceivers (MGTs) are based on the Mindspeed SkyRail™ CMOS technology. Each transceiver runs at rates ranging from 622 Mb/s to 3.125 Gb/s, and includes the entire transceiver support circuitry. Each Rocket I/O transceiver consists of both a digital Physical Coding Sublayer (PCS), as well as an analog Physical Media Attachment (PMA), to provide an integrated SerDes (serializer/deserializer) function.

Award-winning Virtex-II Fabric

The Virtex-II FPGA fabric provides features for high-performance system design that has made it a platform of choice for users. These features include:

- High-performance logic with a wide choice of densities
- Embedded block RAM
- XtremeDSP embedded hardware multiply circuitry
- Advanced digital clock management circuitry
- XCITE technology for digitally controlled impedance matching on I/Os
- Bitstream encryption technology for design protection
- Flexible Select I/O™-Ultra technology for supporting over 20 single-ended and differential I/O signaling standards.

Device	XC2VP2	XC2VP4	XC2VP7	XC2VP20	XC2VP30	XC2VP40	XC2VP50	XC2VP70	XC2VP100	XC2VP125
Logic Cells	3,168	6,768	11,088	20,880	30,816	43,632	53,136	74,448	99,216	125,136
Block RAM (Kbits)	216	504	792	1,584	2,448	3,456	4,176	5,904	7,992	10,008
18x18 Multipliers	12	28	44	88	136	192	232	328	444	556
Digital Clock Management Blocks	4	4	4	8	8	8	8	8	12	12
Configuration Memory (Mbits)	1.31	3.01	4.49	8.21	11.36	15.56	19.02	25.6	33.65	42.78
IBM PowerPC Processors	0	1	1	2	2	2	2	2	2	4
Multi-Gigabit Transceivers	4	4	8	8	8	12*	16*	20	20*	24*
Max Available User I/O	204	348	396	564	692	804	852	996	1164	1200
Package										
FG256	140	140								
FG456	156	248	248							
FF672	204	348	396							
FF896			396	556	556					
FF1152				564	692	692	692			
FF1148*						804	812			
FF1517						804	852	964		
FF1704								996	1040	1040
FF1696*									1164	1200

Table 1- Virtex-II Pro Platform FPGAs Power of Choice (new -5 devices shown in red)

**Note: FF1148 and FF1696 packages support higher user I/O and zero Rocket I/O multi-gigabit transceivers*

Conclusion

The expanded Virtex-II Pro family provides a broad choice of 10 devices in scalable features and speed grades. It marks leadership in every facet, providing more capability than any other competing programmable logic device. In addition, the new pricing delivers programmable system features at programmable logic prices, making embedded PowerPC processors and multi-gigabit transceivers standard features in high-performance programmable devices. For more information on Virtex-II Pro Platform FPGAs and design resources, go to *www.xilinx.com/virtex2pro.* Σ

How to Use Virtex-II FPGAs to Deliver Gigabit Advanced Data Security with Helion Encryption Cores

XPERTS
Xilinx Program for Engineering Resources from Third Parties

Virtex-II Platform FPGAs enabled Helion Technology Ltd. not only to deliver on time but to exceed the client's requirements.

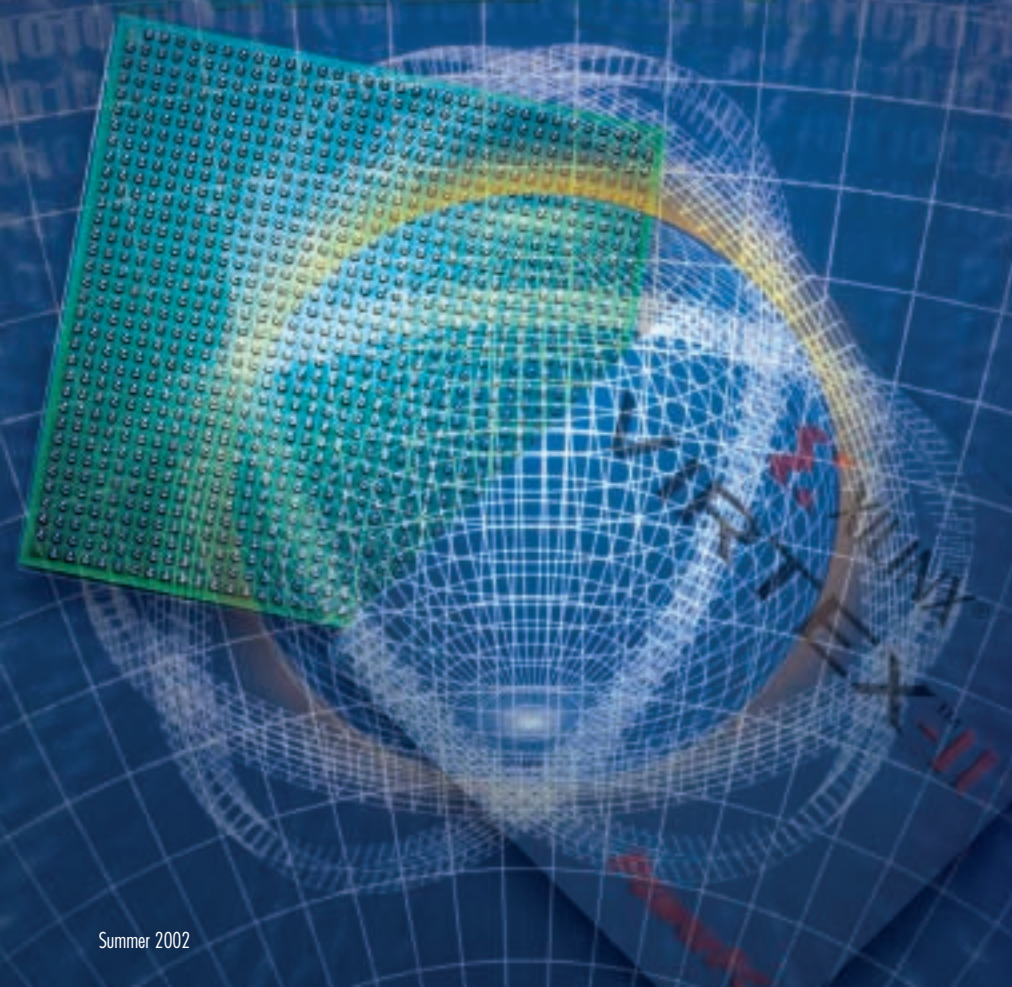

by Simon Cocking
Design Consultant
Helion Technology Limited
simon@heliontech.com

As secure communication data rates increase, software implementations of the key data encryption algorithms present a major system bottleneck. To make matters worse, standard off-the-shelf CPUs and DSPs have failed to keep pace with the computational demands of data encryption algorithms. Besides, CPUs and DSPs just have too many other tasks to perform.

The magnitude of the software performance gap for high-speed data encryption is illustrated by the benchmark results shown in Table 1. These figures are for implementations of the new 128-bit Advanced Encryption Standard (AES) algorithm running on standard CPU and DSP architectures – as compared to how a Virtex™-II Platform FPGA solution based on a Fast AES core from Helion Technology Ltd. can provide the multi-gigabit encryption data throughput required by high-speed interconnect technologies.

The Design Brief

Last year, we at Helion were asked by one of our clients to provide a flexible, high performance FPGA-based data security engine to handle encryption of Internet Protocol (IP) packet-based data for their next-generation network security products. Even though their product was based on a high-performance dual 64-bit MIPS network processor, benchmarking had shown a software encryption solution to be somewhat short of achieving the minimum 600 Mbps throughput our client required.

The system requirement was for a PCI-based engine that could accommodate multiple encryption algorithms: AES and a proprietary algorithm for legacy purposes. The solution had to be low cost, capable of supporting gigabit data rates, and flexible enough to allow for future algorithm changes.

After evaluating the available options, the only solution that met all of the requirements was a PCI card containing the one million-gate Virtex-II XC2V1000 Platform FPGA. The device incorporated a Xilinx LogiCORE™ 32-bit/66-MHz PCI core and Helion-designed encryption cores.

Why Use a Virtex-II FPGA?

From our many years of design experience with leading edge technology, we were well aware of the potential pitfalls, such as inadequate design tool support and unavailable parts. Fortunately, Virtex-II Platform FPGAs use the same EDA toolset as earlier technologies, and the XC2V1000 device was already in production. The reasons for choosing the XC2V1000 instead of the more established Virtex-E family were compelling:

- Costs – The Virtex-II XC2V1000 cost less than half the nearest equivalent Virtex-E(-8) device.
- Performance – Benchmarking with the Helion AES core showed the Virtex-II(-4) FPGA to be 30% faster than the nearest equivalent Virtex-E(-8) device.
- Bigger and wider block RAM – We found 32-bit wide data buffers were more efficiently implemented in Virtex-II block RAM. For example, a 512x32 buffer required only a single Virtex-II block RAM rather than four in Virtex-E technology.
- Enhanced CLBs – The MUXF7 and MUXF8 primitives allowed up to four slices to be combined for fast implementations of any logic function up to eight inputs wide. This led to fewer logic levels and faster critical paths for the very wide logic functions typical of encryption algorithms.
- Digital Clock Manager (DCM) – True clock frequency synthesis allowed the encryption data clock to be tuned to closely match worst-case PAR (place-and-route) timing, optimizing encryption throughput.
- Easier Design Implementation – The raw speed of Virtex-II devices meant no PAR guide file was required for the LogiCORE 66-MHz PCI bus logic.

Encryption Engine Overview

A block diagram of the encryption engine is shown in Figure 1. The only external interface is the LogiCORE 32-bit/66-MHz PCI bus.

All encryption key information and data to be encrypted (plaintext) are stored in PCI system memory and transferred into the FPGA by the direct memory access (DMA) controller. This configuration allows the encryption keys to be easily changed by software – an important security feature.

The DMA read-and-write buffers use block RAM to support the long PCI burst read-and-write transfers needed to achieve optimum performance. Checksum calculation and buffering is also provided for the encrypted data (ciphertext) to off-load this task from the system software, which reads the checksum(s) at the end of each DMA transfer.

The main data path between the DMA buffers is through the encryption cores. The Helion 128-bit fast AES encryption core had already been developed, so it was simply a case of dropping it into the design. However, it was still necessary to develop a high-performance core for the proprietary algorithm.

Encryption Engine Operation

While the encryption engine is inactive (no DMA encrypt transfer in progress), the PCI interface defaults to "target" mode and waits for the system CPU to initiate a DMA transfer. Once a DMA operation is requested, the PCI interface is switched to "master" mode, and the DMA controller initiates the required bus transfers.

Upon completion, the encryption engine asserts the PCI interrupt request and returns the interface to "target" mode so that the system CPU can read the checksums and/or start the next transfer.

Solution	Clock (MHz)	Encryption Data Rate	Comments
TMS320C62XX 32-bit fixed-point DSP	200	112 Mbps	
MIPS-based 64-bit RISC processor	250	392 Mbps	Assumes fully primed cache and CPU fully dedicated to encryption
Pentium III	1,000	464 Mbps	Assumes fully primed cache and CPU fully dedicated to encryption
Helion Fast AES Core	132	1536 Mbps	Virtex-II XC2V1000-4 target

Table 1 - Typical achievable data rates for 128-bit Advanced Encryption Standard solutions

Because the plaintext data stream written into PCI system memory can be fragmented into multiple IP packets, the system software is responsible for creating a control data structure that details the size, location, and number of fragments, as well as the encryption type to use. This structure is transferred from memory into a block RAM in the DMA controller at the start of each encrypt transfer.

The encryption engine can then read in all plaintext fragments, encrypt them using the selected algorithm, and reassemble the ciphertext fragments in memory. In the process of writing ciphertext back to memory, a checksum is calculated for each encrypted fragment, and this value is added to a total checksum of all fragments. These checksums are stored in a small distributed RAM in the FPGA, and then read back by the system CPU at the end of each transfer.

Encryption Engine Performance

Our client was delighted to learn that not only was the final encryption engine design delivered on time, but that it also exceeded the required data throughput for both encryption algorithms by a significant margin.

The proprietary algorithm core is clocked at 66 MHz to achieve a raw encryption throughput of 990 Mbps. The AES core is clocked at 132 MHz to yield a raw encryption throughput of 1.536 Gbps.

However, these raw throughput figures are degraded somewhat in the final system due to bus arbitration overhead in the PCI bridge. In fact, the data throughput of the Helion fast AES core is so high that the maximum available bandwidth on the 32-bit/66-MHz PCI bus acts as a ceiling on its performance.

Frequent changing of the encryption keys will also lead to a reduction in overall data encryption throughput due to the increased transfer load on the PCI bus.

Conclusion

The encryption engine project we've described is a great example of how a combination of third-party intellectual property cores and Virtex-II Platform FPGA technology can yield flexible, high-performance security products in record time.

A large reduction in the effort required to implement such a design relies on the ready availability of intellectual property like the Xilinx LogiCORE PCI products and the range of advanced encryption cores we have developed at Helion. Design blocks like these enable a single engineer to attack a million-gate design and achieve results in a matter of weeks.

Future plans for our encryption engine include the use of a higher bandwidth bus interface, such as HyperTransport™ technology in place of PCI. We also plan the addition of multiple AES cores to provide a full-duplex (encrypt and decrypt), higher performance solution.

The Helion encryption cores allow for even higher throughput: Eight similar Fast AES cores in a larger Virtex-II Platform FPGA could provide encryption at rates in excess of 12 Gbps. This is fast enough to support emerging high-speed interconnect standards, such as OC-192, 10 Gigabit Ethernet, POS-PHY L4, and RapidIO™ Phy.

For more detailed information on Helion and our products and services, please visit the Helion website at *www.heliontech.com*.

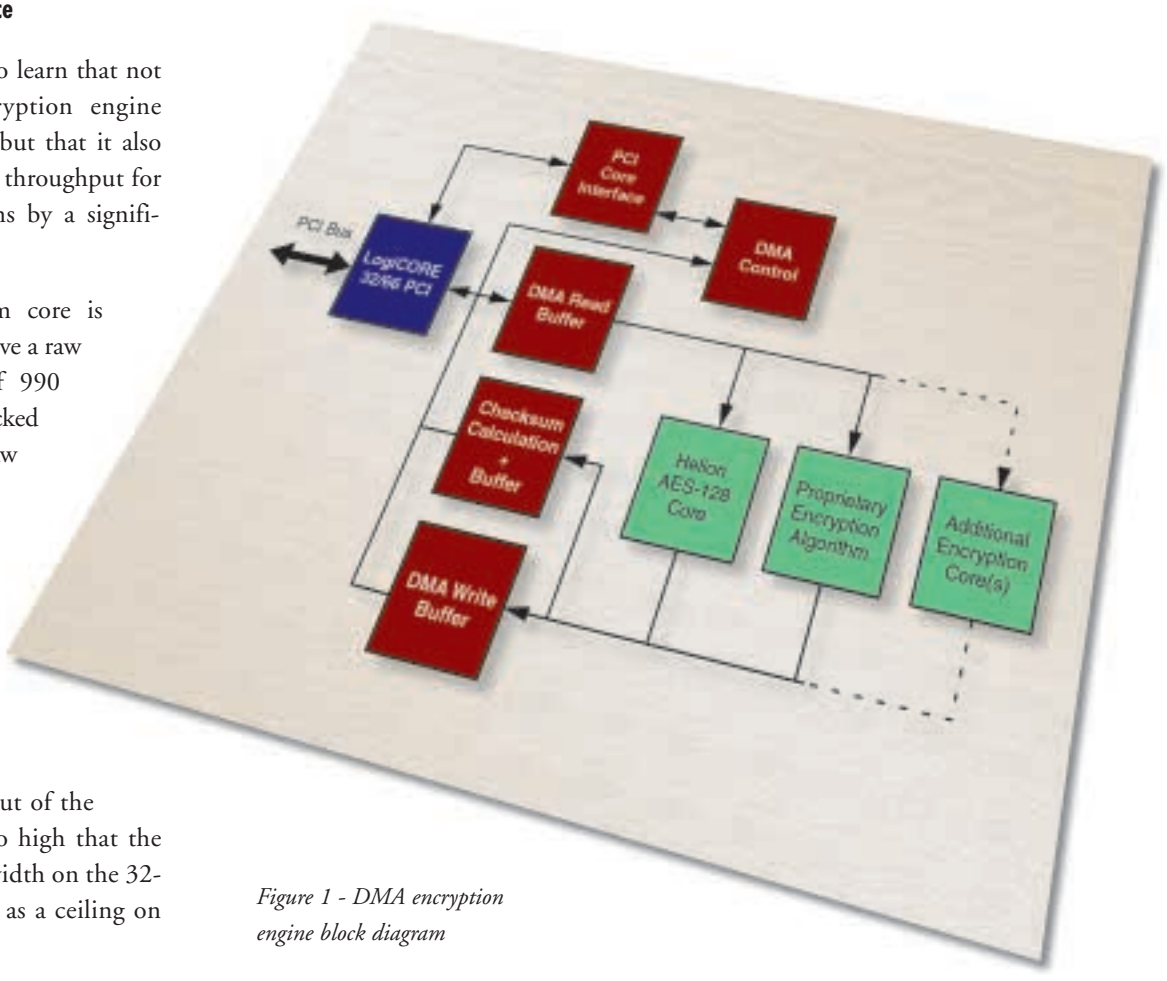

Figure 1 - DMA encryption engine block diagram

CoolRunner-II CPLDs Go All Digital

Xilinx CoolRunner-II RealDigital CPLDs consume less power and offer reprogrammable flexibility with unprecedented design security – without a price premium.

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

When Xilinx asked design engineers what they expected from the new generation of complex programmable logic devices (CPLDs), their responses showed they pretty much wanted it all:

- Advantages derived from state-of-the-art fabrication processes
- More industry standard I/O features
- System security
- Lower power operation
- Higher performance.

Designers required the next generation of CPLDs to increase performance of battery-powered devices. And, they stipulated that the next generation of CPLDs must improve performance without power penalties – and with the features and functions that would meet the expectations of today's discriminating designers and cost-conscious consumers.

In response to the high expectations of top design engineers, Xilinx has created the CoolRunner™-II RealDigital CPLD. The CoolRunner-II CPLD architecture introduces, for the first time, the advantage of both high performance and low power operation. The CoolRunner-II RealDigital CPLD reprogrammable logic solution just may be the defining element of a new era of CPLD-based applications – small form factor, high performance, power sensitive products that are feature rich and protected by multiple levels of security.

CoolRunner-II RealDigital CPLDs

Evolutions of the XPLA (extended programmable logic array) CPLD architecture have led to this ultra low power, ultra high speed, reprogrammable CMOS logic that offers significant improvements over earlier generations. (See sidebar "The Evolution of Low-Power CPLDs.") These capabilities are aiding in dramatic device size reductions and performance increases that are enhancing the look and feel of both handheld and high performance products.

The CoolRunner-II RealDigital CPLD delivers on several design fronts: low power, high performance, design flexibility, and security. Today, designers can choose high performance devices that do not require extra power for speed.

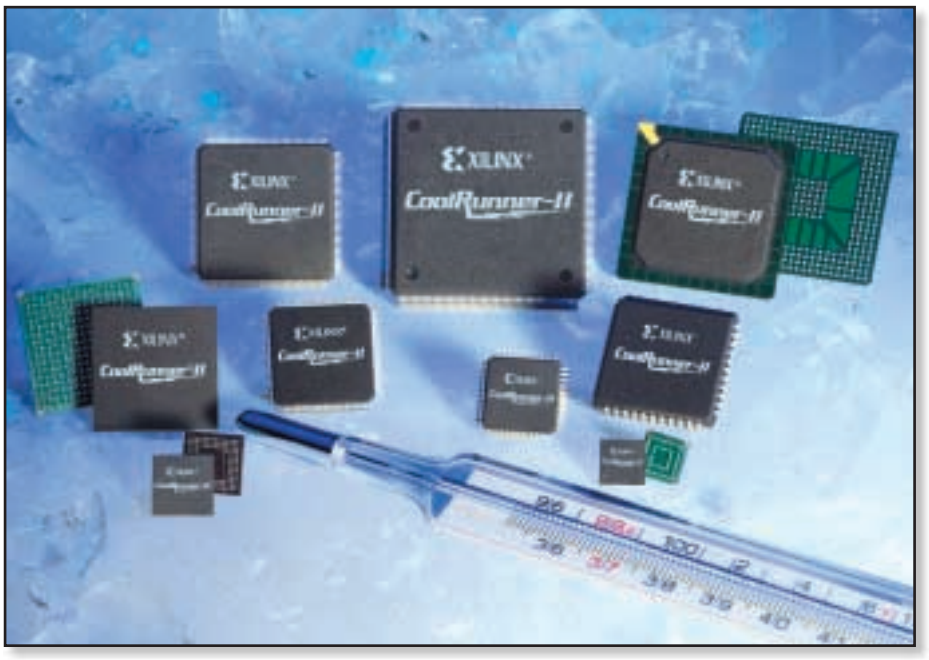

Figure 1 - Xilinx CoolRunner-II CPLD product family

Process Technology

The CoolRunner-II family of CPLDs has benefited from the leading edge of Xilinx FPGA (field programmable gate array) process technology. The original CoolRunner CPLD architecture began on a 0.5-micron process and then migrated to 0.35 micron.

Today, the CoolRunner-II CPLD family is manufactured using a 0.18-micron process technology – the same as the high-end Xilinx Virtex™-II FPGAs. With this ultra high density process technology, and scalable CMOS architecture, it is now possible to lower power consumption even further than before.

Multiple I/O Features

To facilitate its use in multiple system architectures, various input/output (I/O) standards have been incorporated in the CoolRunner-II interface. These include:

- LVTTL
- LVCMOS 33, 25, 18
- SSTL2-1, SSTL3-1
- HSTL-1
- 1.5V I/O.

High-speed transceiver logic (HSTL) I/O standards are common in high-speed memory interfaces. Stub-series terminated logic (SSTL) is associated with circuit designs where buses must be isolated from relatively large stubs. With these added I/O standards, CoolRunner-II CPLDs can easily adapt to a variety of interfaces that other low power processors might not be able to handle. This makes CoolRunner-II CPLDs a great system integration solution for microprocessor-based applications.

100% Digital Core

CoolRunner-II RealDigital CPLDs do not use traditional sense amplifier technology – and that means the historical tradeoff between speed and power has been eliminated. RealDigital technology has enabled Xilinx to offer the industry's lowest power consumption and highest performance on a single device with no price premium.

Low Power Enhancements

Learning that designers of portable products wanted even lower power than CoolRunner XPLA3 CPLDs, Xilinx added architectural features to accommodate tight power budgets. Even with FZP (Fast Zero Power™) technology and voltage reduction, we still had to add two new architectural features to lower overall power consumption in designs for power sensitive applications. Those features are DataGATE and CoolCLOCK.

DataGATE makes it possible to reduce power consumption by reducing unnecessary toggling of inputs when they are not in use. A prime example of this would be a bus interface. These input signals would become active only when information is passed to the CPLD. This eliminates the unnecessary switching from high to low/low to high. As bus interfaces grow, so will the power demand from needless switching of input signals.

CoolCLOCK, another CoolRunner-II low power enhancement, is implemented by dividing, then doubling, a clock signal. Using Global Clock 2 (GCK2) clock input and dividing by 2, then doubling the clock at the macrocell, performance can be maintained while lowering power consumption. The CoolCLOCK method is easily synthesized and is available on 128 macrocell and larger devices.

Schmitt Trigger

When communicating with CPLD design engineers, we discovered they were looking for an easy method to interface to noisy analog components. Thus, CoolRunner-II CPLDs incorporate Schmitt trigger inputs, which are helpful by delaying the switching time on slow rising or falling signals. Schmitt trigger inputs also help eliminate

false switching due to noise spikes. The possibility also exists to form a simple oscillator circuit from these inputs. You must exercise care to assure reliability, but it can save component costs on simple circuit designs.

Security

The loss of intellectual property is always a concern for electronic equipment manufacturers. With CoolRunner-II RealDigital CPLDs, security was enhanced to eliminate the possibility of code theft. The security bits are scattered and affect different modes of operation. If any tampering of the security chain is detected, the device automatically locks down and is erased. Even if the device is de-capped, buried interconnects make it almost impossible to trace security connections without destroying the device. Security details cannot be discussed even under non-disclosure agreements, thus ensuring ultimate protection.

User-Defined Grounds

As voltage levels decrease, noise and ground bounce become larger factors in signal integrity. Now, a few hundred millivolts is the difference between on and off, versus a few volts in legacy devices. We had to take extra care in design and layout to shorten signal paths and minimize potential interference.

Differential signaling (SSTL and HSTL) serves as a good method to reduce noise and increase signal integrity – with the provision that extra pins are available for differential I/O and reference voltage sources.

To simplify the elimination of noise and ground bounce, CoolRunner-II CPLDs also employ software-definable ground pins. These programmable ground pins are an ideal way to reduce noise and ground bounce effects. By using neighboring I/Os as ground pins, CoolRunner-II devices can tolerate a larger amount of potential signal noise.

Voltage Translator Clearinghouse

Due to the availability of various specialty electronic devices, it may be necessary to interface with multiple component voltage levels. With CoolRunner-II RealDigital CPLDs, I/O banking is an easy way to eliminate voltage level translators. "Banking" is the ability to separate function blocks within the CPLD so that these blocks can operate independently at different voltage levels.

For instance, one function block could be I/O compatible with 3.3V devices, and another function block could interface with 1.8V devices. This saves board space and eases voltage level design issues.

Conclusion

From process technology to a 100% digital core to added power reduction features, CoolRunner-II RealDigital devices mark the beginning of a new era for CPLD applications. Many system design concerns – including speed, power, I/O interfaces, clock management, and security – have been consolidated into one reprogrammable logic solution: CoolRunner-II RealDigital CPLDS from Xilinx.

For more information, please visit *www.xilinx.com > Products > CPLDs > Introducing CoolRunner-II RealDigital CPLDs*, or *www.xilinx.com/xlnx/xil_prodcat_product.jsp?title=coolrunner2_page*.

A Short History of Low Power CPLDs

The first reprogrammable logic devices have their roots in bipolar PROM (programmable read-only memory) technology, with added logic capacity and features. The first company to generate customer interest was Signetics, with their introduction of the 82S100 FPLA (field programmable logic array).

Monolithic Memories Inc. attempted to be a second source supplier to the 82S100, but failed due to a process technology mismatch. They instead developed the now famous PAL (programmable array logic). This development by MMI led to aggressive process technology shrinks and increased speeds as logic designers pushed for faster state machine operation.

Philips Semiconductors purchased Signetics, and in 1977, acquired a second source license for the PAL from MMI. Philips continued to work on different architectural enhancements and migrated to a BiCMOS (bipolar CMOS) process. When the first 22V10 PAL from AMD hit the market, it was a greater success than anyone anticipated. From this beginning, devices matured into what are now known as CPLDs. Xilinx entered the CPLD market in the early 1990s by acquiring Plus Logic and introduced the 7200 family of CPLDs.

CMOS CPLD products use power for speed improvements by partially turning on (biasing) transistors and by using sense amp technology. In the past, this was an easy way to promote devices and claim a speed advantage. The by-product, however, was heat.

In 1994, Philips Semiconductors made the move to CMOS CPLDs. With this decision, CPLDs broke away from their BiCMOS PROM roots and entered into the scalable process technology arena. These products reflected innovative approaches in circuit design with an awareness of power consumption by removing the old bipolar style sense amplifier from the circuitry.

Xilinx purchased the CoolRunner line of CPLDs from Philips Semiconductors in 1999 to penetrate the low power CPLD market segment. Today, ultra high density CMOS processes (0.18 micron) make possible reduced size and lower power consumption, which directly reduces heat dissipation.

Consumers' demands for smaller, faster, cheaper, better products are the driving forces in today's competitive markets – and the state-of-the-art CoolRunner-II RealDigital CPLDs meet these demands.

If Disaster Strikes, Xilinx Is Prepared

War, natural disasters, terrorism, economic ruin, and other catastrophes can't be prevented, but their effects can be attenuated. Xilinx has in place a Business Continuity Program designed to minimize the impact on your business.

by Alicia Tripp
Manager, Business Continuity Program
Xilinx, Inc.
alicia.tripp@xilinx.com

Our customers depend on us – and you expect us to provide reliable product and consistent service, even in the event of a disaster or interruption to our business. For these reasons and for the benefit of our stakeholders, we incorporated the Business Continuity Program (BCP) into the Xilinx culture in June 2000. The BCP is a corporation-wide initiative that is supported by our Board of Directors and throughout our employee population.

We began the program by conducting a business impact analysis (BIA) and a threat and risk assessment (TRA) to determine the critical business processes within the company. In addition, the BIA and TRA provided an understanding of the impact of the loss of one or more of the critical functions. Once the analyses were completed, more than 80 departmental recovery plans were developed. The recovery plans are the road maps that will allow us to continue providing products and services to our customers during a major disruption.

Incident management teams (IMTs) are one of the most vital parts of the Business Continuity Program. We have located IMTs strategically throughout the Xilinx enterprise. Each team consists of members of functional business units within the company such as Human Resources, Legal, Purchasing, Security, Facilities, Risk Management, Accounting, and Safety. The IMTs are responsible for handling the event, from the time of a disaster through recovery and restoration activities. The purpose of the IMTs is to protect the health and safety of our employees, to guide the company through a crisis, and to provide a structured organization for overseeing the response. These teams will convene during a major business interruption to assess the situation and determine if a disaster should be declared and whether or not to launch departmental recovery plans.

Incident management teams convened four times in less than 48 hours after the September 11, 2001, terrorist attacks to make plans to assist our employees. We were able to have counselors on site, and travel representatives helped employees traveling on business to schedule return flights home.

The Business Continuity Program has been successfully integrated into our business culture – in large part, because of the top-level support from executive management. From the very beginning, our executive staff has been onboard to make sure the program goals and objectives are achieved.

The BCP includes an executive awareness and ongoing plan development that helps our executive officers to continue to be involved and ready to respond. This executive awareness plan will allow Xilinx to properly manage any incident from the time a disaster is declared through recovery and restoration activities.

Recovery plans and IMTs are exercised on a regular basis. By developing and actively maintaining the Business Continuity Program, we are prepared to limit the effect a long-term business interruption might have on Xilinx – and to mitigate the impact on our customers.

Please contact your local sales representative for additional information.

Even as Designs Grow More Complex, Xilinx Software Improves Your Productivity

Xilinx Integrated Software Environment (ISE) design tools promise to manage design complexity and, at the same time, ensure your logic design flow is easy to use.

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

Your Xilinx programmable logic device offers you more capabilities than ever before, and coming releases of the ISE design tools will help you to handle the pressure that comes from using these complex new features. Come along for a glimpse into the future of ISE.

ISE Delivers a Wealth of Device Features

The most compelling pressure on the logic design flow comes from the added complexity of new capabilities and tool features. Embedded systems, increased clock speeds, verification of high-density designs, new device capabilities, high-speed I/O, a variety of IP (intellectual property) sources, and design reuse – all of these logic trends force more requirements into the design flow.

In 2001, the Virtex™-II FPGA introduced a revolution in clock capabilities with Digital Clock Manager (DCM), and this year, the Virtex-II Pro™ Platform FPGA broke new ground by delivering 3.125-gigabit serial I/O transceivers and embedded IBM PowerPC™ microprocessors. As the demand for more device features continues to grow, engineers must learn all the programming attributes for these new features – and the learning curve escalates with the introduction of each new device feature.

To ease the learning curve, Xilinx ISE 5.1i will offer designers even more interactive architecture and design assistance. In addition to quick and easy dialogs, wizards will automatically step you through configuring advanced device features and inserting those configurations directly into your HDL source code. Upcoming releases of ISE will soften your learning curve, helping to speed design completion.

ISE Enables IP Capture and Design Reuse

As design sizes grow, source management and methods to improve source code productivity must keep pace. In a multi-million-gate design, source code can come from multiple resources, including purchased IP, developed code, and "design reuse" – modules of HDL developed and proved in a prior design.

In upcoming releases of ISE, it will be possible to capture proven IP modules at the floorplan level. This captured IP will not be just a module of code, but it will also have attached relative placement and floorplanning area information that will help speed implementation in later uses.

As time-to-market pressures increase, design change impacts late in the design cycle will become an even larger design challenge. Enhancements to ISE will confine late-cycle design changes to only that portion of the design that is required to change. The remainder of the design will be left intact, thus speeding overall design completion. This technology is enabled by new floorplanning capabilities along HDL hierarchy boundaries that make it easier to define areas of logic.

High-Speed Design Will Push Innovation

High-speed design pressures will continue to push more innovation into logic design and into the board design flow.

Timing analysis and timing constraints have – until now – lived under two different analysis domains, separated at the FPGA pin. The logic designer has looked to the specification and timing requirements to define design closure, while the board designer has picked up the signal at the logic pin and attempted to lay out and analyze the effects that the PCB trace will have on that signal.

Now, however, the line between logic designer and board designer functions will begin to blur as timing constraint languages become more uniform across the logic tools and among the logic-level and board-level tools. Trace analysis packages also will increasingly use more complex pin models provided by the logic tools, but these complex analyses will also become much faster, dumping slower general analog simulation in favor of specialized transmission line trace simulations.

A Higher Level of Abstraction Is Coming

While high-level languages (HLLs) have been explored in logic and systems design, the pain of existing logic design methodologies hasn't been great enough to force a high-level language or associated methodology into common use.

Verilog and VHDL have served the logic design space well for the last two decades, but now logic design sizes are routinely exceeding one million gates. These million-gate-plus designs require large and cumbersome HDL source code and long periods of debugging time. Emerging embedded processor methodologies are also complicating the design process. Moreover, the need to serve two distinct language-driven design flows (hardware and software) is now making HLL design support imperative.

Xilinx is investing time and technology in the industry-wide HLL efforts through partner technologies such as Synopsys® SystemC™ and Celoxica Handel-C – as well as through Forge Compiler, which Xilinx acquired from LavaLogic™ two years ago. The results are leading to optimized performance for FPGA devices, and to better integration of the co-design and co-verification phases of the embedded systems design flow.

Longer-term technology is also being defined to help you "architect" the potential design at the HLL source code level, analyze what modules should be implemented in the logic design flow, and determine what modules are required in the processor design flow to reach optimum system performance.

Conclusion

New device capabilities will continue to push design complexity. The good news is that Xilinx design tools are keeping pace. New releases of ISE software will help you keep your design time down and your time to market fast. For more information on Xilinx ISE logic software, go to *www.xilinx.com/ise* – and watch for the release of ISE 5.1i coming in late summer.

Exploring Hardware/Software Co-Design with Virtex-II Pro FPGAs

To explore the possibilities of hardware/software co-design, the Xilinx Design Services team created a reference design for a real-time operating system (RTOS) for telecommunications – here is an interview with the team.

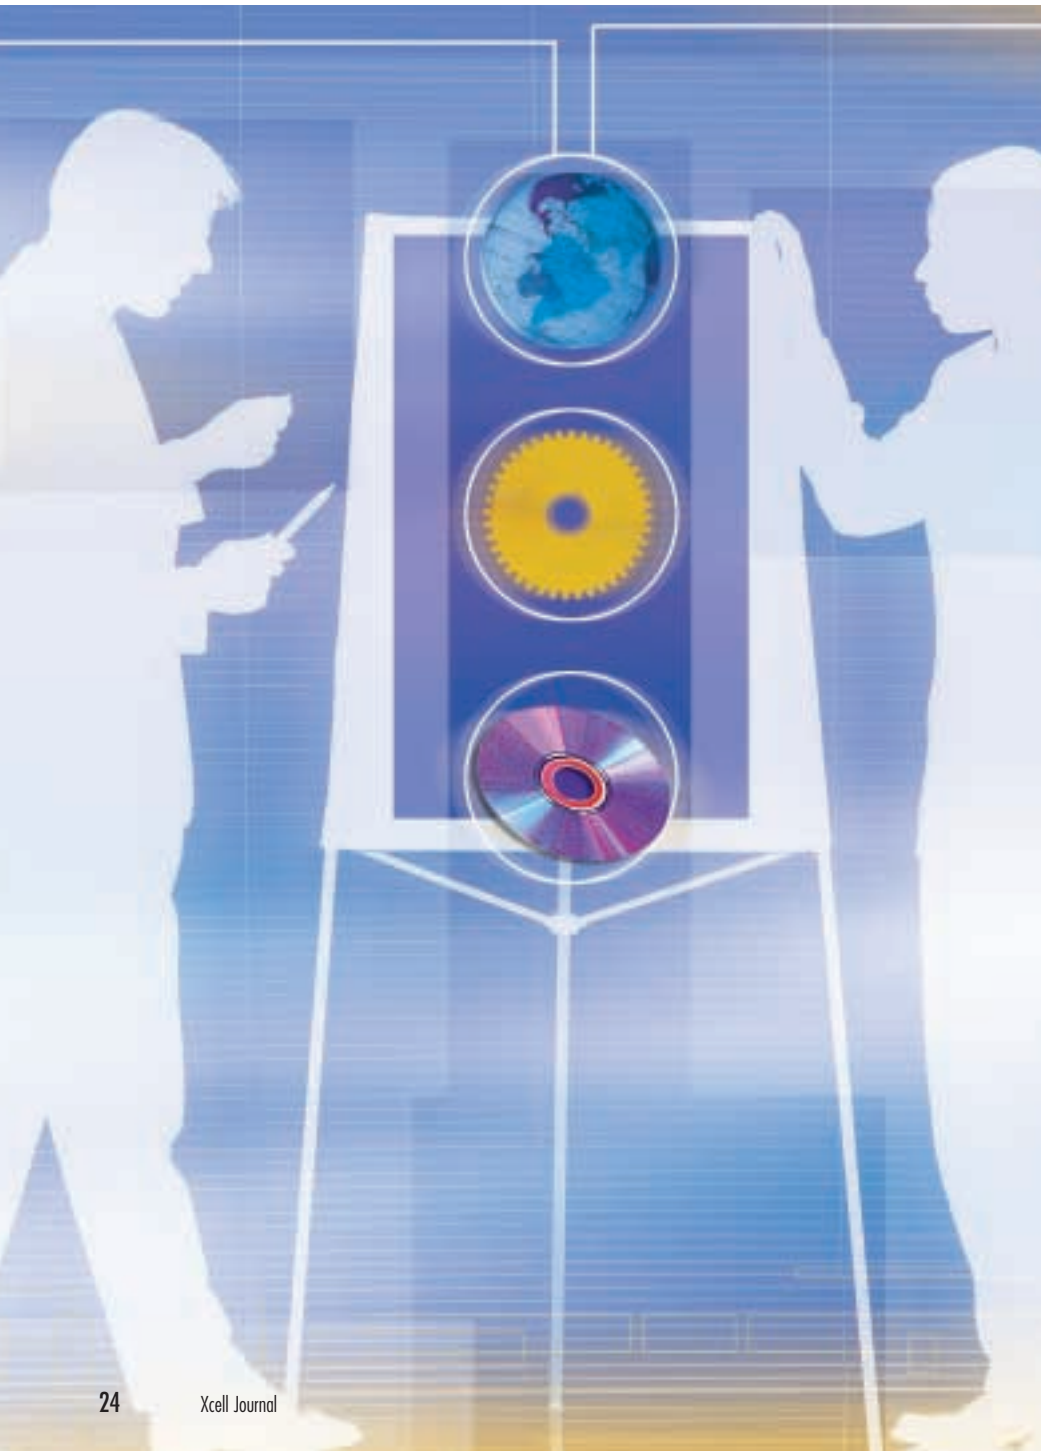

by Jean-Louis Brelet, Manager
APD Technical Market Research
Xilinx, Inc.
jean-louis.brelet@xilinx.com

In this interview, Jean-Louis Brelet, manager of Technical Market Research for the Xilinx Advanced Products Division, interviewed Xilinx Design Services (XDS) software and hardware engineers about their practical experience with an RTOS reference design project. Senior Project Engineer Matthew Roche acted as spokesman for XDS.

What is hardware/software co-design?

Hardware/software co-design generally means that the hardware and software are developed at the same time – some software functions may be implemented in hardware for additional speed, and some hardware functions may be implemented in software to free up logic resources.

What does XDS do?

Xilinx Design Services, located in Dublin, Ireland, is part of the Xilinx Global Services Division (GSD). XDS has two functional groups: software engineers who provide custom software solutions, and hardware engineers who provide custom hardware solutions. The combined team consists of more than 50 engineers.

Typically, we help customers with their full product life cycle, from requirements capture, through design implementation and test, to final acceptance and delivery. We work very closely with our customers, and form a team of experts who best match the job requirements.

What are your previous experiences with embedded software?

We have worked on technologies that include:

- Digital video systems – both set-top boxes and head-ends
- High speed digital network communications systems
- Automotive applications
- Real time operating systems
- Board-level designs.

Due to the nature of design services, we have attained experience from many different development environments and methodologies. We have experience with a wide range of designs, for device-level to system-level applications.

How would you describe the Virtex-II Pro FPGA family from a software perspective?

The introduction of PowerPC™ processors into the Virtex-II Pro™ Platfrom FPGA architecture creates many new possibilities. Now we can design custom co-processors with innovative (and fast) interfaces. Due to the flexibility of the embedded IBM processors, and their surrounding programmable logic, we can offload some parts of our software design into a dedicated hardware implementation that can run much faster.

Previous hardware/software designs have been based around an inflexible hardware/software interface.

Jean-Louis Brelet

When did you start to design for the next-generation Virtex-II Pro Platform FPGA?

As soon as the technology became available, it became apparent that we must evaluate the system in a real project situation – the RTOS reference design. This way we could become aware of the new possibilities offered by the Virtex-II Pro device – and be best placed to contribute our knowledge of the advanced architectures used by our customers.

Project Objectives

What are the main objectives of this project?

The overall and real objective was to gain experience in developing a product in a tightly coupled hardware/software environment. The ability to shift functions from software to hardware, or vice versa, in a very short time frame, creates an environment where the design teams need a much greater knowledge of each other's work.

The smaller and more restricted objective was to identify the speed bottlenecks in the RTOS and remove the bottlenecks through the use of dedicated hardware – effectively implementing large sections of the kernel in hardware.

Before Virtex-II Pro FPGAs, how would you develop a system like this?

The RTOS project is very suited to the Virtex-II-Pro architecture. Given the costs, project time, cycle time, and risk required to develop this project with other technologies – for example, an ASIC – it is unlikely that it would even be undertaken. An ASIC design doesn't offer you the benefit of reprogrammability and flexibility to change/improve or upgrade the system in the field.

What is the expertise required to successfully implement these systems?

Software people must understand the nature of hardware designs and the type of problems encountered by the hardware team. They also must understand the possibilities

Mathew Roche

and capabilities of the hardware.

Likewise, the hardware team must have a good understanding of software and how the application operates. Both teams must have a common language – or a good understanding of each other's language and a willingness to adapt.

Partitioning

What is your approach at the system level to partition embedded hardware and software functions?

We designed and implemented the RTOS reference design project for telecommunications with resource-constrained processors. We built it for state-based SDL (specification and design language) applications.

The application for this RTOS was to be in systems with high numbers of processes with large numbers of context switches. At this level, the overhead of the operating system would be significant and thus, must be minimized.

First, we analyzed the system and identified the time-hungry elements. Our initial performance analysis indicated that processing of the state tables and managing the most heavily used queues (signal queues) were the processor-intensive tasks. The first iterations of the design placed these tasks into hardware with a clearly defined hardware/software interface. This did not completely remove the bottlenecks, but it did have a tendency to move them from the processing into the interface.

Further iterations of the design required much closer integration of the hardware and software teams. This allowed the design of coherent and integrated tasks, which could be completely dealt with in hardware, and required minimal use of the interface.

Having the processor embedded into the FPGA fabric allows a variety of interface techniques, which do not necessarily require the use of the IBM CoreConnect™ proces-

sor local bus. This feature became an important design issue, as the speed of the system was capable of hogging the bus.

In our experience, we have found that there is no distinct division between functions that are suitable for hardware or software. It becomes a trade-off of how fast you want to do something and how many configurable logic blocks you want to spend on doing it. Greater complexity of the function or required speed of operation will bias the design towards software or hardware.

How would you summarize your method for partitioning the system?

We simply integrated the design teams, to avoid a purely hardware- or software-led solution.

How have you split the teamwork?

Most of the software was designed first, because we needed to understand how the system would operate before transferring parts of it into hardware. This gave us the advantage of having functioning prototypes before a full hardware implementation – significantly reducing our theoretical time to market.

At the system design level, the division of labor became less important as the project went on. It was necessary for both sides to have a good understanding of what the other team was doing. Design decisions had to be made with knowledge of both hardware and software. This had some interesting implications for the language and terminology used by the teams.

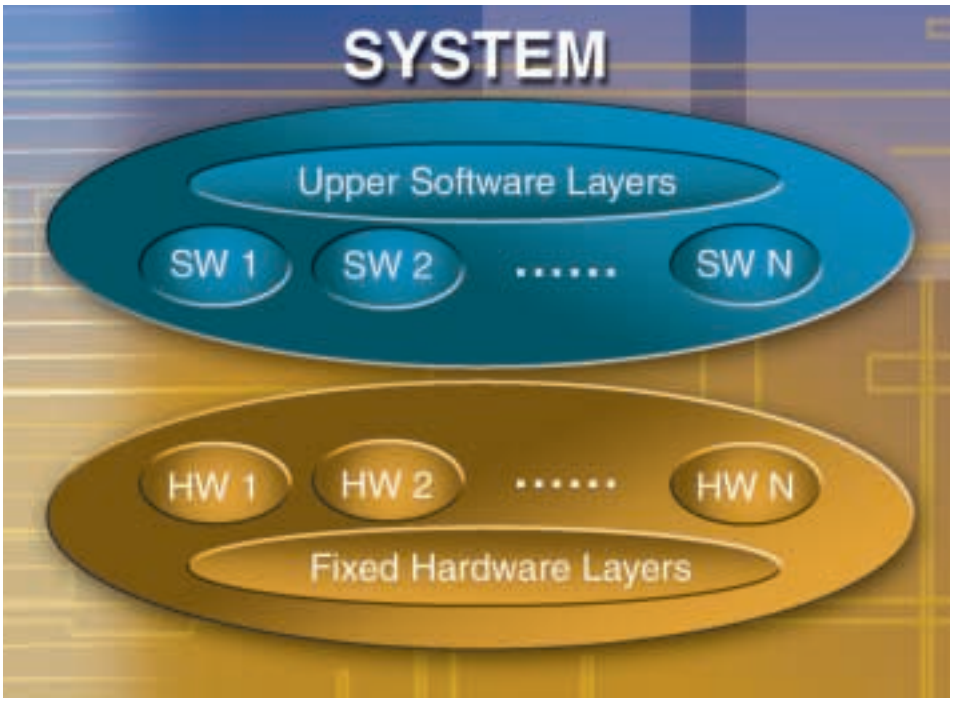

Virtex-II Pro Embedded Software

Is it easier to create embedded software for the Virtex-II Pro fabric?

In the first applications, embedded software development techniques will not necessarily be different for Virtex-II Pro devices – they can be viewed as conventional systems with performance improvements and reduced chip count. However, as experience builds, it will become apparent that traditional design methods do not use the full potential of the device. System partitioning will become part of the architectural design process, and the boundary between hardware and software will become less rigid.

The RTOS project gave us an opportunity to examine the project with system co-design as the goal. Virtex-II-Pro Platform FPGAs allow significantly more integrated designs with all the performance benefits that entails.

What is the most innovative feature of Virtex-II Pro for a software engineer?

Using multiple processors (both IBM PowerPC 405 hard core and Xilinx MicroBlaze™ soft core processors), on a single chip, allow you to dedicate processes to particular processors for critical tasks. You can easily implement interfaces between the different processors, using hardware-based FIFOs, reducing the communications bus overhead.

How easy is it to partition a Virtex-II Pro design?

Partitions are very flexible in the Virtex architecture. Changes were being made to the interface all the way through the design and implementation of our RTOS system. This allowed us to reduce our design time and minimize design roadblocks.

Conclusion

Could you summarize your experience?

Close cooperation between hardware and software teams is critical.

Is this project an example of making the impossible possible?

No. This project does not do anything that has been impossible before. What it does do is make it faster, cheaper, and with reduced risk.

How can Xilinx customers benefit from your expertise?

It is part of our culture to maintain high levels of expertise in advanced and state-of-the-art technological systems and methodologies. As project managers, we are constantly improving our service and quality in dealing with our customers' projects. Increasing competitiveness in the marketplace means that there is a higher emphasis on time to market and quality. XDS continues to improve its skills in order to help our customers achieve these challenging targets.

XDS can be contacted by e-mail at *designservices@xilinx.com* or by visiting *http://support.xilinx.com/xds/*.

Virtex-II Pro Platform FPGA Drives Convergence of Programmable Logic and Embedded Systems

The breakthrough technologies embedded and immersed within the Virtex-II Pro logic fabric deliver the integrated functionality demanded by top digital design engineers.

