

Xcell journal

SOLUTIONS FOR A PROGRAMMABLE WORLD

**At the Heart
of Innovation**

INSIDE

**At the Heart of Consumer
and Automotive Innovation**

Marriage Made in Heaven?

**Scalable and Flexible
In-Vehicle Networking**

**Designing Portable Handsets
Using CoolRunner-II CPLDs**

**A High-Speed Broadcast
Video Connectivity Solution**

XILINX®

www