XPERTS
Xilinx Program for Engineering Resources from Third Parties

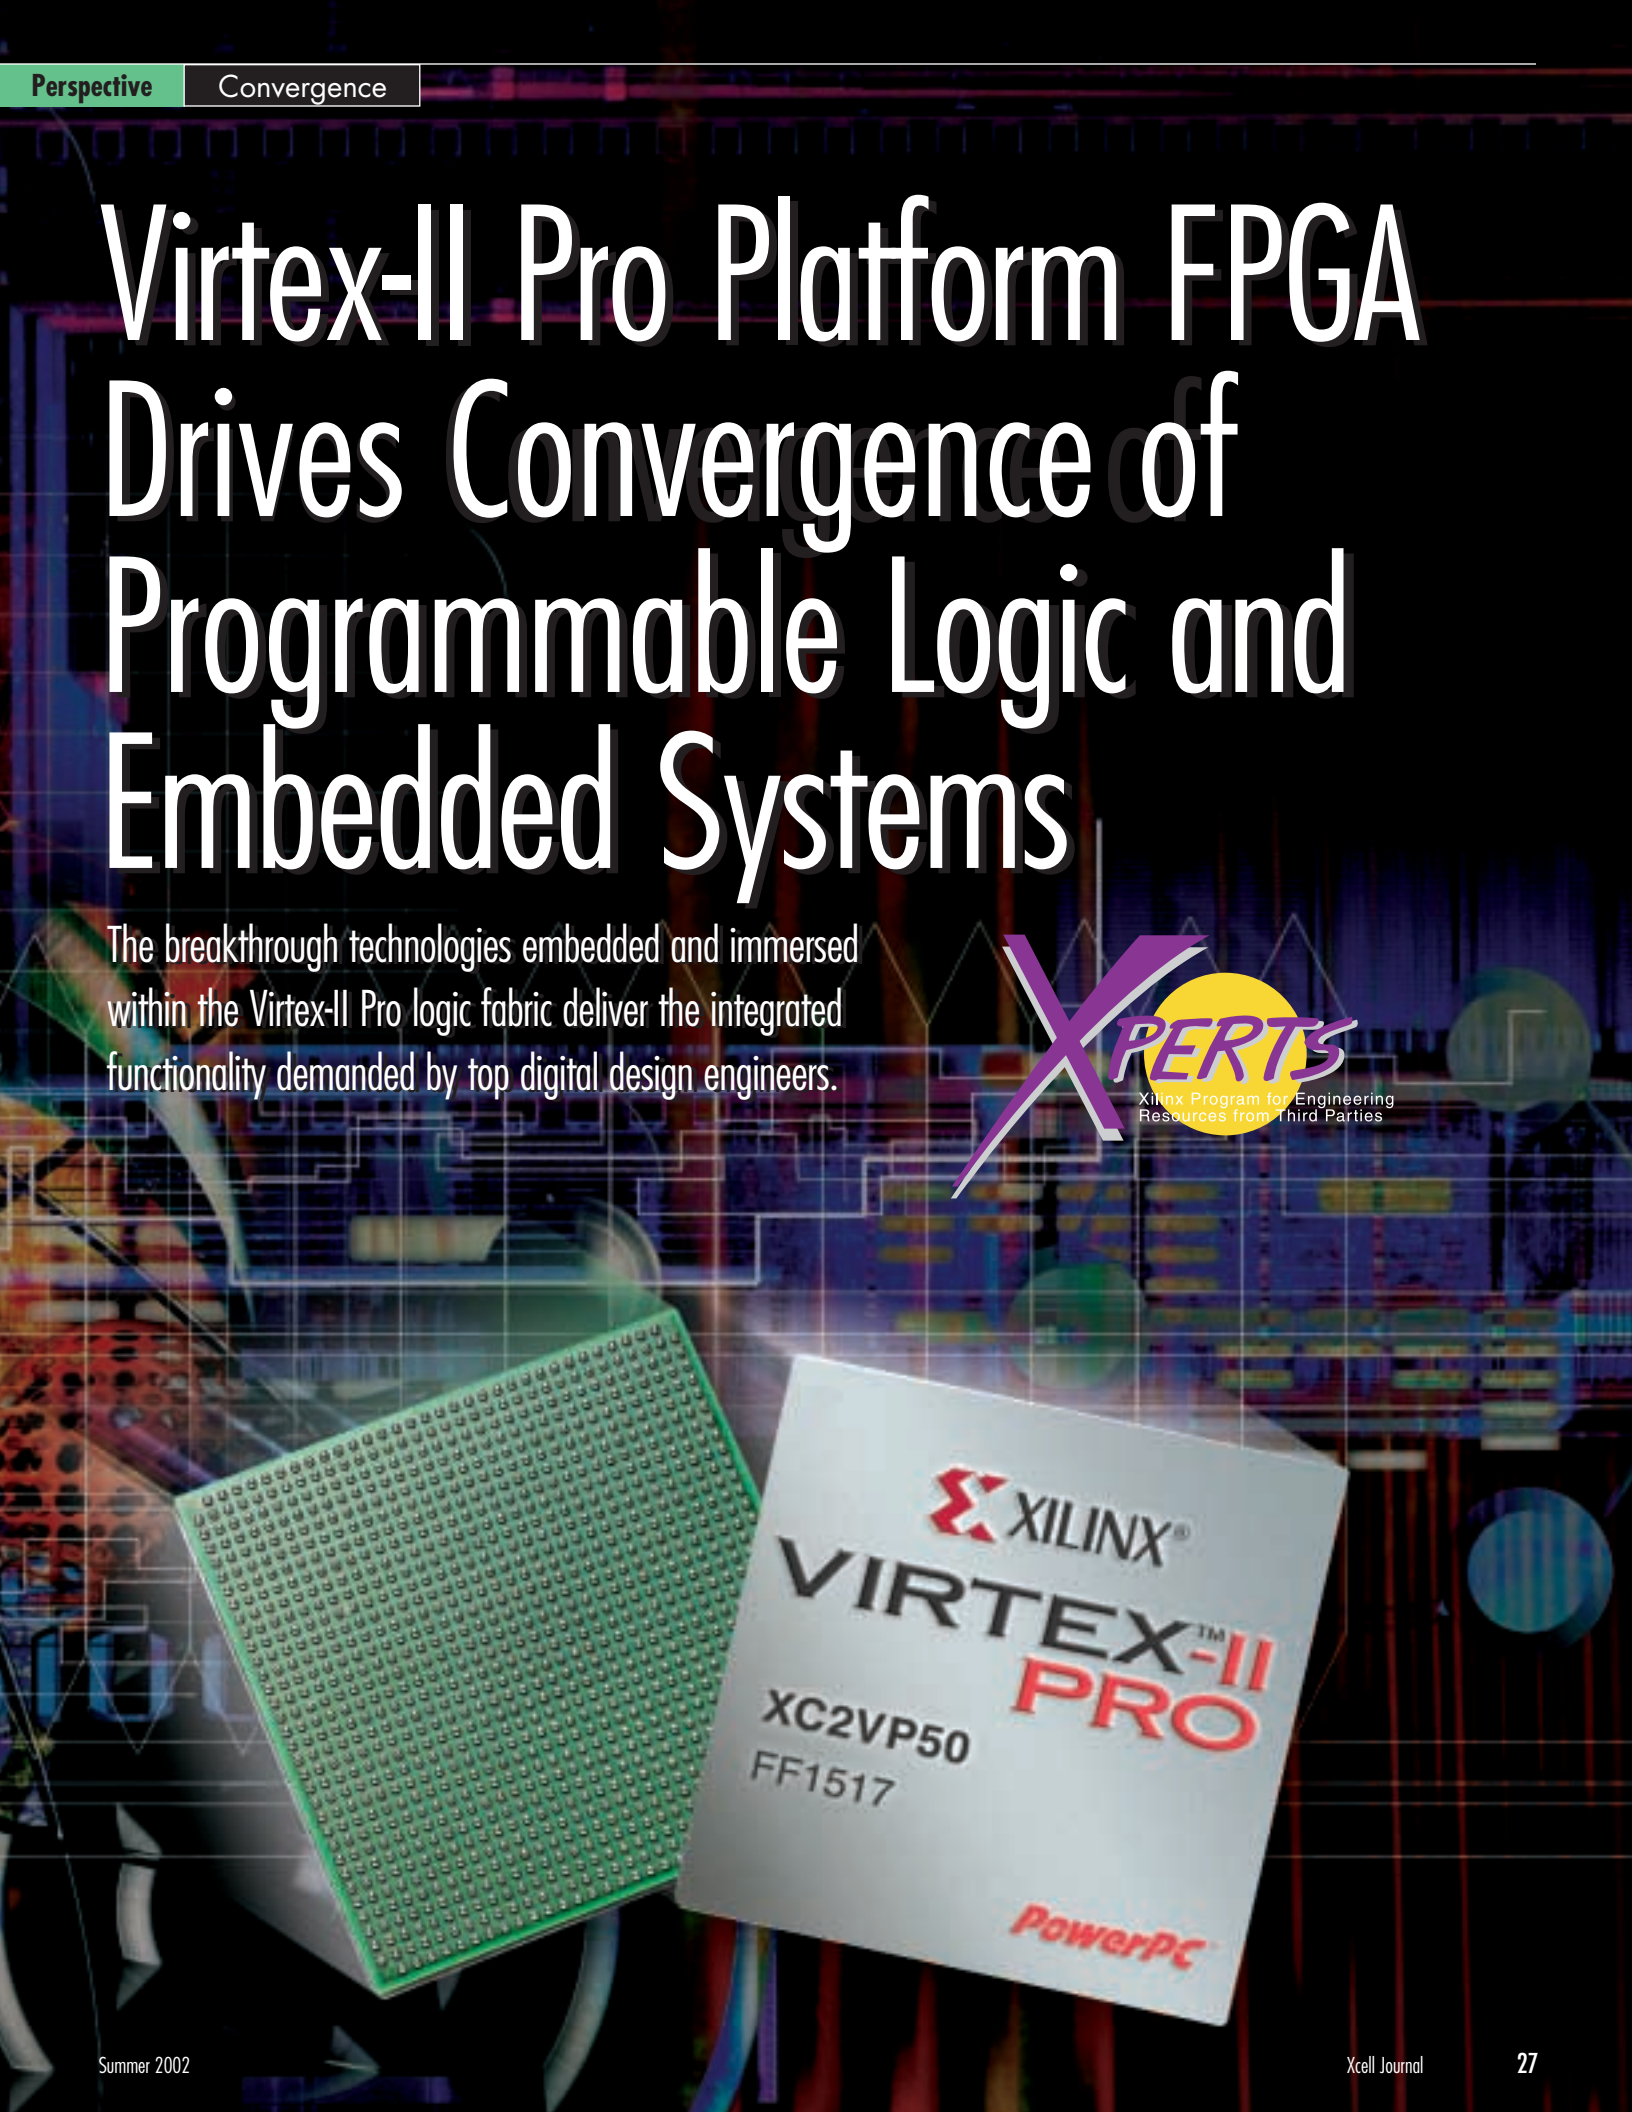

by Michael Worry
CEO, Nuvation Engineering
michael.worry@nuvation.com

Time to market pressures have ruled out designing an ASIC from scratch, but without a custom ASIC, the discrete FPGA, microprocessor, and off-board high-speed serial I/Os will never fit last year's form factor. You just can't fit 16 channels this year on a board designed for four channels last year.

As a Xilinx XPERTS partner and an engineering services company working in the cross disciplines of FPGAs, printed circuit boards, and firmware, we at Nuvation Engineering often encounter these problems of form, fit, and function. We need an onboard industry standard microprocessor, integrated high speed I/Os, off-the-shelf IP blocks, and double last year's gate count.

Xilinx has delivered on all this and more with their introduction of the revolutionary Virtex-II Pro™ Platform FPGA. It starts with the technologically advanced Virtex™-II infrastructure, supporting up to 8 million gates and I/O serial speeds of 3.125 Gbps. Nestled into this digital workshop is the IBM PowerPC™ 405 microprocessor – the number one embedded CPU for networking and telecom infrastructure. The continued convergence of programmable logic with embedded systems has created a new digital workbench with a wealth of integrated platform standards and rapid co-development tools. The Virtex-II Pro Platform FPGA is in a class by itself in the exploding programmable embedded market.

FPGAs and Microprocessors: A Symbiotic Relationship

Fundamentally, microprocessors and FPGAs fulfill different tasks (see Table 1). Programmable logic is needed for wire speed, high bandwidth, custom digital processing, and hardware acceleration. However, that speed and customization comes at a price, both in parts and non-recurring engineering (NRE) costs.

For monitoring and control tasks, it is far cheaper and easier to use a microprocessor. Furthermore, microprocessors have existed longer and have a larger base of engineering familiarity. Thus, many product architectures naturally marry an FPGA with a microprocessor. At Nuvation, our complex designs often start with an FPGA and microprocessor architecture, then we determine which system functions are best implemented in each. Xilinx takes this to the next level by integrating these functions into a single chip.

FPGA	Microprocessor
• Configurable and programmable digital logic	• Fixed processing units accessed by predefined commands
• Instantiate multiple processing units to run tasks in parallel	• Greater knowledge, tools, installation base and engineering familiarity
• Hardware accelerate high speed tasks	• Faster development time
• Scalable	• Lower parts costs
• Coded in HDL	• Coded in firmware such as C

Table 1 - Hardware accelerates high speed tasks

A more traditional design would be a microprocessor integrated with dedicated hardwired logic in an ASIC. However, time to market pressures are demanding platforms with full product cycles in months and integration cycles in hours. Compared to the time for ASIC development and potential respins, FPGAs present a highly appealing alternative.

Because the FPGA is reprogrammable after release, using technologies such as the Blue Iguana™ System, a whole new market of revenue streams and business models is opened. Xilinx has risen to these market pressures by systematically driving down the cost of customized hardware acceleration while producing continued innovation in size, speed, and integrated features.

Virtex-II Pro Features

The Virtex-II Pro Platform FPGA brings together market-leading Xilinx programmable logic with the industry standard PowerPC (PPC) microprocessors and high-speed I/Os. Truly this is the digital sandbox for the 21st century engineer.

Integrate Off-the-Shelf IP

The Virtex-II Pro device offers as many as 8 million system gates. This allows easy and rapid SoC (system-on-chip) integration of existing IP cores without having to painstakingly code a custom implementation of a standard interface to save a few gates.

The Virtex-II Pro FPGA allows you to choose up to four PPC built-in cores. No longer do you have to meticulously calculate how many clock cycles each firmware operation will take to determine the processor infrastructure. If the first PPC becomes overloaded, just change the part number and add another PPC to the Virtex-II Pro FPGA. You don't have to respin the board or, worse yet, you don't have to deal with an increase in the form factor rippling to mechanical changes.

The off-the-shelf features of Virtex-II Pro devices are matched with on-chip memory of up to 4.5 MB. With this onboard memory, you can get high-speed, on-chip bus connects without being bottlenecked by a memory arbiter to off-chip memory.

Flexible Clocking Infrastructure

Any first year student learns the first rule of digital synchronous logic: never gate the clock. Well, Xilinx has packed that up with the vacuum tubes and created an architecture that allows you to switch – glitch-free – among as many as 16 global clocks. This is a huge benefit for sharing system resources among different clock domains. The Xilinx Digital Clock Manager also provides on-chip frequency and phase control, meaning that an off-chip, phase-locked loop is no longer needed.

XCITE On-Chip Impedance Control

After drawing a thousand-pin FPGA, the last thing a design engineer wants to deal with is adding a few hundred termination resistors. Xilinx Controlled Impedance TEchnology (XCITE) eliminates those resistors with dynamic, on-chip, digitally controlled impedance. The savings in board space, schematic time, and layout time are significant.

Integrated High Speed I/O

It seemed not too long ago that it was cutting-edge to support 64-bit, 66 MHz PCIs. Now we are up to PCI-X speeds at 133 MHz. We have support for RapidIO™, POS-PHY4 and SPI-4 interconnect technologies, to name a few. Pick almost any kind of memory you like – ZBT™ SRAM, DDR SDRAM, or QDR SDRAM, for instance – and they can be implemented in the Virtex-II Pro logic fabric. Then we jump into serial I/O space, and the Virtex-II Pro platform can accommodate as many as 16 Rocket I/O™ multi-gigabit transceivers at a whopping 3.125 Gbps. This means applications like OC-48 can integrate SERDES (serializer/deserializers) right into the FPGA – once again, saving board space, cost, and valuable time.

Integrating Customer Hard IP

Hard IP is when digital logic is implemented in hard coded gates, removing the programming capability, but also using much less silicon. Once a design or IP block is stable, it can be converted to hard IP. The PPC 405 microprocessor is currently implemented in the Virtex-II Pro device as hard IP, but Xilinx actually allows customers of sufficient volume to embed their own hard IP. For example, if you wanted to roll out an existing ASIC to provide additional capacity and integrate the microprocessor, the Virtex-II Pro could embed that ASIC as hard IP and provide programmable gates for new functionality.

Integrated Bitstream Security

The Virtex-II Pro device has on-chip storage capacity for a 3DES key, either with a battery or constant power supply. The stored or delivered bitstream is encrypted, and can only be unlocked with the on-chip key. The on-chip key allows Virtex-II Pro devices to be designed into military applications that traditionally have required OTP (one time programmable) solutions to remove any chance of code becoming compromised.

Reconfigurability

The Virtex-II Pro Platform FPGA can be made even more powerful by enabling remote configuration capabilities. Xilinx IRL™ (Internet reconfigurable logic) technology includes a partnership with Wind River Systems Inc. to produce the PAVE Framework – PLD API for VxWorks™ Embedded System Integration Framework. The framework enables Xilinx FPGAs using Wind River's VxWorks real-time operating systems and Tornado™ integrated development environments to be reconfigured during development and after deployment in the field.

Additionally, Xilinx is in partnership with Blue Iguana Networks Inc., a semiconductor company that provides secure, remote, hardware management. Their patented technology enables secure software and hardware upgrades to in-service equipment over the Internet and without downtime. The Virtex-II Pro Platform FPGA is an excellent match for Blue Iguana's technology, as the embedded microprocessor could upgrade the FPGA, or vice versa.

Conclusion

Because microprocessor and FPGA functionalities differ, as applications become more complex and diverse, it is often insufficient to utilize just one or the other. The development of the Xilinx Virtex-II Pro Platform FPGA resulted from the natural integration of FPGAs and microprocessors. In the past, the only solution would have been to utilize multiple, discrete ASICs. However, because of the time to market requirements – in addition to the potential expenses associated with an error – reconfigurable processing has become an ideal alternative.

With digital design engineers facing the pressures of increased flexibility, shorter time to market, and higher clock speeds, they will continue to drive the demand for integrated standard platforms. The Virtex-II Pro Platform FPGA does a masterful job of combining a highly technologically advanced FPGA architecture with the industry standard IBM PowerPC microprocessor. As a member of the Xilinx Early Adopters Program, as well as a Xilinx XPERTS partner, we at Nuvation look forward to implementing this breakthrough technology for the benefits of our clients. For more information, please visit us at *www.nuvation.com*.

Xilinx Offers End-to-End Solutions

For anyone who has been following the history of programmable logic, it's pretty clear that the industry has seen more changes in the course of the past year than it has perhaps in its entire 19 years of its existence.

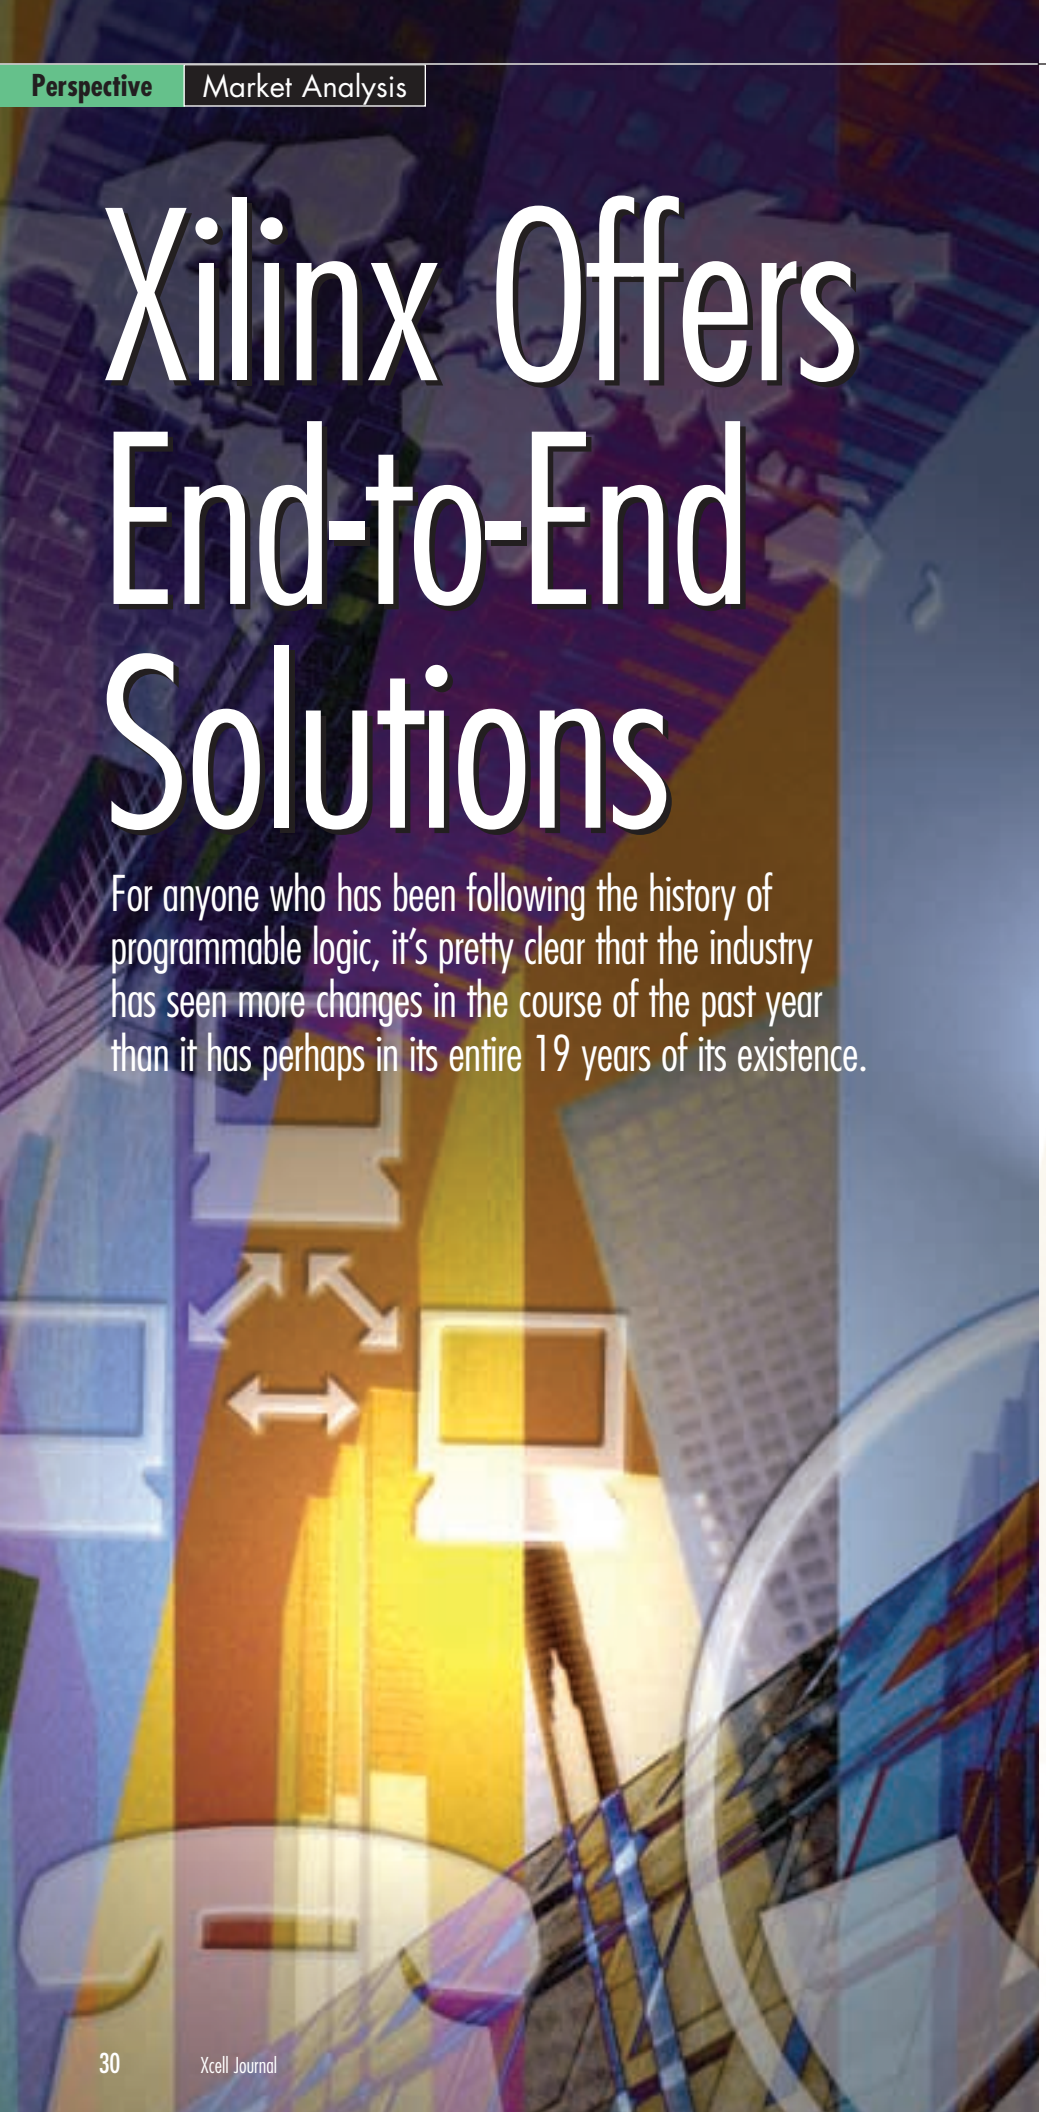

Robert Bielby, Senior Director
Strategic Solutions Marketing
Xilinx, Inc.
robert.bielby@xilinx.com

With the recently introduced Virtex-II Pro™ product family, the much awaited and highly debated immersion of a microprocessor into a programmable logic fabric has finally occurred. With up to four IBM® PowerPC™ 405 microprocessors in the highest density family members, Virtex-II Pro Platform FPGAs deliver more than 2,000 Dhrystone MIPS, a number that is unchallenged in the industry.

Complemented by the integration of multi-gigabit serial I/O signaling technology capable of supporting more than 3 Gbps per I/O pin pair, the Virtex-II Pro family delivers the industry's first truly high-performance Platform FPGAs. These advanced programmable logic devices are unparalleled in their capabilities not only to address a wide range of the most demanding applications – but also to adapt and adopt emerging I/O standards and protocols as they become available.

Advances in semiconductor and programmable logic fabrication have not only made families like the high-end Virtex-II Pro Platform FPGAs possible, but on the other end of the spectrum, technology advances have also enabled the recent delivery of Spartan™-IIE FPGAs and CoolRunner™-II RealDigital CPLDs. The Spartan-IIE family sets a new standard for low cost within the programmable logic industry, and the CoolRunner-II series sets an equivalent benchmark in the industry for ultra low power consumption.

This premiere portfolio of programmable logic solutions gives Xilinx the power to address a wide range of new markets and applications, including consumer electronics and handheld devices. In this issue of Xcell Journal, we discuss some of the markets we've opened and applications we've enabled.

Metropolitan Area Networks

The hottest areas in networking today are metropolitan area networks (MANs) and "edge access" markets. Driven by the explosive growth of data traffic in metropolitan regions, this networking market segment is expanding at a 34% CAGR (compound annual growth rate).

New products and technologies based upon standards proposals that leverage the combined benefits of Ethernet and SONET are now being rushed to market to address this rapidly growing demand for data services. However, the uncertainty surrounding the overall acceptance of these new standards – and the rigorous requirements of supporting line rates in excess of 10 Gbps – have given rise to the need for solutions like Virtex-II Pro Platform FPGAs that offer both flexibility and high performance.

In addition to articles discussing the challenges of developing solutions for the MAN, we are also pleased to present – in conjunction with Avnet Design Services, Avnet Cilicon, and Reed Electronics Group – an industry-wide event: the Metro-Optical Networking Forum to be held July 25 at the Santa Clara Convention Center in California. This free one-day event (which will be simulcast to more than 40 North American venues) is focused entirely on the latest developments in metropolitan and edge access networking. For more information, see the article and ad in this issue of Xcell Journal. To register for this event, visit *www.xilinx.com/metro*.

Consumer Electronics

Consumer electronics continue to grow at a very healthy pace, fueled by the trend towards digitization. Consumers perceive that anything digital is better. This consumer perception started with the introduction of the digital clock. The digital clock was more accurate than the analog clock and over time, more desirable.

Miniaturization drove the development of the digital watch, handheld calculator, and a myriad of other consumer products based upon digital technologies. Naturally, a plethora of new standards and protocols have emerged. Multiple standards, changing protocols, high-performance processing, low power requirements at cost-conscious consumer prices – these are the challenges where the low cost Spartan™ series of FPGAs and the ultra low power CoolRunner™ CPLDs shine.

Digital TV

With a heightened awareness and demonstrated demand for digitized electronics, the television industry logically went about defining a new television format based entirely upon digital technology. The results of high-definition digital TV (DTV) were astounding, but the price tag was prohibitively expensive, and consequently, consumer interest was low. Until just recently, it was hard to pinpoint exactly why consumer interest in DTV languished – was it because of price, or did consumers just think that digital video was not interesting?

Actually, it has only been in the past five years that the story has slowly unraveled. Consumers have, in fact, voted strongly in favor of digital video – but this has come, not in the form of an end-to-end digital system, but rather in the form of a single video component: the digital video disc, or DVD. Never, in the history of any consumer electronics product, has there been such a wide and rapid acceptance of any new medium or format as there has been with the introduction of the DVD. The DVD represents one of the most successful consumer electronics products in history. According to International Data Corp, DVD had penetrated 23% of US households by end of 2001.

The success of the DVD is attributed to several factors: the significant improvement in image quality, the industry-wide standardization that reduced consumer confusion, and high-volume sales, which led to accelerated cost reductions. Another key factor in the widespread acceptance of DVDs was that they can be viewed on existing analog televisions.

Now, as digital video has grown in popularity, it is spawning a host of new digital products – and formats – including plasma displays, personal video recorders, read/write DVDs, and set-top boxes.

DTV has also spawned a widespread concern in the "infotainment" industry over video piracy. This concern is giving rise to a suite of new encryption standards, including 5C, 4C, CPTWG, and TCPA (to mention a few) that are being proposed by a wide range of interested parties, including the motion picture industry.

Conclusion

In this issue of *Xcell*, we discuss some of the exciting Xilinx solutions for consumer electronics, with a special emphasis on digital video applications, such as video image processing, navigating changing standards, and addressing encryption challenges. We hope you enjoy these articles.

FPGAs — Enabling DSP in Real-Time Video Processing

With Xilinx FPGAs you can create DSPs that are much faster than any off-the-shelf stand-alone DSP device.

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

Real-time video processing makes enormous demands on system-level performance requirements and, except for the simplest of functions, is well beyond the capabilities of any general purpose DSP device. Therefore, DSP-based video designs often require several DSP devices to get the necessary throughput, along with the overhead of multiple program and data memory resources. However, designing with programmable logic allows you to implement video signal processing algorithms using parallel processing techniques, giving you the performance you need within a single, highly flexible programmable logic device. By performing video processing functions in real time, you can minimize the need for frame stores and data buffering.

With the industry pushing for higher levels of video quality, and the development of improved compression schemes, system processing rates have increased dramatically. At the same time, we have introduced new programmable logic devices, such as the Xilinx Spartan™-IIE FPGAs, which draw on the architectural heritage of the FPGA devices that are commonly used in professional broadcast equipment today. As FPGA process development follows Moore's Law, our new products can perform the same functionality as their predecessors, with even higher performance and a much lower cost.

By developing systems using our latest cost-effective FPGAs, you can bring unprecedented levels of professional, broadcast quality, video processing into areas of digital video technology such as high-end consumer products, security systems, industrial and machine vision applications, and frame-grabbers.

One driver of this trend, is the combination of networking, broadcast, processing and display technology, in what the industry has termed "digital convergence." The need to send high bandwidth video data over extremely difficult transmission channels, such as wireless, while still maintaining an acceptable quality-of-service (QoS), is extremely difficult. This has lead to wide ranging research in how to improve error correction, compression and image processing technology, much of it based on advanced FPGA technology.

Image Compression/Decompression – DCT/IDCT

The main video compression scheme used in digital video systems today is MPEG2. It can be found at the heart of digital television, set-top boxes, digital satellite systems, high-definition television (HDTV) decoders, DVD players, video conferencing equipment, and Web phones, to name just a few. Raw digital video information invariably has to be compressed so it can be either sent over a suitable transmission channel, or stored onto a suitable medium such as a disc.

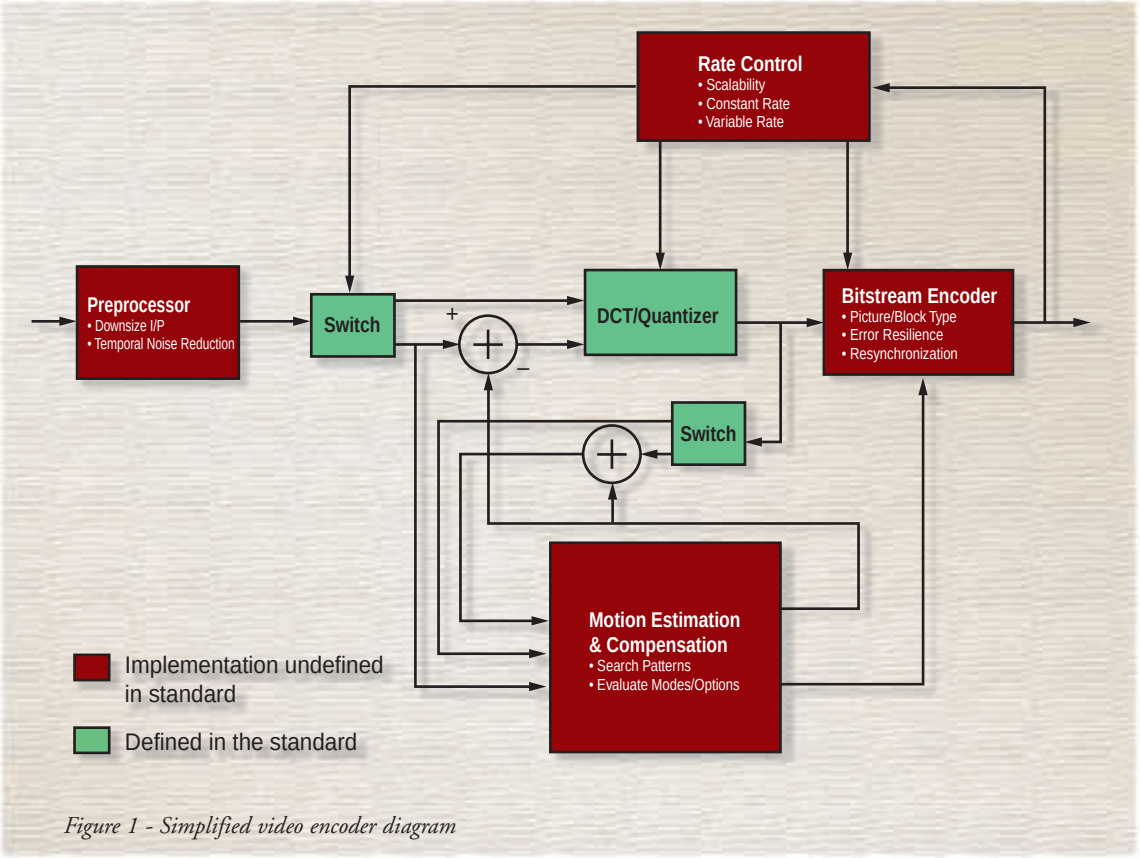

Figure 1 - Simplified video encoder diagram

There are also a number of emerging standards, including most notably MPEG4. Most products based around this technology are still in development, although ramp-up to production is expected shortly. At the heart of the MPEG2 and MPEG4 algorithms is a function called the discrete cosine transform (DCT). The aim of the DCT is to take a square of pixel blocks and remove redundant information that is imperceptible to the viewer. To decompress the data, the inverse discrete cosine transform (IDCT) function is required.

While the DCT section of the MPEG algorithms is standardized and can be implemented very efficiently within FPGAs, MPEG encoding has a lot of blocks that are left undefined. It is within these undefined sections of the standard that a company can truly differentiate its product from its competition and develop its own proprietary algorithms. Many professional MPEG encoders use FPGAs in these sections, such as the motion estimation block, as shown in Figure 1. Because FPGAs are reconfigurable, the equipment can be easily updated to incorporate new algorithms at all stages of development, including in the field after deployment. Companies that rely totally on standard ASSP solutions are limited in their ability to produce products that can make them stand out from the competition, and they therefore run the risk of being seen as just one of a number of similarly specified solutions in the market.

Color Space Conversion

Another important part of a video system is the requirement for color space conversion, a process that defines how an image specified in one color format can be converted into a different color format.

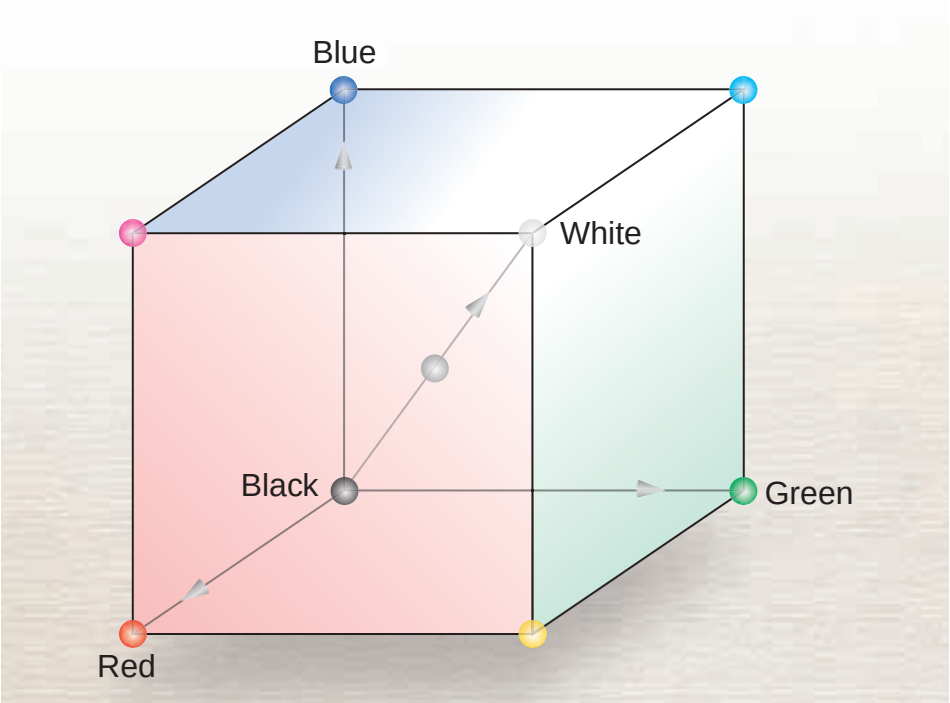

Figure 2 - Color space cube

Receptors in the human eye are only capable of detecting light wavelengths from 400 nm to 700 nm. These receptors are called cones and there are three different types, one for red light, one for green, and one for blue. If a single wavelength of light is observed, the relative responses of these three sensors allow us to discern what we call the color of the light. This phenomenon is extremely useful because it means we can generate a range of colors by simply adding together various proportions of light from just three different wavelengths. The process is known as additive color matching, and is used in color television systems.

It's possible to represent colors of light by plotting the red, green and blue (RGB) components proportions on a 3-dimensional cube, with black at the origin and while at the diagonally opposite corner. The resulting cube is known as the "RGB color space," as shown in Figure 2.

Whether the final display medium is paper, LED, CRT, or plasma displays, the image is always broken down into an array of picture elements or pixels (HDTV, for example, can have 1920 x 1080 pixels). While the mechanics change slightly for each medium, the basic concept is that each pixel displays a proportion of red, green, or blue depending on the voltage signals driven to the display.

Processing an image in RGB format, where each pixel is defined by three 8-bit or 10-bit words corresponding to each primary color, is certainly not the most efficient method. With such a format, every action on a pixel has to be performed on all the red, green, and blue channels. This invariably requires more storage and data bandwidth than other alternative color formats.

To address such issues many broadcast standards, such as the European PAL and North American NTSC television systems, use luminance and color difference video signals. A requirement therefore exists for a mechanism to convert between the different formats, and this is called "color space conversion."

Realizing these circuits in hardware is a relatively simple task, once the coefficients to map from one plane to another are known. One common format conversion is RGB to YCbCr (and conversely YCbCr to RGB). See Figure 3. It has been found that between 60% and 70% of the luminance information (Y) a human eye can detect comes from the green color. The red and blue channels in effect duplicate much of the luminance information, and hence, this duplicate information can be safely removed. The end result is that the image can be represented as signals representing chrominance and luminance. In this format, the luminance is defined as having a range of 16 to 235 in an 8-bit system, and the Cb and Cr signals have a range of 16 to 240, with 128 being equal to 0.

A color in the YCrCb space is converted to the RGB color space using the equations:

$$R' = 1.164\,(Y - 16) + 1.596\,(Cr - 128)$$

$$G' = 1.164\,(Y - 16) - 0.813\,(Cr - 128) - 0.392\,(Cb - 128)$$

$$B' = 1.164\,(Y - 16) + 1.596\,(Cr - 128)$$

R'G'B' are gamma corrected RGB values. A CRT display has a non-linear relationship between signal amplitude and output intensity. By gamma correcting signals before the display, the relationship between received signal amplitude and output intensity can be made linear. The output gain is also limited below certain thresholds to reduce transmission-induced noise in the darker parts of an image.

There are a number of possible implementations – memory, logic, or embedded multipliers – to perform the necessary multiplication functions. It is certainly possible to meet and exceed the 74.25 MHz data rate required for HDTV systems. It is also possible to try different design trade-offs, such as that between system accuracy and design area. For example, for a 3% error in luminance, the design size of a YCrCb-to-RGB color-space-converter can be more than halved. This may be unacceptable in most display products but could be acceptable in other applications such as machine vision or security systems. By using an FPGA, you can tailor the algorithm for the application, thereby maximizing performance, efficiency, or both.

Real-Time Image and Video Processing Functions

Limitations in the performance obtainable with standard DSPs has led to the development of specially designed chips, such as

media processors. However, these devices have often proved to be too inflexible in all but a narrow range of applications, and can suffer from performance bottlenecks. The limitations of a processor-based approach become especially apparent in high-resolution systems, such as HDTV and medical imaging. Fundamentally, a processor solution is restricted in how many cycles can be allocated to each tap of a filter, or each stage of a transform. Once the performance limits have been reached there is often no other way around the problem than to add extra DSP parts.

An FPGA, however, can be custom tailored to provide the maximum efficiency of utilization and performance. It's possible for you to trade off area against speed, and invariably perform a given function at a much lower clock rate than a DSP would require.

For example, Visicom Inc. found that for a median filter application a DSP processor would require 67 cycles to perform the algorithm. An FPGA needed to run at only 25 MHz, because it could realize the function in parallel. For the DSP to match the same performance, it would have to run at over 1.5 GHz. In this particular application, the FPGA solution is some 17 times more powerful than a 100 MHz DSP processor.

There are a wide number of real-time image and video processing functions that are well fitted for implementation in FPGA devices – these include real-time functions such as:

- Image rotation
- Scaling
- Color and hue correction
- Shadow enhancement
- Edge detection
- Histogram functions
- Sharpening
- Median filters
- Blob analysis.

Many of these functions are both application specific and system specific, and are based around core structures such as 2D-FIR filters. Such functions can be implemented quickly using HDL language design or by exploiting the DSP building blocks found in high-level core design tools such as the Xilinx CORE Generator™ software. It's also possible to reduce both design and simulation time by employing a system-level design approach, using products such as The MathWorks' MATLAB™ and Simulink™ software, as well as Xilinx System Generator.

Conclusion

With digital convergence, you need to integrate various standards and requirements into one homogeneous product – this demands flexibility in design and implementation.

FPGA technology is addressing the system requirements of new and emerging video applications, bringing with it the features and signal processing performance that have made it the preferred solution for video and image processing in the professional broadcast equipment market for a number of years. The latest generation of FPGAs can provide the same level of performance and functionality, at a fraction of the cost of the FPGAs that were designed into professional equipment just a few years before.

In comparison to ASSP and chipset solutions, FPGAs offer levels of flexibility that designers demand in today's products, while at the same time maintaining a distinct performance advantage over conventional DSPs.

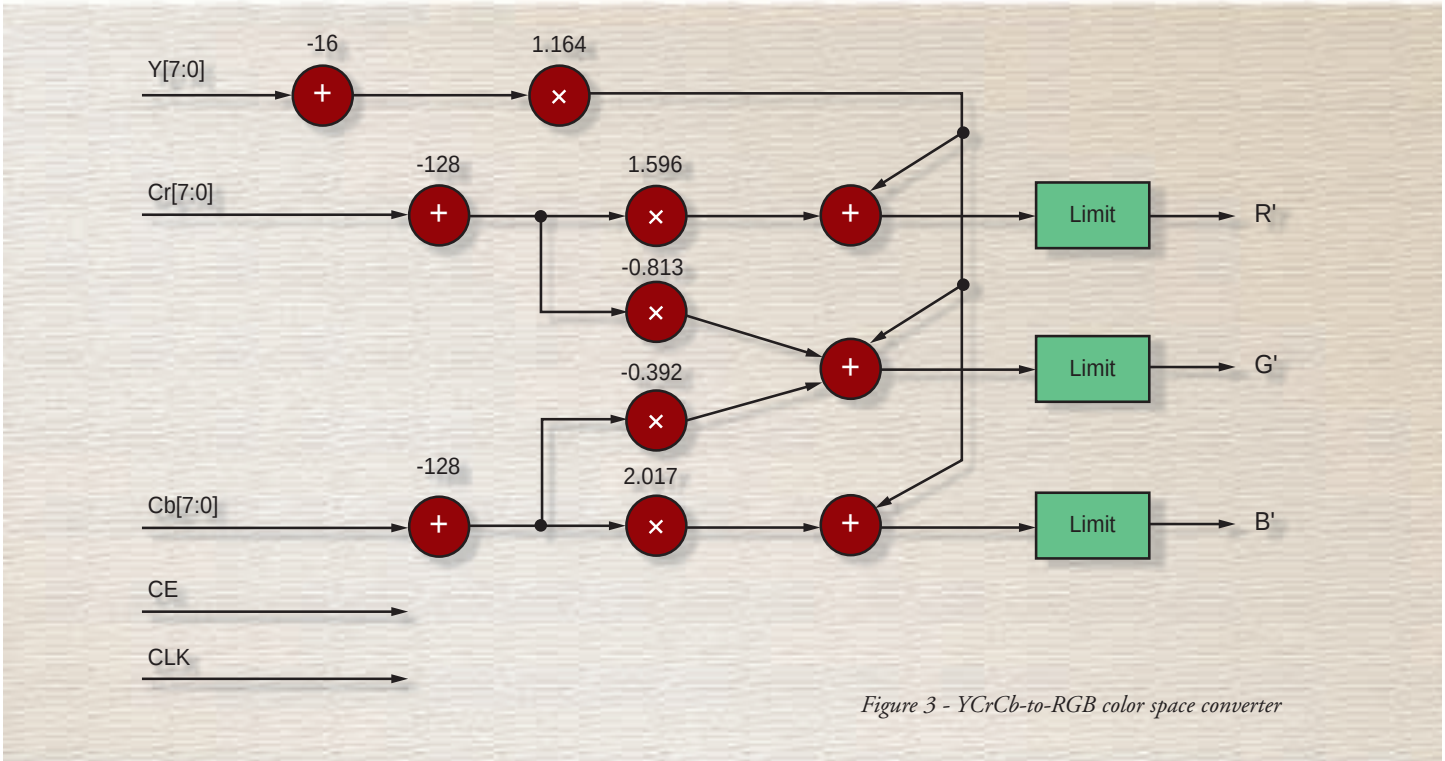

Figure 3 - YCrCb-to-RGB color space converter

FPGAs Provide Flexibility in Digital TV Development

Cost sensitive consumer applications, such as digital TVs, can now benefit from the many advantages of Spartan series programmable logic devices.

by Robert Green, Manager
Strategic Solutions, Digital Audio-Video
Xilinx, Inc.
robert.green@xilinx.com

Broadcasters are very interested in creating new revenue from digital television services. However, a vast range of evolving digital television (DTV) standards in transmission, image processing, interconnectivity, and display interfaces can create considerable barriers to speedy implementation. Yet, the industry promotes digital broadcasting, highlighting better video and sound quality to attract new subscribers from the analog world. Consumer adoption of digital TV is important if broadcasters are to adhere to the schedule of analog switch-off.

Bandwidth is precious; to make the most of it, compression schemes have steadily improved and new algorithms are in development to push the envelope even further. As such, system-processing rates have increased over time, and real-time image processing is now an ideal way of meeting these requirements while removing memory overhead. At the same time, Moore's Law has resulted in low cost programmable logic devices, such as the new Spartan™-IIE FPGAs, that provide the same functionality and performance previously only found in expensive professional broadcast products.

Which Standard Do I Use?

The transition from analog to digital technology is providing opportunities for broadcasters to offer new services, and competition is fierce between OEMs to produce cost-effective systems that are attractive to consumers. As with many technology shifts, this introduces many new proposed standards as companies vie for market leadership. It's possible that the first company to market becomes a de facto standard regardless of the efforts of standards bodies to ensure interoperability and fair competition. Even when standards bodies successfully produce internationally recognized specifications, there are often many versions of a standard that can be adhered to as they try to address many member companies' needs. Add to this the

inevitable revisions to standards during development, and it is easy to see why the flexibility of a fully reprogrammable solution is so compelling.

For example, your return on investment would be much greater if you were able to reprogram your system so that it complies with the latest revision of an emerging forward error correction algorithm. Replacing an aging encryption scheme such as DES with a new, more robust version (AES) is another example.

Table 1 shows a summary of broadcast formats defined by the Advanced Television Standards Committee, and is referred to in the broadcast industry as "Table 3" (from the ATSC A53 specification). There are 36 options that equipment manufacturers need to support in this table alone. There are also regional variations to broadcast standards, with DVB, ISDB, and the proposed DMB used in different parts of the world. Although ASSPs (Application Specific Standard Products) are available, it frequently means that different chips are required for each format.

An FPGA solution, easily supporting data rates exceeding those required by HDTV, would enable all the formats to be supported with one device that is reprogrammed depending on the equipment or operational mode it is used in. Such flexibility in a standard, off-the-shelf component reduces your bill of materials and removes any risk of supply problems from an ASSP provider.

Definition	Lines/Frame	Pixels/Line	Aspect Ratios	Frame Rates
High (HD)	1080	1920	16:9	23.976p, 24p, 29.97p, 29.97i, 30p, 30i
High (HD)	720	1280	16:9	23.976p, 24p, 29.97p, 30p, 59.94p, 60p
Standard (SD)	480	704	4:3, 16:9	23.976p, 24p, 29.97p, 29.97i, 30p, 30i, 59.94p, 60p
Standard (SD)	480	640	16:9	23.976p, 24p, 29.97p, 29.97i, 30p, 30i, 59.94p, 60p

Table 1 - ATSC broadcast formats

Pixels and Performance

With the advent of high-definition broadcasts in many parts of the world, video signal processing requirements have increased dramatically. For example, an HDTV display of 1920 x 1080 resolution, 24-bit pixels, and 30 progressive frames per second, requires a total uncompressed bandwidth of around 1.5 Gbps. Even in areas where high-definition images aren't actually broadcast, HD pictures are still used in all post production stages .

The latest low cost programmable logic devices (the Spartan series) now have multiple I/Os with LVDS (low voltage differential signaling) that support these sorts of data rates. Therefore, even uncompressed video data can be brought on- and off-chip and processed in real time. Performing real-time video processing, even at HDTV rates, allows you to reduce external memory requirements.

Multiple frame stores and data buffers are often used in current digital TV systems because the video signal processor portion of the design becomes a bottleneck. Using the parallel signal processing capabilities of FPGAs means that smaller, or even single, frame stores can be used, and data buffers can be eliminated. The limitations of standard DSPs for real-time image processing has led to the development of specially designed devices. Although performance is greatly improved, these devices have often proved to be too inflexible in all but a narrow range of applications and can still suffer from performance bottlenecks.

FPGAs in Consumer Digital TVs

FPGAs have historically only been found in expensive professional broadcast systems. However, due to low costs, Spartan series FPGAs can be used in high volume consumer products such as digital TVs where set-top box functionality is fully integrated, allowing it to receive digital broadcasts. FPGAs can have many areas of responsibility within digital TV sets, as shown in Figure 1. Interfacing between standard chipsets, as "glue logic," has always been a strong application of FPGAs, but many more image processing tasks (such as color space conversion) and support for network interfaces (such as IEEE 1394) are now possible in low cost programmable devices.

One driver of the trend to offload image-processing tasks to FPGAs is the result of the industry-coined phrase "digital convergence." This arises from the need to send high bandwidth video data over new and extremely difficult transmission channels, such as wireless networks, while still maintaining an acceptable quality-of-service (QoS). This has lead to wide ranging research in how to improve error-correction, compression and image processing technology, much of it based around implementation in an FPGA.

Product Differentiation

By using FPGAs, it is possible for you to differentiate your standard-compliant systems from your competitor's products. With the MPEG-2 compression scheme, for example, it's possible to offload the IDCT (inverse discrete cosine transform) portion of the algorithm from an MPEG processor, to an FPGA, to increase the bandwidth. IDCT (and DCT at the encoder) can be implemented extremely efficiently using FPGAs, and optimized IP cores are readily available to include in MPEG based designs.

By integrating proprietary techniques for other image processing functions alongside defined blocks such as DCT, it's possible to produce a low cost, single chip solution that increases processing bandwidth and gives higher quality images than your competitor. By avoiding systems that rely just

on standard ASSP solutions, you no longer need to be perceived as providing just another product in a number of similarly specified products in the market. Plus, with FPGAs you can easily modify your design at any time to meet new standards or fix bugs – even after your product is in your customers' hands.

Time to Market and Time in Market

As well as enhancing DTV systems and increasing processing performance, FPGAs also get your products to market quicker and keep them generating more revenue once they are in the field. Xilinx FPGAs are based on SRAM technology, which allows you to easily reprogram the device during the development phase. This allows simple debugging of your system, but also enables last minute changes to be made to the product if needed.

Because FPGAs are reprogrammable, there's no waiting for development of a new ASSP chipset to support the latest revision, or a costly and time consuming ASIC re-spin if chip development is done in-house. New ways of programming devices in the field are giving the added benefit of being able to transmit hardware updates via the Internet or some other communications channel. For example, Xilinx IRL™ (Internet Reconfigurable Logic) technology allows you to automatically reconfigure your product, remotely, at any time. The power of this feature is highlighted in situations where first-to-market is key, such as in consumer set-top boxes. It means that boxes can be shipped to customers without necessarily having a complete design – features can be added later.

Conclusion

FPGAs provide both professional and consumer digital broadcast OEMs with real-time image processing capabilities that address the system requirements of new and emerging video applications. Compared to other technologies, FPGAs offer an unrivalled flexibility that enables you to get your products to market quickly. Remote field upgradeability means that systems can be shipped now and features, upgrades, or design fixes added later.

FPGAs also have significant performance advantages over conventional digital signal processors and have the added benefit of allowing you to differentiate your products from the many solutions on the market. Programmable logic is the ideal balance between the flexibility of DSPs and the performance of ASSPs for digital TV development.

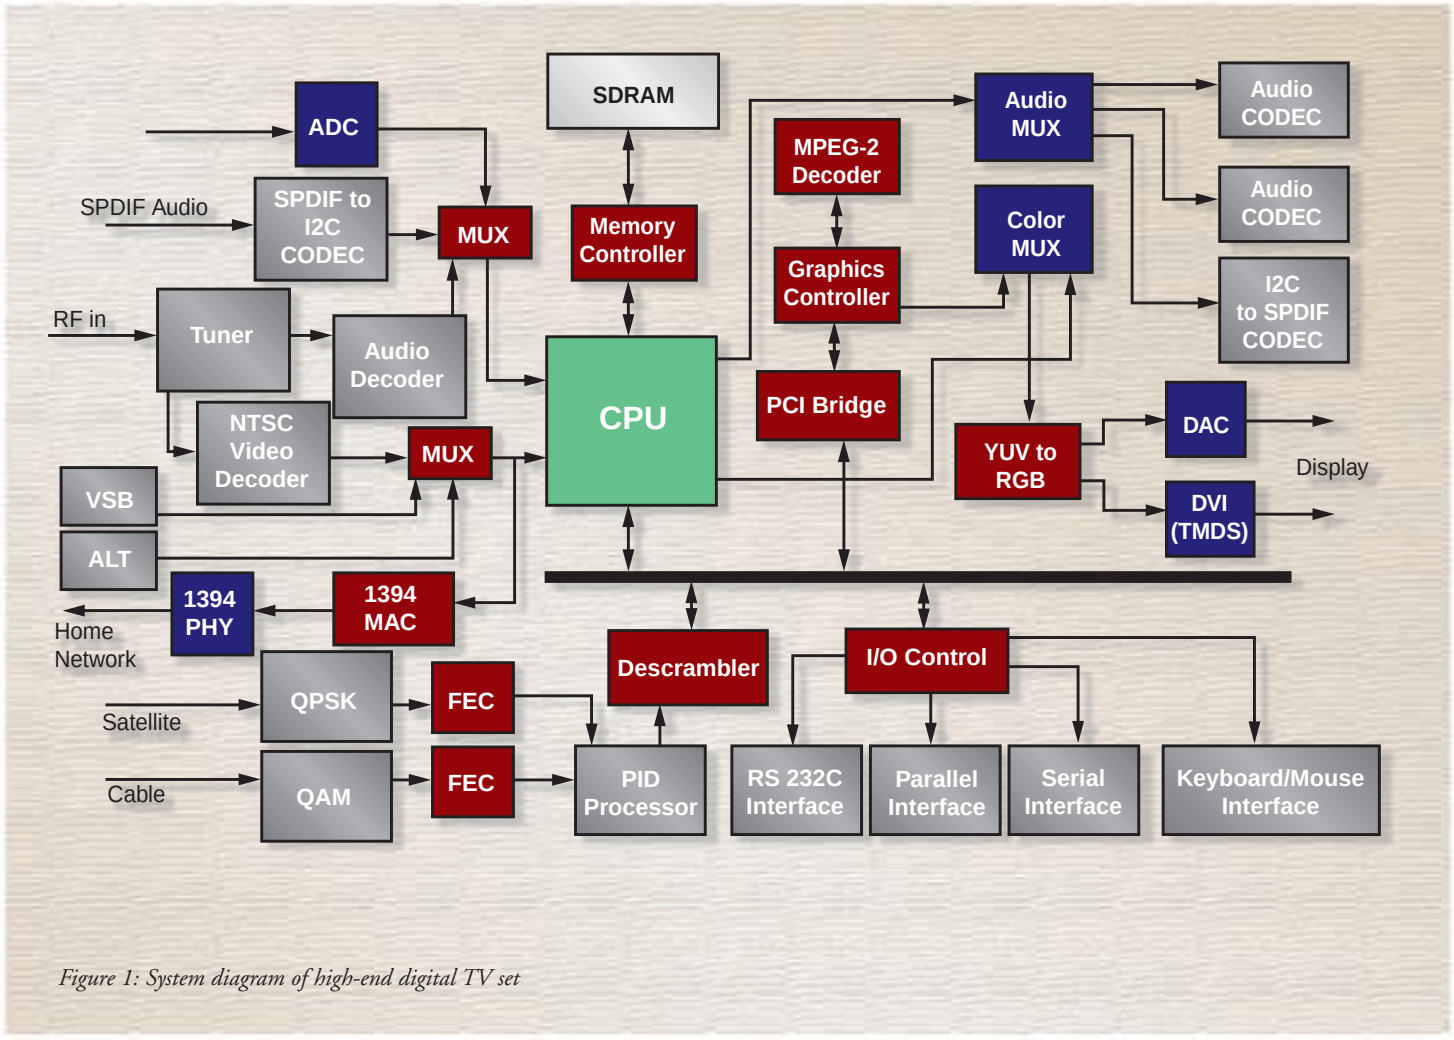

Figure 1: System diagram of high-end digital TV set

Xilinx Technology Can Disable Stolen Cell Phones – and Restore Them When Recovered

Mobile phone theft is on the rise, but cell phones equipped with CoolRunner-II CPLDs, Internet Reconfigurable Logic, and Comprehensive Design Security could remove the incentive to steal.

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

Researchers at the University of California at San Diego have found a way to "blow up" silicon chips with an electrical signal, according to a recent article, "Exploding Chips Could Foil Laptop Thieves," in *NewScientist.com*.

The practical application of self-destroying chips would be the ability to remotely disable stolen laptops, cell phones, PDAs, and other devices using wireless communication and containing confidential information.

According to the *NewScientist.com* article by Duncan Graham-Rowe, when a device is stolen, the victim or service provider could call the device and transmit a self-destruct order to "detonate" a small quantity of gadolinium nitrate. The resulting explosion would destroy the critical chips, rendering the stolen device useless – and irreparable. Unfortunately, information in the device memory might remain intact – and that data could sometimes be more useful than the device itself.

At Xilinx, we have a better idea. With IRL™ (Internet Reconfigurable Logic) technology, Comprehensive Design Security, and CoolRunner™-II CPLDs, we can remotely disable up to four different functions of a stolen device, rendering it inoperable, and with locked-down, tamperproof data. The beauty of the Xilinx

solution is the unit is not damaged in any way – full functionality and data can be easily restored if the device is recovered.

Whose Problem Is It?

Every three minutes in the UK, a mobile phone is stolen, sometimes at knifepoint. This phenomenon is not just confined to the UK. It is a global issue and continuing to rise. For example, in New South Wales, in Australia, nearly 40,000 phones were stolen between October 1999 and September 2000, a 100% rise over the previous year.

But what if the phone hardware could be disabled and effectively rendered useless, then surely the need to steal this much-coveted possession would decline. Wouldn't it? Not necessarily ...

When a mobile phone is stolen, it is relatively easy for the phone operator to disable the Subscriber Identity Module (SIM) card inside the phone so that calls can not be made using that particular number. This is of great benefit to the bill player, but this is not usually why the phone has been stolen. Hackers can install a modified SIM card that fools the phone into believing it is a prepaid, "pay-as-you-go" phone with unlimited credits. This allows the phone to be used indefinitely, free of charge.

What Are Today's Solutions?

In March 2001, the then-UK Home Secretary Jack Straw said that he would meet with the world's mobile phone manufacturers and police to try to stop the growing threat of street robberies that target mobile phones. Among the measures considered were the need for general vigilance, discrete usage, knowledge of security codes to lock the phone, and the use of an International Mobile Equipment Identifier (IMEI) number.

In addition to the SIM card, digital mobile phones connected to the GSM (Global System for Mobile Communications) network possess an IMEI identifier that is associated with the handset itself. (GSM is the dominant cellular system used in Europe and Asia.)

Mobile phone owners are encouraged to note down this number so that if the handset is stolen, then the operator can bar that specific phone from working on their network.

But what about the other networks? The handset can theoretically still be used on other networks if a suitable SIM can be found on the black market.

The need, therefore, is to find a way to completely disable the handset hardware on any cellular system – without blowing up the phone itself. Thus, the phone is useless to the thief, but it can be reactivated and restored if the phone recovered and returned to its rightful owner.

This message was again reinforced in January this year with UK ministers considering whether to introduce legislation that will force networks to introduce the aforementioned anti-theft measures, but, they stipulated, this would be "a last resort."

What Is Tomorrow's Solution?

The average mobile phone thief is becoming more discerning, because in the black market, the newer "smart phones" are becoming the target of choice. To date, the growth of mobile communications has been fuelled by voice calls, but voice is becoming a saturated market, with increasing competition forcing a decline in average revenue per user.

To maintain growth and even sustain current income and profits, cellular providers must introduce new services that will be attractive and beneficial to users, who will then need new types of technically sophisticated handsets from the providers' manufacturers.

In order to capture the market for replacement cellular handsets, the next wave of mobile phone handsets must be smarter, lighter, and last longer on one charge. These new smart phones are reinvigorating the flagging mobile market by providing must-have functionality. The cell phone has now moved from being just a humble voice communications device to a combined PDA, Internet access device, MP3 player, and games console.

For instance, Nokia recently unveiled its latest "lifestyle" function phone, the 5510. The phone includes an integrated digital music player (with 64 MB memory) for ACC (Advanced Audio Coding) and MP3 files, a full keyboard to facilitate better messaging functionality, and a handful of embedded games.

Contrary to recent times, the full cost of smart phone will be borne by the customer – not subsidized as a free loss leader by a wireless service provider, who intends to recover the cost of the hardware with inflated air-time rates.

These smart phones are more costly for the owner and/or network operator to replace if stolen. This cost risk has made the mobile phone companies rethink how their products are designed. The hardware of handset must have the built-in capability to be remotely disabled if stolen – but still traceable via GPS (global positioning system) or similar geographic locator systems.

Once recovered, the phones must be able to be reactivated to full functionality when returned to their owners. A phone with blown-up chips isn't worth recovering. To be a viable consumer product, the disabling component of a stolen cellular phone must be small, lightweight, low cost, low power – and capable of fully restoring the functionality of the phone upon recovery.

The Xilinx Solution

With typical forward thinking, Xilinx has already developed the tools to combat phone theft – IRL technology, Comprehensive Security Design, and the CoolRunner-II CPLD.

If a CoolRunner-II CPLD were used to perform the keypad interfacing in a handset, then if it were stolen, the keypad could be disabled remotely by the mobile phone operator using Xilinx IRL technology (Figure 1). When returned to its owner, the handset could then be reprogrammed again via an IRL bitstream to enable the keypad again. (IRL technology enables the remote upgrading or programming process of CPLD or FPGA hardware over any kind of network, including wireless.)

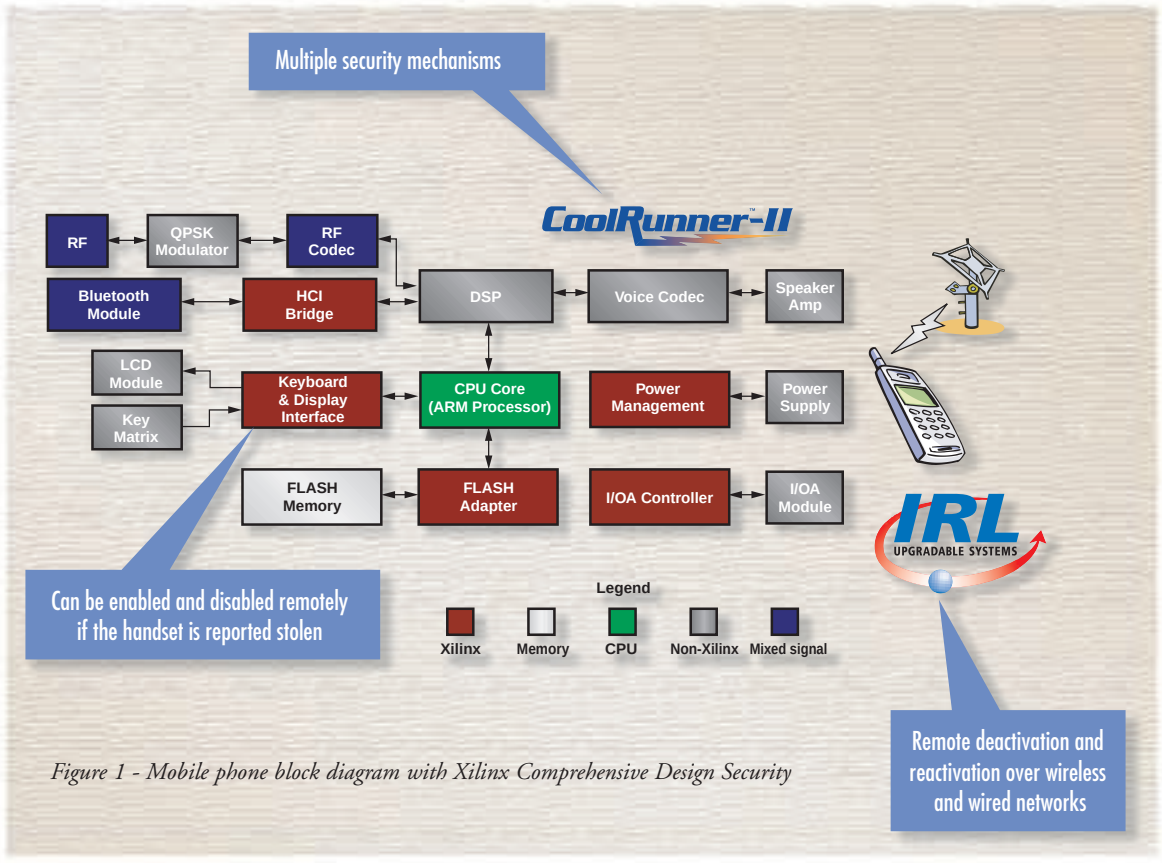

Figure 1 - Mobile phone block diagram with Xilinx Comprehensive Design Security

It is virtually impossible for the CoolRunner-II device (Figure 2) to be reactivated or for data to be retrieved from the handset by the thief or hacker. Xilinx Comprehensive Design Security provides an unprecedented four aspects of design security:

- Prevention of accidental/purposeful overwriting or read back of the configuration pattern
- Blocking visual or electrical detection of the configuration pattern
- Automatic device lockdown in response to electrical or laser tampering
- Physical implementation of the protection scheme that is virtually undetectable.

These design security aspects are not only buried within the layers of the device, but they are also scattered throughout the die to make their detection impossible. Designers can now design with the knowledge they are using the best design security available on any CPLD in the industry.

The enhanced features of CoolRunner-II also elevate the CPLD from an ASIC fix to that of an ASIC replacement. New features include clock doubling and clock division for lower power consumption, various I/O standards support for level translation, selectable Schmitt trigger inputs to solve signal integrity challenges, and bus hold.

Figure 2 - Multiple security functions embedded in CoolRunner-II CPLD

Conclusion

By utilizing a combination of IRL technology, Comprehensive Security Design, and ultra-low power CoolRunner-II CPLDs to remotely disable the hardware within a stolen mobile phone, phone operators and manufacturers can remove the incentive for theft – and improve the odds of recovery with geographic locator technology.

Furthermore, the triple-threat security technology of Xilinx-protected products can be extended to any consumer appliance, including laptops, set-top boxes for cable TV, PDAs, and other devices with wireless or wired network connectivity. As the word gets out, consumers will look for the Xilinx brand in the products they buy – and thieves will look the other way.

Customize Virtual Private Networks with Spartan-IIE FPGA Solutions

Spartan-IIE FPGA solutions can help you to customize your virtual private network for highest performance, greatest security, and lowest cost.

by Amit Dhir
Manager, Strategic Solutions
Optical/Wired Networks
Xilinx, Inc.
amit.dhir@xilinx.com

The worldwide VPN (virtual private network) market will reach $32 billion by 2003, up from $3 billion in 2000, according to Infonetics Research. VPNs provide highly secure, temporary, point-to-point "tunnels" through the public Internet. These tunnels are created on demand and are removed when a session ends. VPNs offer you the following advantages:

- Lower operating costs
- Extendable networks to respond to changing business demands
- Anywhere, anytime access to e-mail, intranet, and other shared applications.

Essentially, VPNs are private networks deployed over public networks that provide similar levels of privacy, security, quality of service (QoS), reliability, prioritization, end-to-end management, and manageability as local area networks (LANs).

With Spartan™-IIE FPGAs, the right software, and Internet gateways, you get secure data encryption and packet authentication even over the wide-open public Internet.

Furthermore, compared to leased line networks, Xilinx-enabled VPNs are:

- Simpler to set up and easier to administer
- More dynamic and extendable
- Less expensive to create and deploy
- More accessible anywhere, anytime.

How Do VPNs Work?

VPNs use the bandwidth of public, packet-routed networks – typically the Internet or a service provider's backbone network (Internet Protocol, frame relay, or ATM). Corporations have company headquarters, remote offices, mobile workers, independent contractors, and business partners connected to a network service provider's local points of presence (POP). Remote users and the company's LAN(s) connect to the provider's network using dial-up, DSL, cable, ISDN, T1/T3, and wireless.

The key to VPNs is "tunneling" – the practice of repackaging data from one network to another. At the originating end of a tunneled transmission, a data packet is "wrapped" or encapsulated with new header information that allows an intermediary network to reorganize and deliver the packet. At the terminating end, the terminal protocol "wrapper" is removed, and the original packet is delivered to the destination.

Tunneling does not ensure privacy or security – just delivery. To secure a tunneled transmission against interception and tampering, all traffic must be encrypted. Therefore, VPNs must include additional functional features to enhance transmission security, ensure quality of service (QoS), and protect the VPN perimeter with a firewall.

Security and privacy features include tunneling itself, data encryption, packet authentication, firewalls, and user identification. Tunneling uses IPSec (Internet Protocol Security), and encryption uses DES (Data Encryption Standard), Triple DES, and Diffie-Hellman cryptographic algorithms.

VPN QoS and bandwidth management make possible delivery of high transmission quality, mission-critical applications (such as financial reporting and order processing), and real-time voice/video applications (such as distance learning and videoconferencing). Packets are tagged with priority and time sensitivity markers to allow traffic to be routed based on delivery priorities. Tagged packets take precedence over bandwidth-consuming applications (such as Web surfing).

Network management features simplify the addition, deletion, or changing of users; software upgrades; and policy management for security and QoS assurance. These features make scalability and interoperability possible.

Xilinx Spartan-IIE FPGAs – Providing Encryption with Flexibility

VPN gateways secure Internet-based communication, perform user identification, authenticate data integrity and confidentiality, and reject all non-tunneled IPSec traffic. Encryption and authentication lie at the heart of VPN products, and speeding up the processing of encryption algorithms boosts the performance and scale of VPN solutions.

Spartan-IIE FPGAs programmed with proprietary and/or standard encryption cores provide flexibility without compromising cost or performance – and reducing time to market.

Implementing encryption in programmable hardware significantly improves performance over software solutions. By using parallel computing, FPGAs encode and decode larger transmission blocks more effectively. The encrypted key can be changed within the FPGA fabric, and if a key is ever broken, the Xilinx Spartan-IIE solution can be reconfigured instantly with new algorithms.

Spartan-IIE FPGAs give you the power to to choose and optimize just the right feature set and intellectual property cores you need to implement your designs. You can integrate functionality, such as PCI, memory controllers, and other components, within the Spartan-IIE device to customize your product and reduce costs. This flexibility comes at a significantly lower cost compared to ASSPs (application specific standard products).

Conclusion

VPNs improve the productivity of remote workers and hence, their organizations. VPNs promote flexible work styles, extend workplaces beyond office walls, connect remote offices with the headquarters, and foster competitive advantages with strategic partners – all at reduced costs, compared to other options.

Using the Internet and/or service provider backbones to transfer confidential data requires state-of-the-art encryption and privacy solutions. Spartan-IIE FPGA-based encryption and security give you the scalability and flexibility you need to implement the perfect virtual private network for your enterprise.

Leveraging Bluetooth Technology to Build Secure and Robust Wireless Communications

Programmable logic is used to design wireless products that address the needs of data security and integrity.

by Mamoon Hamid, Manager
Strategic Solutions, Consumer/Home Networking
Xilinx, Inc.
mamoon.hamid@xilinx.com

Although Bluetooth™ wireless technology promises to be a ubiquitous, low cost solution – perfect for a very wide range of data communication – there are two major areas that further need to be addressed: data security and data integrity. These two areas will serve to limit range of applications where Bluetooth can be deployed.

Bluetooth Data Security

The Bluetooth specification has defined a set of security mechanisms to provide a very basic level of protection for short-range wireless communication. Each Bluetooth device is required to implement key management, authentication, and encryption. The Bluetooth Generic Access Profile (GAP) divides security into three different modes. In addition, Bluetooth technology implements a frequency-hopping scheme that also serves as an additional security measure against eavesdroppers.

Key Management

Bluetooth security uses several different keys in higher layer software to ensure secure transmissions. Figure 1 shows a block diagram for key management.

Device Authentication

Authentication is a critical security component of Bluetooth systems, allowing a domain of trust between an ad-hoc network of Bluetooth devices. Bluetooth authentication is based on a challenge-response scheme that secures the connection between two devices. When a device is not authenticated, a period of time must pass before a new attempt for authentication can be made.

Encryption

Bluetooth Encryption is used to ensure the privacy of a connection. Prior to encryption, an authenticated link must be in place. Encryption scrambles the data payload of Bluetooth packets using an 8-bit to 128-bit key, however, the Bluetooth access code and the packet header are never encrypted. Payload encryption depends on the strength of encryption desired and geographical regulations in the countries in which the product is being deployed.

Physical Layer Data Security – Frequency Hopping Spread Spectrum

In addition to the other security features built into Bluetooth equipment, the frequency-hopping scheme employed for Bluetooth communication (Frequency Hopping Spread Spectrum or FHSS) also serves as mechanism to make eavesdropping extremely difficult.

The Bluetooth radio operates in the 2.4 GHz ISM band. In North America and most of Europe, Bluetooth radio uses a frequency range from 2.402 GHz to 2.480 GHz, with this band divided into 79, 1-MHz subchannels. In frequency hopping, a data signal is modulated with a narrowband carrier signal that hops from frequency to frequency as a function of time.

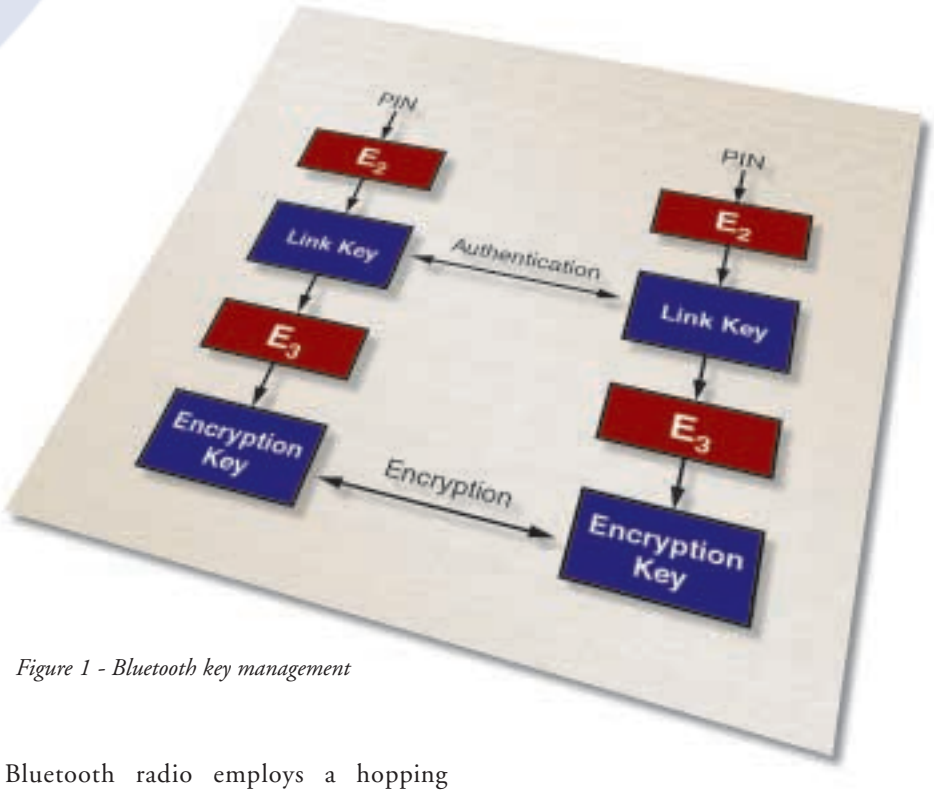

Figure 1 - Bluetooth key management

Bluetooth radio employs a hopping sequence of 1600 hops per second.

FHSS relies on frequency diversity to combat interference. If the radio encounters interference on one frequency, the radio will retransmit the signal on a subsequent hop on another frequency. The aggregate interference will be very low, resulting in little or no bit errors.

Bluetooth Data Integrity – Forward Error Correction (FEC)

Forward error correction reduces the number of retransmissions for a data payload. However, the drawback is that it can tremendously reduce the achievable throughput. There are three types of forward error correction that are implemented by Bluetooth technology:

- 1/3 rate FEC – Each header bit is repeated three times. It is primarily used for masking errors in the header, which contains vital link information.
- 2/3 rate FEC – Every ten information bits are encoded into a 15-bit code word. This type of error correction is primarily used for sending data packets in the least vulnerable way.
- Automatic Repeat reQuest (ARQ) – Used to transmit and retransmit data packets, among other things, until the reception of those packets is acknowledged from the recipient. The recipient returns the acknowledgement in the header sent back to the sender of the initial packet.

Data Whitening

All the header and payload information is scrambled with data whitening bits before transmission occurs, so that redundant 0's or 1's are eliminated from the transmissions. This can be a problem with analog data signals received by the baseband processor because they need to be classified as 0 or 1. Because there is no reference point such as DC for these signals, you must rely on the last few signal transmissions for calibration purposes. Any long sequential string of 0's or 1's may cause calibration to fail. Hence data whitening is used to scramble these signals so that the probability of long strings of 0's or 1's is significantly reduced.

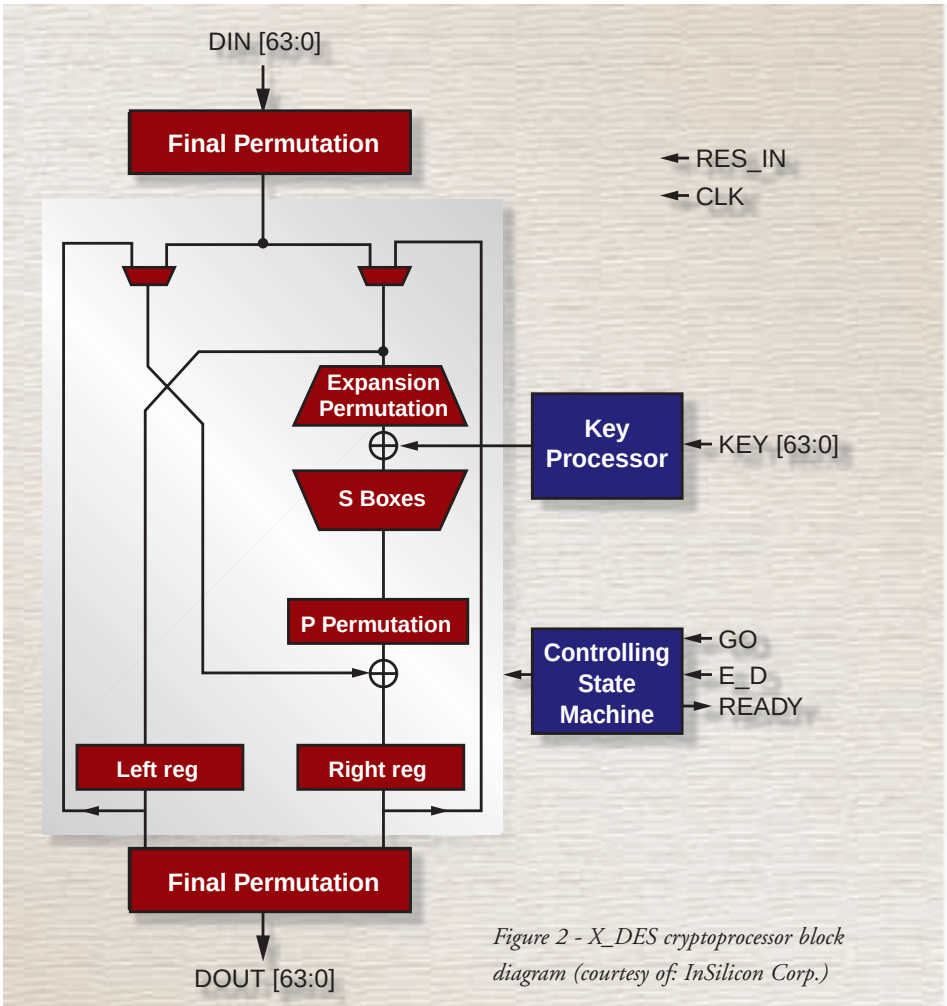

Figure 2 - X_DES cryptoprocessor block diagram (courtesy of: InSilicon Corp.)

Programmable Solutions for Advanced Data Security

Bluetooth data security is inadequate for most applications where a high level of confidentiality is of the utmost importance. It seems to be adequate for smaller applications, but any sensitive or otherwise problematic data should not be transmitted via Bluetooth. For example, the encryption scheme used in Bluetooth shows some weaknesses. In one possible scenario, an attacker can obtain the encryption key that secures communication between two devices and can then eavesdrop on their messages. The attacker also could claim to be one of the devices and insert false messages. One way to avoid this, Lucent Technologies says, is for the users to employ lengthy personal identification numbers to increase the difficulty of discovering the encryption keys. This again involves manually entering personal identification numbers (PINs). This can quickly become a nuisance; it is a tedious process to secure the connection.

Another way to overcome these security issues would be to use a more robust encryption algorithm, such as Data Encryption Standard (DES), or even Triple DES, or Advanced Encryption Standard (AES), instead of the stream cipher E0 algorithm. DES is a block cipher, which essentially means that the cipher encrypts blocks of data all at once and then proceeds to the next block. In DES, the original message is divided into fixed length blocks of 64 bits, enciphered using a 56-bit secret key into a 64-bit encrypted message by means of permutation and substitution. A block diagram of DES is shown in Figure 2.

Unlike the Bluetooth stream cipher algorithm, it has been mathematically proved that the block ciphers are completely secure. The DES block cipher is highly random, nonlinear, and produces encrypted text that functionally depends on every bit of the original message and the key. The DES has more than 72 quadrillion (72 x 1015) possible encryption keys that can be used. For each given message, the key is chosen at random, from among this enormous number of keys. The DES algorithm is widely used and is considered very dependable, however, an even more secure variant, Triple DES, which applies DES three times with different keys, in succession is also available.

Both of these encryption algorithms – DES and Triple DES – can be implemented using low-cost programmable logic devices and readily available intellectual property (IP) for a higher level of encryption. Today, you can buy 100,000 system gate Spartan™ FPGAs in volume, off the shelf, and ready to go, for just $10 per unit. These FPGAs also allow additional functionality, such as advanced error correction, to be added to a design as well. Hence, these devices enable drastic system-level cost reductions.

Another encryption algorithm that is gaining rapid acceptance (due its superior robustness and optimal implementation in hardware and software) is AES, also called Rijndael. AES is a 128-bit block cipher with selectable 128, 192, or 256-bit key lengths. AES encryption and decryption blocks can be implemented individually or as a whole in programmable logic. In addition, the key generator can be parameterized based on key length. A programmable logic-based implementation is modular and also allows you to choose a hardware implementation based on the level of security desired. Shown in Figure 3 is a block diagram of AES. The cost on a silicon basis for this encryption core with a 128-bit key generator is less than $2 today. Variations can be implemented depending on the end application security requirements.

Software-based encryption solutions offer high flexibility, but low performance. Hardware-based encryption solutions provide high performance, yet provide very little flexibility once designed. An encryption solution based on programmable logic provides the best of both worlds – you get high levels of flexibility as well as high performance.

A more robust encryption algorithm will allow Bluetooth technology to be safely deployed in a wide range of applications where security is of the utmost importance. This includes:

- Electronic financial transactions such as ATM or Smart Cards
- Secure e-commerce transactions
- Secure office communications
- Secure video surveillance system
- Digital set-top boxes
- HDTV
- Other consumer electronics devices.

Programmable Solutions for Advanced Data Integrity

The Bluetooth specification defines measures that are effective for preventing unintentional errors from being transmitted using data whitening and error checking. However, a particular stream of data, which would consistently cause transmission errors, is theoretically possible, but highly unlikely in normal transactions.

There is still a risk of error-prone transmissions in more hostile environments, such as industrial locations, office buildings, airports, and metropolitan public transportation. In locations like these, noise and interference from other multiple Bluetooth devices, wireless networks, phone systems, or other electronic devices operating on the same frequency band can cause problems. Again, programmable logic and advanced error-correction IP can be used to facilitate error-free communication. A robust forward error correction technique, such as turbo convolution coding, is an alternative for a hostile communications environment.

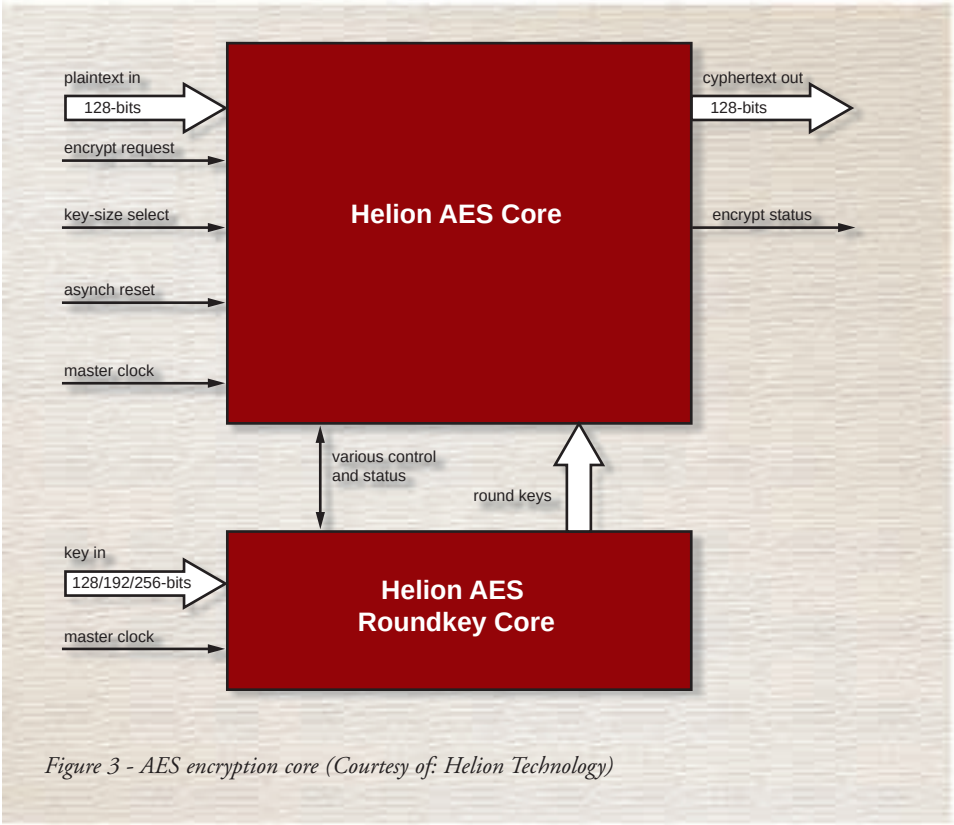

Figure 3 - AES encryption core (Courtesy of: Helion Technology)

Turbo coding is an advanced forward error correction algorithm. It is a standard component in third-generation (3G) wireless communication systems, like those employing wideband code division multiple access. The principle of turbo coding is that the encoder generates a data stream consisting of two independently encoded data streams and a single un-encoded data stream. The two parity streams are weakly correlated due to the interleaving. In the turbo decoder, the two parity streams are separately decoded with soft decision outputs, referred to as "extrinsic" information.

The strength of the turbo decoding results from the sharing of the extrinsic information during a number of iterations. The extrinsic information is passed from one parity decoding step to the other, from one iteration to the other. IP for such a turbo convolutional codec is readily available and is easily implemented on low cost programmable logic devices such as the Xilinx Spartan series. This solution can ensure the integrity of Bluetooth data transmissions in hostile and error-prone environments.

Conclusion

Bluetooth wireless technology is changing the way we communicate. However, it still faces challenges on the level of data security and integrity for secure applications. Low-cost programmable logic devices are ideally suited for advanced data security and integrity implementation. This will be a necessity for those who wish to implement higher levels of security into applications where security is of the utmost importance. Also, applications that will be in operation in hostile environments will benefit from a customizable programmable logic implementation. Programmable logic is an excellent solution for integrating custom Bluetooth implementations, such as these, into embedded systems.

IEEE 1394 and HAVi Are the Leading Technologies for Wired Home Networking

Home networking promises to be one of the hottest new markets in consumer electronics – and the combination of IEEE 1394 and HAVi technologies show the most promise to become the standard for home networking.

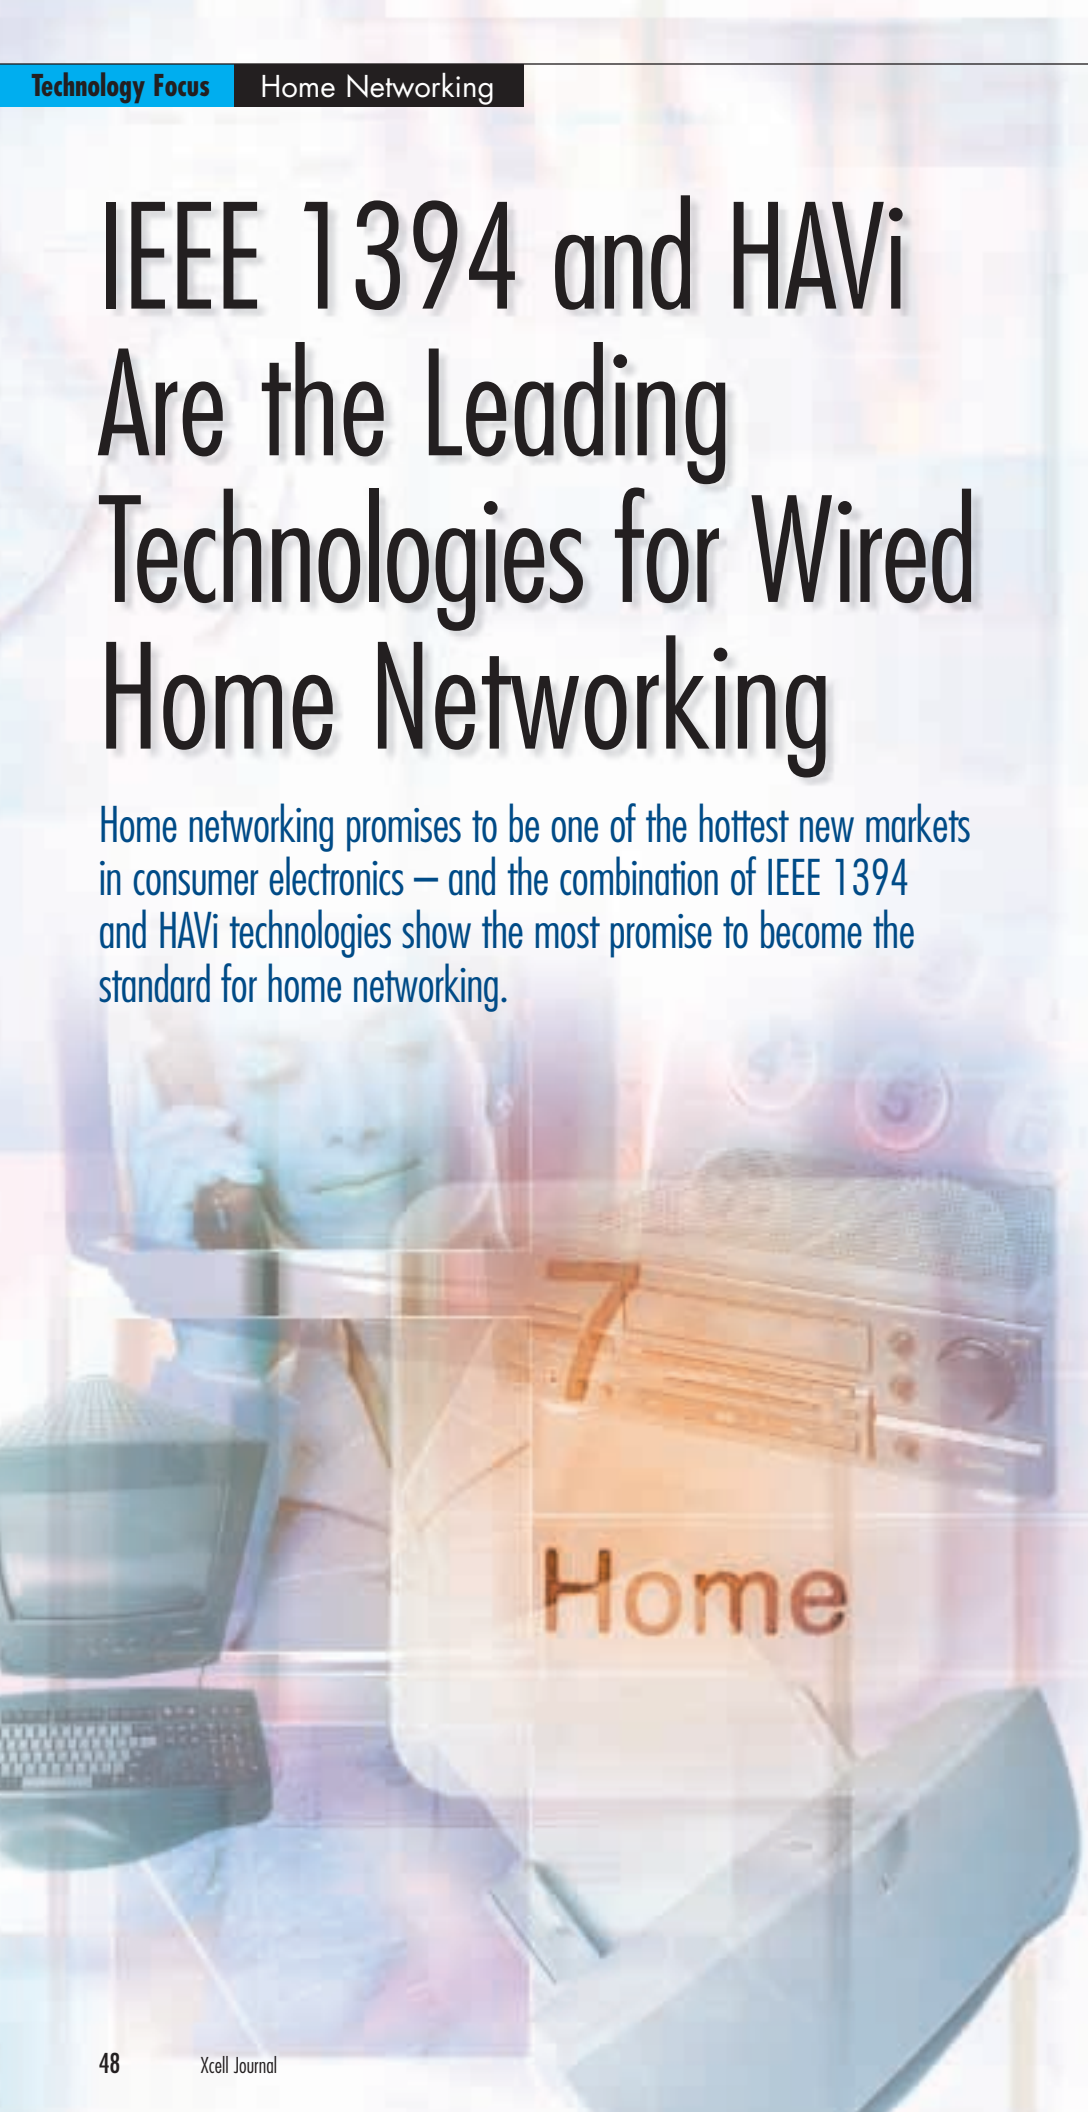

by Amit Dhir, Manager
Strategic Solutions, Optical/Wired Networks
Xilinx, Inc.
amit.dhir@xilinx.com

The consumer world is buzzing with news and standards to network home appliances, PCs, peripherals, and other consumer devices. A number of technologies are vying to be the next technology to network all appliances in the home. Some of these 22 technologies include HomeRF™, Bluetooth™, HiperLAN2, IEEE 802.11, Ethernet, HomePNA™, USB 2.0, and IEEE 1394.

Although each technology presents unique pros and cons, the technologies that require wires do provide one critical capability – the ability to deliver reliable high-speed voice, data, and video transfer. Home networking is incomplete without the ability to transfer voice, data, and video together.

Ethernet, USB 1.1/2.0, and IEEE 1394 top the list of these "new wires" technologies:

- Ethernet is a low-cost technology, which uses a protocol ideal for data traffic.
- USB 1.1 and its more recent specification USB 2.0 are primarily PC-centric – networking PCs and PC peripherals.
- IEEE 1394 is the technology that holds the most promise to be the perfect home networking technology that will network audio, video, and PC equipment.

IEEE 1394 Backgrounder

The move from analog to digital technology has enabled the sharing of video, Internet, data, and audio. Consumers are constantly searching for faster, cheaper, more reliable, and easier ways of transferring and sharing information. While information sharing has been happening at the enterprise level for the last few years, the home is just seeing its introduction. This phenomenon, known as home networking, allows the interconnection of PCs, consumer equipment, and communications, and the distribution of infotainment among these appliances.

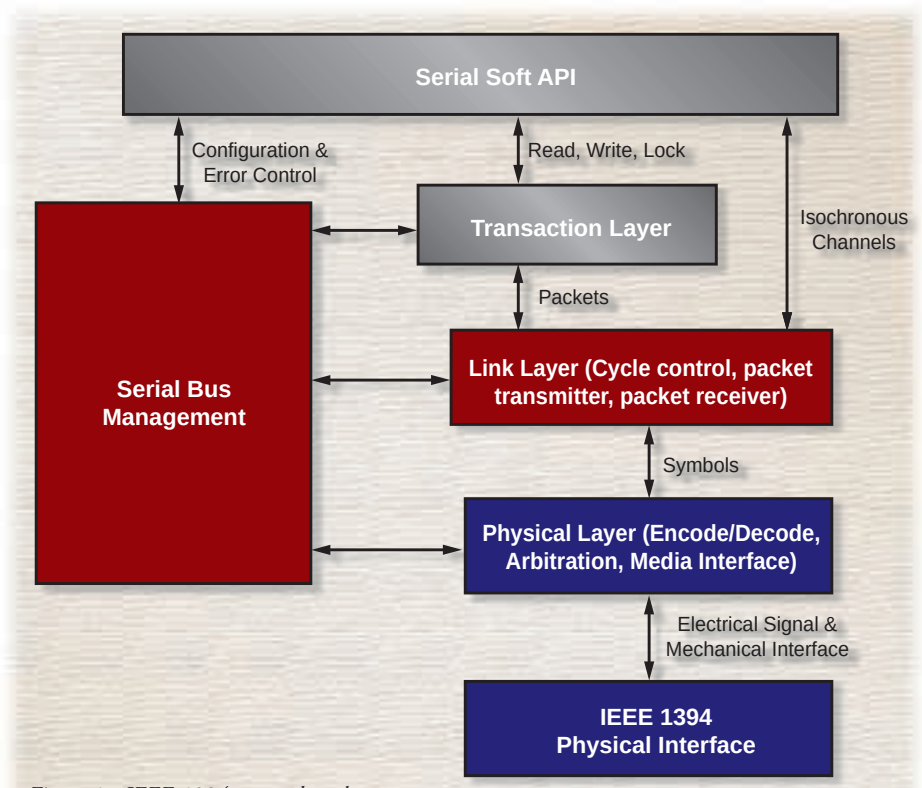

Figure 1 - IEEE 1394 protocol stack

The convergence of voice, data, and video in the home will happen only when seamless, high-speed communication becomes readily available. IEEE 1394 is one such interconnection technology that will enable connection of these devices throughout the home. IEEE 1394, also known as 1394, FireWire™ or iLink™, is a versatile, high-speed, and inexpensive method of interconnecting a variety of consumer electronic devices (such as digital TV, set-top boxes, and home theater equipment), PCs, and PC peripherals (such as scanners and printers).

The FireWire bus standard, originally created by Apple Computer, was born out of the need for a low cost, consumer-oriented connection between digital-video recorders and PCs. FireWire technology grew into a standard – called IEEE 1394 – for low cost, high data rate connections. In 1994, the 1394 Trade Association (1394ta) was formed to support and promote the adoption of the IEEE 1394 standard. In 1995, the 1394ta formally released the 1394 specification. Further revisions like 1394a in 1998 and 1394b in 1999, were introduced. The 1394b standard is fully backward compatible with the current 1394 and 1394a specifications. Each revision of 1394 has added features, performance, and capabilities.

IEEE 1394 Architecture

The components that form a 1394-based home network include the actual protocol itself, the cabling system, and the architectural design of the network. Similar to other high-speed networking systems, IEEE 1394 adopts a layered approach to transmitting data across a physical medium. The four layers used by the IEEE 1394 protocol are shown in Figure 1.

Physical Layer

The physical layer provides the electrical and mechanical connection between the 1394 appliance (connector) and the cable itself. Besides the actual data transmission and reception tasks, the physical layer also provides arbitration to insure all devices have fair access to the bus. As its physical media, 1394 requires optical fiber or high-grade copper wiring between the appliances. The 1394 physical layer is physically point-to-point and logically a bus (each node is a repeater).

The physical layer transmits the unstructured, raw, bit stream over a physical medium, and describes the electrical, mechanical, and functional interface to the carrier. It provides the linking to the upper sessions via signaling and the initialization and arbitration services necessary to assure that only one node at a time is sending data. The physical layer of the 1394 protocol includes electrical signaling, mechanical connectors and cabling, arbitration mechanisms, serial coding and decoding of the data being transferred or received, and transfer speed detection.

Link Layer

The link layer takes the raw data from the physical layer and formats it into two types of recognizable 1394 packets – asynchronous and isochronous.

Asynchronous data transfer is the conventional transmit-acknowledgment protocol and puts the emphasis on guaranteed delivery of data, with less emphasis on guaranteed timing.

Isochronous data transfer is a real-time guaranteed-bandwidth protocol for just-in-time delivery of information. It puts the emphasis on the guaranteed timing of the data and less emphasis on delivery. Isochronous transfers are always broadcast in a one-to-one or one-to-many fashion. No error correction or retransmission is available for isochronous transfers.

Transaction Layer

The third layer in the IEEE 1394 protocol is called the transaction layer and is responsible for managing the commands that are executed across the home network. It supports the asynchronous protocol write, read, and lock commands. A write sends data from the originator to the receiver, and a read returns the data to the originator.

Serial Bus Management

The fourth and final logical grouping of functions is responsible for the overall configuration control of the serial bus. It controls the serial bus in the form of optimizing arbitration timing, guarantee of adequate electrical power for all devices on the bus, assignment of which 1394 device is the cycle master, assignment of isochronous channel ID, and basic notification of errors. The bus management is built upon IEEE 1212 standard register architecture.

Benefits of IEEE 1394

- Broad industry and standards bodies support
- Low cost
- High and scalable speeds
- Plug and play
- Nonproprietary.

IEEE 1394 is an enabling technology for connecting multimedia devices, such as:

- Digital camcorders and VCRs
- Satellite modems
- Set-top boxes
- Digital TV
- PCs
- DVD players
- Gaming consoles
- Home theater
- Musical synthesizers/samplers with digital audio capabilities
  - Digital audio tape (DAT) recorders
  - Mixers
  - Hard-disk recorders
- Video editors.

HAVi Backgrounder

HAVi (Home Audio/Video interoperability) is an industry initiative started by Sony and Philips in 1996. Since then six other companies have joined – Thomson, Hitachi, Toshiba, Matsushita, Sharp, and Grundig.

HAVi adopted the IEEE 1394 bus standard as the underlying network technology for the HAVi protocols and for the transport of real-time audio/video (A/V) streams. Using 1394 as the bus standard (shown in Figure 2) provides benefits, such as high-speed, flexible connectivity, and the ability to link up to 63 appliances together. No other interconnection technologies, such as wireless, HomePlug™, HomePNA, and USB, are capable of distributing high-speed video applications.

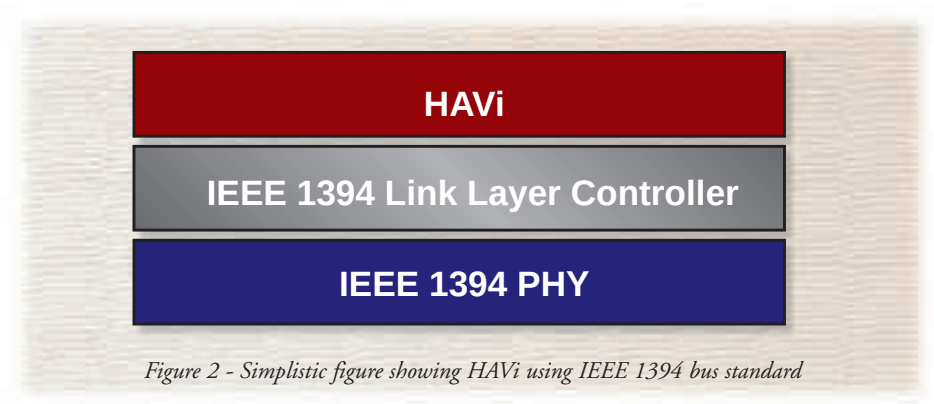

Figure 2 - Simplistic figure showing HAVi using IEEE 1394 bus standard

The HAVi middleware architecture is an open (nonproprietary), lightweight, and platform-independent specification that allows development of home networking applications. It does not, however, address home networking functions such as controlling the lights or monitoring the climate within the house. HAVi specifically focuses on the transfer of digital A/V content between in-home digital appliances, as well as the processing (rendering, recording, and playback) of this content by HAVi-enabled appliances.

The HAVi middleware system is independent of any particular operating system or CPU and can be implemented on a range of hardware platforms, including digital products such as cable modems, set-top boxes, integrated TVs, Internet TVs, or intelligent storage appliances for A/V content.

Today in the world of analog consumer electronic appliances, there exist a number of proprietary solutions for interoperability among appliances from one brand or vendor. In the upcoming world of digital technologies, HAVi extends this networking model by allowing communication among consumer electronic appliances from multiple brands in the home. The HAVi middleware architecture specifies a set of APIs (application programming interfaces) that allow consumer electronic manufacturers and software engineering companies to develop applications for IEEE 1394-based home networks.

One of the main reasons HAVi selected IEEE 1394 over other transmission protocols is because of its support for isochronous communications. HAVi comprises software elements that facilitate the interoperability among different brands of entertainment appliances within the house. Interoperability is an industry term that refers to the ability of an application running on an in-home appliance to detect and use the functionality of other appliances that are connected to the home network.

The underlying structure for a home network based on HAVi technologies is a peer-to-peer network, where all appliances can talk to and interact with each other. HAVi has been designed to allow the incremental addition of new appliances, which will most likely result in a number of interconnected clusters of appliances. Typically, there will be several clusters of networks in the home, with one per floor or per room. Over time these clusters will be connected with technologies such as 1394 long or wireless 1394.

Benefits of HAVi

A HAVi-compliant appliance offers a number of advantages, which include:

- Automatic detection
- Automatic registration of added appliance to the HAVi network
- Automatic software upgrades
- Manageability
- Brand independence
- Hot plug-and-play
- Legacy appliance support.

Xilinx Programmable Solutions Enable IEEE 1394/HAVi Products

The IEEE 1394 technology consists of a physical layer for encoding-decoding, arbitration, a medium interface, and provides an electrical signal and mechanical interface. The link layer provides cycle control, packet transmits, packet receives, CRC, and provides the host and application interface.

As shown in Figure 3, programmable logic solutions provide complete link layer functionality with the ability to connect to multiple interfaces such as PCI, USB, and proprietary audio-video buses. The advantage of programmability is realized when there is a proprietary application interface. However, in an IEEE 1394 system, the Spartan™-II FPGA provides system interface and other ASSP functionalities.

With the 1394 specification still continuing to evolve, having the link layer controller programmed in a FPGA provides the ability to reprogram the FPGA with the latest 1394 revision.

Supporting different products requires the support and interface to different interfaces. For example, using 1394 in a PC requires an interface to PCI, PCMCIA, and other proprietary interfaces. Programmable solutions are ideal for providing this interface, because developing ASSPs for these applications is relatively expensive. Also, the decreasing cost of programmable logic solutions makes them ideal for IEEE 1394 and HAVi-based products.

Conclusion

The digital home continues to evolve and smarter appliances continue to populate the home. These smarter appliances and the need for sharing broadband data, voice, and video are pushing the need for home networking. Although several technologies exist, the technology that provides high-speed and reliable delivery of voice, data, and video will win.

IEEE 1394 is one such home networking technology that provides both high-speed and reliable delivery. The proliferation of 1394 as the A/V standard will be accelerated through the use of HAVi as its middleware solution to connect disparate devices, thus providing a complete solution to the consumer.

However, the specification continues to evolve and 1394 products need to coexist with other home networking and system interface technologies, thus requiring products that interface between these technologies. With the continuing reduction in prices of programmable logic solutions, these are ideal solutions to provide interfaces between 1394 and other technologies such as PCI, USB, PCI-X, SCSI, and HomePNA.

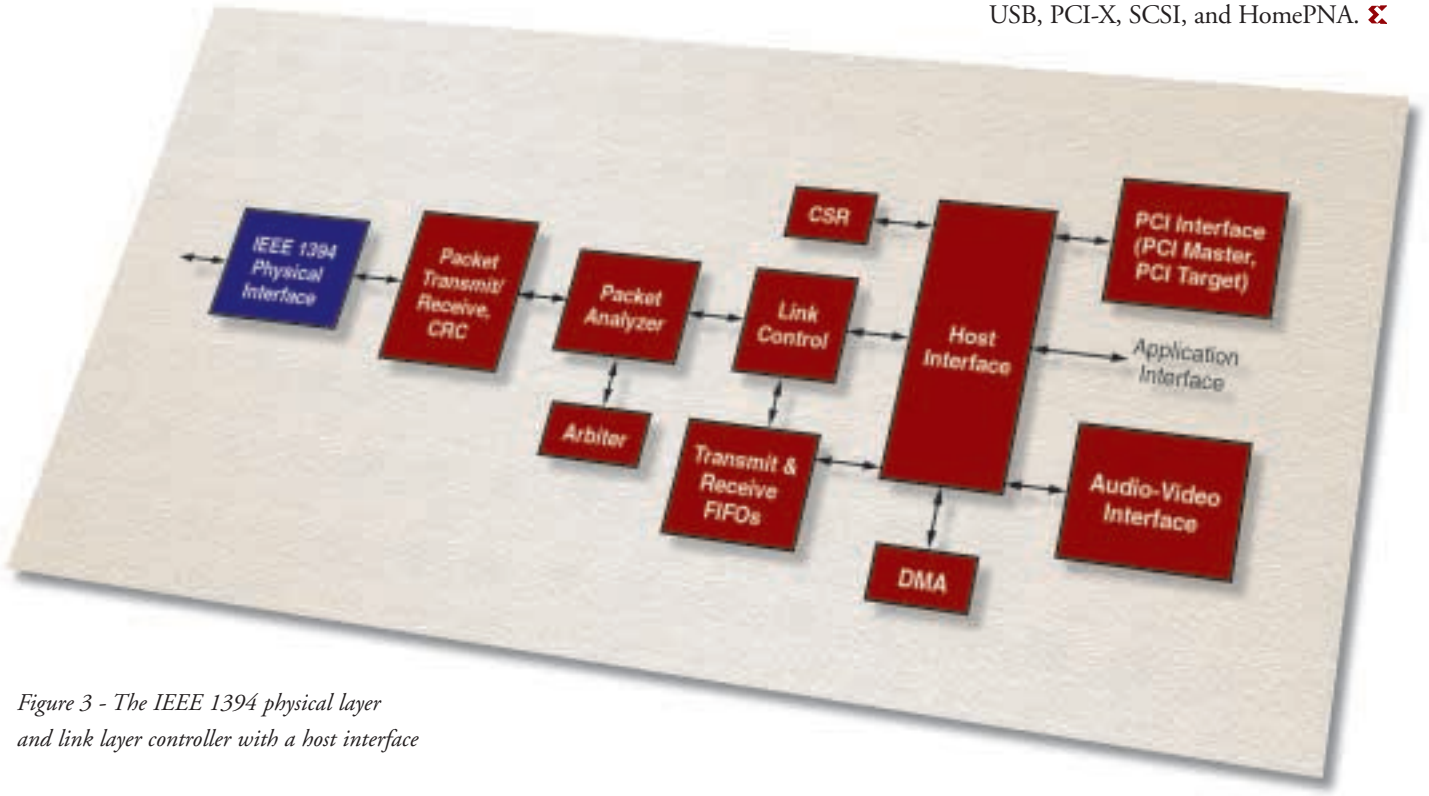

Figure 3 - The IEEE 1394 physical layer and link layer controller with a host interface

You Can Take It with You: On the Road with Xilinx

Xilinx products and technology are putting office technology and functionality into next-generation automobiles.

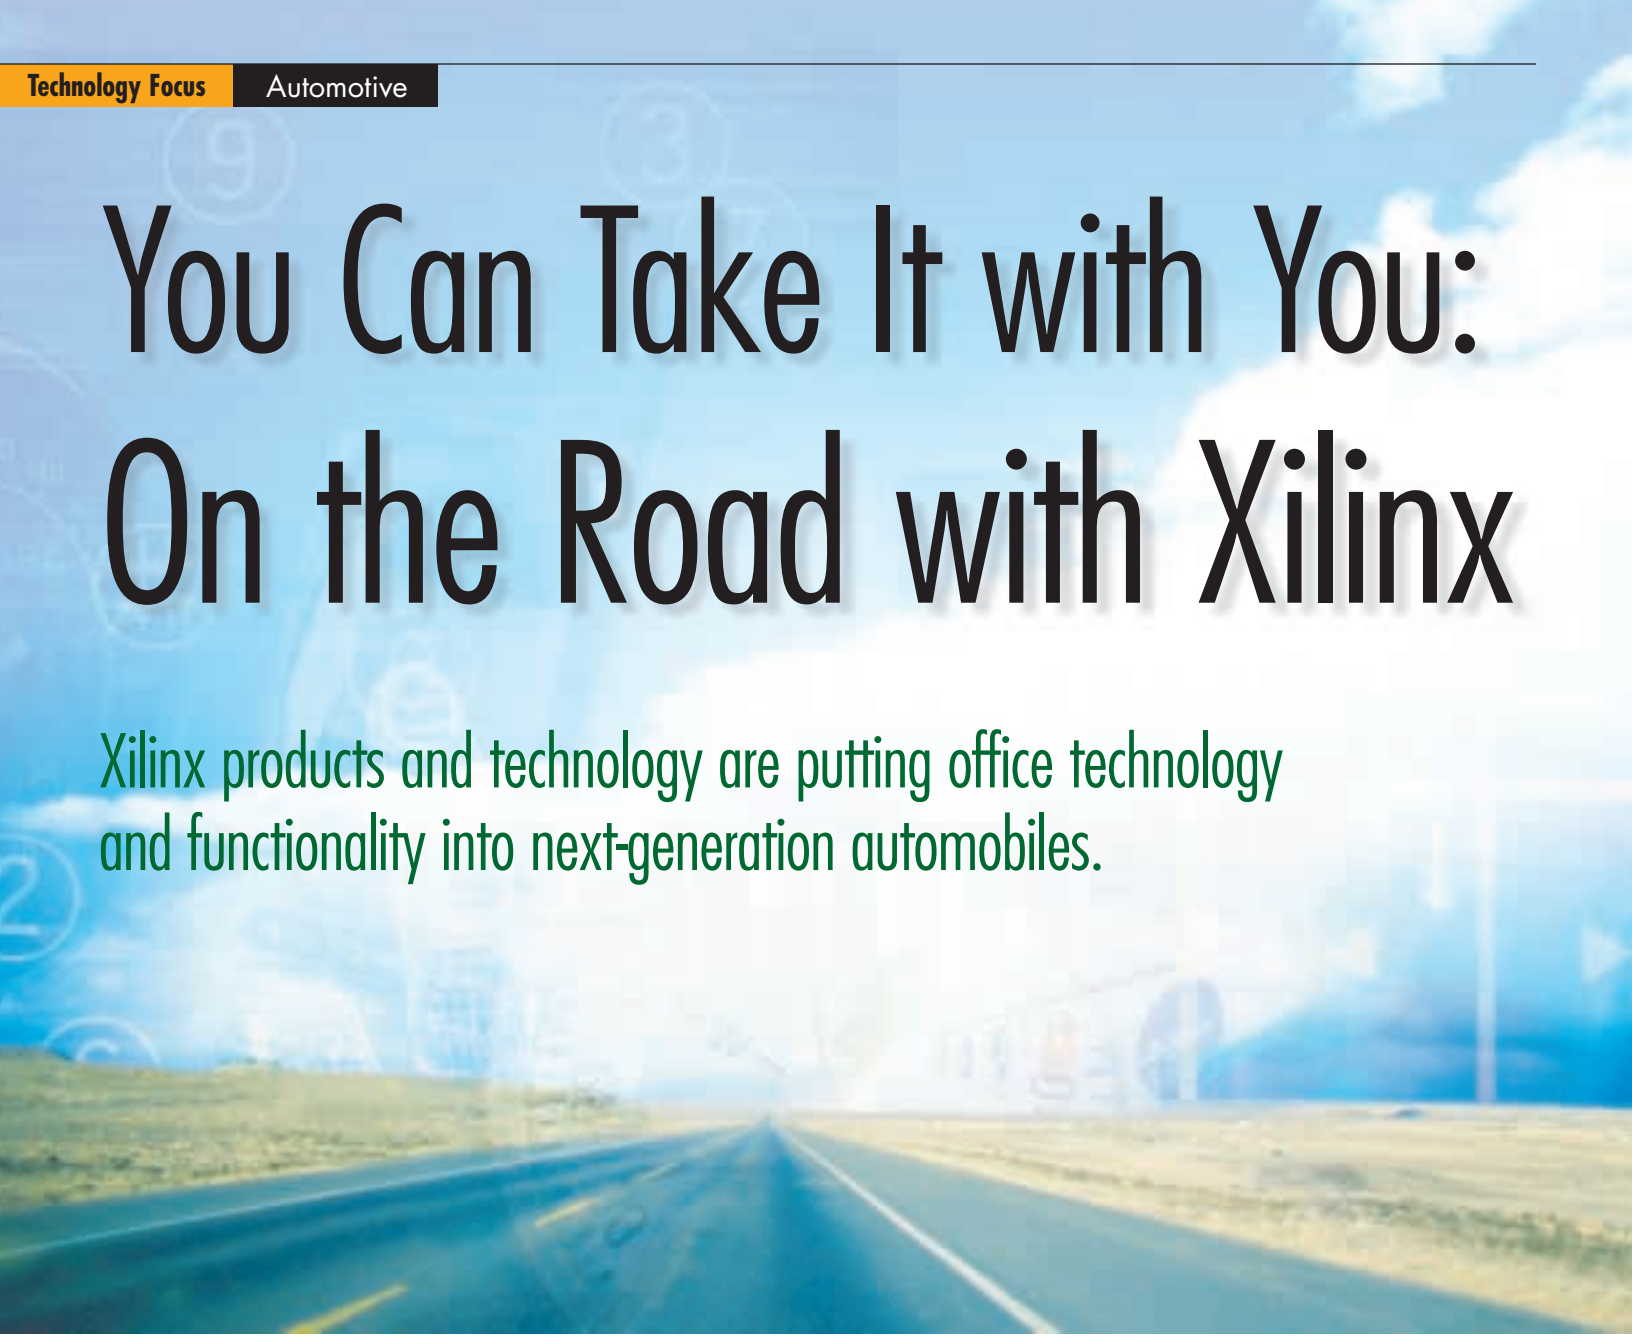

by Karen Parnell, Manager
Automotive Product Marketing
Xilinx, Inc.
karen.parnell@xilinx.com

Today's consumers demand the comforts of home and the productivity of the office when they drive. They also want the latest in information and safety systems for their cars. The resulting convergence requires interface standards and protocols to interconnect and interoperate. These standards and protocols are still emerging, and will get even more complicated before being standardized. The result is a manufacturing dilemma. Should manufacturers make an educated guess now as to which standard will prevail, and risk making the wrong choice? Or should they wait for the standards to be fixed, and risk getting left behind?

Xilinx has the answer. With Xilinx IQ programmable logic solutions for automotive applications, manufacturers can win the time-to-market battle risk free by using the ability to reconfigure their products to accommodate any standard – past, present, or future. Reconfigurable logic in production allows manufacturers to reconfigure the units in-car to implement new hardware-based features. Xilinx high-volume IQ FPGA and CPLD devices are cost-effective solutions that retain the traditional PLD time to market advantage. As a result, many leading telematics and infotainment product manufacturers use Xilinx programmable logic devices. These manufacturers recognize the added flexibility and time to market benefits that programmable logic solutions provide.

Driving Your Office

Information Age professionals need to be able to work from wherever they happen to be – from home, from hotel rooms, from airports, from 36,000 feet over the Atlantic. They need to take their offices with them.

The laptop computer was a critical step in the development of offices-on-the-go, but now workers are demanding even more. One result has been the in-car office. The automobile is no longer just a way to get from Point A to Point B. Now it can include GPS navigation with intelligent route finding, integrated PDA functions, built-in mobile communications, entertainment systems, and even in-dash personal computers.

Automotive electronic designers face the same challenges as designers of consumer

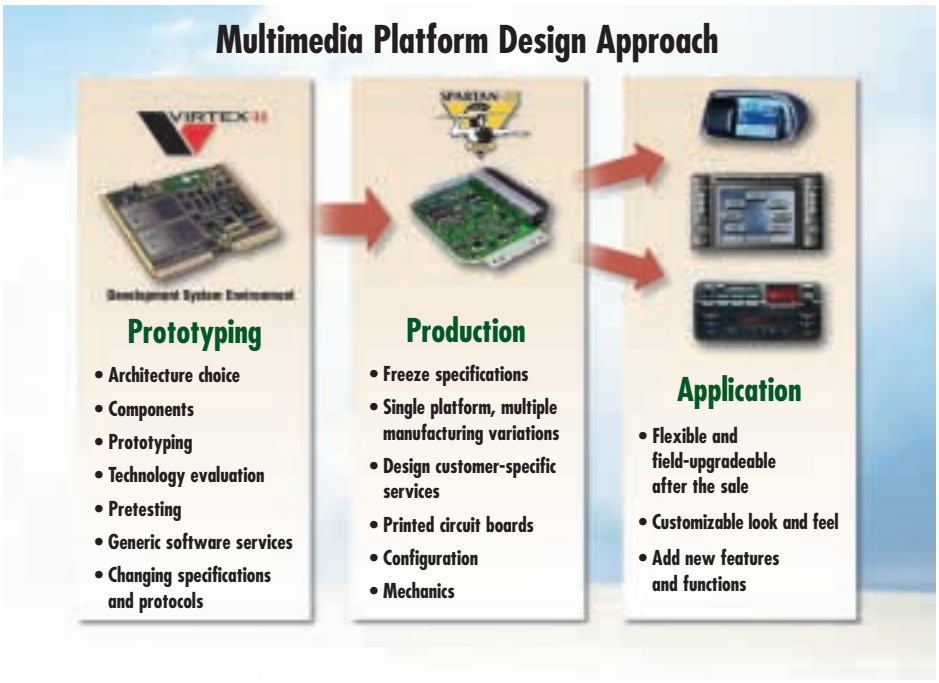

Figure 1 - Automotive multimedia platform design approach

equipment. They must grapple with time to market pressures, cost restrictions, and a profusion of new standards and protocols. They also face added challenges like a constrained design area, and the need for a simple user interface so drivers can keep their eyes on the road.

As one example of the new technology, the latest BMW 7 Series automobiles allow the driver to operate the car's heating and air conditioning, entertainment systems, wireless communications devices, navigation instruments, and PC applications, all on a single display, by means of a joy-stick type of device.

Xilinx programmable logic devices enable automotive designers to provide their customers with a more informative, productive, safe, and entertaining in-car environment. By using the flexibility, time to market advantage, and after-sale reconfigurability of Xilinx devices, manufacturers can be first to market and best in class long after the sale.

The Automotive Multimedia Platform

As more information and entertainment systems are added to automobiles, there's an inevitable conflict among the bewildering array of standards and protocols being tested, including Bluetooth™, BlueCAN, MOST (media-oriented system transport), FireWire™, CAN (controller area network), TTP (time triggered protocol), and Flex Ray™ technologies.

Designers of in-car multimedia systems may soon have to provide for traffic information systems, Internet and Web access, electronic game consoles, MPEG music download capability, digital radio reception, and mobile commerce services.

Designers must ensure that the car's multimedia system can communicate with the automobile's other devices. For example, the car should automatically detect a new mobile phone and connect it seamlessly to the car's communications network. This demand for automatic connectivity may soon extend to PDAs, portable PCs, MP3 players, and other personal portable electronic equipment.

Figure 2 - Acunia automotive development platform with Xilinx Spartan-II FPGA and CoolRunner™ CPLD onboard

As in-car and consumer functions converge – and new automotive and consumer standards and protocols emerge – designers have begun to prototype multimedia platforms that can provide as much, or as little, functionality as required. The best way to create such multimedia platforms is to design reconfigurable hardware. Engineers can program reconfigurable hardware late in the production flow to provide custom functionality on a standard hardware platform. At the same time, they can configure the hardware to accommodate new standards.

Figure 1 shows the automotive multimedia platform design approach. This concept allows upgrades throughout the life of the car. Designers can implement these upgrades remotely by means of wireless communications and Internet connectivity. Xilinx Internet Reconfigurable Logic (IRL™) technology makes this remote upgrading process possible.

With Xilinx IRL technology, designers can upgrade, modify, and fix systems long after the car leaves the dealership. For example, an engineer could remotely add a new MP3 player to an in-car multimedia system, or upgrade the system with the latest protocol. This same technology can even disable the multimedia unit if it is stolen, then re-enable it when it is returned to the rightful owner.

Multimedia System Design Flow Using FPGAs

The multimedia platform should ideally be based on one Human Machine Interface (HMI). This HMI allows the user to access all functions by means of a touch screen. If the functions are implemented in software and reconfigurable hardware, the manufacturer or dealer can upgrade the system even after the buyer has driven the car away.

The new way of developing in-car systems is to prototype with FPGAs in a generic development environment. The designer can develop the elements quickly and easily without fixing specifications. This initial prototyping phase can be realized using Virtex™-II Platform FPGAs. At this early stage, the designer can try, test, and debug the different standards, protocols, and functions by using the headroom provided by a large FPGA.

As the design firms up, the specifications are frozen. The designer can then select the specific standard, protocol, or function from the many tested. This move from prototype to production enables the designer to optimize the design and fit it into smaller, lower cost FPGAs, such as Spartan™-IIE programmable logic devices. This migration from Virtex-II™ to Spartan-IIE FPGAs still allows for future system upgradability via IRL connectivity.

Once the device is in production, the engineer can use the FPGA as an aid in total printed circuit board testing with JTAG techniques. If necessary, the engineer can also tweak and enhance the design even at this late stage.

Recognizing the need for intelligent prototyping platforms for use by in-car multimedia designers, leading next-generation telematics technology providers such as Acunia have produced the Xingu™ telematics platform shown in Figure 2.

Conclusion

The final stage is the look and feel of the product. Each car manufacturer can customize the product to fit into a specific dashboard (or fascia). All of the production multimedia units are built up around the standard FPGA-based platform. The designer can program this standard platform with its personality late in the production flow to accommodate last-minute design changes or end-user preferences.

For more information, visit these websites:

Xilinx Automotive IQ Solutions: *www.xilinx.com/automotive/*

Acunia: *www.acunia.com*

BMW 7 Series Sedan: *www.bmw.com/bmwe/products/automobiles/7er/sedan/*

Xilinx IQ Solutions – Creating Automotive Intelligence

To address the needs of automotive telematics designers, Xilinx has created a new family of devices with an extended industrial temperature range option. This new "IQ" family consists of existing Xilinx industrial grade (I) FPGAs and CPLDs with the addition of a new extended temperature grade (Q), available for selected devices. The new IQ product grade (-40°C to +125°C ambient for CPLDs and junction for FPGAs) is ideal for automotive and industrial applications. See Table 1.

The wide range of device density and package combinations enable you to deliver cost effective, high performance, flexible solutions to meet your most extreme application needs. See Table 2.

Temperature Grade/Range °C			
Products Group	C	I	Q
FPGA	T_J = 0 to +85	T_J = -40 to +100	T_J = -40 to +125
CPLD	T_A = 0 to +70	T_A = -40 to +85	T_A = -40 to +125

Table 1 - IQ temperature range

With Xilinx IQ devices, you can design flexibility into your application and get your product to market faster than ever before. Because many new standards are evolving – such as the MOST (media-oriented system transport) and FlexRay™ in-car bussing protocols – you need the flexibility to quickly modify your designs at any time. With our Internet Reconfigurable Logic (IRL™) capability, you can remotely and automatically modify your designs, in the field, after your product has left the factory, and even after the sale.

Device Family	Availability
Spartan-XL (3.3V)	NOW - XCS05XL, XCS10XL, XCS20XL, XCS30XL, XCS40XL
XC9500XL (3.3V)	NOW - XC9536XL, XC9572XL
CoolRunner XPLA3 (3.3V)	Q3 - 02
Spartan-II (2.5V)	Q3 - 02
CoolRunner-II (1.8V)	Q1 - 03
Spartan-IIE (1.8V)	Q1 - 03

Table 2 - Device availability schedule

Q Grade Ordering Information

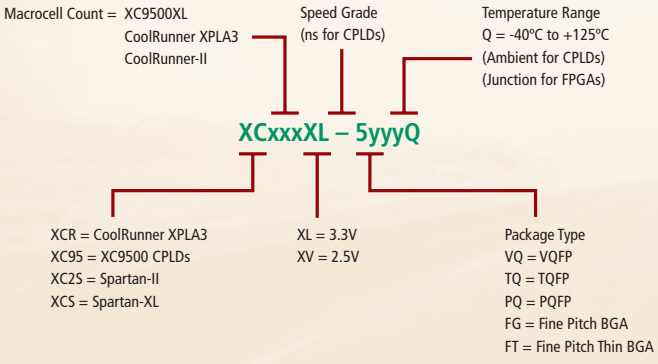

By combining our latest IQ programmable logic devices with our solutions infrastructure of high productivity software, IP cores, design services, and customer education, you can develop advanced, highly flexible products, faster than ever before.

For more information, full product selector guide, data sheets, white paper, and more, go to *www.xilinx.com/automotive/.*

How to Build an Efficient FIR Filter Using System Generator

Xilinx System Generator 2.2 software enables high-level modeling and implementation of DSP systems in Xilinx Platform FPGAs to outperform traditional DSP processors.

by Jim Hwang, Senior Manager
DSP Software and Design Methodologies
Xilinx, Inc.
jim.hwang@xilinx.com

Platform FPGAs have become key components for implementing high-performance digital signal processing (DSP) systems, especially in digital communications, video, and image processing applications. FPGAs have memory bandwidth that far exceeds that of microprocessors and DSP processors running at clock rates two to ten times faster, and unlike processors, Platform FPGAs possess the ability to implement highly parallel custom signal processing architectures.

One of the main impediments to wider adoption of FPGAs for signal processing has been the relative unfamiliarity with FPGA technology within the DSP community. What is needed are tools that speak the language of signal processing engineers – for example, MATLAB™ language and Simulink™ design tools from The MathWorks, Inc. – while also supporting FPGA designers who already know how to achieve the highest performance circuit using the least number of FPGA resources.

The new Xilinx System Generator 2.2 software meets the needs of both DSP engineers and FPGA designers. System Generator enables high-level modeling and implementation of DSP systems in the Virtex-II Pro™, Virtex™-II, Virtex, and Spartan™-II families of FPGAs. New in the 2.2 release is an increased capability to generate hardware that approaches the efficiency of hand-crafted modules, both in terms of performance and resource usage. In this article, we describe how System Generator can be used to create a custom FIR (finite impulse response) filter, an operation that lies at the heart of most signal processing systems. We demonstrate that in addition to accessing high-level Xilinx LogiCORE™ modules, you can use your own FPGA knowledge to advantage in customizing a filter data path to minimize resources while achieving high performance. You will see how System Generator provides a wide range of options for implementing signal processing systems.

Xilinx System Generator

The Xilinx System Generator is a state-of-the-art "on-ramp" that allows you to move a DSP algorithm into a Xilinx FPGA quickly and easily. System Generator extends the capabilities of the Simulink system-level simulation environment with bit and cycle-true modeling of an FPGA circuit. System Generator simultaneously provides access to key features in the FPGA fabric, including SRL16E shift register logic, distributed and block memory, and embedded multipliers.

System Generator includes a Simulink library of functional blocks for building DSP, arithmetic, and digital logic circuits. These polymorphic blocks compute their output types based on their inputs, although alternatively, you can specify their quantized output types explicitly. You can combine Xilinx blocks with MATLAB and Simulink blocks to create a realistic test bench and to analyze data computed by your model. The high level of abstraction provided by System Generator greatly simplifies algorithm development and verification.

In addition to a system-level modeling library, System Generator includes a code generator that automatically generates a synthesizable VHDL netlist from your Simulink model. This netlist includes IP (intellectual property) blocks that have been carefully designed for high performance and density in Xilinx FPGAs. System Generator also creates project and constraint files to assist implementation using the Xilinx Foundation™ ISE 4.2i and soon-to-be-released ISE 5.1i tools, as well as the major synthesis tools.

Building an FIR Filter

As a simple but instructive example, let's consider how System Generator can be used to create a parametric finite impulse response (FIR) filter. An N-tap FIR filter is defined by its impulse response, a length N sequence of filter coefficients: $h_0, h_1, \ldots, h_{N-1}$. If $x_0, x_1, x_2, \ldots$, is a sequence of input values, where by convention, we define $x_i = 0$ for $i < 0$, the filter output sequence $y_0, y_1, y_2, \ldots$ is defined by the convolution sum

$$y_n = \sum_{i=0}^{N-1} h_i x_{n-i}$$

That is, the filter output at time n is computed by accumulating a sum of products of the filter coefficients with the N most recent input samples.

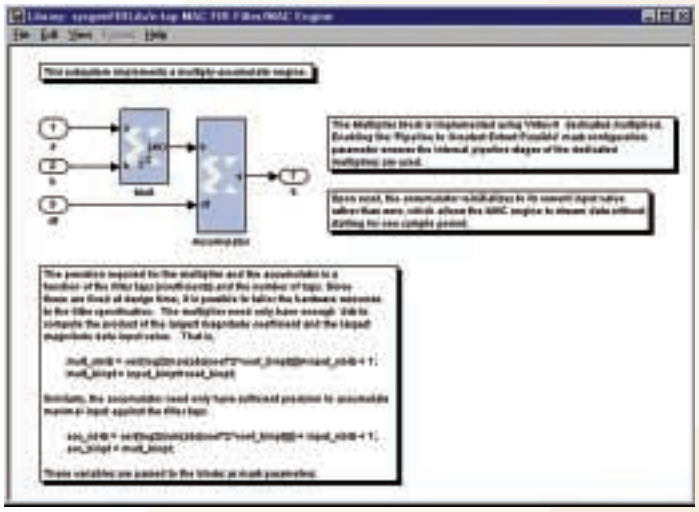

Figure 1 - Multiply-accumulate engine

In practice, all numbers must be represented with a finite number of bits. With a traditional processor, numeric data are typically represented as 8, 16, or 32-bit integers, or in a floating-point representation. In contrast, in an FPGA, we have no such word length limitations. We can create a custom data path processor having an arithmetic precision tailored to the application. System Generator supports this capability by providing an arbitrary precision fixed-point data type. Each block allows us to specify its output precision and the policy for handling quantization and overflow. We can model the system in the Simulink environment under a number of scenarios, and analyze the data to ensure exactly the right precision for the application.

Mapping the Algorithm onto an FPGA

You can implement an FIR filter in an FPGA in many ways. A versatile approach that maps well onto an FPGA employs a multiply-accumulate (MAC) engine to compute the sum of products. As shown in Figure 1, a MAC unit is easily constructed in System Generator using the multiplier and accumulator blocks. The multiplier can be implemented either in the logic fabric or, for Virtex-II family FPGAs, using dedicated 18-bit x18-bit embedded multipliers. System Generator ensures the underlying IP core provides an efficient implementation. Note that upon reset, the accumulator reinitializes to its current input value rather than zero, to avoid a one-clock cycle stall.

Customizing the Data Path

Although Simulink provides a graphical block editor, System Generator should not be mistaken for a "schematic capture" tool. System Generator models are fully customizable and executable in Simulink without recompilation. (In fact, some blocks can be reconfigured during simulation.) Xilinx blocks support Simulink's data type propagation capability, and provide extensive error checking on their parameters and usage. Because System Generator is seamlessly integrated with the Simulink tool suite, you can customize Xilinx blocks in ways that are impossible in schematic and other visual tools.

For example, in our System Generator model we can specify the arithmetic precision of the blocks in the data path using MATLAB expressions, making it possible to minimize the hardware used, and still avoid the possibility of overflow. For a filter with k-bit coefficients and m-bit input, we know that the output

$$|y_n| = \left|\sum_{i=0}^{N-1} h_i x_{n-i}\right| \le \sum_{i=0}^{N-1} |h_i x_{n-i}|$$
$$\le \sum_{i=0}^{N-1} |h_i 2^m| = 2^m \sum_{i=0}^{N-1} |h_i|$$

so the accumulator requires no more than $m + \lceil \log_2 \sum |h_i| \rceil$ bits. This is in practice considerably fewer than the $m + k + \lceil \log_2 N \rceil$ bits implied by the input and coefficient precision. The desired accumulator width is, of course, readily expressed in the MATLAB language. (In Figure 1, the binary point in the fixed-point data is also taken into account.) When you know the input values have limited dynamic range, it may be possible to further tighten the bond. The accumulator width we just derived is a function of the input precision and the MATLAB array that stores the filter coefficients, so if you want to change the input precision or change the filter itself, the same model will "right-size" the data path without any modification. The implications of this ability to use the MATLAB interpreter to customize your System Generator model should not be underestimated – it is unrivaled by any other design flow.

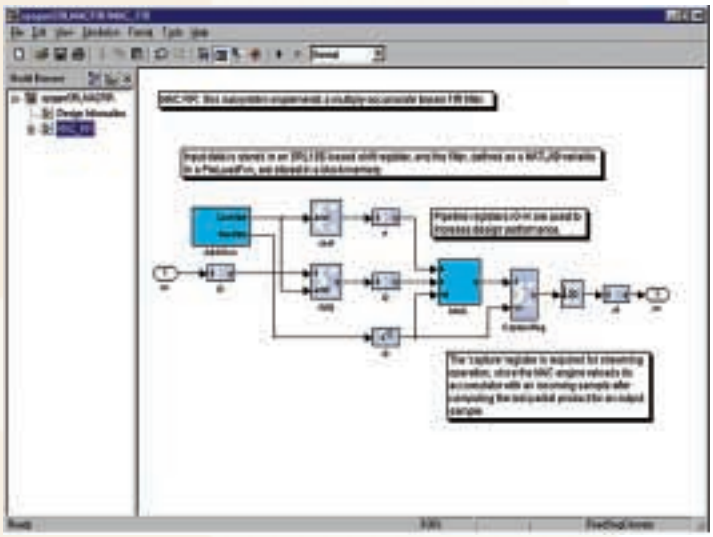

Figure 2 - A multiply-accumulate (MAC)-based FIR filter

Completing the FIR Filter Data Path

As shown in Figure 2, the input data buffer is implemented as an SRL16E-based addressable shift register, and the filter coefficient buffer is implemented as a block memory. (Both are supplied as handcrafted Xilinx LogiCORE™ algorithms.) By storing the filter coefficients in reverse order in memory, the same address counter can be used to drive both buffers.

Because there are N multiply-accumulate operations per input sample, the filter must run internally at N times the data rate to supply a continuous data stream. The capture register on the output of the MAC is used to latch the accumulated sum of products, and its output is down-sampled by N to match the input data rate. This simple filter architecture is quite compact and efficient. A 64-tap non-symmetric filter with 12-bit coefficients and data requires only 110 slices in a Virtex-II XC2V250-6 FPGA, and it runs at 195 MHz, or 3 Msps (ISE 4.2i, production speeds files 1.96).

System Generator provides many ways to tailor the implementation of a design, and changes are tracked automatically and transparently. In the FIR filter, you might choose to use dedicated embedded multipliers for the MAC engine, coupled with a block memory (BRAM) for the coefficient data. (It is natural to do so, because these resources are juxtaposed in the FPGA – although both can also be implemented efficiently in the logic fabric.) Filter throughput can be increased significantly by employing additional MAC engines. System Generator makes this a straightforward extension of our example. When you switch the implementation strategy, latencies are automatically adjusted as necessary in the System Generator model to match the hardware behavior.

Other Implementation Options

Of course, you do not have to build filters from scratch. There is an FIR filter block in the System Generator library that employs distributed arithmetic (DA) to map the computation into the FPGA. This is often the best way to implement a filter, because the underlying IP core supports many efficient architectures, including fully parallel DA, fully serial DA, half band, polyphase interpolators and decimators, and more.

At the same time, there are occasions when you will want to design a custom filter, for example, to exploit symmetries, take advantage of the embedded Virtex-II multipliers, or to build adaptive filters. While providing a high-level simulation and modeling capability, System Generator also allows you to apply your knowledge of the FPGA to map your algorithm onto specific resources. As we saw in our example, you can often parameterize the design using MATLAB functions, so that changing parameters automatically gives you an appropriately customized implementation.

Conclusion

With Virtex/Virtex-II family FPGAs, handcrafted IP LogiCORE algorithms, and System Generator software, Xilinx is rapidly changing the way people think about DSP. Designing a custom DSP data path processor in an FPGA has never been easier.

This article barely scratches the surface of the capabilities of the Xilinx System Generator. The System Generator 2.2 release includes a number of useful annotated demonstration designs, including a QAM demodulator, concatenated codec for digital video broadcast, discrete wavelet transform, and several extensions of the FIR filter discussed in this article, to name but a few. These demonstration designs can be used as-is, or provide a starting point for you to create your own applications. A full-featured, free, 90-day evaluation version of System Generator 2.2 is available for download from the Xilinx website at *www.xilinx.com/systemgenerator_dsp.*

For more information on MATLAB and Simulink software, visit *www.mathworks.com.* Σ

Use Virtex FPGAs and Xilinx Software Tools to Reduce Line Echoes in PSTN and Packet Networks

By using Xilinx software and Virtex FPGAs with DSP capabilities, eInfochips Ltd. brought its line-echo canceller to market faster and better than a traditional DSP chip solution.

by Nilesh Ranpura
Project Manager, ASIC Group
eInfochips ltd
nilesh@einfochips.com

Line echoes occur in public switched telephone networks (PSTN) at so-called "hybrid" points, where 2-wire and 4-wire copper lines are interconnected at the central office. Due to electric current leakage in the hybrid, a part of the signal energy is reflected back to the source of the signal, which causes an echo on the phone line. Likewise, in packet networks, the delays associated with processing the voice stream can also produce an echo on the line.

At eInfochips Ltd., we have developed an echo canceller using Virtex™ FPGAs. We chose Virtex devices because they give us high-speed performance and an edge in area-to-cost ratio over traditional ASIC solutions.

Echo cancellers are advanced signal processing algorithms running on DSPs or fixed-function ASICs that remove line echoes in voice networks. An adaptive FIR filter is used to predict the echo from the history of the transmitted signal.

Echo cancellers can be integrated into a phone system or packet network along with other voice-processing functions, such as subtracter, double-talk detector, nonlinear processor, narrow band signal detector, PCM encoder/decoder, offset null filter, and controller/processor interface, as shown in Figure 1.

Virtex FPGAs: The Obvious Choice

We started our project with the goal of creating a single-channel echo canceller capable of handling an echo tail length of 128 ms. We had the option of using a DSP chip for algorithm implementation, but during the course of development, we discovered we could deploy a Virtex FPGA-based solution with DSP capabilities to achieve echo cancellation of up to 128 ms.

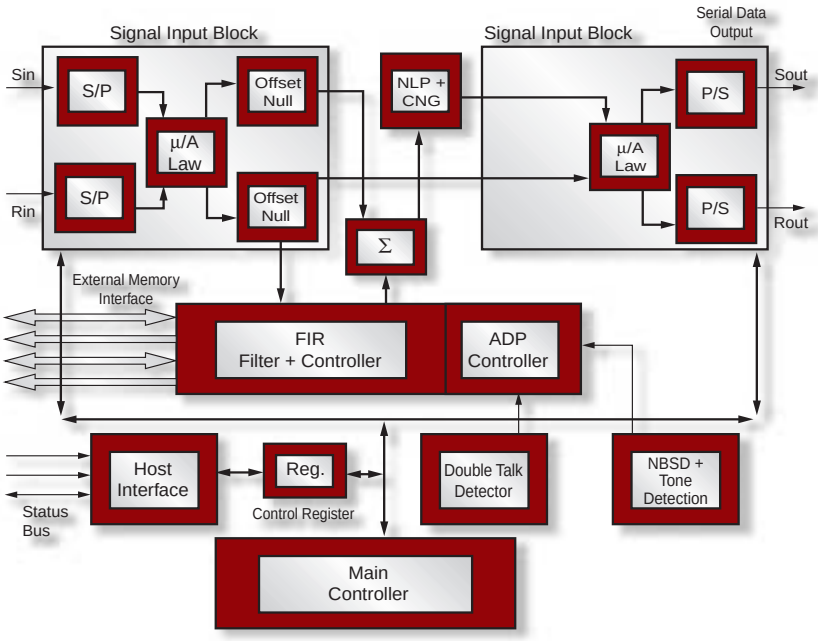

Figure 1 - Echo canceller logic block diagram

The choice was obvious. Not only do Virtex FPGAs have high-speed digital signal processing functions, but they also include extensive flexibility in using block RAM and LUTs as RAM and ROM.

We implemented the following major blocks as a part of the data path design in the FPGA:

- Adaptive FIR filter for estimating the echo signal
- Programmable double-talk detection threshold
- Nonlinear processor with adaptive suppression threshold for removing the residual echo signals
- Offset null filtering of PCM channels
- Narrow band signal detector conforming to ITU-T G.165 requirements for preventing the divergence of the adaptive filter coefficients
- Selectable μ/A law companding, PCM coding, and sign magnitude.

The core was verified using Industry Standard Vectors compliant with the G.165 standard.

Conclusion

Xilinx Virtex FPGAs support high density and speed with extensive EDA tools support on a wide spectrum of features. The combined solution of Virtex hardware (in this case, the Virtex XCV800 FPGA), associated Xilinx Foundation™ Series software, and the Xilinx ChipScope™ integrated logic analysis tool saved us invaluable time to market.

As a result, of implementing a total Xilinx Virtex solution, we were able to be among the first to bring our echo canceller solution into the emerging areas of Internet gateway, ATM gateway, XDSL technology, cable modem, and wireless telephony. For more information visit *www.einfochips.com* or write *info@einfochips.com.*

SystemIO Solution Expands with Bandwidth Demands

The Xilinx Platform FPGA SystemIO solutions solve the emerging interface challenges for high-speed system interconnectivity.

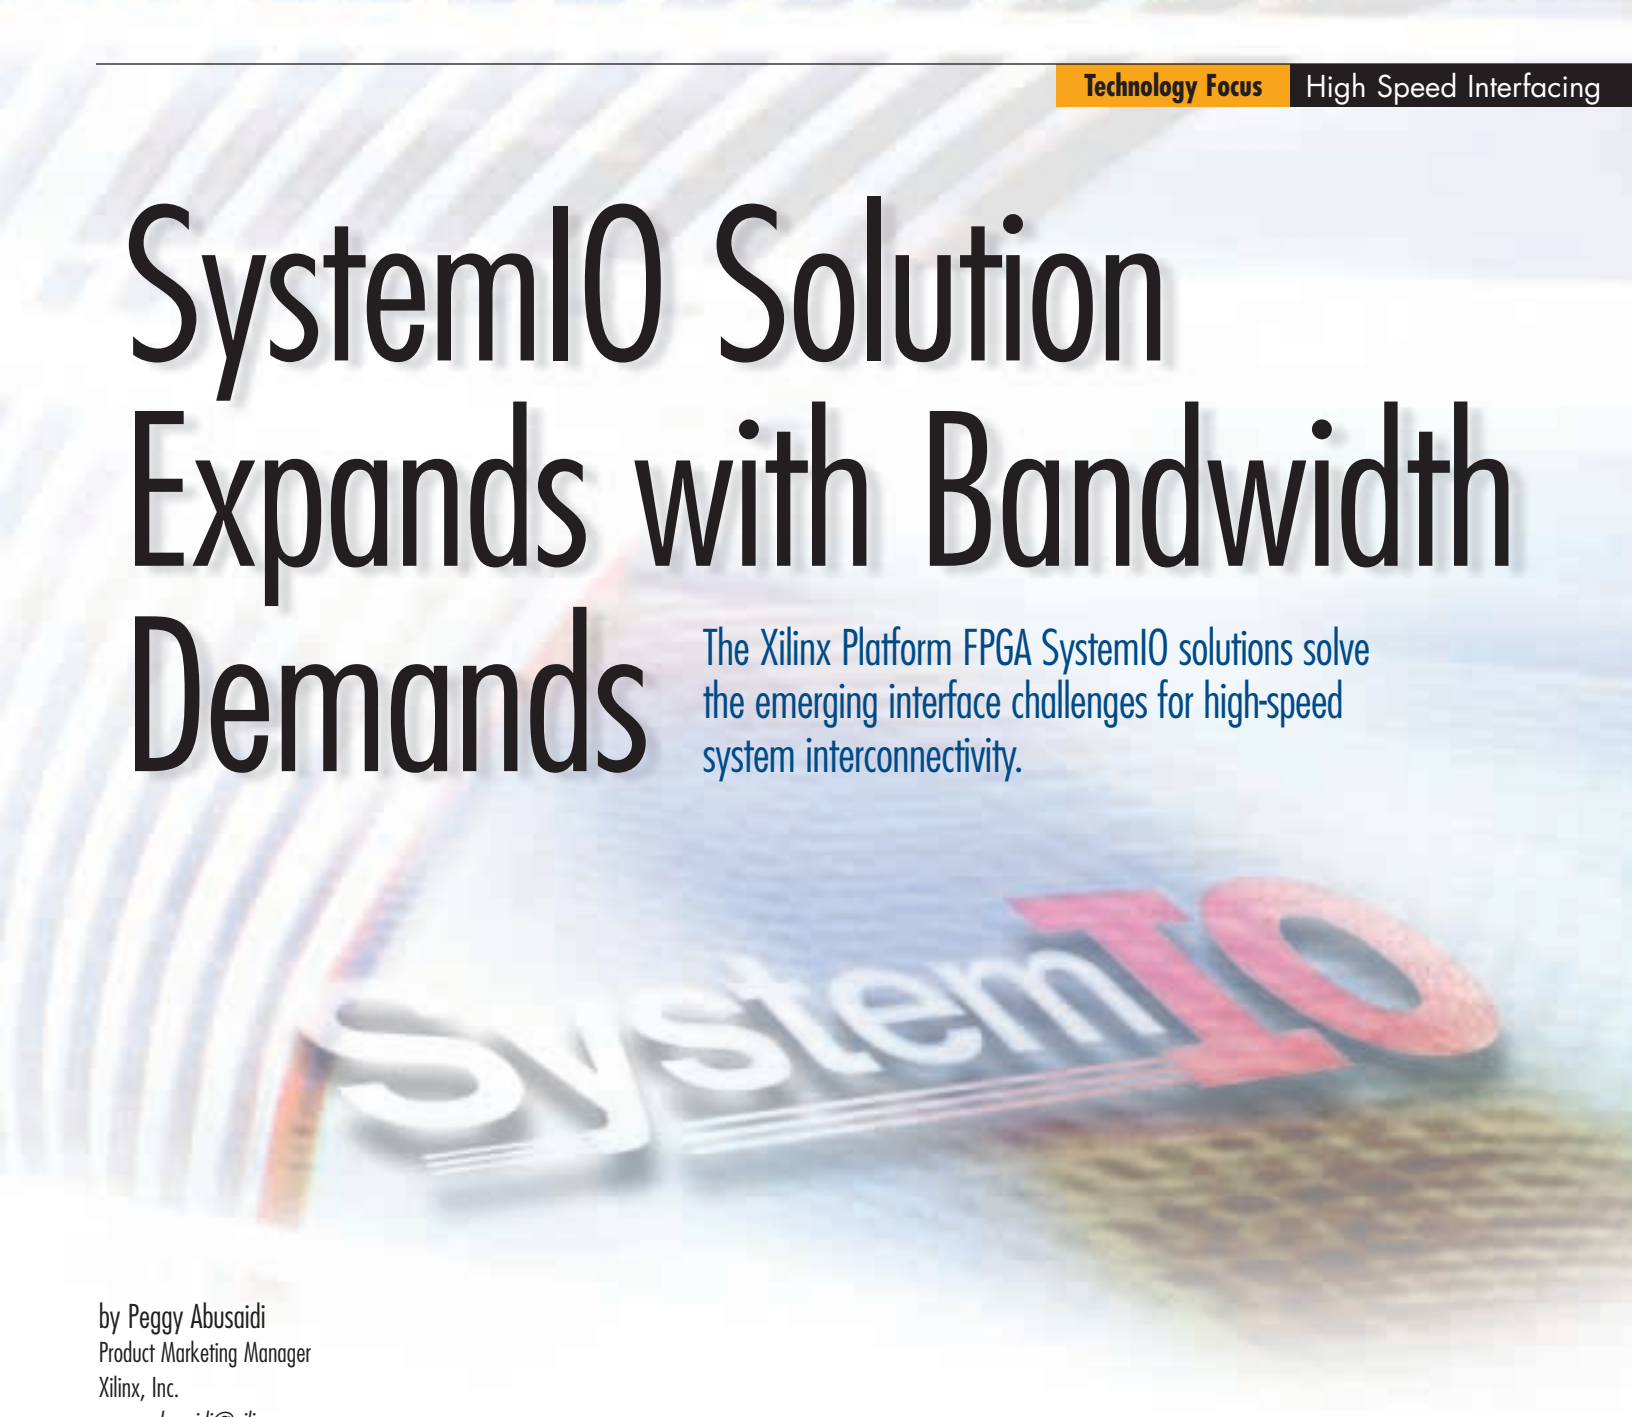

by Peggy Abusaidi
Product Marketing Manager
Xilinx, Inc.
peggy.abusaidi@xilinx.com

Traditional system interfaces, such as the existing PCI and VME parallel bus schemes, cannot keep up with today's ever-growing bandwidth requirements. Therefore, RapidIO™, HyperTransport™, InfiniBand™, PCI Express™ (aka 3GIO), and other high-speed bus interconnect technologies have been developed to open the I/O bottleneck. Figure 1 illustrates the variety of Internet applications that drive the need for more bandwidth.

How do you choose the right I/O interface standards for your systems now – and how do you keep current with the evolving I/O standards in the future? Moreover, how do you meet your time-to-market and cost goals with minimal risk when the interface standards keep changing?

As shown in Figure 2, Xilinx provides a comprehensive list of Platform FPGA SystemIO Solutions to help you solve the I/O bottleneck problem. SystemIO gives you the ability to implement designs using any of the latest I/O standards. By using a Xilinx Platform FPGA with a SystemIO solution, you can accelerate your time to market and be assured that your designs will remain current even as I/O standards evolve and specifications change. In addition, SystemIO solutions enable you to "future-proof" your products by allowing to you update the code in the FPGA to bring your product up to compliance – even after the product has been deployed in the field.

Solving the Bandwidth Problem

Today, most computer, embedded processing, and telecommunications equipment are parallel bus-oriented. However, as device performance increases, and demand for bandwidth rises, these multi-drop bus structures are reaching their performance limits. In most cases, these buses can only process one module at a time. This results in idle time for the subsystems waiting for bus access, thus decreasing overall system performance.

The industry's response has been the development of a host of new interconnection schemes, some with data transfer rates exceeding 10 Gbps. The new proposed standards have backers like Intel, AMD, Microsoft, Hewlett-Packard (Compaq), Dell, Xilinx, and Sun Microsystems. In addition to the aforementioned RapidIO, PCI Express, HyperTransport, and InfiniBand technologies, other standards already in use, or evolving toward higher levels of performance, include 10-gigabit attachment unit interface (XAUI), POS-PHY Level 4, Gigabit Ethernet, and various optical ATM formats like OC-12, OC-48, and OC-192. These new or evolving standards cover all architectural environments, such as motherboards in chip-to-chip communication; backplanes for communications between subsystems; and storage, local, and wide area networks (SANs, LANs, and WANs).

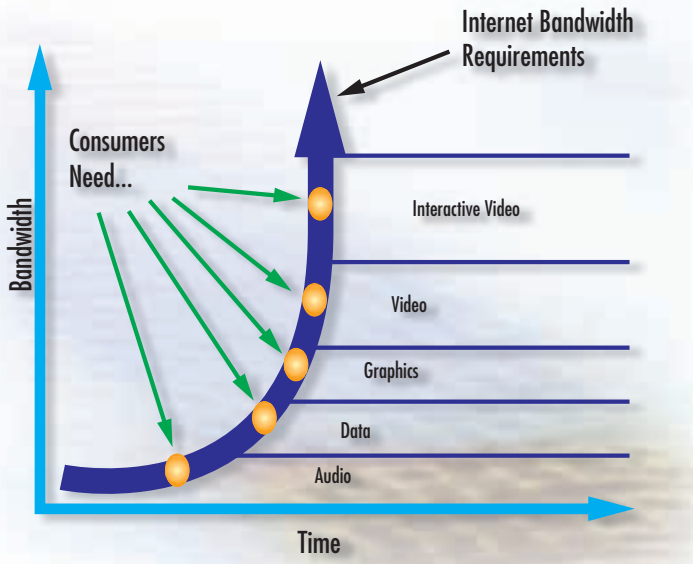

Figure 1- Internet bandwidth trend

Parallel bus standards are divided into two general categories:

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

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

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

The SystemIO Solution

The Xilinx Platform FPGA SystemIO Solution uses the unique Virtex™-II SelectI/O™-Ultra blocks to provide the fastest and most flexible electrical interfaces available. Each user I/O pin is individually programmable for 19 single-ended I/O standards or six differential I/O standards, including LVDS, SSTL, HSTL, and GTL+. The SystemIO solution is capable of delivering 840 Mbps double data rate (DDR) and 644 MHz single data rate (SDR) LVDS performance. Furthermore, any two I/O pins can be used as a differential pair, providing maximum board layout flexibility and reducing overall system cost by saving both component and board layer costs.

Using Virtex-II 840 Mbps LVDS performance and an abundance of memory and logic resources, you can easily create designs using 1GE (Gigabit Ethernet) and 10GE MAC, PCI and PCI-X, POS-PHY Levels 3 and 4, RapidIO, HyperTransport, and Flexbus 4 protocols. Free reference designs for implementing interfaces, such as SFI-4, XSBI, XGMII, and CSIX are available from the Xilinx website. Table 1 shows the Platform FPGA SystemIO Solutions Summary, including both IP cores and reference designs.

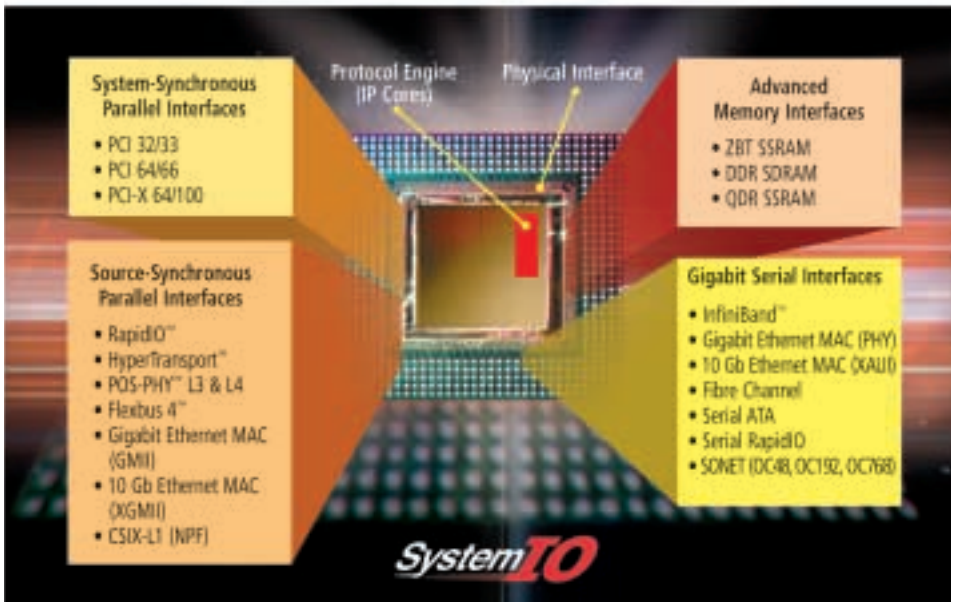

Figure 2 - Xilinx Platform FPGA SystemIO Solutions

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

With our recent introduction of the Virtex-II Pro™ Platform FPGA, we have incorporated Mindspeed's SkyRail™ architecture to embed up to 16 Rocket I/O™ 3.125 Gbps multi-gigabit transceivers (MGTs), the highest number of MGTs in a single programmable device. This breakthrough technology supports all the popular emerging serial I/O standards, including the PCI Express, InfiniBand, XAUI, and Fibre Channel protocols.

The combination of SkyRail architecture and SelectI/O-Ultra technology supplies a comprehensive list of popular parallel and serial interface IP cores in the overall SystemIO solution. Additionally, both 1GE and 10GE MAC IP cores are available with both parallel and serial interface options.

Conclusion

Platform FPGA SystemIO solutions enable Xilinx to provide you with the ultimate connectivity platform to connect and bridge interface requirements for your next-generation data communication systems.

A terabit system can simultaneously contain multiple interface requirements, including both emerging serial I/O protocols, as well as established parallel I/O standards. The new Virtex-II Pro Platform FPGA is ideally suited to help you address interface challenges by providing Rocket I/O MGTs for serial connectivity and SelectI/O-Ultra technology for parallel connectivity.

Combining LogiCORE IP cores for protocols and data processing, along with design tools and third-party partnerships, Xilinx provides the complete system connectivity solution for today's – and tomorrow's – high speed designs.

IP Core	Standard Compliance	Aggregate Bandwidth	Performance	SelectI/O Bus	Availability
POS-PHY L3	OIF-SPI3-01.0 Saturn POS-PHY L3	2.48 Gbps	104 MHz	32b 3.3V CMOS	Now
POS-PHY L4	OIF-SPI4-02.0 Saturn POS-PHY L4	12.8 Gbps	800 Mbps per pair 400 MHz DDR	16b LVDS	Now
Flexbus 4	OIF-SPI4-01.0 AMCC Flexbus 4	12.8 Gbps	200 MHz	64b HSTL	Now
Flexbus 4 to POS-PHY L4 Bridge	OIF-SPI4-01.0 AMCC Flexbus 4 OIF-SPI4-02.0 Saturn POS-PHY L4	11.2 Gbps	200 MHz (Flexbus 4 side) 350 MHz DDR (PL4 side)	64b HSTL (Flexbus 4 side) 16b LVDS (PL4 side)	Now
1GE MAC w/ GMII	IEEE 802.3-2000	1 Gbps	125 MHz	8b 3.3V GMII	Now
1GE MAC w/ PCS-PMA	IEEE 802.3-2000	1.25 Gbps	1.25 Gbps	1 channel of Rocket I/O Transceiver	Now
10GE MAC w/ XGMII	Designed to IEEE P802.3ae draft 4.1	10 Gbps	156.25 MHz DDR	32b XGMII HSTL	Now
10GE MAC w/ XAUI	Designed to IEEE P802.3ae draft 4.1	12.5 Gbps	3.125 Gbps per channel	4 channels of 3.125 Gbps Rocket I/O Transceivers	Now
RapidIO PHY	RapidIO Interconnect Specification v1.1	8 Gbps	500 Mbps per pair 250 MHz DDR	8b LVDS	Now
PCI64	PCI Spec V2.3	264 -528 MBps	33/66 MHz	64b 5/3.3V PCI	Now
PCI32	PCI Spec V2.3	132 -264 MBps	33/66 MHz	32b 5/3.3V PCI	Now
PCI-X 64/100	PCI-X Spec V1.0a	800 MBps	66/100 MHz	64b 3.3V PCI-X	Now
HyperTransport Single-Ended Slave	HyperTransport V1.01a	6.4 Gbps	800 Mbps per pair 400 MHz DDR	8b HT I/O	Q3 '02
CSIX Reference Design	CSIX-L1	6.4 Gbps	200 MHz	32b HSTL	Now
XGMII Reference Design	Designed to IEEE P802.3ae draft 4.1	10 Gbps	156.25 MHz DDR	32b XGMII HSTL	Now

Table 1 - Summary of Platform FPGA SystemIO solutions

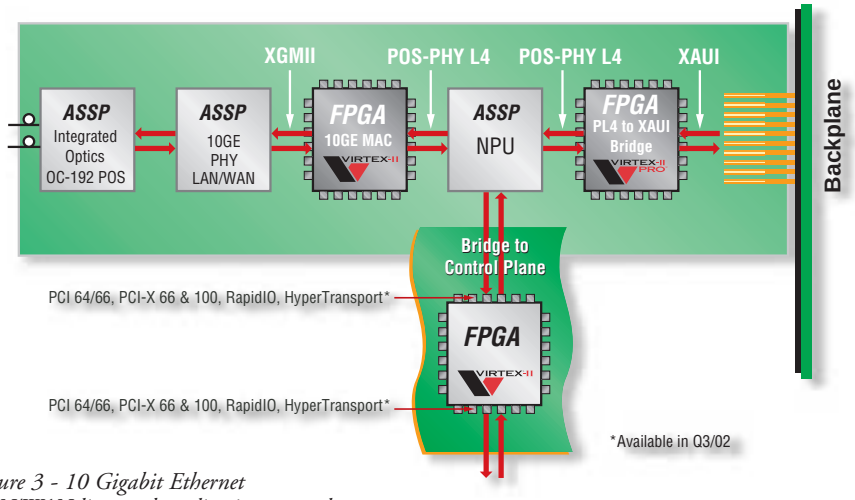

Figure 3 - 10 Gigabit Ethernet LAN/WAN line card application example

Develop MicroBlaze Applications with Three Flexible Hardware Evaluation Platforms

MicroBlaze Development Kits from Memec Design demonstrate the versatility of the MicroBlaze soft processor core in a variety of Xilinx FPGAs.

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

Running at 125 MHz, the Xilinx MicroBlaze™ 32-bit soft processor core is the industry's fastest soft processing solution. The MicroBlaze processor delivers a true 32-bit processor, critical for building complex systems for the networking, telecommunication, data communication, embedded, and consumer markets. The soft processor features RISC architecture with Harvard-style separate 32-bit instruction and data busses, which run at full speed to execute programs and access data from both on-chip and external memory. With 900 logic cells and 82 D-MIPS, the MicroBlaze processor meets the utilization, performance, and cost targets that most FPGA designers require.

To support the development of MicroBlaze-based applications, Memec Design has introduced three different MicroBlaze Development Kits (MDKs) to accelerate design, prototype, and evaluation cycles. The three MDKs allow you to quickly prototype your MicroBlaze processor and peripherals in real-world hardware environments. With functioning hardware platforms based on Xilinx FPGAs, you can easily verify design concepts, system interfaces, and real-time functionality.

Three Families – Three Kits

The Virtex™-II, Spartan™-IIE, and Spartan-II FPGA families are the logical targets for MicroBlaze applications. Therefore, Memec offers specially designed standalone system boards for a device from each of these families.

Additionally, each MDK offers an expansion module that contains common processor peripheral interfaces and memory, a prototype module for creating custom circuits, the Xilinx MicroBlaze processor license, software tools, and a power supply.

- The Virtex-II platform is based on the one million-gate XC2V1000 device (Figure 1). The system board includes 2M x 16 DDR SDRAM and a 16-bit LVDS Tx/Rx port.
- The Spartan-IIE system board is based on the 300K-gate XC2S300E device (Figure 2). This board offers a 10-bit LVDS Tx/Rx port and 82 general purpose I/Os.
- The Spartan-II kit uses the 200K-gate XC2S200 device (Figure 3), which resides on a 32-bit PCI interface, along with 2M x 32 SDRAM. The PCI interface presents an interesting configuration by allowing the MicroBlaze processor to sit on the backend of the PCI bus.

All three platforms include the P160 expansion slot, an ISP PROM, on-board power regulation, and other user support circuits.

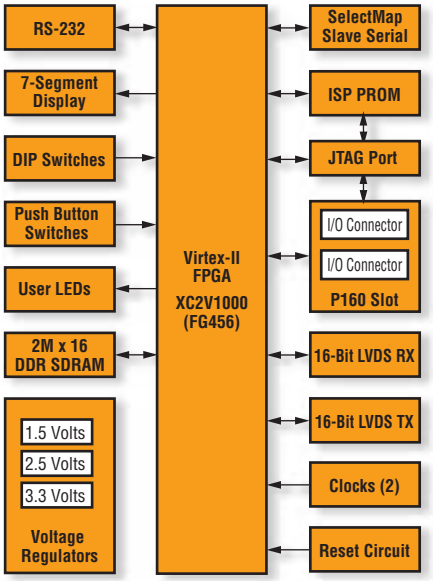

Figure 1 - Virtex-II system board block diagram

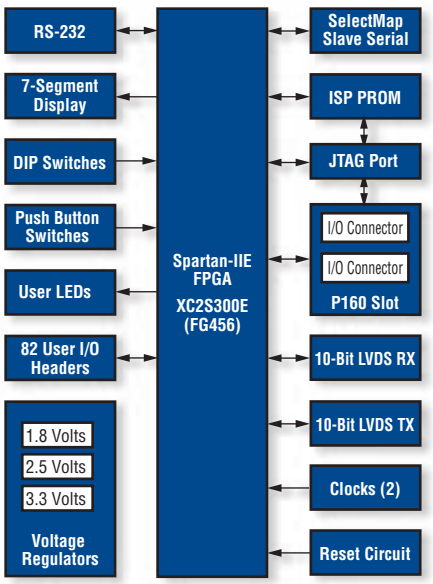

Figure 2 - Spartan-IIE system board block diagram

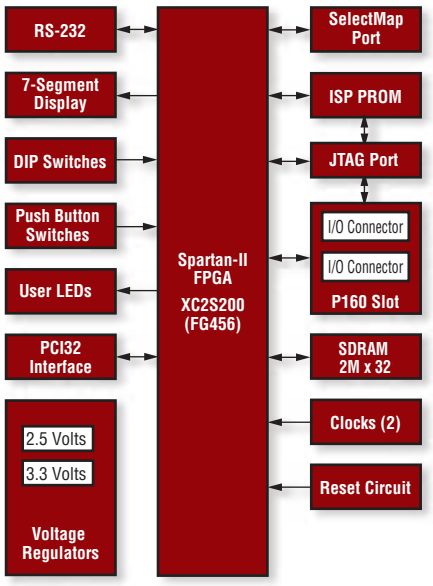

Figure 3 - Spartan-II system board block diagram

P160 Expansion

The P160 expansion slot provides the flexibility you need in a processor and IP development environment. The expansion slot allows you to easily plug custom peripheral modules into the system board and interface to the FPGA device.

The P160 slot supplies 110 user-defined I/O signals from the FPGA to the user application circuit on the expansion card. By adding a P160 module, you can tailor your development environment and ease hardware verification.

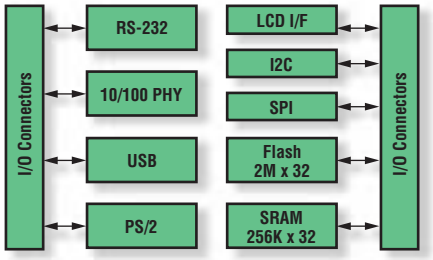

Figure 4 - P160 communications module block diagram

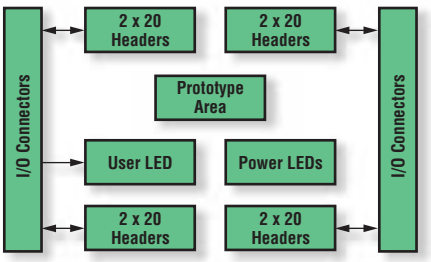

Figure 5 - P160 prototype module block diagram

Two P160 modules are included in the MDKs:

- The P160 communication module (Figure 4) provides interfaces for 10/100 Ethernet, RS-232, USB, PS/2 keyboard, LCD module, I^2C, and SPI connections. The communication module also includes 2M x 32 flash and 256K x 32 SRAM for off-chip memory expansion by way of the IBM CoreConnect™ on-chip peripheral bus architecture.
- The P160 prototype module (Figure 5) offers general purpose expansion headers for all 110 user I/O signals, as well as a prototype area for building custom circuits.

Conclusion

As MicroBlaze implementations move into the mainstream, you will need easy to use, flexible, low-cost hardware platforms to verify and validate your design concepts. Supporting Virtex-II, Spartan-IIE, and Spartan-II devices, each Memec Design MDK bundles the essential tools you need to explore MicroBlaze-based designs.

The Virtex-II MDK is priced at $795. The Spartan-IIE and Spartan-II MDKs are available for $695. Call 888-488-4133, ext. 235, or go to *www.insight-electronics.com/microblaze* for more information or to order a MicroBlaze Development Kit.

Xyron ZOTS IP Core Uses Virtex-II FPGA to Demonstrate RTOS Breakthrough

Xyron Semiconductor's zero overhead task switch technology gives a 40 MHz FPGA the real-time video rendering power of a GHz desktop microprocessor.

by William Dress
Director of Strategic Development
Xyron Semiconductor
wdress@xyronsemi.com

Xyron Semiconductor's patented "zero overhead task switch" (ZOTS™) technology promises to revolutionize both hardware and software design for embedded systems applications. Essentially, we remove the task-switching and interrupt control decision-making off the core microprocessor – which dramatically improves microprocessor efficiency and provides end results similar to increasing a microprocessor's clock speed by several orders of magnitude.

Our technological breakthrough is novel and radical. Never before has RTOS task management been implemented in hardware in such a manner. We knew we could attract customers if we could just effectively demonstrate the technology. Fortunately, the functionality of the Xilinx Virtex™-II FPGA gave us the solution.

Show and Tell

We chose to demonstrate our technology by building on a public domain 32-bit RISC architecture. By introducing a simple, seven-task simultaneous environment containing real-time video, stored video, and real-time audio, we have been able to demonstrate a fully functional multimedia microprocessor that handles interrupts on a pixel-by-pixel basis. With Xyronium™ technology embedded into a 32-bit RISC processor as a firm IP core in a Xilinx Virtex-II XC2V1000 FPGA, our demonstration board runs at only 40 MHz. Nevertheless, the board processes full-motion, full-frame video and stereo CD-quality audio – with enough performance power remaining to manipulate video in real time in a moving picture-in-picture synchronized with the audio amplitude (Figure 1). To see an interactive Macromedia® Flash™ demonstration, go to *www.xyronsemi.com/xdemo.shtml*.

A similarly burdened standalone desktop microprocessor programmed with the same RTOS capability and without video, audio, or memory support architecture would need to run at over 1.0 GHz to accomplish the same real-time tasks.

Figure 1 - Special effects in real-time video

Figure 2 - Xyron Semiconductor Edison development board incorporates a Xilinx Virtex-II XC2V1000 FPGA.

Show and Sell

The Xyron ZOTS approach is so revolutionary that we anticipated the initial market acceptance might be difficult. Potential customers might have difficulty believing you can implement traditional RTOS software task management functions in hardware – and realize significant microprocessor performance improvements as a result. Therefore, we created a powerful standalone, single-board computer that showcased our technology so users could see for themselves (Figure 2). Multiple video, audio, and user interfaces allow our customers the flexibility to design for any application in a familiar Xilinx FPGA environment.

Using a Virtex-II XC2V1000 FPGA, the development board is an ideal environment to not only illustrate the power of the ZOTS innovation, but also to implement these powerful new technologies in real-world applications. Once the Xyron ZOTS firm IP microprocessor core is inserted into the fabric of a Virtex-II device, there is ample room to create custom circuitry to support specific design requirements.

Conclusion

The path from demonstrating a radical new microprocessor technology to providing a practical development environment was optimized with the Virtex-II product family. The process only took a few months with a small design team.

In all fairness, we explored other options, but they were not nearly as efficient. Without the functionality of the Virtex-II FPGA, it would not have been possible to construct such a robust development board as the Edison™ Development Board. That's why Xyron Semiconductor has become a partner in the Xilinx AllianceCORE™ program.

In a classic win-win situation, combining Xyron technology with Xilinx products has given us the ability to realize standard-cell, dedicated IC performance in a flexible FPGA environment. The Virtex-II product offering has enabled Xyron Semiconductor to take its first steps toward becoming a full-fledged microprocessor company.

For more information on Xyron technology, please visit *www.xyronsemi.com*. To purchase a development board, go to *www.xyronsemi.net*.

Elantec Semiconductor Has Power Supply Solutions for Virtex-II Pro Platform FPGAs

High efficiency and fewer components give you high reliability while reducing printed circuit board real estate.

by Art Stryer
Senior Field Applications Engineer
Elantec Semiconductor, Inc.
astryer@elantec.com

Virtex-II Pro™ series Platform FPGAs use separate supply voltages for the core circuitry and I/O interface power supplies. Typically, the I/O interface requires 2.5V and the core circuitry requires 1.8V. As a designer, you need to minimize board space and maximize the reliability of the power supplies.

Elantec Semiconductor Inc. has highly integrated solutions to fulfill your local power needs. Elantec's family of integrated FET DC:DC converters provides optimal solutions for Virtex-II Pro designs. These power products are unique synchronous buck converters with integrated FETs and internal current sensing. These features, as well as other embedded functions, enable high efficiency and higher frequencies, which lead to smaller inductors. These features and functions require fewer external components, giving you the advantage of minimum board space. Additionally, higher reliability is achieved through lower die temperature.

Integrated FET DC:DC Converters

When you are designing a local power supply for a Virtex-II Pro system, several factors allow you to achieve a small PC board area. Elantec power products include the following features:

- Synchronous DC:DC switching regulators running at higher frequencies
  - Up to 1 MHz switching frequency
- Higher frequencies, allowing use of smaller external capacitors and inductors
  - Steady peak-to-peak ripple voltage with fewer external parts
- Higher integration, resulting in fewer external components
  - No external FETs
- Reduced parts count.

Embedded Functions Reduce ICs

When designing a local power supply for a Virtex-II Pro system, you need certain features for a complete solution. To achieve these complete solutions usually requires additional ICs. Elantec parts, however, feature embedded functions, including:

- Current mode control
  - On-chip current sensing feedback for internal PWM controller
- Over-current protection
  - Cycle-by-cycle current sensing and limiting
- Over-temperature protection
  - Junction temperature monitor circuit on die
  - Control circuits with hysteresis for automatic supply restart
- Adjustable switching frequency
  - User-selectable frequency to avoid switching interference
- Soft start
  - Internal power-up control
- Power sequencing
  - Sequential activation of multiple power supplies without additional control ICs
- Simultaneous power tracking with controlled voltage out.

High Reliability Components

The following list of functions in Elantec power products improve the reliability of your power supply:

- Synchronous converter
  - Both upper and lower power switches on-chip
  - Up to 95% efficiency
- Integrated Power FETs
  - Reduced power dissipation versus external FETs with resistive connections
- Advanced packages
  - Optional low die temperature, small HTSSOP.

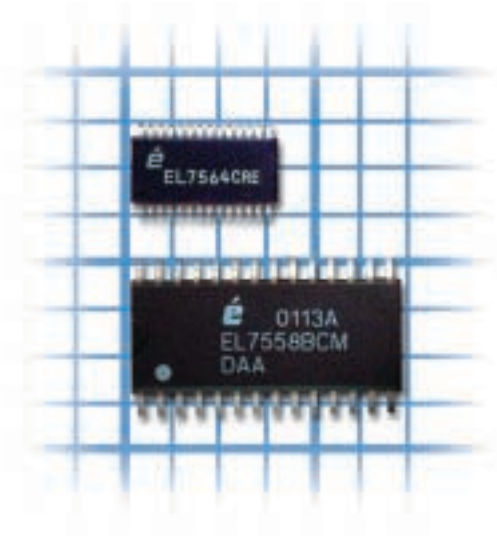

Figure 1 - Integrated FET 4A and 8A DC:DC converter packages

Complete DC:DC Converter Family

Different designs will have varying requirements for voltage and current, switching frequency, and available board space. The Elantec family of DC:DC products meets Virtex-II Pro local system needs by providing:

- Continuous output currents
  - Products to choose from with currents up to 1A, 2A, 4A, 6A, and 8A
- Flexible output voltages
  - Adjustable voltages from 1.0V to 3.3V available from each product
- Multiple package types
  - HTSSOP, QSOP, MSOP, HSOP, and SOL available.

Table 1 shows a comprehensive summary of Elantec power products. Figure 1 shows packages and board area space for 4A and 8A products.

Conclusion

Design engineers engaged in Virtex-II Pro board designs need to have power supplies that provide a complete solution for local power in their systems. Elantec DC:DC converters have all of the features and embedded functions to get the job done. For data sheets, application briefs, and additional information, visit Elantec online at *www.elantec.com*, or send an e-mail to *tech@elantec.com*.

Part No.	Type	Supply Voltage	Output Voltage	Output Current	Package
EL7558BC	Buck	4.5V to 5.5V	1.0V to 3.8V	8A	HSOP28
EL7556BC	Buck	4.5V to 5.5V	1.0V to 3.8V	6A	SO28
EL7564C	Buck	4.5V to 5.5V	1.0V to 3.8V	4A	HTSSOP28
EL7563C	Buck	3.0V to 3.6V	1.0V to 2.5V	4A	HTSSOP28
EL7562C	Buck	4.5V to 5.5V	1.0V to 3.8V	2A	QSOP16
EL7551C	Buck	4.5V to 5.5V	1.0V to 3.8V	1A	QSOP16
EL7512C	Boost	2.0V to 14V	5.0V to 16V	0.2A to 0.5A	MSOP10

Table 1 - Elantec power products

Nallatech's DIME-II Modular Architecture Goes Pro

Nallatech's second-generation modular DIME-II architecture has been designed to maximize Virtex-II Pro Platform FPGA performance and bandwidth.

by Dr. Malachy Devlin
Chief Technology Officer
Nallatech
m.devlin@nallatech.com

Typical data handling and processing systems require a combination of high-performance capabilities. Until now, these capabilities could only be achieved through combining various programmable logic architectures and processor technologies to handle complex control algorithms.

The release of the Xilinx Virtex-II Pro™ Platform FPGAs, however, now gives us a single, integrated package for manipulating complex control algorithms. As many as four IBM® PowerPC™ 405 processors can be embedded deeply in the Virtex-II Pro programmable logic fabric.

In order to create final systems based on Virtex-II Pro devices – and to harness maximum power from these devices – appropriate hardware architecture is critical. The DIME™ and DIME-II™ architectures are open modular standards from Nallatech that create the framework for scalable systems. A small, dedicated FPGA-based system can be easily scaled to create high-end high-speed DSP applications through the DIME-II plug-and-play capability. This architecture provides support for processor integration within FPGA-based systems, including the

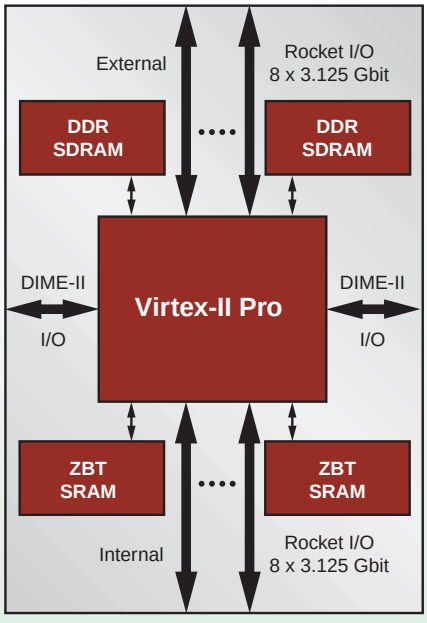

Figure 1 - BenPRO block diagram

use of ICE (in-circuit emulation) debuggers, remote booting, and configuration.

BenPRO Module

The BenPRO™ module from Nallatech provides maximum scalability and maximum communications bandwidth. The BenPRO™ module will be available with all Virtex-II Pro devices that have embedded PowerPC processors. In this configuration, a single DIME-II module can have up to four hard processors and up to four million system gates of logic for computational and interfacing tasks.

By utilizing up to 16 Rocket I/O™ transceivers on a single Virtex-II Pro, and standard I/O interfaces offered by DIME-II technology, you can get an I/O bandwidth in the order of 100 gigabits per second

To support onboard data and program storage, four high-speed memory banks are built into the BenPRO module. As shown in Figure 1, DDR SDRAM and ZBT SRAM enable suitable memory architectures to be designed for tight integration with the algorithm or platform architecture required.

Controlling Virtex-II Pro Systems

Nallatech's FUSE™ (field upgradeable systems environment) software supplies the reconfigurable computing operating system for FPGA-based systems. The powerful API (application programming interface) enables rapid development of systems and applications based on FPGAs, either with or without, embedded processors. FUSE system software is portable and supports multiple operating systems, such as Linux, VXWorks®, and Windows® systems.

Additionally, FUSE software supports a wide range of languages, including C/C++, Java™, MATLAB™ and DIMEscript™, Nallatech's own dedicated language for controlling reconfigurable computing environments.

FUSE software also supports multiple embedded processors, including the IBM PowerPC hard processor and the MicroBlaze™ soft processor from Xilinx – both of which can be embedded in one Virtex-II Pro Platform FPGA. The power of FUSE system software provides for high-speed booting and data communication, as well as system control. The FUSE capability allows full integration with high-level design tools from Xilinx, industry-standard protocols, and user application software. Figure 2 shows the relationship among FUSE, DIME-II, and Virtex-II Pro components.

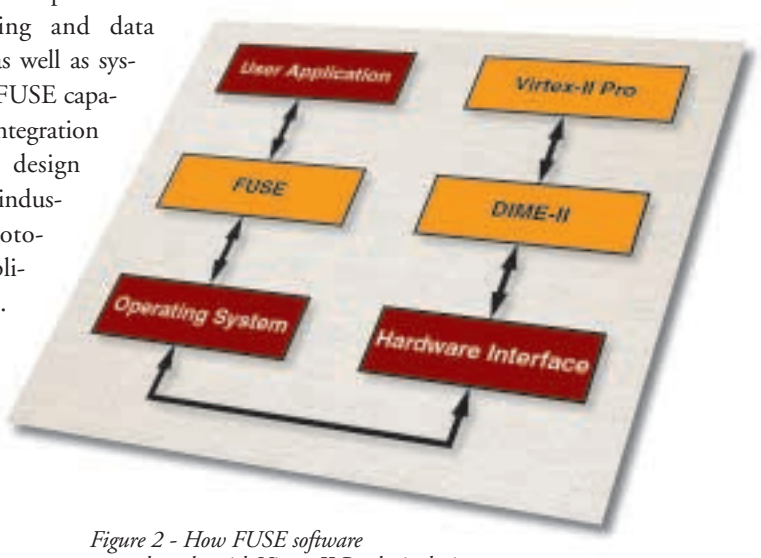

Figure 2 - How FUSE software control works with Virtex-II Pro logic devices

Distributed Parallel Processing

Although a single Virtex-II Pro device can have as many as four processors, DIME-II technology can scale systems far beyond this. For example, DIME-II interconnect technology will enable you to tightly couple four BenPRO modules on a BenERA™ cPCI card, thus providing 16 PowerPC processors on a single card, as illustrated in Figure 3. With FUSE system control software, you can expand connectivity to multiple cPCI cards, leading to a virtually unlimited array of processors within the overall system.

To benefit from this capability, the tools need to be in place for handling this distributed parallel processing system. Nallatech is enhancing its system design products to deliver the tools and communications infrastructure to support distributed parallel processing. Once in place, the combined DIME-II and Virtex-II Pro technologies will enable you to create a diverse range of system topologies that can take advantage of the high communications bandwidths now available – and in the future.

Conclusion

Today, system designers require suitable software tools, development platforms, and modular hardware to efficiently create complex high-performance systems. With DIME-II hardware system architecture and the FUSE reconfigurable computing operating system, you can not only harness the full power of Virtex-II Pro Platform FPGAs, but you can also use Nallatech's advanced engineering to bundle multiple devices within a system – taking system designs to unprecedented levels of complexity and functionality.

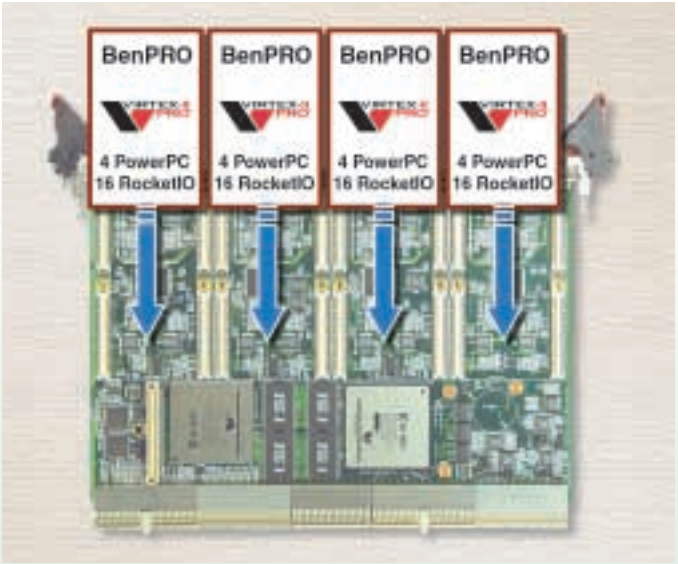

Figure 3 - BenERA cPCI card with four Virtex-II Pro FPGAs on BenPRO modules

Platform FPGAs Take on ASIC SOCs

Here are seven good reasons why Platform FPGAs provide a superior design environment and faster time to market than ASIC SOCs.

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

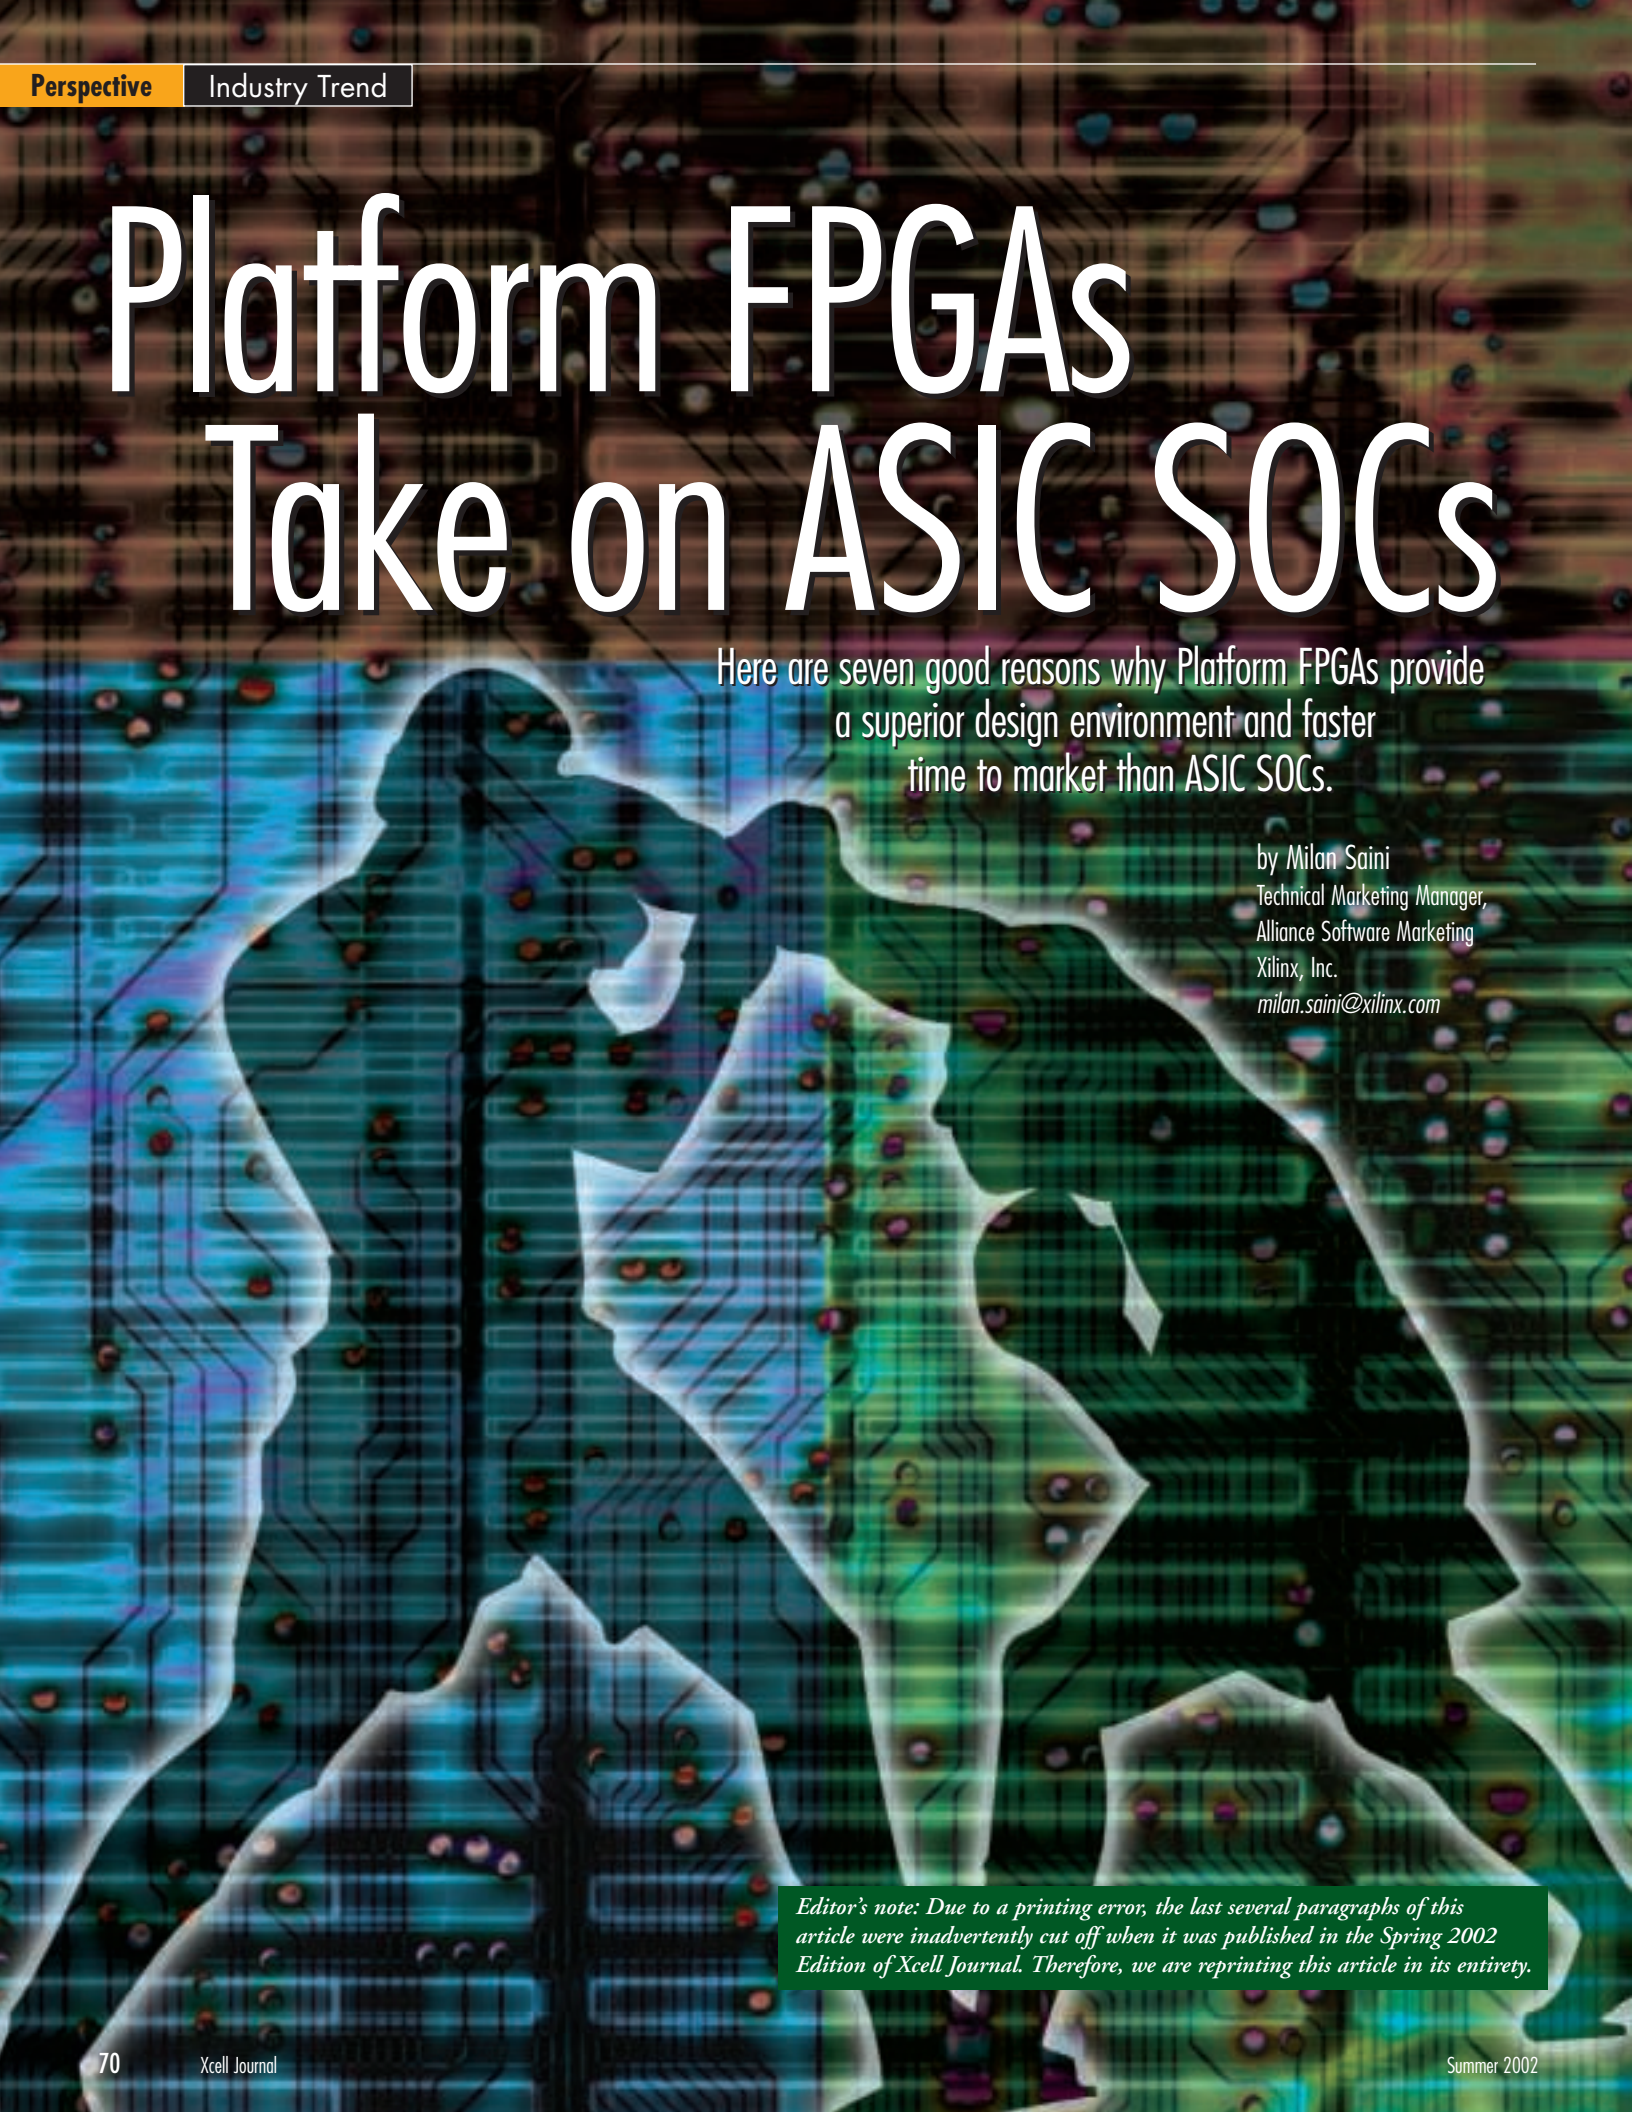

Editor's note: Due to a printing error, the last several paragraphs of this article were inadvertently cut off when it was published in the Spring 2002 Edition of Xcell Journal. Therefore, we are reprinting this article in its entirety.

The system-on-a-chip (SOC) market has experienced steady and consistent growth. Dataquest estimates roughly half of all ASIC design starts are SOC based. That percentage is projected to reach 80% by 2005. There are several reasons for this clear shift in design methodology. Some of the obvious advantages include greater component integration, increased speed (Logic <-> Processor), lower packaging and test costs, and increased overall system reliability. All of these combined can potentially make significant contributions to achieving the often elusive but always important goal of accelerated time to market.

Programmable logic device vendors have entered the SOC solution space with the introduction of a new class of devices known as Platform FPGAs – with the most recent addition being the Virtex-II Pro™ Platform FPGA. These devices offer the same level of integration as ASIC SOCs, but in contrast, Platform FPGAs facilitate the development of a wide range of applications on the same chip. This article focuses on seven of the key advantages of the Platform FPGA approach.

#1 – Pre-Engineered Platform

Platform FPGAs integrate several fixed and predetermined blocks of hard IP (intellectual property) components (system elements) within the programmable fabric. Notable among these are high-performance RISC CPUs, multi-gigabit and high speed I/Os, block RAM, system clock management, and dedicated DSP processing hardware. This powerful assembly and harmonious blend of components creates a cohesive system design environment. This environment offers unprecedented flexibility and performance, thus enabling the deployment of a wide range of applications.

The critical technological breakthrough is in the ability to tightly interface the various elements into the programmable fabric. Without this tight integration, much of the speed benefits could not be realized. The fact that the choice of system elements is already made greatly simplifies the design and development process. A fixed architecture is particularly beneficial for software tool and IP providers in allowing them to deliver better value, customization, and architecture-optimized solutions.

The assembly of the various hybrid IP blocks in an ASIC adds substantial complexity and hardship to the users' design and development environment, because of a variety of issues relating to tool and IP interoperability, physical layout, timing, and system verification. For Platform FPGAs, on the other hand, it is much easier to tailor and optimize components – such as silicon, software, support, and IP – because FPGAs represent a fixed and pre-engineered target.

Summary: A fixed, pre-engineered but programmable FPGA solution offers a more productive and efficient development environment from both the software and silicon perspective.

#2 – Process Technology

PLD vendors have been able to extract great value and benefits from Moore's Law and shrinking device geometries. While the majority of ASIC design starts are at or higher than 180 nm, FPGAs have raced ahead to bring the cost and performance advantages of 130 nm to its customers. This ability to rapidly migrate to the leading edge of technology is essential to delivering the performance, capacity, and integration that is necessary to challenge and displace the current established platforms of system design.

For instance, FPGAs now make it entirely feasible to build systems of up to 2M (ASIC) gates with CPU(s) running at 400 MHz, serial I/O channels at 3.125 Gbps CLB fabric-switching at 300 MHz, and the entire system clocking over 150 MHz. This surpasses the projected sweet spot of ASIC SOC designs.

Looking at it in pure economic terms, the costs for a typical mask set for a 130 nm ASIC run into the $700K+ range. That raises the bar for entry into the ASIC space, making the Platform FPGA an even more attractive option for a growing percentage of all SOC designs.

Summary: FPGA silicon is best-in-class in process and engineering, thereby delivering best-in-class system performance.

#3 – Software Tools and Methodology

It is well known that EDA designer productivity for ASICs is lagging behind recent silicon advances. The ability to create a productive design and debug environment is absolutely essential to the success of any silicon platform. Therefore, software plays a critical role in not only making the platform easy to use and work with, but also in extracting the most performance and device utilization out of the silicon.

The goal is to insulate the user from having to learn extraneous details about the platform, yet providing empowerment and control when and where it is needed. To this end, FPGA vendors have created customized and user-friendly processor system generator tools that aid in instantiation, initialization, and configuration of all the various system component blocks. In addition, these software tools automate otherwise manual and error-prone tasks, such as the interconnections among the processor, its peripherals, and buses.

Design entry is engineered to tightly couple the HW/SW domains. Such engineer-

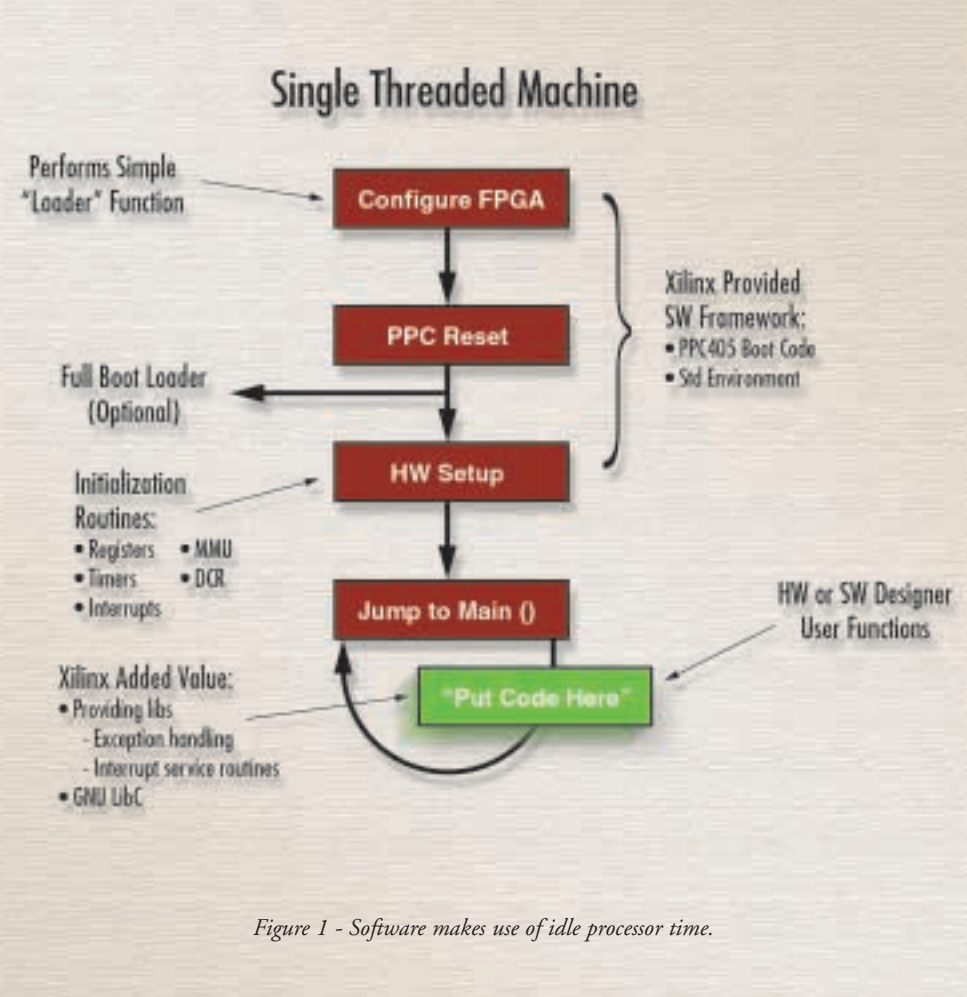

Figure 1 - Software makes use of idle processor time.

ing leads to not only simplified and easy-to-use design flows, but it ensures reliability and robustness from the very start by performing design rule and data consistency checks across the two domains.

Two examples of the advantages of cross-domain HW/SW co-design include:

- Automatic generation of device drivers and header files for SW engineers once a particular block is instantiated by the HW designer
- Tools to automatically populate SW binary code into appropriate FPGA memory bitstreams.

Summary: Software sells silicon.

#4 – Advanced Debug

The importance of finding problems early cannot be overstated. Yet, up-front ASIC verification is extremely designer- and computer-resource intensive. Compared to systems on a board, SOC limits visibility into the internal nodes of the system, making the task of verification and debug more challenging than ever.

The critical part includes the verification of complex interactions between the application software and the custom-designed peripheral hardware. Traditional HDL-centric verification and debug techniques can no longer deal with the rising complexity of system designs.

Consider a typical application, such as MPEG A/V decoding, where a large number of simulation cycles are required to complete a small sequence of frames. Co-verification tools in this case would either take an impractically long time to complete or validate only a mere fraction of the software code, falling far short of what it takes to find problems in the HW/SW interface.

FPGAs overcome this problem in large part by being able to provide access to real or near real targets at a very early stage in the design cycle. Among other things, this means that SW engineers using FPGAs can quickly and easily sort out logic and design flaws by targeting real silicon. The engineers do not have to rely on inefficient ASIC-centric techniques like co-verification or writing stub code. Application software can be debugged at system speeds with full hardware and software register access and control.

Additionally, the FPGA fabric allows for construction of highly customized and value-added cores to enable powerful real-time, on-chip debug capability. Some examples of such instrumentation include:

- Logic and bus analyzer functions
- Bus protocol compliance monitors
- Memory buffers for debug and trace port data
- Cross-domain triggers and breakpoints
- Hardware run control of the CPU
- HW/SW time synchronization logic.

Furthermore, a single cable is able to perform multiple functions, like debugging hardware, debugging software, as well as programming the FPGAs. This greatly simplifies the lab setup making it much easier to exploit the debug advantages.

Summary: Platform FPGAs offer a clearer, more cohesive, and overall more effective debug strategy. Specifically, Platform FPGAs offer up-front silicon access along with unprecedented visibility and control of the processor and its peripherals residing in the programmable fabric.

#5 – Top Tier Partnerships and Vendor Tools Support

In extending the concept of traditional programmable logic to Platform FPGAs, certain critical technologies have had to be developed or acquired. One of most exciting aspects brought forth by FPGA vendors has been to successfully forge

strategic partnerships with vendors holding various system technology components. Driven largely by the successes of the FPGA business models, leading vendors have been eager to partner in fulfilling the vision of building powerful programmable systems platforms.

Areas of cooperation include partnerships in the form of cutting-edge process technology, powerful mainstream processing elements – such as RISC CPU, high-speed I/O – and other components deemed of significance to a system design platform. IBM, Conexant Systems, Inc., Wind River Systems, and other high profile vendors are currently engaged in such strategic partnerships. What this means to you is there is no need for these partners to negotiate licensing, royalty, and integration issues with individual vendors, thereby greatly reducing your engineering, management, and accounting overhead.

The appeal and draw of FPGAs has caught the attention of independent SW tool vendors. Increasing numbers of vendors are able to sustain business models selling to FPGA customers. Several new vendors are setting up shop to write custom tools to help enhance and exploit the unique capabilities of Platform FPGAs.

Summary: The FPGA business model has attracted top-tier silicon and IP vendors to forge strategic partnerships to create powerful system design platforms. Software vendors are able to financially justify investment in research and development leading to a continuous stream of an increasing number of innovative solutions.

#6 – Application Space: Co-Design Flexibility

Programmable HW combined with processors on a single chip softens the HW/SW design boundary. By using hardware-assisted architectural exploration, designers can optimally search for the right HW and SW partitioning, which leads to the increased probability of being able to meet performance and area targets. Sequential computing, exception handling, and control functions, for instance, programmed in HW could be implemented in SW running on the processor to save silicon.

On the other hand, SW structures and algorithms – which can be broken into parallel, non-blocking processes – can achieve significant speed and data throughput improvements by implementing them in HW. In fact, software engineers represent a new and emerging market for Platform FPGAs. Aided by design tools, it is now easier than ever for SW engineers to explore concepts of software acceleration via HW. Both hardware and firmware are reprogrammable and field upgradeable, which enables the development and deployment of several product generations from one base. This translates into a much broader applicability than ASICs and ASSPs.

In addition to being suitable for building complex embedded applications, HW engineers can utilize an otherwise idle processor to run relevant portions of their design algorithms or control logic. As Figure 1 shows, the CPU can serve as a simple microcontroller running a single-threaded application. PLD vendors add value by providing software device drivers and library functions for rapid implementation without requiring the HW engineer to have detailed knowledge of SW practices.

Summary: The Platform FPGAs provide the most flexible co-design platform by enabling dynamic HW-SW design partitioning.

#7 – Risk Management

When compared to Platform FPGAs, ASIC SOCs present a huge design risk. With more variables and issues to worry about, and with limited debug capabilities when mistakes are found, ASIC designers are forced to either resolve problems suboptimally in software, or in the worst case scenario, they are forced to respin the silicon at great cost and loss of time to market.

Reprogrammability comes up big here. Programmable platforms allow early access to silicon. Engineers can validate performance and functionality in real silicon, leading to greater confidence in the reliability of the completed system. The programmability of the platform allows itself to be debugged and upgraded even after the system has been deployed. This helps promote and protect the time-to-market advantage by alleviating a large part of the risk.

Summary: Programmable Platform FPGAs provide better control over the life cycle management of products by minimizing the cost and time penalty for design errors and specification changes.

Conclusion

In an era when SOCs continue to dominate mainstream design, Platform FPGAs are emerging to take a share of the spotlight from ASICs. While ASIC SOCs have and will continue to be strong in certain segments – like low power, small form factor, and high-volume, cost-sensitive consumer electronic gadgets – an increasing range of other infrastructure applications in areas such as networking, telecommunications, industrial electronics, and data storage have compelling reasons to move to a programmable platform. Σ

CoolRunner-II CPLD Family

DS090 (v1.0) January 3, 2002 **Advance Product Specification**

Features

- Optimized for 1.8V systems
  - Industry's fastest low power CPLD
  - Static Icc of less than 100 microamps at all times
  - Densities from 32 to 512 macrocells
- Industry's best 0.18 micron CMOS CPLD
  - Optimized architecture for effective logic synthesis
  - Multi-voltage I/O operation — 1.5V to 3.3V
- Advanced system features
  - Fastest in system programming
    - 1.8V ISP using IEEE 1532 (JTAG) interface
  - IEEE1149.1 JTAG Boundary Scan Test
  - Optional Schmitt trigger input (per pin)
  - Unsurpassed low power management
  - FZP 100% CMOS product term generation
  - DataGATE external signal control
  - Flexible clocking modes
    - Optional DualEDGE triggered registers
    - Clock divider (÷ 2,4,6,8,10,12,14,16)
    - CoolCLOCK
  - Global signal options with macrocell control
    - Multiple global clocks with phase selection per macrocell
    - Multiple global output enables
    - Global set/reset
  - Abundant product term clocks, output enables and set/resets
  - Efficient control term clocks, output enables and set/resets for each macrocell and shared across function blocks
  - Advanced design security
  - Open-drain output option for Wired-OR and LED drive
  - Optional bus-hold or weak pullup on selected I/O pins
  - Optional configurable grounds on unused I/Os
  - Mixed I/O voltages compatible with 1.5V, 1.8V, 2.5V, and 3.3V logic levels on all parts
    - SSTL2-1,SSTL3-1, and HSTL-1 on 128 macrocell and denser devices
  - PLA architecture
    - Superior pinout retention
    - 100% product term routability across function block
  - Hot pluggable
  - Wide package availability including fine pitch:
    - Chip Scale Package (CSP) BGA, Fine Line BGA, TQFP, PQFP, VQFP, and PLCC packages
  - Design entry/verification using Xilinx and industry standard CAE tools
  - Free software support for all densities using Xilinx WebPACK™ or WebFITTER™ tools
  - Industry leading nonvolatile 0.18 micron CMOS process
  - Guaranteed 1,000 program/erase cycles
  - Guaranteed 20 year data retention

Family Overview

Xilinx CoolRunner™-II CPLDs deliver the high speed and ease of use associated with the XC9500/XL/XV CPLD family with the extremely low power versatility of the XPLA3™ family in a single CPLD. This means that the exact same parts can be used for high-speed data communications/computing systems and leading edge portable products, with the added benefit of In System Programming. Low power consumption and high-speed operation are combined into a single family that is easy to use and cost effective. Xilinx patented Fast Zero Power™ (FZP) architecture inherently delivers extremely low power performance without the need for any special design measures. Clocking techniques and other power saving features extend the users' power budget. The design features are supported starting with Xilinx ISE 4.1i, WebFITTER, and ISE WebPACK.

Table 1 shows the macrocell capacity and key timing parameters for the CoolRunner-II CPLD family.

Table 1: **CoolRunner-II CPLD Family Parameters**

	XC2C32	XC2C64	XC2C128	XC2C256	XC2C384	XC2C512
Macrocells	32	64	128	256	384	512
Max I/O	33	64	100	184	240	270
T_{PD} (ns)	3.5	4.0	4.5	5.0	5.5	6.0
T_{SU} (ns)	1.7	2.0	2.1	2.2	2.3	2.4
T_{CO} (ns)	2.8	3.0	3.4	3.8	4.2	4.6
F_{SYSTEM} (MHz)	303	270	244	222	204	189

Table 2 shows the CoolRunner-II CPLD package offering with corresponding I/O count. All packages are surface mount, with over half of them being ball-grid technologies. The ultra tiny packages permit maximum functional capacity in the smallest possible area. The CMOS technology used in CoolRunner-II CPLDs generates minimal heat, allowing the use of tiny packages during high-speed operation.

There are at least two densities present in each package with three in the VQ100 (100-pin 1.0mm QFP) and TQ144 (144-pin 1.4mm QFP), and in the FT256 (256-ball 1.0mm spacing FLBGA). The FT256 is particularly important for slim dimensioned portable products with mid- to high-density logic requirements.

Table 2: **CoolRunner-II CPLD Family Packages and I/O Count**

	XC2C32	XC2C64	XC2C128	XC2C256	XC2C384	XC2C512
PC44	33	33	-	-	-	-
VQ44	33	33	-	-	-	-
CP56	33	45	-	-	-	-
VQ100	-	64	80	80	-	-
CP132	-	-	100	106	-	-
TQ144	-	-	100	118	118	-
PQ208	-	-	-	173	173	173
FT256	-	-	-	184	212	212
FG324	-	-	-	-	240	270

Table 3 details the distribution of advanced features across the CoolRunner-II CPLD family. The family has uniform basic features with advanced features included in densities where they are most useful. For example, it is very unlikely that four I/O banks are needed on 32 and 64 macrocell parts, but very likely they are for 384 and 512 macrocell parts. The I/O banks are groupings of I/O pins using any one of a subset of compatible voltage standards that share the same V_{CCIO} level. (See Table 4 for a summary of CoolRunner-II I/O standards.) The clock division capability is less efficient on small parts, but more useful and likely to be used on larger ones. DataGATE, an ability to block and latch inputs to save power, is valuable in larger parts, but brings marginal benefit to small parts.

Table 3: **CoolRunner-II CPLD Family Features**

	XC2C32	XC2C64	XC2C128	XC2C256	XC2C384	XC2C512
IEEE 1532	✓	✓	✓	✓	✓	✓
I/O banks	1	1	2	2	4	4
Clock division	-	-	✓	✓	✓	✓
Clock doubling	✓	✓	✓	✓	✓	✓
DataGATE	-	-	✓	✓	✓	✓
LVTTL	✓	✓	✓	✓	✓	✓
LVCMOS33, 25, 18, and 1.5V I/O	✓	✓	✓	✓	✓	✓
SSTL2-1	-	-	✓	✓	✓	✓
SSTL3-1	-	-	✓	✓	✓	✓
HSTL-1	-	-	✓	✓	✓	✓
Configurable ground	✓	✓	✓	✓	✓	✓
Quadruple data security	✓	✓	✓	✓	✓	✓
Open drain outputs	✓	✓	✓	✓	✓	✓
Hot plugging	✓	✓	✓	✓	✓	✓

Architecture Description

CoolRunner-II CPLD is a highly uniform family of fast, low power CPLDs. The underlying architecture is a traditional CPLD architecture combining macrocells into Function Blocks (FBs) interconnected with a global routing matrix, the Xilinx Advanced Interconnect Matrix (AIM). The Function Blocks use a Programmable Logic Array (PLA) configuration which allows all product tems to be routed and shared among any of the macrocells of the FB. Design software can efficiently synthesize and optimize logic that is subsequently fit to the FBs and connected with the ability to utilize a very high percentage of device resources. Design changes are easily and automatically managed by the software, which exploits the 100% routability of the Programmable Logic Array within each FB. This extremely robust building block delivers the industry's highest pinout retention, under very broad design conditions. The architecture will be explained by expanding the detail as we discuss the underlying Function Blocks, logic and interconnect.

The design software automatically manages these device resources so that users can express their designs using completely generic constructs without knowledge of these architectural details. More advanced users can take advantage of these details to more thoroughly understand the software's choices and direct its results.

Figure 1 shows the high-level architecture whereby Function Blocks attach to pins and interconnect to each other within the internal interconnect matrix. Each FB contains 16 macrocells. The BSC path is the JTAG Boundary Scan Control path. The BSC and ISP block has the JTAG controller and In-System Programming Circuits.

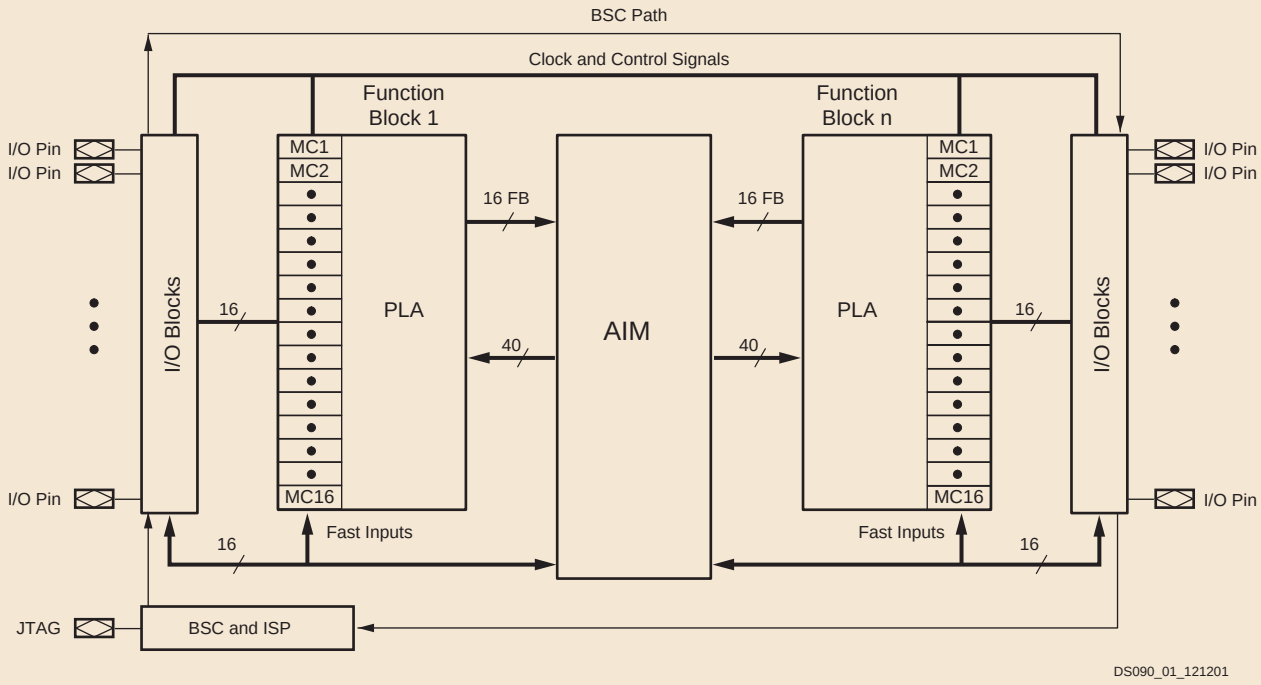

Figure 1: **CoolRunner-II CPLD Architecture**

Function Block

The CoolRunner-II CPLD Function Blocks contain 16 macrocells, with 40 entry sites for signals to arrive for logic creation and connection. The internal logic engine is a 56 product term PLA. All Function Blocks, regardless of the number contained in the device, are identical. For a high-level view of the Function Block, see Figure 2.

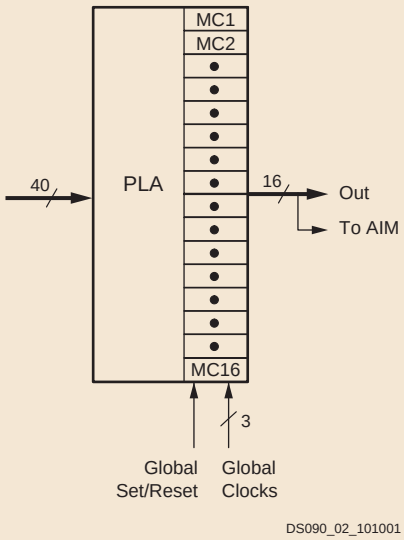

Figure 2: **CoolRunner-II CPLD Function Block**

At the high level, it is seen that the product terms (p-terms) reside in a programmable logic array (PLA). This structure is extremely flexible, and very robust when compared to fixed or cascaded product term function blocks.

Classic CPLDs typically have a few product terms available for a high-speed path to a given macrocell. They rely on capturing unused p-terms from neighboring macrocells to expand their product term tally, when needed. The result of this architecture is a variable timing model and the possibility of stranding unusable logic within the FB.

The PLA is different — and better. First, any product term can be attached to any OR gate inside the FB macrocell(s). Second, any logic function can have as many p-terms as needed attached to it within the FB, to an upper limit of 56. Third, product terms can be re-used at multiple macrocell OR functions so that within a FB, a particular logical product need only be created once, but can be re-used up to 16 times within the FB. Naturally, this plays well with the fitting software, which identifies product terms that can be shared.

The software places as many of those functions as it can into FBs, so it happens for free. There is no need to force macrocell functions to be adjacent or any other restriction save residing in the same FB, which is handled by the software. Functions need not share a common clock, common set/reset or common output enable to take full advantage of the PLA. Also, every product term arrives with the same time delay incurred. There are no cascade time adders for putting more product terms in the FB. When the FB product term budget is reached, there is a small interconnect timing penalty to route signals to another FB to continue creating logic. Xilinx design software handles all this automatically.

Macrocell

The CoolRunner-II CPLD macrocell is extremely efficient and streamlined for logic creation. Users can develop sum of product (SOP) logic expressions that comprise up to 40 inputs and span 56 product terms within a single function block. The macrocell can further combine the SOP expression into an XOR gate with another single p-term expression. The resulting logic expression's polarity is also selectable. As well, the logic function can be pure combinatorial or registered, with the storage element operating selectably as a D or T flip-flop, or transparent latch. Available at each macrocell are independent selections of global, function block level or local p-term derived clocks, sets, resets, and output enables. Each macrocell flip-flop is configurable for either single edge or DualEDGE clocking, providing either double data rate capability or the ability to distribute a slower clock (thereby saving power). For single edge clocking or latching, either clock polarity may be selected per macrocell. CoolRunner-II macrocell details are shown in Figure 3. Note that in Figure 3, standard logic symbols are used except the trapezoidal multiplexers have input selection from statically programmed configuration select lines (not shown). Xilinx application note XAPP376 gives a detailed explanation of how logic is created in the CoolRunner-II CPLD family.

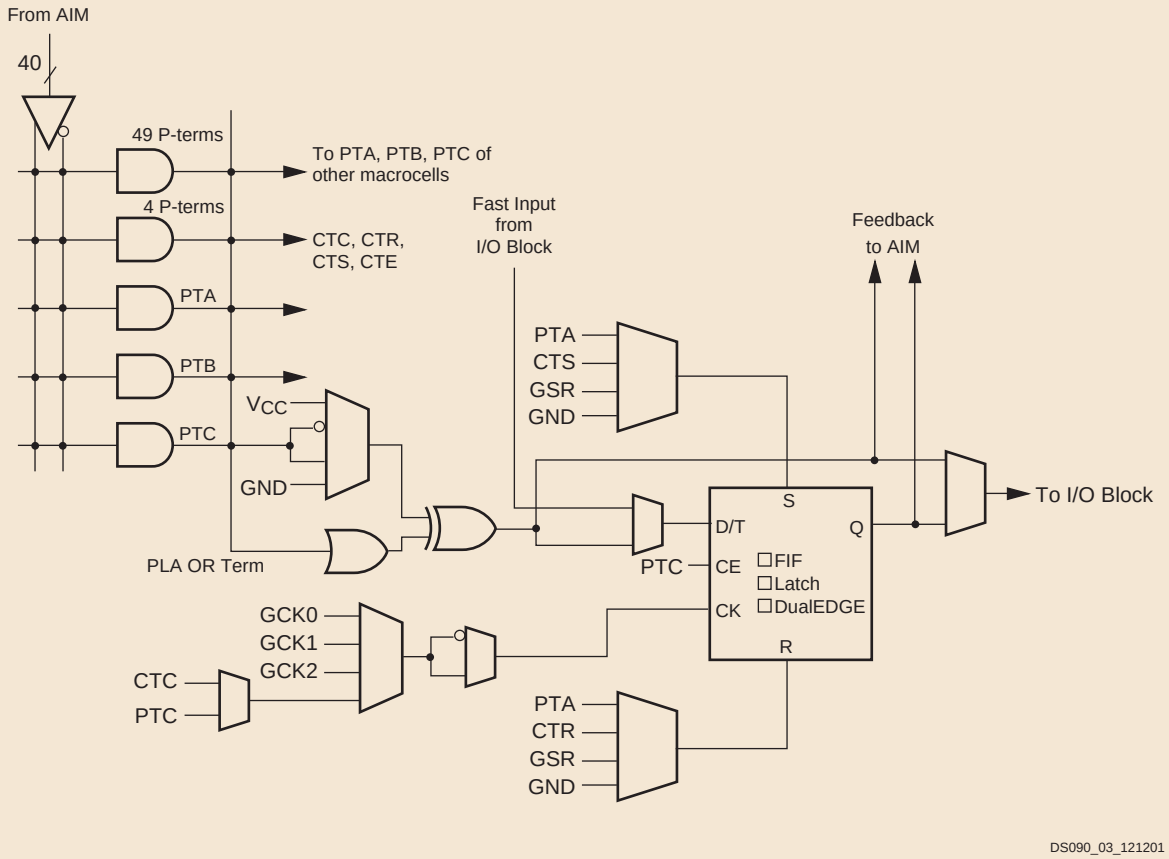

Figure 3: **CoolRunner-II CPLD Macrocell**

When configured as a D-type flip-flop, each macrocell has an optional clock enable signal permitting state hold while a clock runs freely. Note that Control Terms (CT) are available to be shared for key functions within the FB, and are generally used whenever the exact same logic function would be repeatedly created at multiple macrocells. The CT product terms are available for FB clocking (CTC), FB asynchronous set (CTS), FB asynchronous reset (CTR), and FB output enable (CTE).

Any macrocell flip-flop can be configured as an input register or latch, which takes in the signal from the macrocell's I/O pin, and directly drives the AIM. The macrocell combinational functionality is retained for use as a buried logic node if needed.

Advanced Interconnect Matrix (AIM)

The Advanced Interconnect Matrix is a highly connected low power rapid switch. The AIM is directed by the software to deliver up to a set of 40 signals to each FB for the creation of logic. Results from all FB macrocells, as well as, all pin inputs circulate back through the AIM for additional connection available to all other FBs as dictated by the designsoftware. The AIM minimizes both propagation delay and power as it makes attachments to the various FBs.

I/O Block

I/O blocks are primarily transceivers. However, each I/O is either automatically compliant with standard voltage ranges or can be programmed to become so.

In addition to voltage levels, each input can selectively arrive through Schmitt-trigger inputs. This adds a small time delay, but substantially reduces noise on that input pin. Hysteresis also allows easy generation of external clock circuits. The Schmitt-trigger path is best seen in Figure 4.

Outputs can be directly driven, 3-stated or open-drain configured. A choice of slow or fast slew rate output signal is also available. Table 4 summarizes various supported voltage standards associated with specific part capacities. All inputs and disabled outputs are voltage tolerant up to 3.3V.

Figure 4 details the I/O pin, where it is noted that the inputs requiring comparison to an external reference voltage (SSTL2-1, SSTL3-1, HSTL-1) are available on the 128 macrocell and denser parts.

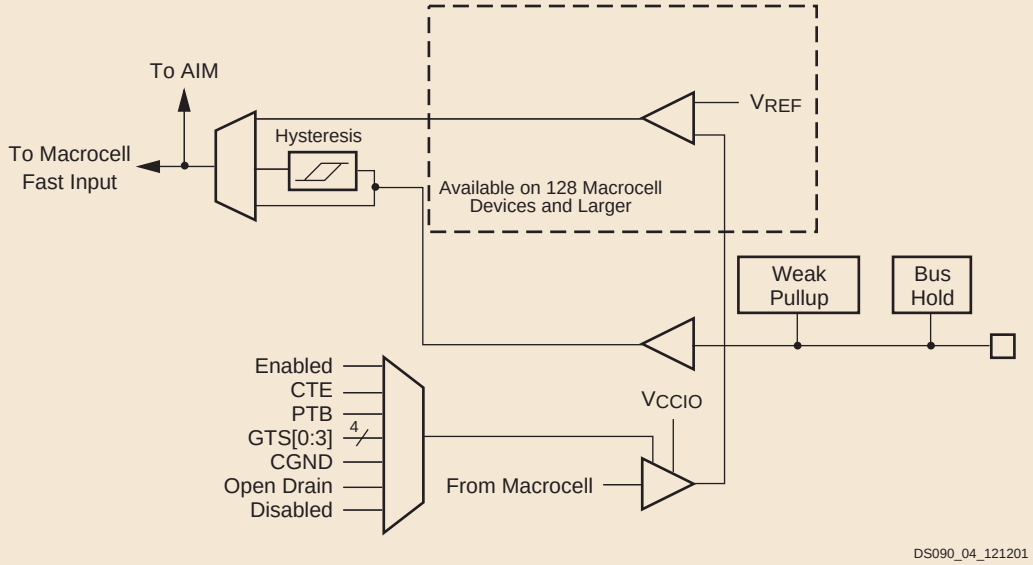

Figure 4: **CoolRunner-II CPLD I/O Block Diagram**

Table 4 summarizes the single ended I/O standard support and shows which standards require V_{REF} values and board termination. V_{REF} detail is given in specific data sheets.

Table 4: **CoolRunner-II CPLD I/O Standard Summary**

I/O Standard	V_{CCIO}	Input V_{REF}	Board Termination Voltage (V_{TT})
LVTTL	3.3	N/A	N/A
LVCMOS33	3.3	N/A	N/A
LVCMOS25	2.5	N/A	N/A
LVCMOS18	1.8	N/A	N/A
1.5V I/O	1.5	N/A	N/A
HSTL-1	1.5	0.75	0.75
SSTL2-1	2.5	1.25	1.25
SSTL3-1	3.3	1.5	1.5

Output Banking

CPLDs are widely used as voltage interface translators. To that end, the output pins are grouped in large banks. The smallest parts are not banked, so all signals will have the same output swing for 32 and 64 macrocell parts. The medium parts (128 and 256 macrocell) support two output banks. With two, the outputs will switch to one of two selected output voltage levels, unless both banks are set to the same voltage. The larger parts (384 and 512 macrocell) support four output banks split evenly. They can support groupings of one, two, three or four separate output voltage levels. This kind of flexibility permits easy interfacing to 3.3V, 2.5V, 1.8V, and 1.5V in a single part.

DataGATE

Low power is the hallmark of CMOS technology. Other CPLD families use a sense amplifier approach to creating product terms, which always has a residual current component being drawn. This residual current can be several hundred milliamps, making them unusable in portable systems. CoolRunner-II CPLDs use standard CMOS methods to create the CPLD architecture and deliver the corresponding low current consumption, without doing any special tricks. However, sometimes designers would like to reduce their system current even more by selectively disabling circuitry not being used.

The patented DataGATE technology was developed to permit a straightforward approach to additional power reduction. Each I/O pin has a series switch that can block the arrival of free running signals that are not of interest. Signals that serve no use may increase power consumption, and can be disabled. Users are free to do their design, then choose sections to participate in the DataGATE function. DataGATE is a logic function that drives an assertion rail threaded through the medium and high-density CoolRunner-II CPLD parts. Designers can select inputs to be blocked under the control of the DataGATE function, effectively blocking controlled switching signals so they do not drive internal chip capacitances. Output signals that do not switch, are held by the bus hold feature. Any set of input pins can be chosen to participate in the DataGATE function. Figure 5 shows the familiar CMOS I_{CC} versus switching frequency graph. With DataGATE, designers can approach zero power, should they choose to, in their designs

Figure 6 shows how DataGATE basically works. One I/O pin drives the DataGATE Assertion Rail. It can have any desired logic function on it. It can be as simple as mapping an input pin to the DataGATE function or as complex as a

counter or state machine output driving the DataGATE I/O pin through a macrocell. When the DataGATE rail is asserted low, any pass transistor switch attached to it is blocked. Note that each pin has the ability to attach to the AIM through a DataGATE pass transistor, and thus be blocked. A latch automatically captures the state of the pin when it becomes blocked. The DataGATE Assertion Rail threads throughout all possible I/Os, so each can participate if chosen. Note that one macrocell is singled out to drive the rail, and that macrocell is exposed to the outside world through a pin, for inspection. If DataGATE is not needed, this pin is an ordinary I/O.

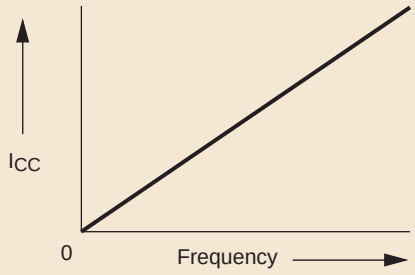

Figure 5: **CMOS I_{CC} vs. Switching Frequency Curve**

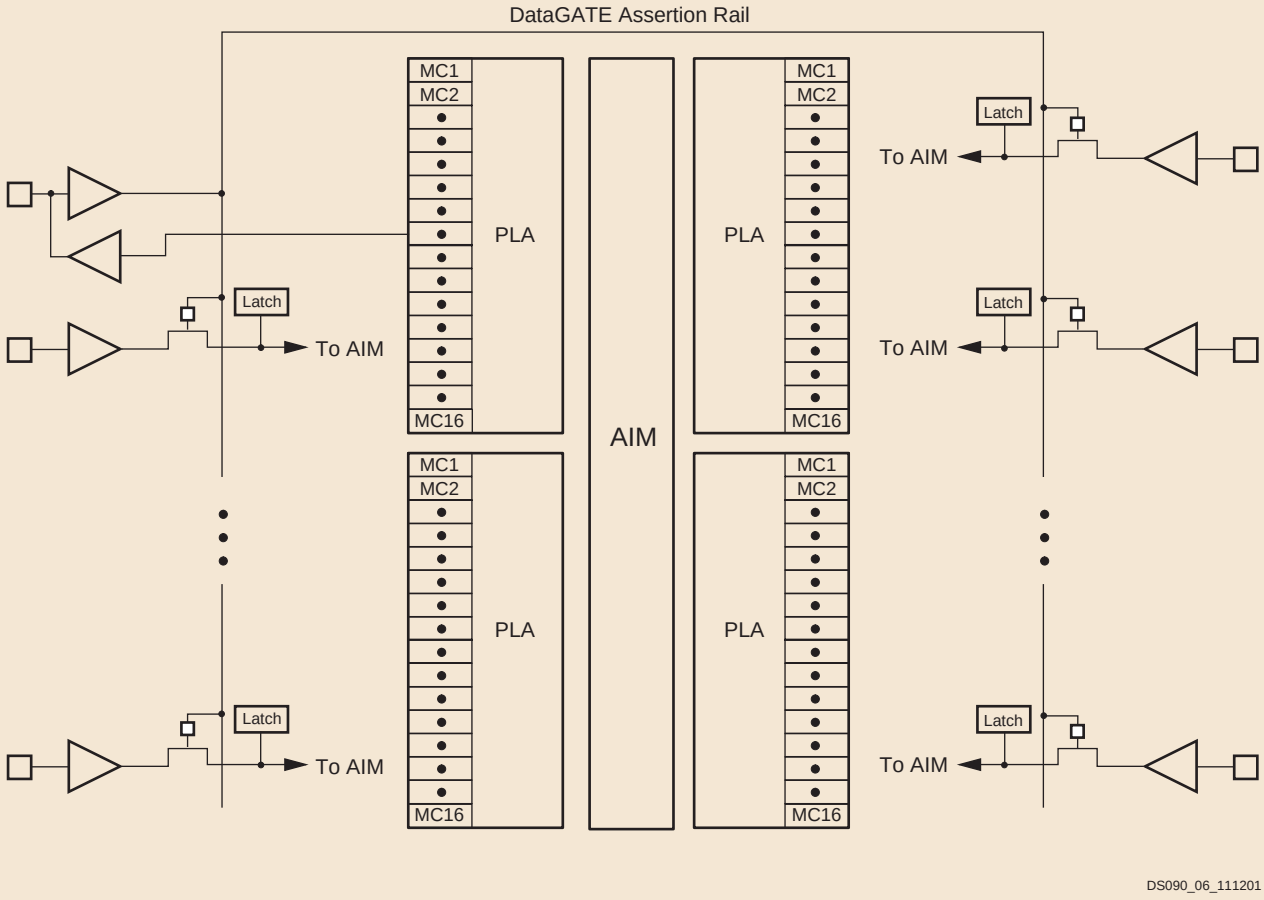

Figure 6: **DataGATE Architecture (output drivers not shown)**

Global Signals

Global signals, clocks (GCK), sets/resets (GSR) and output enables (GTS), are designed to strongly resemble each other. This approach enables design software to make the best utilization of their capabilities. Each global capability is supplemented by a corresponding product term version. Figure 7 shows the common structure of the global signal trees. The pin input is buffered, then drives multiple internal global signal traces to deliver low skew and reduce loading delays. The DataGATE assertion rail is also a global signal.

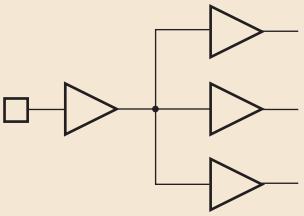

DS090_07_101001

Figure 7: **Global Clocks (GCK), Sets/Resets (GSR) and Output Enables (GTS)**

Additional Clock Options: Division, DualEDGE, and CoolCLOCK

Division

Circuitry has been included in the CoolRunner-II CPLD architecture to divide one externally supplied global clock by standard values. Division by 2,4,6,8,10, 12, 14 and 16 are the options (see Figure 8). This capability is supplied on the GCK2 pin. The resulting clock produced will be 50% duty cycle for all possible divisions. Note that a Synchronous Reset is included to guarantee no runt clocks can get through to the global clock nets. Note that again, the signal is buffered and driven to multiple traces with minimal loading and skew.

DualEDGE

Each macrocell has the ability to double its input clock switching frequency. Figure 9 shows the macrocell flip-flop with the DualEDGE option (doubled clock) at each macrocell. The source to double can be a control term clock, a product term clock or one of the available global clocks. The ability to switch on both clock edges is vital for a number of synchronous memory interface applications as well as certain double data rate I/O applications.

CoolCLOCK

In addition to the DualEDGE flip-flop, additional power savings can be had by combining the clock division circuitry with the DualEDGE circuitry. This capability is called CoolCLOCK and is designed to reduce clocking power within the CPLD. Because the clock net can be an appreciable power drain, the clock power can be reduced by driving the net at half frequency, then doubling the clock rate using DualEDGE triggering at the macrocells. Figure 10 shows how CoolCLOCK is created by internal clock cascading with the divider and DualEDGE flip-flop working together.

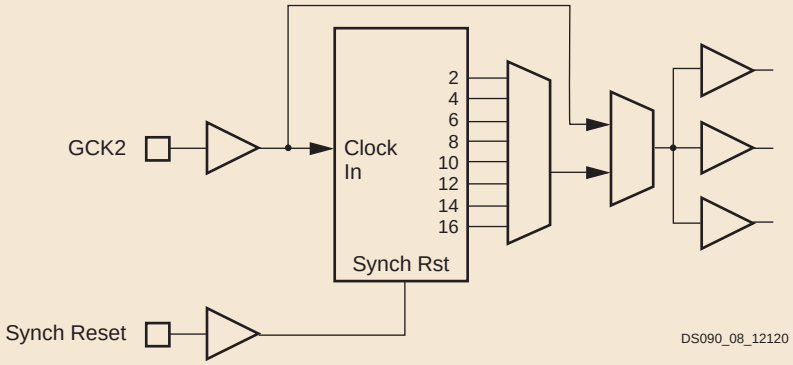

DS090_08_12120

Figure 8: **Clock Division Circuitry for GCK2**

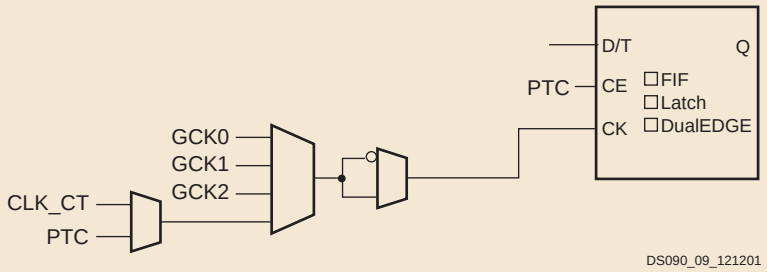

Figure 9: **Macrocell Clock Chain with DualEDGE Option Shown**

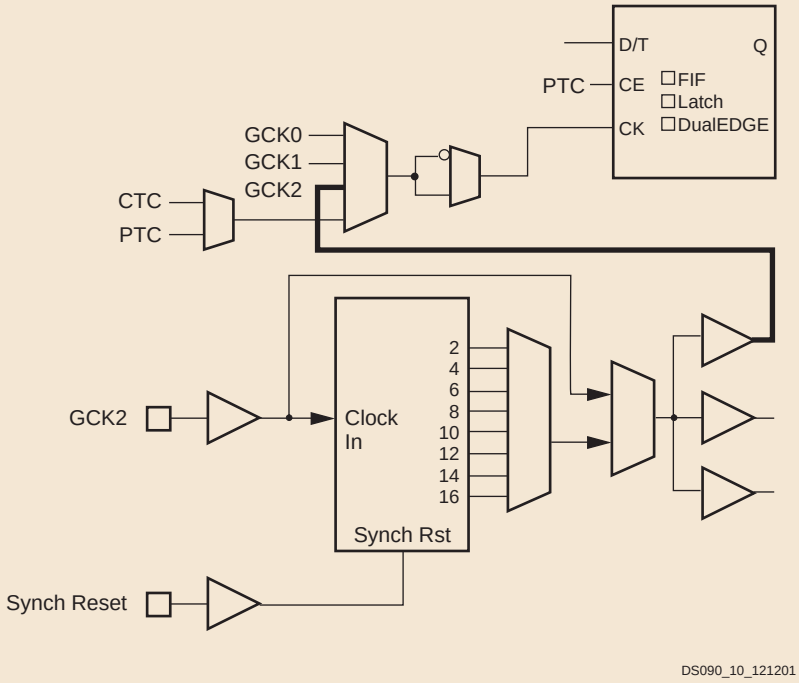

Figure 10: **CoolCLOCK Created by Cascading Clock Divider and DualEDGE Option**

Design Security

Designs can be secured during programming to prevent either accidental overwriting or pattern theft via readback. Four independent levels of security are provided on-chip, eliminating any electrical or visual detection of configuration patterns. These security bits can be reset only by erasing the entire device. Additional detail is omitted intentionally.

Timing Model

Figure 11 shows the CoolRunner-II CPLD timing model. It represents one aspect of the overall architecture from a timing viewpoint. Each little block is a time delay that a signal will incur if the signal passes through such a resource. Timing reports are created by tallying the incremental signal delays as signals progress within the CPLD. Software creates the timing reports after a design has been mapped onto the specific part, and knows the specific delay values for a given speed grade. Equations for the higher level timing values (i.e., T_{PD} and F_{SYSTEM}) are available. Table 5 summarizes the individual parameters and provides a brief definition of their associated functions. Xilinx application note XAPP375 details the CoolRunner-II CPLD family timing with several examples.

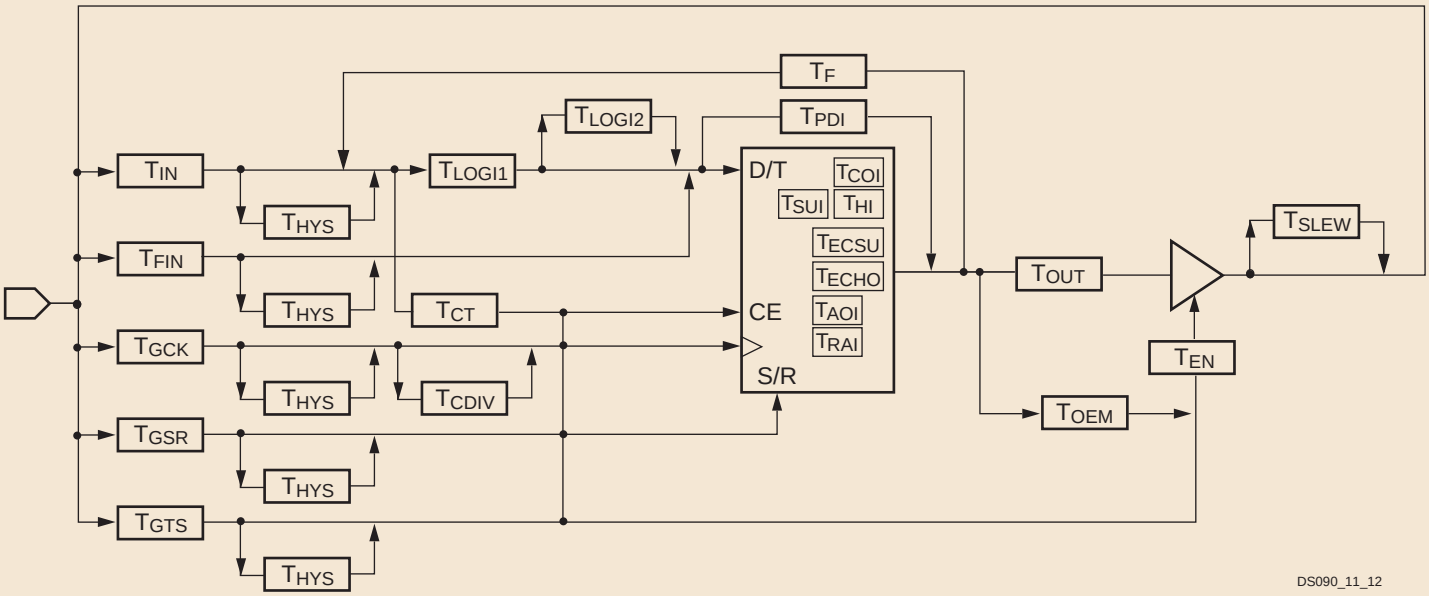

Figure 11: **CoolRunner-II CPLD Timing Model**

Table 5: **Timing Parameter Definitions**

Symbol	Parameter
Buffer Delays	
T_{IN}	Input Buffer Delay
T_{FIN}	Fast data register input delay
T_{GCK}	Global clock (GCK) buffer delay
T_{GSR}	Global set/reset (GSR) buffer delay
T_{GTS}	Global output enable (GTS) buffer delay
T_{OUT}	Output buffer delay
T_{EN}	Output buffer enable/disable delay
T_{SLEW}	Output buffer slew rate control delay
P-term Delays	
T_{CT}	Control Term delay (single PT or FB-CT)
T_{LOGI1}	Single P-term logic delay
T_{LOGI2}	Multiple P-term logic delay adder

Table 5: **Timing Parameter Definitions *(Continued)***

Symbol	Parameter
Macrocell Delays	
T_{PDI}	Macro cell input to output valid
T_{SUI}	Macro register setup before clock
T_{HI}	Macro register hold after clock
T_{ECSU}	Macro register enable clock setup time
T_{ECHO}	Macro register enable clock hold time
T_{COI}	Macro register clock to output valid
T_{RAI}	Macro register set/reset recovery before clock
T_{AOI}	Macro register set/reset to output valid
T_{CDIV}	Clock divider delay adder
T_{HYS}	Hysteresis selection delay adder
Feedback Delays	
T_F	Feedback delay
T_{OEM}	Macrocell to Global OE delay

In System Programming

All CoolRunner-II CPLD parts are 1.8V in system programmable. This means they derive their programming voltage and currents from the 1.8V V_{CC} (internal supply voltage) pins on the part. The V_{CCIO} pins do not participate in this operation, as they may assume another voltage ranging as high as 3.3V down to 1.5V. A 1.8V V_{CC} is required to properly operate the internal state machines and charge pumps that reside within the CPLD to do the nonvolatile programming operations. Xilinx software is provided to deliver the bit-stream to the CPLD and drive the appropriate IEEE 1532 protocol. To that end, there is a set of IEEE 1532 commands that are supported in the CoolRunner-II CPLD parts. Programming times are less than one second for 32 to 256 macrocell parts. Programming times are less than four seconds for 384 and 512 macrocell parts.

JTAG Instructions

Table 6 shows the commands available to users. These same commands may be used by third party ATE products, as well. The internal controllers can operate as fast as 66 MHz. The JTAG interface buffers are powered by a dedicated power pin, V_{CCAUX} , which is independent of all other supply pins.

Table 6: **JTAG Instructions**

Code	Instruction	Description
00000000	EXTEST	Force boundary scan data onto outputs
00000011	PRELOAD	Latch macrocell data into boundary scan cells
11111111	BYPASS	Insert bypass register between TDI and TDO
00000010	INTEST	Force boundary scan data onto inputs and feedbacks
00000001	IDCODE	Read IDCODE
11111101	USERCODE	Read USERCODE
11111100	HIGHZ	Force output into high impedance state
11111010	CLAMP	Latch present output state

Power-Up Characteristics

CoolRunner-II CPLD parts must operate under the demands of both the high-speed and the portable market places, therefore, they must support hot plugging for the high-speed world and tolerate most any power sequence to its various voltage pins. They must also not draw excessivecurrent during power-up initialization. To those ends, the general behavior is summarized as follows:

1. I/O pins are disabled until the end of power-up.
2. As supply rises, configuration bits transfer from nonvolatile memory to SRAM cells.
3. As power up completes, the outputs become as configured (input, output, or I/O).
4. The process takes less than 25 ms and draws less than 5 mA of current.

I/O Banking

CoolRunner-II CPLD 32 and 64 macrocell parts support a single V_{CCIO} rail that can range from 3.3V down to 1.5V operation. Two V_{CCIO} rails are supported on the 128 and 256 macrocell parts where outputs on each rail can independently range from 3.3V down to 1.5V operation. Four V_{CCIO} rails are supported on the 384 and 512 macrocell parts with each rail independently supporting any voltage between 3.3V and 1.5V. The V_{CC} (internal supply voltage) for a CoolRunner-II CPLD must be maintained within 1.8V ±5% for correct speed operation and proper in system programming.

Mixed Voltage, Power Sequencing, and Hot Plugging

As mentioned in I/O Banking, CoolRunner-II CPLD parts support mixed voltage I/O signals where signals within the same bank can range from 3.3V down to 1.5V. The power applied to the V_{CCIO} and V_{CC} pins can occur in any order and the CoolRunner-II CPLD will not be damaged. CoolRunner-II CPLDs can reside on boards where the board is inserted into a “live” connector (hot plugged) and the parts will be well-behaved as if powering up in a standard way.

Development System Support

Xilinx CoolRunner-II CPLDs are supported by all configurations of Xilinx standard release development software as well as the freely available WebFITTER and ISE WebPACK software available from www.xilinx.com. Third party development tools include synthesis tools from Cadence, Exemplar, Mentor Graphics, Synplicity, and Synopsys.

ATE Support

Third party ATE development support is available for both programming and board/chip level testing. Vendors providing this support include Hewlett-Packard, GenRad, and Teradyne. Other third party providers are expected to deliver solutions in the future.

Absolute Maximum Ratings(1)

Symbol	Parameter	Min.	Max.	Unit
V_{CC}	Supply voltage(2) relative to GND	–0.5	2.0	V
V_I	Input voltage(3) relative to GND	–0.5	4.0	V
TJ	Maximum junction temperature	–40	125	˚C
T_{STR}	Storage temperature	–65	150	˚C

Notes:
1. Stresses above those listed may cause malfunction or permanent damage to the device. This is a stress rating only. Functional operation at these or any other condition above those indicated in the operational and programming specification is not implied.
2. The chip supply voltage should rise monotonically.
3. Maximum DC undershoot below GND must be limited to either 0.5V or 10 mA, whichever is easier to achieve. During transitions, the device pins may undershoot to –2.0V or overshoot to 3.5 V, provided this over- or undershoot lasts less than 10 ns and with the forcing current being limited to 200 mA. The I/O voltage may never exceed 4.0V.

Quality and Reliability Parameters

Symbol	Parameter	Min	Max	Units
T_{DR}	Data retention	20	-	Years
N_{PE}	Program/erase cycles (Endurance)	1,000	-	Cycles
V_{ESD}	Electrostatic discharge	2,000	-	Volts

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
01/03/02	1.0	Initial Xilinx release.

Xilinx FPGA Product Selection Matrix

	Device	Rocket I/O™ Transceiver Blocks	PowerPC™ Processor Blocks	CLB (1 CLB = 4 slices = Max 128 bits) Array Row x Col	Slices	Maximum Distributed RAM Kbits	18 x 18 Bit Multiplier Blocks	Block SelectRAM 18 Kbit Blocks	Max Block RAM Kbits	DCMs	Max I/O Pads
VIRTEX-II PRO™	XC2VP2	4	0	16 x 22	1,408	44	12	12	216	4	204
	XC2VP4	4	1	40 x 22	3,008	94	28	28	504	4	348
	XC2VP7	8	1	40 x 34	4,928	154	44	44	792	4	396
	XC2VP20	8	2	56 x 46	9,280	290	88	88	1,584	8	564
	XC2VP50	16	4	88 x 70	22,592	706	216	216	3,888	8	852

	CLB Resources						BLK RAM		CLK Resources					I/O Features				Speed				
	System Gates (see note 1)	CLB Array (Row X Col)	Number of Slices	Logic Cells (see note 2)	CLB Flip-Flops	Max. Distributed RAM Bits	# Block RAM Blocks	Block RAM Bits	# Dedicated Multipliers	DLL Frequency (min/max)	# DLL's	Frequency Synthesis	Phase Shift	Digitally Controlled Impedance	Number of Differential I/O Pairs	Max. I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Serial PROM Family		Config. Memory (Bits)
Virtex-II Family — 1.5 Volt																		**.15um Eight Layer Metal Process**				
XC2V40	40K	8 x 8	256	576	512	8K	4	72K	4	24/420	4	DCM	DCM	YES	44	88	LDT-25, LVPECL-33,	-4 -5 -6	-4 -5	ISP	OTP	0.4M
XC2V80	80K	16 x8	512	1,152	1,024	16K	8	144K	8	24/420	4	DCM	DCM	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	24/420	8	DCM	DCM	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	24/420	8	DCM	DCM	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	24/420	8	DCM	DCM	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	24/420	8	DCM	DCM	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	24/420	8	DCM	DCM	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	24/420	12	DCM	DCM	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	24/420	12	DCM	DCM	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	24/420	12	DCM	DCM	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	24/420	12	DCM	DCM	YES	554	1108	SSTL3 I, SSTL3 II	-4 -5	-4			29.1M
Virtex-E Family — 1.8 Volt																		**.18um Six Layer Metal Process**				
XCV50E	72K	16 x 24	768	1,728	1,536	24K	16	64K	NA	25/350	8	YES	YES	NA	88	176		-6 -7 -8	-6 -7	ISP	OTP	0.6M
XCV100E	128K	20 x 30	1,200	2,700	2,400	37.5K	20	80K	NA	25/350	8	YES	YES	NA	98	196		-6 -7 -8	-6 -7			0.9M
XCV200E	306K	28 x 42	2,352	5,292	4,704	73.5K	28	112K	NA	25/350	8	YES	YES	NA	142	284	LVTTL, LVCMOS2,	-6 -7 -8	-6 -7			1.45M
XCV300E	412K	32 x 48	3,072	6,912	6,144	96K	32	128K	NA	25/350	8	YES	YES	NA	158	316	LVCMOS18, PCI33,	-6 -7 -8	-6 -7			1.88M
XCV400E	570K	40 x 60	4,800	10,800	9,600	150K	40	160K	NA	25/350	8	YES	YES	NA	202	404	PCI66, GTL, GTL+,	-6 -7 -8	-6 -7			2.7M
XCV600E	986K	48 x 72	6,912	15,552	13,824	216K	72	288K	NA	25/350	8	YES	YES	NA	256	512	HSTL I, HSTL III, HSTL IV,	-6 -7 -8	-6 -7			3.97M
XCV1000E	1,569K	64 x 96	12,288	27,648	24,576	384K	96	384K	NA	25/350	8	YES	YES	NA	330	660	SSTL3 I, SSTL3 II,	-6 -7 -8	-6 -7			6.6M
XCV1600E	2,188K	72 x 180	25,920	34,992	51,840	486K	144	576K	NA	25/350	8	YES	YES	NA	362	724	SSTL21, SSTL211, BLVDS,	-6 -7 -8	-6 -7			8.4M
XCV2000E	2,542K	80 x 120	19,200	43,200	38,400	600K	160	640K	NA	25/350	8	YES	YES	NA	402	804	LVDS, LVPECL	-6 -7 -8	-6 -7			10.2M
XCV2600E	3,264K	92 x 138	25,392	57,132	50,784	793.5K	184	736K	NA	25/350	8	YES	YES	NA	402	804		-6 -7 -8	-6 -7			13M
XCV3200E	4,074K	104 x 156	32,448	73,008	64,896	1014K	208	832K	NA	25/350	8	YES	YES	NA	402	804		-6 -7 -8	-6 -7			16.3M
Virtex-EM Family — 1.8 Volt																		**.18um Six Layer Metal Process**				
XCV405E	1.31M	40 x 60	4,800	10,800	9,600	150K	140	560K	NA	25/350	8	YES	YES	NA	202	404	Same As	-6 -7 -8	-6 -7	ISP	OTP	3.43M
XCV812E	2.54M	56 x 84	9,408	21,168	18,816	294K	280	1120K	NA	25/350	8	YES	YES	NA	278	556	Virtex-E	-6 -7 -8	-6 -7			6.52M
Spartan-IIE Family — 1.8 Volt																		**.18um Six Layer Metal Process**				
XC2S50E	50K	16 x 24	768	1,728	1,536	24K	8	32K	NA	25/320	4	YES	YES	NA	84	182	LVTTL,LVCMOS2, LVCMOS18,	-6 -7	-6	ISP	OTP	0.6M
XC2S100E	100K	20 x 30	1,200	2,700	2,400	37.5K	10	40K	NA	25/320	4	YES	YES	NA	86	202	PCI33, PCI66, GTL, GTL+,	-6 -7	-6			0.9M
XC2S150E	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	NA	25/320	4	YES	YES	NA	114	263	HSTL I, HSTL III, HSTL IV, SSTL3 I,	-6 -7	-6			1.1M
XC2S200E	200K	28 x 42	2,352	5,292	4,704	73.5K	14	56K	NA	25/320	4	YES	YES	NA	120	289	SSTL3 II, SSTL2 I,SSTL2 II, AGP-2X,	-6 -7	-6			1.4M
XC2S300E	300K	32 x 48	3,072	6,912	6,144	96K	16	64K	NA	25/320	4	YES	YES	NA	120	329	CTT, LVDS, BLVDS, LVPECL	-6 -7	-6			1.9M
Spartan-II Family — 2.5 Volt																		**.22/.18um Six Layer Metal Process**				
XC2S15	15K	8 x 12	192	432	384	6K	4	16K	NA	25/200	4	YES	YES	NA	NA	86	LVTTL, LVCMOS2,	-5 -6	-5	ISP	OTP	0.2M
XC2S30	30K	12 x 18	432	972	864	13.5K	6	24K	NA	25/200	4	YES	YES	NA	NA	132	PCI33 (3.3V & 5V),	-5 -6	-5			0.4M
XC2S50	50K	16 x 24	768	1,728	1,536	24K	8	32K	NA	25/200	4	YES	YES	NA	NA	176	PCI66 (3.3V), GTL, GTL+,	-5 -6	-5			0.6M
XC2S100	100K	20 x 30	1,200	2,700	2,400	37.5K	10	40K	NA	25/200	4	YES	YES	NA	NA	196	HSTL I, HSTL III, HSTL IV,	-5 -6	-5			0.8M
XC2S150	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	NA	25/200	4	YES	YES	NA	NA	260	SSTL3 I, SSTL3 II, SSTL2 I,	-5 -6	-5			1.1M
XC2S200	200K	28 x 42	2,352	5,292	4,704	73.5K	14	56K	NA	25/200	4	YES	YES	NA	NA	284	SSTL2 II, AGP-2X, CTT	-5 -6	-5			1.4M

Note: 1. System Gates include 20-30% of CLBs used as RAM
2. A Logic Cell is defined as a 4 input LUT and a register
DCM – Digital Clock Management

Important: Verify all Data with Device Data Sheet

Xilinx Spartan FPGAs

		CLB Resources					BLK RAM		CLK Resources					I/O Features				Speed				
	System Gates (see note 1)	CLB Array (Row X Col)	Number of Slices	Logic Cells (see note 2)	CLB Flip-Flops	Max. Distributed RAM Bits	# Block RAM	Block RAM (kbits)	# Dedicated Multipliers	DLL Frequency (min/max)	# DLL's	Frequency Synthesis	Phase Shift	Digitally Controlled Impedance	Number of Differential I/O Pairs	Max. I/O	I/O Standards	Commercial Speed Grades (slowest to fastest)	Industrial Speed Grades (slowest to fastest)	Serial PROM Family		Config. Memory (Bits)
Spartan-IIE Family — 1.8 Volt																						.18um Six Layer Metal Process
XC2S50E	50K	16 x 24	768	1,728	1,536	24K	8	32K	NA	25/320	4	YES	YES	NA	84	182	LVTTL,LVCMOS2, LVCMOS18,	-6 -7	-6	ISP	OTP	0.6M
XC2S100E	100K	20 x 30	1,200	2,700	2,400	37.5K	10	40K	NA	25/320	4	YES	YES	NA	86	202	PCI33, PCI66, GTL, GTL+,	-6 -7	-6			0.9M
XC2S150E	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	NA	25/320	4	YES	YES	NA	114	263	HSTL I, HSTL III, HSTL IV, SSTL3 I,	-6 -7	-6			1.1M
XC2S200E	200K	28 x 42	2,352	5,292	4,704	73.5K	14	56K	NA	25/320	4	YES	YES	NA	120	289	SSTL3 II, SSTL2 I,SSTL2 II, AGP-2X,	-6 -7	-6			1.4M
XC2S300E	300K	32 x 48	3,072	6,912	6,144	96K	16	64K	NA	25/320	4	YES	YES	NA	120	329	CTT, LVDS, BLVDS, LVPECL	-6 -7	-6			1.9M
Spartan-II Family — 2.5 Volt																						.22/.18um Six Layer Metal Process
XC2S15	15K	8 x 12	192	432	384	6K	4	16K	NA	25/200	4	YES	YES	NA	NA	86	LVTTL, LVCMOS2,	-5 -6	-5	ISP	OTP	0.2M
XC2S30	30K	12 x 18	432	972	864	13.5K	6	24K	NA	25/200	4	YES	YES	NA	NA	132	PCI33 (3.3V & 5V),	-5 -6	-5			0.4M
XC2S50	50K	16 x 24	768	1,728	1,536	24K	8	32K	NA	25/200	4	YES	YES	NA	NA	176	PCI66 (3.3V), GTL, GTL+,	-5 -6	-5			0.6M
XC2S100	100K	20 x 30	1,200	2,700	2,400	37.5K	10	40K	NA	25/200	4	YES	YES	NA	NA	196	HSTL I, HSTL III, HSTL IV,	-5 -6	-5			0.8M
XC2S150	150K	24 x 36	1,728	3,888	3,456	54K	12	48K	NA	25/200	4	YES	YES	NA	NA	260	SSTL3 I, SSTL3 II, SSTL2 I,	-5 -6	-5			1.1M
XC2S200	200K	28 x 42	2,352	5,292	4,704	73.5K	14	56K	NA	25/200	4	YES	YES	NA	NA	284	SSTL2 II, AGP-2X, CTT	-5 -6	-5			1.4M

PACKAGE OPTIONS AND USER I/O

			Spartan-IIE (1.8V)					Spartan-II (2.5V)					
			XC2S50E	XC2S100E	XC2S150E	XC2S200E	XC2S300E	XC2S15	XC2S30	XC2S50	XC2S100	XC2S150	XC2S200
Pins	**Body Size**	**I/O's**	182	202	263	289	329	86	132	176	196	260	284
PQFP Packages (PQ)													
208	28 x 28 mm		146	146	146	146	146		132	140	140	140	140
VQFP Packages (VQ)													
100	14 x 14 mm							60	60				
TQFP Packages (TQ)													
144	20 x 20 mm		102	102				86	92	92	92		
Chip Scale Packages — wire-bond chip-scale BGA (0.8 mm ball spacing)													
144	12 x 12 mm							86	92				
FGA Packages (FT) — wire-bond fine-pitch thin BGA (1.0 mm ball spacing)													
256	17 x 17 mm		182	182	182	182	182						
FGA Packages (FG) — wire-bond fine-pitch BGA (1.0 mm ball spacing)													
256	17 x 17 mm									176	176	176	176
456	23 x 23 mm			202	263	289	329				196	260	284

Note: 1. System Gates include 20-30% of CLBs used as RAM
2. A Logic Cell is defined as a 4 input LUT and a register
DCM – Digital Clock Management
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.

Xilinx FPGA Package Options and User I/O

VIRTEX-II PRO

Package	Pitch (mm)	Size (mm)	User Available I/Os: XC2VP2	XC2VP4	XC2VP7	XC2VP20	XC2VP50
FG256	1.00	17 x 17	140	140			
FG456	1.00	23 x 23	156	248	248		
FF672	1.00	27 x 27	204	348	396		
FF896	1.00	31 x 31			396	556	
FF1152	1.00	35 x 35				564	692
FF1517	1.00	40 x 40					852
BF957	1.27	40 x 40				564	584

			Virtex-II (1.5V)											Virtex-E (1.8V)											V-EM		Spartan-IIE (1.8V)					Spartan-II (2.5V)					
			XC2V40	XC2V80	XC2V250	XC2V500	XC2V1000	XC2V1500	XC2V2000	XC2V3000	XC2V4000	XC2V6000	XC2V8000	XCV50E	XCV100E	XCV200E	XCV300E	XCV400E	XCV600E	XCV1000E	XCV1600E	XCV2000E	XCV2600E	XCV3200E	XCV405E	XCV812E	XC2S50E	XC2S100E	XC2S150E	XC2S200E	XC2S300E	XC2S15	XC2S30	XC2S50	XC2S100	XC2S150	XC2S200
Pins	Body Size	I/O's	88	120	200	264	432	528	624	720	912	1104	1296	176	176	284	316	404	512	660	724	804	804	804	404	556	182	202	263	289	329	86	132	176	196	260	284
PQFP Packages (PQ)																																					
208	28 x 28 mm																										146	146	146	146	146		132	140	140	140	140
240	32 x 32 mm													158	158	158	158	158	158	158																	
HQFP Packages (HQ)																																					
240	32 x 32 mm														158				158																		
VQFP Packages (VQ)																																					
100	14 x 14 mm																															60	60				
TQFP Packages (TQ)																																					
144	20 x 20 mm																										102	102				86	92	92	92		
Chip Scale Packages — wire-bond chip-scale BGA (0.8 mm ball spacing)																																					
144	12 x 12 mm		88	92	92									94	94	94																86	92				
BGA Packages (BG) — wire-bond standard BGA (1.27 mm ball spacing)																																					
352	40 x 40 mm														196	260	260																				
432	40 x 40 mm																316	316	316																		
560	42.5 x 42.5 mm																	404	404	404	404	404			404	404											
575	31 x 31 mm						328	392	408																												
728	35 x 35 mm								456	516																											
FGA Packages (FT) — wire-bond fine-pitch thin BGA (1.0 mm ball spacing)																																					
256	17 x 17 mm																										182	182	182	182	182						
FGA Packages (FG) — wire-bond fine-pitch BGA (1.0 mm ball spacing)																																					
256	17 x 17 mm		88	120	172	172	172							176	176	176	176																176	176	176	176	
456	23 x 23 mm				200	264	324									284	312											202	263	289	329				196	260	284
676	27 x 27 mm							392	456	484								404	444						404												
680	40 x 40 mm																		512	512	512	512															
860	42.5 x 42.5 mm																			660	660	660															
900	31 x 31 mm																		512	660	700					556											
1156	35 x 35 mm																			660	724	804	804	804													
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																								

Note: 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.

Xilinx CPLD Product Selection Matrix

PRODUCT SELECTION MATRIX

						I/O Features		Speed			Clocking	
	System Gates	Macrocells	Product terms per Macrocell	Input Voltage Compatible	Output Voltage Compatible	Max. I/O	I/O Banking	Min. Pin-to-Pin Logic Delay (ns)	Commercial Speed Grades (fastest to slowest)	Industrial Speed Grades (fastest to slowest)	Global Clocks	Product Term Clocks per Function Block
CoolRunner-II Family — 1.8 Volt												
XC2C32	750	32	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	33	1	3.5	-3 -4 -6	-6	3	17
XC2C64	1500	64	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	64	1	4	-4 -5 -7	-7	3	17
XC2C128	3000	128	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	100	2	4.5	-4 -6 -7	-7	3	17
XC2C256	6000	256	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	184	2	5	-5 -6 -7	-7	3	17
XC2C384	9000	384	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	240	4	6	-6 -7 -10	-10	3	17
XC2C512	12000	512	40	1.5/1.8/2.5/3.3	1.5/1.8/2.5/3.3	270	4	6	-6 -7 -10	-10	3	17
CoolRunner XPLA3 Family — 3.3 Volt												
XCR3032XL	750	32	48	3.3/5	3.3	36		5	-5 -7 -10	-7 -10	4	16
XCR3064XL	1500	64	48	3.3/5	3.3	68		6	-6 -7 -10	-7 -10	4	16
XCR3128XL	3000	128	48	3.3/5	3.3	108		6	-6 -7 -10	-7 -10	4	16
XCR3256XL	6000	256	48	3.3/5	3.3	164		7.5	-7 -10 -12	-10 -12	4	16
XCR3384XL	9000	384	48	3.3/5	3.3	220		7.5	-7 -10 -12	-10 -12	4	16
XCR3512XL	12000	512	48	3.3/5	3.3	260		7.5	-7 -10 -12	-10 -12	4	16
XC9500XV Family — 2.5 Volt												
XC9536XV	800	36	90	2.5/3.3	1.8/2.5/3.3	36	1	5	-5 -7	-7	3	18
XC9572XV	1600	72	90	2.5/3.3	1.8/2.5/3.3	72	1	5	-5 -7	-7	3	18
XC95144XV	3200	144	90	2.5/3.3	1.8/2.5/3.3	117	2	5	-5 -7	-7	3	18
XC95288XV	6400	288	90	2.5/3.3	1.8/2.5/3.3	192	4	6	-6 -7 -10	-7 -10	3	18
XC9500XL Family — 3.3 Volt												
XC9536XL	800	36	90	2.5/3.3/5	2.5/3.3	36		5	-5 -7 -10	-7 -10	3	18
XC9572XL	1600	72	90	2.5/3.3/5	2.5/3.3	72		5	-5 -7 -10	-7 -10	3	18
XC95144XL	3200	144	90	2.5/3.3/5	2.5/3.3	117		5	-5 -7 -10	-7 -10	3	18
XC95288XL	6400	288	90	2.5/3.3/5	2.5/3.3	192		6	-6 -7 -10	-7 -10	3	18

PACKAGE OPTIONS AND USER I/O

		CoolRunner-II						CoolRunner XPLA3						XC9500XV				XC9500XL			
	Body Size	XC2C32	XC2C64	XC2C128	XC2C256	XC2C384	XC2C512	XCR3032XL	XCR3064XL	XCR3128XL	XCR3256XL	XCR3384XL	XCR3512XL	XC9536XV	XC9572XV	XC95144XV	XC95288XV	XC9536XL	XC9572XL	XC95144XL	XC95288XL
PLCC Packages (PC)																					
44	17.5 x 17.5 mm	33	33					36	36					34	34			34	34		
PQFP Packages (PQ)																					
208	28 x 28 mm				173	173	173				164	172	180				168				168
VQFP Packages (VQ)																					
44	12 x 12 mm	33	33					36	36					34	34			34	34		
64	12 x 12 mm																	36	52		
100	16 x 16 mm		64	80	80				68	84											
TQFP Packages (TQ)																					
100	14 x 14 mm														72	81			72	81	
144	20 x 20 mm			100	118	118				108	120	118*				117	117			117	117
Chip Scale Packages (CP) — wire-bond chip-scale BGA (0.5 mm ball spacing)																					
56	6 X 6 mm	33	45						48												
132	8 X 8 mm			100	106																
Chip Scale Packages (CS) — wire-bond chip-scale BGA (0.8 mm ball spacing)																					
48	7 x 7mm							36	40					36	38			36	38		
144	12 x 12 mm									108						117				117	
280	16 x 16 mm										164						192				192
BGA Packages (BG) — wire-bond standard BGA (1.27 mm ball spacing)																					
256	27 x 27 mm																				192
FGA Packages (FT) — wire-bond fine-pitch thin BGA (1.0 mm ball spacing)																					
256	17 x 17 mm				184	212	212				164	212	212								
FBGA Packages (FG) — wire-bond Fineline BGA (1.0 mm ball spacing)																					
256	17 x 17 mm																192				192
324	23 x 23 mm					240	270					220	260								

*JTAG pins and port enable are not pin compatible in this package for this member of the family.
Important: Verify all Data with Device Data Sheet and Product Availability with your local Xilinx Rep

Xilinx Software

	Feature	ISE *WebPACK*	ISE *BaseX*	ISE *Foundation*	ISE *Alliance*
	Platform	PC	PC	PC/UNIX	PC/UNIX
Device Support	Virtex™Series	Up to 300K (Virtex-E and Virtex-II)	Up to 300K	ALL	ALL
	Spartan™ II/IIE Families	ALL	ALL	ALL	ALL
	Spartan™ /XL Families	No	Implementation Only*	Implementation Only*	ALL
	XC4000™ Series	No	Implementation Only*	Implementation Only*	ALL
	CoolRunner XPLA3 / CoolRunner™-II Families	ALL	ALL	ALL	ALL
	XC9500™ Series	ALL	ALL	ALL	ALL
Design Planning	Modular Design	No	Sold as an Option	Sold as an Option	Sold as an Option
	Educational Services	Yes	Yes	Yes	Yes
	Design Services	Sold as an Option	Sold as an Option	Sold as an Option	Sold as an Option
	Support Services	Web Only	Yes	Yes	Yes
Design Entry	Schematic Editor (Gate & Block Level HDL)	Yes	Yes	PC Only	No
	HDL Editor	Yes	Yes	Yes	Yes
	State Diagram Editor	Yes	Yes	PC Only	No
	CORE Generator	No	Yes	Yes	Yes
	System Generator for DSP	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	No	Integration Kit	Integration Kit (PC Only)	Integration Kit (PC Only)
	Leonardo Spectrum	Integration Kit	Integration Kit	Integration Kit	Integration Kit
	Synopsys FPGA Compiler II	EDIF Only	EDIF Only	EDIF Only	EDIF Only
	ABEL	CPLD	CPLD	CPLD (PC Only)	No
Implementation Tools	iMPACT	Yes	Yes	Yes	Yes
	FloorPlanner	Yes	Yes	Yes	Yes
	Xilinx Constraints Editor	Yes	Yes	Yes	Yes
	Timing Driven Place & Route	Yes	Yes	Yes	Yes
	Timing Improvement Wizard	No	Yes	Yes	Yes
Board Level Integration	IBIS Models	Yes	Yes	Yes	Yes
	STAMP Models	Yes	Yes	Yes	Yes
	LMG SmartModels	Yes	Yes	Yes	Yes
	HSPICE Models**	Yes	Yes	Yes	Yes
Verification	HDL Bencher™	Yes	Yes	PC Only	No
	ModelSim Xilinx Edition (MXE II)	ModelSim XE II Starter***	ModelSim XE II Starter***	ModelSim XE II Starter***	ModelSim XE II Starter***
	Static Timing Analyzer	Yes	Yes	Yes	Yes
	ChipScope	No	Sold as an Option	Sold as an Option	Sold as an Option
	FPGA Editor with Probe	No	Yes	Yes	Yes
	ChipViewer	Yes	Yes	Yes	Yes
	XPower (Power Analysis)	Yes	Yes	Yes	Yes
	3rd Party Simulator Support	Yes	Yes	Yes	Yes
IP/CORE	For more information on the complete list of Xilinx IP products, visit the Xilinx IP Center at http://www.xilinx.com/ipcenter				

* For Spartan, SpartanXL, and 4K device families, ISE supports the implementation by reading in gate-level EDIF netlist.

** 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.

Xilinx Configuration Storage Solutions

System ACE	Memory Density	Number of Components	Min board space	Compression	FPGA Config. Mode	Multiple Designs	Software Storage	Removable	IRL Hooks	Max Config. Speed	Non-Volatile Media
System ACE CF	up to 8 Gbit	2	25 cm^2	No	JTAG	Unlimited	Yes	Yes	Yes	30 Mbit/sec	CompactFlash
System ACE MPM	16 Mbit 32 Mbit 64 Mbit	1	12.25 cm^2	Yes	SelectMAP (up to 4 FPGA) Slave-Serial (up to 8 FPGA chains)	Up to 8	No	No	Yes	152 Mbit/sec	AMD Flash Memory
System ACE SC	16 Mbit 32 Mbit 64 Mbit	3	Custom	Yes	SelectMAP (up to 4 FPGA) Slave-Serial (up to 8 FPGA chains)	Up to 8	No	No	Yes	152 Mbit/sec	AMD Flash memory

PROM

	Density	PD8	VO8	SO20	PC20	PC44	VQ44	Core Voltage	I/O Voltage 2.5V	I/O Voltage 3.3V
In-System Programming (ISP) Configuration PROMs										
XC18V256	256Kb			Y	Y		Y	3.3V	Y	Y
XC18V512	512Kb			Y	Y		Y	3.3V	Y	Y
XC18V01	1Mb			Y	Y		Y	3.3V	Y	Y
XC18V02	2Mb					Y	Y	3.3V	Y	Y
XC18V04	4Mb					Y	Y	3.3V	Y	Y
One-Time Programmable (OTP) Configuration PROMs										
XC17V01	1.6Mb		Y	Y	Y			3.3V	Y	Y
XC17V02	2Mb				Y	Y	Y	3.3V	Y	Y
XC17V04	4Mb				Y	Y	Y	3.3V	Y	Y
XC17V08	8Mb					Y	Y	3.3V	Y	Y
XC17V16	16Mb					Y	Y	3.3V	Y	Y

	FPGA	PD8	VO8	SO20	PC20	PC44	VQ44	Core Voltage	I/O Voltage 2.5V	I/O Voltage 3.3V
OTP Configuration PROMs for Spartan-II/IIE										
XC17S50A	XC2S50E	Y	Y	Y				3.3V	Y	Y
XC17S100A	XC2S100E	Y	Y	Y				3.3V	Y	Y
XC17S200A	XC2S150E	Y	Y				Y	3.3V	Y	Y
XC17S200A	XC2S200E	Y	Y				Y	3.3V	Y	Y
XC17S300A	XC2S300E						Y	3.3V	Y	Y
XC17S15A	XC2S15	Y	Y	Y				3.3V	Y	Y
XC17S30A	XC2S30	Y	Y	Y				3.3V	Y	Y
XC17S50A	XC2S50	Y	Y	Y				3.3V	Y	Y
XC17S100A	XC2S100	Y	Y	Y				3.3V	Y	Y
XC17S150A	XC2S150	Y	Y	Y				3.3V	Y	Y
XC17S200A	XC2S200	Y	Y				Y	3.3V	Y	Y

Xilinx Home Page
http://www.xilinx.com

Xilinx Online Support
http://www.xilinx.com/support/support.htm

Xilinx IP Center
http://www.xilinx.com/ipcenter/index.htm

Xilinx Education Center
http://www.xilinx.com/support/education-home.htm

Xilinx Tutorial Center
http://www.xilinx.com/support/techsup/tutorials/index.htm

Xilinx WebPACK
http://www.xilinx.com/sxpresso/webpack.htm

Xilinx IP Selection Guide

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Implementation Example			Key Features	Application Examples
										Occ	MHz	Device		
Communication & Networking														
1 Gigabit Ethernet Full Duplex MAC	Alcatel Technology Licensing Group	AllianceCORE		V-II						46%	31	XC2V1500-5	Single/multi-mode fiber optics; 802.3x full duples flow control; 10-bit SERDES for GMII; Auto-negotiation for 1000BASE-X	NICs, routers, switches, hubs
1 Gigabit Ethernet MAC	Xilinx	LogiCORE	V-IIP	V-II	V-E		S-IIE			23%	125 (XGMII) or 1.25 Gbps	XC2V1000-4	IEEE 802.3-2000 compliant, support 8-bit GMII interface or integrated PCS/PMA interface, supporting 1000BASE-X application	GbE Network Interface Cards (NICs), Edge switches and terabit routers — packet based line cards, iSCSI line cards, PL3 to Gb Ethernet and other bridges
10 Gigabit Ethernet MAC	Xilinx	LogiCORE	V-IIP	V-II						25%	156.25 DDR for XGMII or 4 channels of 3.125Gbps Rocket I/O transceivers	XC2V3000-5	Designed to IEEE 802.3ae, version D4.1 support both 32-bit XGMII parallel interface or XAUI interface, supports 10 GBASE-X, WAN/LAN functionality. Statistics gathering	Layer 2 swtiches/hubs, test equipment, bridge to POS PHY4, iSCSI line cards
10/100 Ethernet MAC	Alcatel Technology Licensing Group	AllianceCORE		V-II				S-II		13%	31	XC2V1500-5	Single/multi-mode fiber optics; 10/100 MII PHY, 10Base-T, 100Base-T/TX/FX/4; RMON & Etherstats	NICs, routers, switches, hubs, printers
10/100 Ethernet MAC	Paxonet Communications	AllianceCORE				V		S-II	S	45%	50	XCV150-4	IEEE 802.3 compliant RMON, MIBs stats, MII support	Ethernet switched, hub, NICS
10/100 Ethernet MAC	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II				Virtex-II	Interfaces through CoreConnect Bus	Networking, communications, processor apps10/100 Ethernet MAC
3G FEC Package	Xilinx	LogiCORE		V-II	V-E	V							Viterbi Decoder, Turbo Codec, Convolutional Encoder	3G Wireless Infrastructure
8b/10b Decoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		0.4%	160	XC2V40-5	Industry std 8b/10b en/decode for serial data transmission	Physical layer of Fiber Channel
8b/10b Encoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		0.4%	160	XC2V40-5	Industry std 8b/10b en/decode for serial data transmission	Physical layer of Fiber Channel
ADPCM, 16 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II		V				89%	16	XCV150-6	Supports G.721 G.723 G.726 G.726a G.727 G.727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM, 32 channel	Xilinx	LogiCORE		V-II	V-E	V		S-II		62%	25	XC2V500	G.726, G.727, 32 duplex channels	DECT, VOIP, Wireless local loop, DSLAM, PBX
ADPCM, 64 channel	Xilinx	LogiCORE		V-II	V-E	V		S-II		61%	27	XC2V500	G.726, G.727, 64 duplex channels	DECT, VOIP, Wireless local loop, DSLAM, PBX10/100 Ethernet MAC
ADPCM, 256 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E					66%	30	XCV400E-8	Supports G.721 G.723 G.726 G.726a G.727 G.727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM, 512 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II		V				49%	6	XCV400E-8		
ADPCM, 768 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II		V				89%	50	XC2V500-5	Supports G.721 G.723 G.726 G.726a G.727 G.727a, u-law, a-law	DECT, VOIP, cordless telephony
ADPCM, 1024 channel	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E			S-II	S	89%	50	XC2V500-5	Supports G.721 G.723 G.726 G.726a G.727 G.727a, u-law, a-law	DECT, VOIP, cordless telephony
AES Decryption core family	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V		S-II		36%	103	XC2V250-5		eCommerce, Banking, Video phones, PDA, satellite communications
AES Encryption Core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		16%	131	XC2V250-5	Supports ECB, OFB, CFB, CBC modes; Supports 128, 192 and 256-bit keys	Transaction and secure communications; surveillance, storage and embedded systems10/100 Ethernet MAC
AES Encryption core family	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V		S-II		26%	134	XC2V250-5		eCommerce, Banking, Video phones, PDA, satellite communications
Bit Stream Analyzer and Data Extractor	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	71%	67	XCV50-6	Data syntax analysis of IP, MPEG, ATM	ATM, IP, MPEG10/100 Ethernet MAC
Bluetooth Baseband Processor	NewLogic GmbH	AllianceCORE		V-II		V						XC2V1000-4	Compliant to Bluetooth v1.1, BQB qualified software for L2CAP, LHP, HC1, voice support	Bluetooth applications
Bluetooth Hardware Baseband Controller	Wipro, Ltd.	AllianceCORE		V-II		V				69%	25	XC2V1500-4		
Cell Assembler	Paxonet Communications	AllianceCORE				V			S		60	XC4005XL-1	Octet wide operation, HEC computation, cell scrambling	ATM adapter cards, routers, switches
Cell Delineation	Paxonet Communications	AllianceCORE				V			S		40	XC4010XL-9	Octet wide operation, HEC verification, cell scrambling	ATM adapter cards, rtouters, switches
Convolutional Encoder	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	2%	144	XCV50-6	code rate, gen. vectors, CMSTR length customizable	Error correction
Convolutional Encoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		10%	26	XC2V40-6	k from 3 to 9, puncturing from 2/3 to 12/13	3G base stations, broadcast, wireless LAN, cable modem, xDSL, satellite com, uwave
CRC10 Generator and Verifier	Paxonet Communications	AllianceCORE				V			S	22%	20	XCS30-4	Separate generator and verifier blocks, compatible with ITU-T I.363 for AAL3/AAL4	ATM, SONET, and Ethernet10/100 Ethernet MAC
CRC32 Generator and Verifier	Paxonet Communications	AllianceCORE				V			S	44%	29	XCS30-4	Separate generator and verifier blocks, compat with ITU-T I.363 for AAL5	ATM, SONET, and Ethernet
DES3 Encryption	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		15%	167	XC2V1000-5		
DES and 3DES Cryptoprocessor	inSilicon Corporation	AllianceCORE				V		S-II	S		48	XC2S150-6	Compliant with ANSI X9.52, 128-bit key or two independent 64-bit keys	Secure communication, data storage
DES and 3DES encryption engine	Memec Core	AllianceCORE							S	79%	25	XCS20-4	NIST certified, supports EBC, CBC, CFB, and OFB	Secure communication, data storage
DES Cryptoprocessor	inSilicon Corporation	AllianceCORE				V		S-II	S		94	XC2S100-6	NIST certified, supports ECB, CBC, CFB, and OFB	Secure communication, data storage1 Gigabit Ethernet MAC
DES Encryption	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		93%	204	XC2V40-5		
Distributed Sample Descrambler	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	14%	74	XCV50-6	Compliant with ITU-T I.432. Parameterizable data width, cell & header length	ATM PHY layer
Distributed Sample Scrambler	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	9%	104	XCV50-6	Compliant with ITU-T I.432 scrambler. Parameterizable data width, cell length, header length	ATM PHY layer
DVB Satellite Modulator	Memec Core	AllianceCORE				V					45-70	XCV50-4	Conforms to ETSI EN 300 421 v1.1.2, selectable convolutional code rate	Digital broadcast, microwave transmitter
Flexbus 4 Interface Core, 16-Channel	Xilinx	LogiCORE		V-II						31%	200	XC2V3000 FG676-5	OIF SPI-4 Phase 1 and Flexbus4 compliant. Fully HW interoperable with AMCC OC-192 framers.	Line card: terabit routers & optical switches
Flexbus 4 Interface Core, 1-Channel	Xilinx	LogiCORE		V-II						12%	200	XC2V1000 FG456-5	OIF SPI-4 Phase 1 and Flexbus4 compliant. Fully HW interoperable with AMCC OC-192 framers.	Line card: terabit routers & optical switches
Flexbus 4 Interface Core, 4-Channel	Xilinx	LogiCORE		V-II						27%	200	XC2V1000 FG456-5	OIF SPI-4 Phase 1 and Flexbus4 compliant. Fully HW interoperable with AMCC OC-192 framers.	Line card: terabit routers & optical switches
HDLC Controller Core, 1-Channel	Xilinx	LogiCORE		V-II	V-E	V		S-II		15%	115	XC2V250	16/32-bit frame seq, 8/16-bit addr insert/delete, flag/zerop insert/detect	X.25, POS, cable modems, frame relay switches, video conferencing over ISDN
HDLC Controller, 1-Channel	Memec Core	AllianceCORE				V		S-II		95%	77	XC2S15-5	16/32-bit frame seq, 8/16-bit addr insert/delete, flag/zerop insert/detection	X.25, Frame Relay, B/D-Channel
HDLC Controller Core, 32-Channel	Xilinx	LogiCORE		V-II	V-E	V		S-II		34%	81	XC2V250	32 full duplex, CRC-16/32, 8/16-bit address insertion/deletion	X.25, POS, cable modems, frame relay switches, video conferencing over ISDN
HDLC-PPP for Packet over SONET	Paxonet Communications	AllianceCORE				V		S-II		75%	80	XC2S150-6	RFC1619 (IP&IPX) POS, 16/32 bit FCS generation and verification, stats	Bridges, switches, WAN links
Interleaver/De-interleaver	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		30%	208	XC2V40-6	Block & convolutional, width up to 256 bits, 256 branches	Broadcast, wireless LAN, cable modem, xDSL, satellite com,uwave nets, digital TV, CDMA2000
Interleaver/Deinterleaver	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	21%	73	XCV50-6	Block & convolutional support. param features. 3GPP, UMTS, GSM, DVB compliant	Channel coding in telecom/wireless, broadcast
Inverse Multiplexer for ATM	ModelWare, Inc.	AllianceCORE		V-II	V-E					100%	31	XC2V1000-4	Supports up to 32 links/32 groups; UTOPIA L2 PHY & ATM I/Fs, supports IDCR	Access systems, multi-service switches, DSLAMs, Basestation controllers
Network Processor	IP Semiconductors A/S	AllianceCORE		V-II	V-E						80	XC2V1500-5	Total Solution requires this core + SPEEDAnalyzer ASIC, 2.5 Gbps full duplex wire speed; network processor (NPV), low power an low device count than competing network processors.	Networking, edge and access, Switches and routers
Noisy Transmission Channel Model	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	23%	100	XCV50-6	Programmable noise generation profile	Noise emulation in transmission channel
PCM Codec, G.711	Xilinx	LogiCORE			V-E	V		S-II		12%	44	XCV50	µ-Law, ITU G.711, EBI for A-Law	Digital telephony, DECT, T1 & E1 Links
PCM Compressor, G.711	Xilinx	LogiCORE			V-E	V		S-II		7%	44	XCV50	µ-Law, ITU G.711, EBI for A-Law	Digital telephony, DECT, T1 & E1 Links
PCM Expander, G.711	Xilinx	LogiCORE			V-E	V		S-II		6%	57	XCV50	Digital telephony, DECT, T1 & E1	Digital telephony, DECT, T1 & E1 Links
POS-PHY L3 Link Layer Interface, 16-Ch	Xilinx	LogiCORE		V-II						40%	104	XC2V1000 FG456-4	OIF SPI-3 compliant. Fully HW interoperable with PMC-Sierra OC-48 framers.	Line cards, iSCSI cards, gigabit routers and switches
POS-PHY L3 Link Layer Interface, 1-Ch	Xilinx	LogiCORE	V-IIP	V-II	VE	V	S-IIE	S-II		6%	200	XC2V1000 FG456-4	OIF SPI-3 compliant. Fully HW interoperable with PMC-Sierra OC-48 framers.	Line cards, iSCSI cards, gigabit routers and switches
POS-PHY L3 Link Layer Interface, 2-Ch	Xilinx	LogiCORE	V-IIP	V-II	VE	V	S-IIE	S-II		55%	200	XCV50E-8	OIF SPI-3 compliant. Fully HW interoperable with PMC-Sierra OC-48 framers.	Line cards, iSCSI cards, gigabit routers and switches
POS-PHY L3 Link Layer Interface, 4-Ch	Xilinx	LogiCORE	V-IIP	V-II	VE	V	S-IIE	S-II		15%	200	XC2V1000 FG456-4	OIF SPI-3 compliant. Fully HW interoperable with PMC-Sierra OC-48 framers.	Line cards, iSCSI cards, gigabit routers and switches

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Implementation Example Occ	MHz	Device	Key Features	Application Examples
Communication & Networking (continued)														
POS-PHY L3 Physical Layer Interface	Xilinx	LogiCORE			VE					52%	104	XCV50E-8	OIF SPI-3 compliant.	Line cards, iSCSI cards, gigabit routers and switches
POS-PHY L4 Multi-Channel Interface	Xilinx	LogiCORE		V-II						29%	350 DDR	XC2V3000 FG676-5	OIF SPI-4 Phase 2 compliant. Fully HW interoperable with PMC-Sierra and Mindspeed OC-192 framers.	Line cards, switches, routers and optical switches
POS-PHY L4 to Flexbus 4 Bridge	Xilinx	LogiCORE		V-II						35%	200 on FB4, 350 DDR on PL4, 175 internal	XC2V3000 FG676-5	OIF SPI-4 Phase 1& 2 compliant. Fully HW interoperable with AMCC, PMC-Sierra and Mindspeed OC-192 framers.	
POS-PHY Level 3 Link Layer Interface Core, 48-Ch	Xilinx	LogiCORE		V-II						33%	104	XC2V6000 FF1152-4	OIF SPI-3 compliant. Fully HW interoperable with PMC-Sierra OC-48 framers.	Line cards, iSCSI cards, gigabit routers and switches
Reed-Solomon Decoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II	S	40%	98	XC2V250-6	Std or custom coding, 3-12 bit symbol width, up to 4095 symbols, error & erasure decoding	Broadcast, wireless LAN, cable modem, xDSL, satellite com,uwave nets, digital TV
Reed-Solomon Decoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II		V				97%	61	XC2V500-5	Parameterizable, RTL available	Error correction, wireless, DSL
Reed Solomon Decoder	Amphion Semiconductor Ltd.	AllianceCORE				V			S	50%	50	XCV100-4		Error correction
Reed Solomon Decoder	Memec Core	AllianceCORE				V		S-II	S	83%	73	XCV50-6	Customizable, >580 Mbps	Error correction
Reed-Solomon Encoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II	S	42%	180	XC2V40-6	Std or cust coding, 3-12 bit width, up to 4095 symbols with 256 check symb.	Broadcast, wireless LAN, cable modem, xDSL, satellite com,uwave nets, digital TV
Reed Solomon Encoder	Amphion Semiconductor Ltd.	AllianceCORE				V			S	11%	82	XCV50-4		Error correction
Reed Solomon Encoder	Memec Core	AllianceCORE				V		S-II	S	12%	113	XCV50-6	Customizable, > 900 Mbps	Error Correction
SDLC Controller	CAST, Inc.	AllianceCORE		V-II		V				11%	158	XC2V1000-5	Like Intel 8XC152 Global Serial Channel, Serial Comm., HDLC apps, telecom	Embedded systems, professional audio, video
SHA-1 Encryption processor	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		24%	79	XC2V500-6	SHA-1 algorithm compliant	Secure comms, video suveillance, data storage, financial transactions
T1 Deframer	Xilinx	LogiCORE			V-E	V		S-II		15%	54	XC2S150		ISDN PRA links, mux equip, satellite com, digital PABX, high-speed computer links
T1 Framer	Virtual IP Group	AllianceCORE											D4, ESF, SLC-96 formats. For XC4000.	DSI trunk, PBX I/F
T1/E1 Framer	Xilinx	LogiCORE			V-E	V		S-II		7%	72	XC2S150		ISDN PRA links, mux equip, satellite com, digital PABX, high-speed computer links
Turbo Convolutional Decoder - 3GPP Compliant	Xilinx	LogiCORE		V-II	V-E	V				80%	40	XC2V500	3GPP specs, 2 Mbps, BER=10-6 for 1.5dB SNR	3G Wireless Infrastructure
Turbo Decoder, 3GPP	SysOnChip, Inc.	AllianceCORE		V-II		V					66	XC2V500-5	3GPP/UMTS compliant, IMT-2000, 2Mbps data rate	Error correction, wireless
Turbo Decoder, DVB-RCS	iCoding Technology, Inc.	AllianceCORE		V-II	V-E	V				44%	71	XC2V2000-5	DVB-RCS compliant, 9Mbps, data rate, switchable code rates and frame sizes	Error correction, wireless, DVB, Satellite data link
Turbo Decoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II		V				70%	65	XC2V2000-5	3GPP/UMTS compliant, >2Mbps data rate	Error correction, wireless
Turbo Convolutional Encoder - 3GPP Compliant	Xilinx	LogiCORE		V-II	V-E	V				65%	60	XC2V250	Compliant w/ 3GPP, puncturing	3G Wireless Infrastructure
Turbo Encoder, DVB-RCS	iCoding Technology, Inc.	AllianceCORE		V-II	V-E	V				2%	69	XC2V2000-5	DVB-RCS compliant, 9Mbps, data rate, switchable code rates and frame sizes	Error correction, wireless, DVB, Satellite data link
Turbo Encoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II		V		S-II		48%	120	XC2V80-5	3GPP/UMTS compliant, upto 4 interleaver laws	Error correction, wireless
UTOPIA level 2 slave inteface	Paxonet Communications	AllianceCORE											Cell handshake in SPHY mode, 8/16 bit operation, internal FIFO, detects runt cells	ATM PHY layer
UTOPIA Level-2 PHY Side RX Interface	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	8%	53	XCV50-6	Protocol conversion from Pb (RACE BLNT) to UTOPIA L2, 8/16 bit operation	ATM PHY layer
UTOPIA Level-2 PHY Side TX Interface	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	10%	61	XCV50-6	Protocol conversion from UTOPIA L2 Pb (RACE BLNT), 8/16 bit operation	ATM PHY layer
UTOPIA Level-3 ATM Receiver	inSilicon Corporation	AllianceCORE				V		S-II	S					
UTOPIA Level-3 ATM Transmitter	inSilicon Corporation	AllianceCORE				V		S-II	S					
UTOPIA Level-3 PHY Receiver	inSilicon Corporation	AllianceCORE				V		S-II	S					
UTOPIA Level-3 PHY Transmitter	inSilicon Corporation	AllianceCORE				V		S-II	S					
UTOPIA Master	Paxonet Communications	AllianceCORE				V		S-II	S				SPHY, MPHY, HEC processing, round robin polling, ind. transmitter receiver	ATM PHY layer
UTOPIA Slave	Paxonet Communications	AllianceCORE				V		S-II	S	26%	79	XCV50-4	Cell handshake in SPHY mode, 8/16 bit operation, 32 bit FIFO interface, detects runt cells	ATM PHY layer
Viterbi Decoder, IEEE 802-compatible	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		35%	157	XC2V500-5	Parameterizable source code with constraint length(k)=7, G0=171, G1=133; or G0=133, G1=17, includes BestState Logic, ability to change code rate and traceback depth on the fly, supports Trellis Coded Modulation, IEEE802.11a/16a compatible, reaches OC3 (155Mbps) and higher	L/MMDS, broadcast equip, wireless LAN, cable modem, xDSL, sat com, uwave nets
Viterbi Decoder, General Purpose	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		38%	128	XC2V500-5	Puncturing, serial & parallel architecture, dynamic rate change, parameterized constraint length, soft/hard decision with programmable number of soft bits, dual rate decoder, erasure pins for external puncturing, compatible with standards such as DVB ETS, 3GPP2, IEEE802.16, Hiperlan, Intelsat IESS-308/309	3G base stations, broadcast, wireless LAN, cable modem, xDSL, satellite com, uwave, CDMA2000
Viterbi Decoder	Telecom Italia Lab S.p.A.	AllianceCORE		V-II	V-E			S-II		74%	91	XC2V500-5	Radix-2/radix4 architectures, BER, depuncturing. Code rate, contraint length parameterizable	Data transmission, wireless
Digital Signal Processing														
1024-Point Complex FFT/IFFT	Xilinx	LogiCORE			V-E	V								
1024-Point Complex FFT/IFFT for Virtex-II	Xilinx	LogiCORE		V-II						62%	41us, 100MHz	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
16-Point Complex FFT/IFFT	Xilinx	LogiCORE		V-II	V-E	V								
16-Point Complex FFT/IFFT for Virtex-II	Xilinx	LogiCORE		V-II						37%	123ns 130MHz	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
256-Point Complex FFT/IFFT	Xilinx	LogiCORE		V-II	V-E	V								
32-Point Complex FFT/IFFT	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		29%	110	XC2V500-6	Complex FFT, Forward and Inverse transform. Supports bit precisions from 2-32 bits. Embedded memory	
64-, 256-, 1024-Point Complex FFT/IFFT	Xilinx	LogiCORE	V-IIP	V-II						32%	140	XC2V1500-6	64-,256-,1024-point programmable point size, forward and inverse transform, 16-bit complex data, 18-bit phase factors, 2's complement, built-in memories, programmable data scaling, 140 MHz, 1024-point transform — 7.31 us, 256-point transform — 1.83 us, 64-point transform - 0.46 us	Error correction
64-Point Complex FFT/IFFT	Xilinx	LogiCORE			V-E	V								
64-Point Complex FFT/IFFT for Virtex-II	Xilinx	LogiCORE		V-II						38%	1.9us, 100MHz	XC2V500	16 bit complex data, 2's comp, forward and inverse transform	
Bit Correlator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					4096 taps, serial/parallel input, 4096 bits width	
Cascaded Integrator Comb (CIC) Filter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					32 bits data width, rate change from 8 to 16384	
Comb Filter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Cascaded Integrator Comb Filter	
CORDIC	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					Polar to rectangular, rectangular to polar, sin & cos, sinh & cosh, atan & atanh, square root	Digital receivers
Digital Down Converter (DDC)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		47%	116	XC2V500-5	Configurable datapath comprising mixer, DDS, optional CIC filter, optional polyphase decimators, 25 to 108db DDS spurious free dynamic range	Wireless & wireline communication systems such as software defined radios, digital receivers, cable modems, BPSK, QPSK, QAM demodulators, spread spectrum communication systems, CDMA2000 & 3G basestations
Direct Digital Synthesizer (DDS)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		12%	245	XC2V80-6	8-65K samples, 32-bits output precision, phase dithering/offset	
Distributed Arithmetic (DA) FIR Filter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					32-bit input/coeff width, 1024 taps, 1-8 chan, polyphase, online coeff reload	
Dual-Channel Numerically Controlled Oscillator	Xilinx	LogiCORE			V-E	V	S-IIE	S-II	S				Not recommended for new designs. Suggested replacement: Direct Digital Synthesizer	
LFSR, Linear Feedback Shift Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					168 input widths, SRL16/register implementation	
MAC FIR	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		16%		XC2V250	Single rate, Polyphase Decimator, Polyphase Interpolator	3G base stations, wireless communications, image filtering
Numerically Controlled Oscillator	Xilinx	LogiCORE			V-E	V	S-IIE	S-II	S				Not recommended for new designs. Suggested replacement: Direct Digital Synthesizer	

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Implementation Example Occ	MHz	Device	Key Features	Application Examples
Digital Signal Processing (continued)														
Parallel Distributed Arithmetic FIR Filter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Arithmetic FIR Filter	
Serial Distributed Arithmetic FIR Filter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Arithmetic FIR Filter	
Time-Skew Buffer, Nonsymmetric 16-Deep	Xilinx	LogiCORE							S				Not recommended for new designs.	
Time-Skew Buffer, Nonsymmetric 32-Deep	Xilinx	LogiCORE							S				Not recommended for new designs.	
Time-Skew Buffer, Symmetric 16-Deep	Xilinx	LogiCORE							S					
TMS32025 DSP Processor core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		66%	63	XC2V500-5		
Math Functions														
1s and 2s Complement	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Bus Gate or Twos Complementer	
Accumulator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256s bit wide	
Adder Subtracter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256s bit wide	
Constant Coefficient Multiplier	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Multiplier	
Constant Coefficient Multiplier - Pipelined	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Multiplier	
Dynamic Constant Coefficient Multiplier	Xilinx	LogiCORE			V-E	V							Not recommended for new designs. Suggested replacement: Multiplier	
Floating Point Adder	Digital Core Design	AllianceCORE		V-II		V		S-II			66	XC2V250-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
Floating Point Divider	Digital Core Design	AllianceCORE		V-II		V		S-II			53	XC2V250-5	Full IEEE-754 compliance, 15 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
Floating Point Multiplier	Digital Core Design	AllianceCORE		V-II	V-E	V		S-II			74	XC2V250-5	Full IEEE-754 compliance, 7 pipelines, 32x32 mult, Single precision real format support	DSP, Math, Arithmetic apps.
Floating Point Square Comparator	Digital Core Design	AllianceCORE		V-II		V		S-II			91	XC2V80-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps.
Floating Point Square Root Operator	Digital Core Design	AllianceCORE		V-II	V-E			S-II			66	XC2V250-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
Floating Point to Integer Converter	Digital Core Design	AllianceCORE		V-II		V		S-II			66	XC2V250-5	Full IEEE-754 compliance, 4 pipelines, Single precision real format support	DSP, Math, Arithmetic apps
Integer to Floating Point Converter	Digital Core Design	AllianceCORE		V-II		V		S-II			73	XC2V250-5	Full IEEE-754 compliance, double word input, 2 pipelines, Single precision real output	DSP, Math, Arithmetic apps
Integrator	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Cascaded Integrator Comb Filter	
Multiply Accumulator (MAC)	Xilinx	LogiCORE	V-IIP	V-II		V	S-IIE	S-II					Input width up to 32 bits, 65-bit accumulator, truncation rounding	
Multiply Generator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					64-bit input data width, constant, reloadable or variable inputs, parallel/sequential implementation	
Parallel Multipliers - Area Optimized	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Multiplier	
Pipelined Divider	Xilinx	LogiCORE		V-II	V-E	V	S-IIE	S-II	S				32-bit input data width, multiple clock per output	
Registered Adder	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Adder Subtracter	
Registered Loadable Adder	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Adder Subtracter	
Registered Loadable Subtracter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Adder Subtracter	
Registered Scaled Adder	Xilinx	LogiCORE							S				Not recommended for new designs.	
Registered Serial Adder	Xilinx	LogiCORE							S				Not recommended for new designs.	
Registered Subtracter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Adder Subtracter	
Scaled-by-One-Half Accumulator	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Accumulator	
Sine Cosine Look Up Table	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II	S	12%	270	XC2V40-6	3-10 bit in, 4-32 bit out, distributed/block ROM	
Square Root	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: CORDIC	
Twos Complementer	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					Input width up to 256 bits	
Variable Parallel Virtex Multiplier	Xilinx	LogiCORE			V-E	V	S-IIE	S-II					Not recommended for new designs. Suggested replacement: Multiplier	
Memories & Storage Elements														
ATM Utopia Level 2	Xilinx	LogiCORE	V-IIP											
Block Memory, Dual-Port	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits, 2-13K words	
Block Memory, Single-Port	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits, 2-128K words	
Content Addressable Memory (CAM)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-512 bits, 2-10K words, SRL16	
CAM, for Internet Protocol	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	19%	49	XCV50-6		IP routers
Distributed Memory	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-1024 bit, 16-65536 word, RAM/ROM/SRL16, opt output regs and pipelining	
FIFO, Asynchronous	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits, 15-65535 words, DRAM or BRAM, independent I/O clock domains	
FIFO, Synchronous	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Synchronous FIFO supporting Spartan-II/Spartan-IIE	
FIFO, Synchronous	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits, 16-256 words, distributed/block RAM	
Pipelined Delay Element	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Memory using SRL16 based memory type	
Registered ROM	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Memory	
Registered Single Port RAM	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Distributed Memory	
Microprocessors, Controllers & Peripherals														
10/100 Ethernet MAC	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II				Virtex-II	Interfaces through CoreConnect Bus	Networking, communications, processor apps
16450 UART	Virtual IP Group	AllianceCORE							S		60	XC2S50-6	Independently controlled transmit, receive and data interrupts. 16X clock.	Serial data applications, modems
16550 UART w/FIFOs	Virtual IP Group	AllianceCORE							S		16	XCS20-4	Prog. Data width, parity, stop bits. 16X internal clock, FIFO mode, false start bit detection	Serial data applications, modems
16550 UART w/FIFOs & synch interface	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		57%	87	XC2V80-5	Supports 16450 and 16550a synchronous s/w; programmable serial interface	Serial and modem applications
16-bit proprietary RISC Processor	Loarant Corporation	AllianceCORE		V-II	V-E	V		S-II			91 MHz	XC2V500-5	44 opcodes, 64-K word data, program, Harvard arch.	Control functions, State machines, Coprocessor
16-Word-Deep Registered Look Up Table	Xilinx	LogiCORE							S					
2901 Microprocessor Slice	CAST, Inc.	AllianceCORE				V		S-II		19%	36	XC2S50-6	Eight function ALU, 4 status flags- Carry, Overflow, Zero and Negative	Simple microcontroller applications
2910A Microprogram Controller	CAST, Inc.	AllianceCORE				V		S-II	S	12%	63	XCV50-6	Based on AMD 2910a	High-speed bit slice design
32-bit Java Processor	Digital Communications Technologies, Ltd.	AllianceCORE		V-II		V		S-II		33%	40	XC2V1000-5	32bit data, 24 bit address, 3 Stage pipeline, Java/C dev. tools	Internet appliance, industrial control, HAVi multimedia, set top boxes
68000 compatible microprocessor	CAST, Inc.	AllianceCORE		V-II		V				89%	32	XC2V500-5	MC68000 Compatible	Embedded systems, professional audio, video
80186 compatible processor	eInfochips Pvt. Ltd.	AllianceCORE		V-II		V				77%	49	XC2V1000-5	i80186 compatible plus enhanced mode	Industrial, wireless, communications, embedded
8051 compatible microcontroller	CAST, Inc.	AllianceCORE				V		S-II		52%	68	XCV200E-8	80C31 instruction set, 8 bit ALU, 8 bit control, 32 bit I/O ports, two 16 bit timer/counters, SFR I/F	Embedded systems, telecom

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Implementation Example			Key Features	Application Examples
										Occ	MHz	Device		
Microprocessors, Controllers & Peripherals (continued)														
8051 compatible microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			73	XC2V250-5	80C31 instruction set, RISC architecture 6.7X faster than standard 8051	Embedded systems, telecom, video
8051 compatible microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			90	XC2V250-5	80C31 instruction set, high speed multiplier, RISC architecture 6.7X faster than standard 8051	Embedded systems, telecom, video
8051 compatible microcontroller	Digital Core Design	AllianceCORE				V		S-II	S	33%	29.8	XCV300-6	12X faster (average) and code compatible wrt legacy 8051, verification bus monitor, SFR interface	Telecom, industrial, high speed control
8051 compatible microcontroller	Dolphin Integration	AllianceCORE		V-II	V-E			S-II		39%	38	XC2V1000-5	8X faster (average) and code compatible wrt legacy 8051, verification bus monitor, SFR interface, DSP focused	DSP, Telecom, industrial, high speed control
80515 High-speed 8-bit RISC Microcontroller	CAST, Inc.	AllianceCORE				V				90%	42	XCV200E-8	RISC implementation, 8 bit ALU, 8 bit control, 32 bit I/O, 16 bit timer/counters, SFR I/F, ext. memory I/F	High speed embedded systems, audio, video
8052 compatible microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			71	XC2V250-5	80C31 instruction set, high speed multiplication and division, RISC architecture 6.7X faster than standard 8051	Embedded systems, telecom, video
80530 8-bit compact microcontroller	CAST, Inc.	AllianceCORE				V		S-II		88%	51	XC2S150-6	32 bit I/O, 3 counters, interrupt controller, SFR interface, dual data pointer	Low cost embedded systems, telecom
80530 8-bit Microcontroller	CAST, Inc.	AllianceCORE				V				81%	66	XCV200E-8	32 bit I/O, 3 counters, 27-bit watchdog timer, 3-priority interrupt controller, SFR interface	Embedded systems, telecom
80C51 compatible RISC microcontroller	CAST, Inc.	AllianceCORE			V-E	V		S-II		75%	34	XC2S150-6	12X faster, SRF I/F	Embedded systems
8237 Programmable DMA Controller	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		40%	34	XC2S100-6	4 independent DMA channels (expandable), software DM	Microprocessor based systems
8250 UART	Memec Core	AllianceCORE							S	58%	10	XCS10-4	DC to 625K baud	Serial communications
8254 programmable interval timer/counter core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		74%	68	XC2V80-5	Status feedback, counter latch, Square wave mode, 6 counter modes, binary/BCD count, LSB/MSB R/W	Event counterr, baud rate generator
8254 Programmable Timer	Virtual IP Group	AllianceCORE				V		S-II	S					
8254 programmable timer/counter	eInfochips Pvt. Ltd.	AllianceCORE		V-II				S-II		11%	51	XC2V1000-4	Status red, Six prog, counter modes, Intel8254 like	Processor I/O interface
8255 programmable I/O controller	eInfochips Pvt. Ltd.	AllianceCORE		V-II				S-II		2%	175	XC2V1000-5	Three 8-bit parallel ports, 24 programmable IO lines, 8-bit bidi data bus	Processor I/O interface
8255 Programmable Peripheral Interface	Memec Core	AllianceCORE							S	64%	8	XCS05-4	Bit set/reset support	Embedded systems
8255 Programmable Peripheral Interface	Virtual IP Group	AllianceCORE				V		S-II	S		227	XCV50E-8	Three 8-bit peripheral ports, 24 programmable I/O lines, 8-bit bidi data bus	Processor I/O interface
8255A peripheral interface	CAST, Inc.	AllianceCORE				V		S-II		10%	227	XCV50E-8	Three 8-bit peripheral ports, 24 programmable IO lines, 8-bit bidi data bus	Processor I/O interface
8256 Multifunction Microprocessor Support Controller	Memec Core	AllianceCORE							S	89%	10	XCS20-4	Baud rate generator for 13 common baud rates, parallel I/O ports, prog. timer/counters	Communication, embedded systems
8259 Programmable Interrupt Controller	Virtual IP Group	AllianceCORE							S			XC2S50-6	8 vectored priority interrupts, all 8259/A modes programmable- e.g., special mask, buffer	Real-time interrupt based uP designs
8259A programmable interrupt controller	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		90%	142	XC2V40-5	8 vectored priority interrupts, all 8259/A modes programmable- e.g., special mask, buffer	Real-time interrupt based uP designs
8279 Programmable Keyboard Display Interface	Memec Core	AllianceCORE							S	46%	8	XCS20-4	8 char keyboard FIFO, 2-key lockout, n-key rollover, 4-16 char display	Embedded systems interface
ARC 32-bit Configurable RISC Processor	ARC International plc	AllianceCORE			V-E			S-II		89%	37	XC2S150-6	4 stage pipeline, 16 single cycle instructions10, 3 interrupt exception levels, 24 bit stack pointer	32 bit processing, DSP
Configurable Java Processor Core	Derivation Systems, Inc.	AllianceCORE		V-II		V				38%	20	XC2V1000-5	32b data/address optional DES	Internet appliance, industrial control
DDR SDRAM Controller	Memec Core	AllianceCORE		V-II	V-E			S-II		7%	133	XC2V1000-5	DDR SDRAM burst length support for 2,4,8 per access, supports data 16,32, 64, 72.	Digital video, embedded computing, networking
Generic compact UART	Memec Core	AllianceCORE		V-II	V-E	V	S-IIE	S-II	S	15%	73	XC2S50E-6	UART and baud rate generator	Serial data communication
IIC	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II				Virtex-II	Interfaces through OPB to MicroBlaze/PPC405	Networking, communications, processor apps
Memory Tests BRAM	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II						
Memory Tests DDR	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II						
Memory Tests SRAM	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II						
Memory Tests ZBT	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II						
MicroBlaze Soft RISC Processor	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	Soft RISC Processor, 82 D-MIPS, 125 MHz, 900 LUTs in Virtex-II	Networking, communications
OPB Arbiter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB BRAM	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB Bus Structure	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					CoreConnect Bus (OPB)	Processor applications
OPB Flash	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB GPIO	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB Interrupt Controller	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB IPIF Address Decode	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					CoreConnect Bus (OPB), Interface to custom IP	PowerPC embedded system design
OPB IPIF DMA	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					CoreConnect Bus (OPB), Interface to custom IP	PowerPC embedded system design
OPB IPIF Interrupt Controller	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					CoreConnect Bus (OPB), Interface to custom IP	PowerPC embedded system design
OPB IPIF Master/Slave Attachment	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					CoreConnect Bus (OPB), Interface to custom IP	PowerPC embedded system design
OPB IPIF Read/Write Packet FIFO	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					CoreConnect Bus (OPB), Interface to custom IP	PowerPC embedded system design
OPB IPIF Scatter/Gather	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					CoreConnect Bus (OPB), Interface to custom IP	PowerPC embedded system design
OPB Memory Interface (Flash, SRAM)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB SRAM	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB Timebase/WDT	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB Timer/Counter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB UART (16450, 16550)	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	Interfaces through OPB to MicroBlaze/PPC405	Processor applications
OPB UART Lite	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II			125	XC2V80-5	Bundled in the Development Kit, CoreConnect Bus (OPB)	Processor applications
OPB ZBT	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					CoreConnect Bus (OPB)	Processor applications
PIC125x fast RISC microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			126	XC2V80-5	PIC 12c4x like, 2X faster, 12-bit wide instruction set, 33 instructions	Embedded systems, telecom, audio and video
PIC1655x fast RISC microcontroller	Digital Core Design	AllianceCORE		V-II		V		S-II			140	XC2V80-5	S/W compatible with PIC16C55X, 14-bit instruction set, 35 instructions	Embedded systems, telecom, audio and video
PIC165X compatible microcontroller	CAST, Inc.	AllianceCORE		V-II		V		S-II		61%	128	XC2V80-5	Microchip 16C5X PIC like	Embedded systems, telecom
PIC165x fast RISC microcontroller	Digital Core Design	AllianceCORE		V-II	V-E	V		S-II			126	XC2V80-5	PIC 12c4x like, 2X faster, 12-bit wide instruction set, 33 instructions	Embedded systems, telecom, audio and video
PLB Arbiter	Xilinx	LogiCORE	V-IIP										CoreConnect Bus (PLB)	PowerPC embedded system design
PLB BRAM	Xilinx	LogiCORE	V-IIP										CoreConnect Bus (PLB)	PowerPC embedded system design
PLB Flash	Xilinx	LogiCORE	V-IIP										CoreConnect Bus (PLB)	PowerPC embedded system design

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Implementation Example			Key Features	Application Examples
										Occ	MHz	Device		
Microprocessors, Controllers & Peripherals (continued)														
PLB SRAM	Xilinx	LogiCORE	V-IIP										CoreConnect Bus (PLB)	PowerPC embedded system design
PLB ZBT	Xilinx	LogiCORE	V-IIP										CoreConnect Bus (PLB)	PowerPC embedded system design
PLB<->OPB Bridge	Xilinx	LogiCORE	V-IIP										CoreConnect Bus (PLB)	PowerPC embedded system design
PowerPC Bus Master	Eureka Technology	AllianceCORE				V				3%	80	XCV400-6		
PowerPC Bus Slave	Eureka Technology	AllianceCORE			V-E					10%	80	XCV400E-8		
PPC405 Boot Code	Xilinx	LogiCORE	V-IIP										CoreConnect Bus (PLB)	PowerPC embedded system design
SDRAM Controller, 200 MHz	Rapid Prototypes, Inc.	AllianceCORE				V		S-II						
SDRAM Controller	Eureka Technology	AllianceCORE				V				37%	91	XCV50-6		
SDRAM Controller	NMI Electronics, Ltd.	AllianceCORE				V		S-II			137	XCV50-6	SDRAM refresh, customizable	Embedded systems using SDRAMs
SPI	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II				Virtex-II	Interfaces through OPB to MicroBlaze/PPC405	Networking, communications, processor apps
System Reset for PPC405	Xilinx	LogiCORE	V-IIP										Used to reset PowerPC based system	
V8-uRISC 8-bit RISC Microprocessor	ARC International plc	AllianceCORE									16	XC4000E	Proprietary 8-bit processor, 8 bit ALU, 16 bit stack pointer, 33 opcodes, 4 addr. Modes, 2 user opcodes	Embedded systems, 8-bit processing apps.
VxWorks Board Support Package (BSP)	Xilinx	LogiCORE	V-IIP										Interfaces HW and WindRiver RTOS	Embedded system design
VxWorks Support	Xilinx	LogiCORE	V-IIP										Interfaces HW and WindRiver RTOS	Embedded system design
Z80 compatible programmable timer/counter core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		9%	167	XC2V500-5		Programmable frequency divider, Pulse counter, pulse generator, interrupt controller
Z80 peripheral I/O controller core	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		49%	76	XC2S50-6		Programmable , dual-port device, keyboards, printers, paper table readers
Z80CPU Microprocessor	CAST, Inc.	AllianceCORE		V-II		V		S-II		56%	72 MHz	XC2V500-5	Zilog Z80 compatible, 8-bit processor	Embedded systems, Communications
Standard Bus Interfaces														
Arbiter	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II	S	24%	33	XCV50-6	Two priority classes - strong/weak, access counters	General purpose bus arbitration
IIC	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II				Virtex-II	Interfaces through OPB to MicroBlaze/PPC405	Networking, communications, processor apps
I2C Bus Controller	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		21%	103	XC2S100-6	Philips I2C 1.1; supports master tx/rx and slave tx/rx modes	Embeded microcontroller and communications
I2C Bus Controller Master	Digital Core Design	AllianceCORE		V-II	V-E			S-II	S		143	XC2V40-5	I2C-like, multi master, fast/std. modes	Embedded systems
I2C Bus Controller Slave	Digital Core Design	AllianceCORE		V-II	V-E			S-II	S		157	XC2V40-5	I2C-like, Slave	Embedded
I2C Bus Controller Slave Base	Digital Core Design	AllianceCORE		V-II	V-E			S-II	S		187	XC2V40-5	I2C-like, Slave	Embedded Systems
I2C Two-Wire Serial Interface Master-Only	Memec Core	AllianceCORE		V-II	V-E	V	S-IIE	S-II		16%	83	XC2S50E-6	I2C-like, multi master fast/std. modes	Embedded microprocessor systems, I2C peripherals
I2C Two-Wire Serial Interface Master-Slave	Memec Core	AllianceCORE		V-II	V-E	V	S-IIE	S-II		28%	108	XC2V80-4	I2C-like, multi master fast/std. modes	Embedded microprocessor systems, I2C peripherals
PCI 64-bit/66-MHz master/target interface	Eureka Technology	AllianceCORE		V-II	V-E					22%	66	XC2V1000-5		
PCI host bridge	Eureka Technology	AllianceCORE		V-II	V-E					22%	33	XC2V1000-5		
PCI32 Interface Design Kit (DO-DI-PCI32-DKT)	Xilinx	LogiCORE		V-II	V-E	V	S-IIE	S-II		6%	66	XC2V1000 FG456-5	Includes PCI32 board, driver development kit, and customer education 3-day training class for US & Canada locations	PC boards, CPCI, Embedded, hiperf video, gb ethernet
PCI32 Interface, IP Only (DO-DI-PCI32-IP)	Xilinx	LogiCORE		V-II	V-E	V	S-IIE	S-II		6%	66	XC2V1000 FG456-5	v2.3 compiliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC add-in boards, CPCI, Embedded
PCI32 Single-Use License for Spartan (DO-DI-PCI32-SP)	Xilinx	LogiCORE					S-IIE	S-II		12%	66	XC2S200 PQ208-6	v2.3 compiliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC add-in boards, CPCI, Embedded
PCI64 & PCI32, IP Only (DO-DI-PCI-AL)	Xilinx	LogiCORE		V-II	V-E	V	S-IIE	S-II		6 - 7%	66	XC2V1000 FG456-5	v2.3 compliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC boards,CPCI,Embedded,hiperf video,gb ethernet
PCI64 Interface Design Kit (DO-DI-PCI64-DKT)	Xilinx	LogiCORE		V-II	V-E	V	S-IIE	S-II		7%	66	XC2V1000 FG456-5	v2.3 compliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly, customer education 3-day training class for US & Canada locations	PC boards, CPCI, Embedded, hiperf video, gb ethernet
PCI64 Interface, IP Only (DO-DI-PCI64-IP)	Xilinx	LogiCORE		V-II	V-E	V	S-IIE	S-II		7%	66	XC2V1000 FG456-5	v2.3 compiliant, assured PCI timing, 3.3/5-V, 0-waitstate, CPCI hot swap friendly	PC boards,CPCI,Embedded,hiperf video,gb ethernet
PCI-X 64/100 Interface for Virtex-II (DO-DI-PCIX64-VE). Includes PCI 64 bit interface at 33 MHz	Xilinx	LogiCORE		V-II						30%	100	XC2V1000 FG456-5	PCI-X 1.0a comp, 64/32-bit, 66 MHz PCI-X initiator and target IF, PCI 2.3 comp, 64/32-bit, 33 MHz PCI initiator and target IF, 3.3 V PCI-X at 33-66 MHz, 3.3 V PCI at 0-33 MHz	Server,Embedded,gb ethernet,U320 SCSI,Fibre Ch,RAID cntl,graphics
PCI-X 64/66 Interface for Virtex-E (DO-DI-PCIX64-VE). Includes PCI 64 bit interface at 33 MHz	Xilinx	LogiCORE		V-II						30%	100	XC2V1000 FG456-5	PCI-X 1.0a comp, 64/32-bit, 66 MHz PCI-X initiator and target IF, PCI 2.3 comp, 64/32-bit, 33 MHz PCI initiator and target IF, 3.3 V PCI-X at 33-66 MHz, 3.3 V PCI at 0-33 MHz	Comm systems, SAN, clustered servers, Ultra 3 SCSI/Fibre Ch RAID, multi-port Gb
RapidIO 8-bit port LP-LVDS Phy LayerXilinx (DO-DI-RIO8-PHY)	LogiCORE		V-II							24%	250 FG456-5	XC2V1000	RapidIO Interconnect v1.1 compliant, verified with Motorola's RapidIO bus functional model v1.4	Routers, switches, backplane, control plane, data path, embedded sys, high speed interface to memory and encryption engines, high end video
Serial Protocol Interface (SPI) Slave motive radio	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE			14%	80	XC2S50E-6		Embedded microprocessor boards, and SOCs, audio/video, home and auto-
USB 1.1 Device Controller	Memec Core	AllianceCORE		V-II	V-E			S-II			12	XC2V1000-5	Compliant with USB1.1 spec., Supports VCI bus, Performs CRC, Supports 1.5 Mbps & 12 Mbps	Scanners, Printers, Handhelds, Mass Storage
USB 1.1 Function Controller	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		92%	48	XC2S200-6	USB 1.1; up to 31 endpoints; suspend and resume power mgmt; remote wake-up	Embedded systems, communications
Video & Image Processing														
2D discrete/inverse cosine transform	Barco-Silex	AllianceCORE		V-II	V-E	V		S-II		77%	133	XC2V250-5		Picture and video, archiving, digital television compression and transmission, teleconference
2D Forward Discrete Cosine Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		42%	83	XC2V500-5		
2D Inverse Discrete Cosine Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		45%	95	XC2V500-5		
Block-based 2D Discrete Wavelet Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		62%	52	XC2V250-5		
Combined 2D Forward/Inverse Discrete Cosine Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		65%	95	XC2V500-5		
Combined 2D Forward/Inverse Discrete Wavelet Transform	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		55%	95	XC2V500-5		
Compact Video Controller	Xylon d.o.o.	AllianceCORE		V-II		V		S-II			88	XC2V250-4	Single & double panel, LCD/CRT support, 4 gray, 256 colors	Video phone, Set-top box, PDA display
DCT/IDCT Forward/Inverse Discrete Cosine Transform, 1-D	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					8-32 pt FDCT, IDCT with 8-24 bits for coeff & input	JPEG, MPEG, H.261, H.263
DCT/IDCT Forward/Inverse Discrete Cosine Transform, 2-D	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II		32%	140	XC2V1000-5	8x8 parameterized FDCT, IDCT & IEEE 1180-1990 compliant IDCT	JPEG, MPEG, H.261, H.263
Discrete Cosine Transform	eInfochips Pvt. Ltd.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		35%	83	XC2V1000-5	8-bit input/12-bit output precision; 76 clock cycle latency	Video coding and security, medical imaging, scanners, copiers, digital still cameras
Fast JPEG color image decoder	Barco-Silex	AllianceCORE		V-II		V				78%	56	XC2V1000-4	Conforms to ISO/IEC Baseline 10918-1, color, multi-scan, Gray-Scale	Video editing, digital camera, scanners
Fast JPEG gray scale image decoder	Barco-Silex	AllianceCORE		V-II	V-E	V				68%	73	XC2V1000-4	Conforms to ISO/IEC Baseline 10918-1, Gray-Scale	Video editing, digital camera, scanners
FIDCT Forward/Inverse Discrete Cosine Transform	Telecom Italia Lab S.p.A.	AllianceCORE				V		S-II		77%	78	XCV200-6	DCT for 8X8, 16X16, IDCT IEEE1180-1990 compliant	JPEG, MPEG, H.26X
Huffman Decoder	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		22%	25	XC2V1000-5		

Function	Vendor Name	IP Type	Virtex-II Pro	Virtex-II	Virtex-E	Virtex	Spartan-IIE	Spartan-II	Spartan	Implementation Example			Key Features	Application Examples
										Occ	MHz	Device		
Video & Image Processing (continued)														
JPEG encoder/decoder	inSilicon Corporation	AllianceCORE				V		S-II	S		20	XCV400E-8	Conforms to ISO/IEC Baseline 10918-1, 4 quantization tables, 4 Huffman tables. Stallable	Video editing, digital camera, scanners
Line-based programmable forward DWT	CAST, Inc.	AllianceCORE		V-II	V-E	V		S-II		145%	51	XC2V250-5		
Longitudinal Time Code Generator	Deltatec S.A.	AllianceCORE				V				36%	80	XC2S15-5	SMTPE/EBU compliant, PAL/NTSC, lock-on external video reference	Audio/Video recording and editing equipment
Motion JPEG Codec core V1.0	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V					40	Virtex-II		
Motion JPEG Decoder core V1.0	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V					42	Virtex-II		
Motion JPEG Encoder core V2.0	Amphion Semiconductor Ltd.	AllianceCORE		V-II	V-E	V					70	Virtex-II		
RGB2YCrCb Color Space Converter	CAST, Inc.	AllianceCORE		V-II	V-E	V	S-IIE	S-II		29%	96	XC2S50E-7	8-bit I/O, 10-bit coeff, 13-bit internal precision; 5-cycle latency; fully synchronous.	Digital RGB to TV output conversion, image filtering, machine vision, still and video image processing.
RGB2YCrCb Color Space Converter	Perigee, LLC	AllianceCORE				V		S-II	S		202	XCV100E-8	One clock cycle throughput	HDTV, real time TV output modulation
YCrCb2RGB Color Space Converter	Perigee, LLC	AllianceCORE				V		S-II	S			XCV100E-8	One clock cycle throughput	HDTV, real time video
Basic Elements														
Binary Counter	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					2-256 bits output width	
Binary Decoder	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					2-256 bits output width	
Bit Bus Gate	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Bit Gate	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Bit Multiplexer	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
BUFE-based Multiplexer Slice	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
BUFT-based Multiplexer Slice	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Bus Gate	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Bus Multiplexer	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					IO widths up to 256 bits	
Comparator	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
FD-based Parallel Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
FD-based Shift Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-64 bits wide	
Four-Input MUX	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Bit Multiplexer or Bus Multiplexer	
LD-based Parallel Latch	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide	
Parallel-to-Serial Converter	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: FD-based Shift Register	
RAM-based Shift Register	Xilinx	LogiCORE	V-IIP	V-II	V-E	V	S-IIE	S-II					1-256 bits wide, 1024 words deep	
Register	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: FD-based Parallel Register	
Three-Input MUX	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Bit Multiplexer or Bus Multiplexer	
Two-Input MUX	Xilinx	LogiCORE							S				Not recommended for new designs. Suggested replacement: Bit Multiplexer or Bus Multiplexer	

For the most up-to-date information, go to
www.partner.xilinx.com/common/coresolutions/ip/referenceguide

www.partner.xilinx.com/common/coresolutions/ip/referenceguide

Xilinx Global Services

Xilinx Global Services

Xilinx Design Services

- Provide 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.

Xilinx Platinum Technical Services

- Senior Applications Engineers
- Dedicated Toll Free Number*
- Priority Case Resolution
- Proactive Status Updates
- Ten Education Credits
- Electronic Newsletter
- Formal Escalation Process
- Service Packs and Software Updates
- Application Engineer to Customer Ratio, 2x Gold Level

*Available in US only

Part Number	Product Description	Duration	Availability
Educational Services			
FPGA13000-4-ILT	Fundamentals of FPGA Design (v4.x)	8	Now
FPGA16000-4-ILT	ISE Design Entry (v4.x) — Instructor-Led course	8	Now
FPGA23000-4-ILT	Designing for Performance (v4.x)	16	Now
LANG1100-3-ILT	VHDL — Introduction to VHDL (v3.x) — Instructor-Led Course	24	Now
LANG11000-4-ILT	VHDL — Introduction to VHDL (v4.x) — Instructor-Led Course	24	November
LANG12000-4-ILT	VERILOG — Introduction to Verilog (v4.x) — Instructor-Led Course	24	Now
PCI1800-3-ILT	PCI CORE Basics (v3.x) — Instructor-Led Course	8	Now
PCI18000-4-ILT	PCI CORE Basics (v4.x) — Instructor-Led Course	8	December
PCI2800-3-ILT	PCI — Designing a PCI System (v3.x) — Instructor-Led Course	16	Now
PC128000-4-ILT	PCI — Designing a PCI System (v4.x) — Instructor-Led Course	16	December
PCI2900-3-ILT	PCI-X — Designing for PCI-X (v3.x)	16	Now
PCI129000-4-ILT	PCI-X — Designing for PCI-X (v4.x)	16	December
ASIC1500-3-ILT	FPGA Design for the ASIC User (v3.x) — Instructor-Led Course	24	Now
ASIC25000-4-ILT	FPGA Design for the ASIC User (v4.x) — Instructor-Led Course	24	November
DSP2000-3-ILT	DSP Implementation Techniques for Xilinx FPGAs	24	November
PROMO-5002-4-LEL	E-Series II (v4)	10	October
TC-1CR	One Training Credit	1	Now
TC-OS-1DY	One Day of On-Site Training	8	Now
TC-OS-2DY	Two days of On-Site Training	16	Now
TC-OS-3DY	Three days of On-Site Training	24	Now
TC-OS-4DY	Four days of On-Site Training	32	Now
TC-OS-5DY	Five days of On-Site Training	40	Now
Platinum Technical Services			
SC-PLAT-SVC-10	1 Seat Platinum Technical Service w/10 education credits		Now
SC-PLAT-SITE-50	Platinum Technical Service site license up to 50 customers		Now
SC-PLAT-SITE-100	Platinum Technical Service site license for 51-100 customers		Now
SC-PLAT-SITE-150	Platinum Technical Service site license for 101-150 customers		Now
Design Services			
DC-DES-SERV	Design Services Contract		

Call Education Services

North America: 877-XLX-CLASS (877-959-2527)
http://support.xilinx.com/support/training/training.htm

Europe: +44-870-7350-548
eurotraining@xilinx.com

Japan: +81-03-5321-7750
http://support.xilinx.co.jp/support/education-home.htm

Asia Pacific: http://support.xilinx.com/support/education-home.htm

Call Design Services

North America & Asia:
Richard Foder: 408-626-4256

Europe:
Alex Hillier: +44-870-7350-516
Martina Finnerty: +353-1-4032469

Call Technical Services

North America: 800-255-7778 or 408-879-5199
http://support.xilinx.com

United Kingdom: +44-870-7350-610
eurosupport@xilinx.com

Japan: +81-03-5321-7750
jhotline@xilinx.com

Other Locations:
http://support.xilinx.com/support/services/contact_info.htm

For digital video, Spartan®-IIE is the clear choice.

The new Spartan-IIE FPGAs deliver over 20 billion MACs/s, giving you the DSP performance required for high-resolution digital images. Together with advanced signaling standards including LVDS, HSTL and SSTL, you've got the speed you need for digital video processing.

From Picture to Pixel, the Complete Solution

Through capture, process and display, Spartan-IIE FPGAs bring the total solution into focus. Encryption IP, operating at high speed video rates, offers a new level of security protecting every pixel. And the desktop programmability of Spartan-IIE FPGAs solves your time-to-market pressures as design cycles get shorter.

The eSP Web Portal . . . a Unique Designer's Resource

The Xilinx eSP web portal gives you all the support you could wish for as standards continue to rapidly change. Simply log on to **xilinx.com/esp** and you'll get immediate access to reference boards, system block diagrams, a full spectrum of technology & standards tutorials, and much more. You'll soon see why eSP is one of the most successful websites of its kind.

Visit ***www.xilinx.com/spartan2e/dvt*** today and find out why Spartan-IIE FPGAs are the clear choice for digital video design.

XILINX®

The Programmable Logic Company℠

www.xilinx.com/spartan2e/dvt

FORTUNE 2001 100 BEST COMPANIES TO WORK FOR

Introducing The Platform for Programmable Systems.

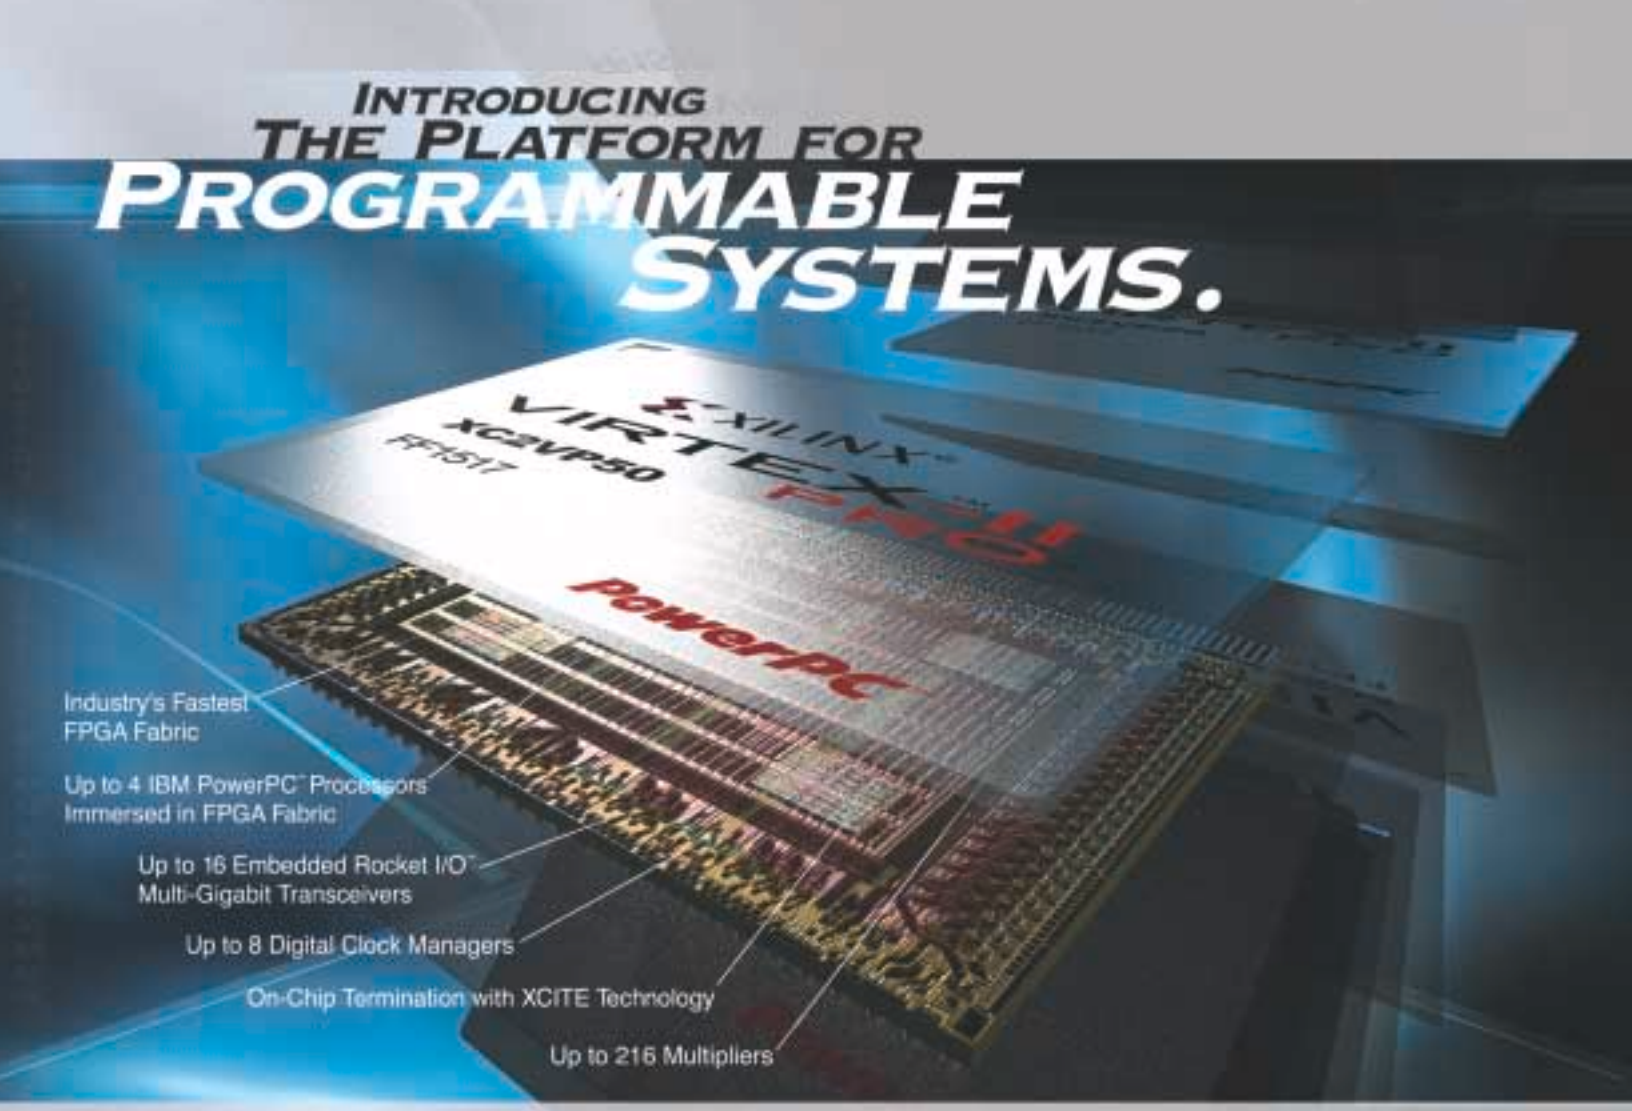

VIRTEX-II PRO™

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

The Power of Xtreme Processing

Each IBM PowerPC runs at 300+ MHz and 420 Dhrystone MIPS, and is supported by IBM CoreConnect™ bus technology. In addition, the FPGA fabric enables ultra-fast hardware processing, such as TeraMACs/s DSP applications. And Xilinx's unique IP-Immersion™ architecture gives you the power of high-performance processors, and easy integration of soft IP. A new performance standard now enables hardware accelerated processing and multiple processing in a single off-the-shelf device.

IBM technology PowerPC processor™

The Ultimate Connectivity Platform

The first programmable device to combine embedded processors along with Rocket I/O™ 3.125 Gbps transceivers, the Virtex-II Pro series supports 10 Gigabit Ethernet with XAUI, 3GIO, SerialATA, InfiniBand, you name it. And our SelectI/O™- Ultra supports 840 Mbps LVDS and other parallel methodologies. With 40+Gbps serial bandwidth, market challenges become market solutions.

With so many features on a single device such as microprocessors, multi-gigabit transceivers, highest density on-chip memory and on-chip termination, the benefits include simplified board layout, a reduced bill of materials, and unbeatable time to market.

Industry-Leading Tools from Wind River and Xilinx

Wind River's embedded tools are the premier support for real-time microprocessor and logic designs. And Xilinx's lightning-fast ISE 4.2i software provides the most comprehensive, easy-to-use development system available.

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

XILINX®
The Programmable Logic Company℠

www.xilinx.com/virtex2pro

FORTUNE 2002 100 BEST COMPANIES TO WORK FOR

PN 0010651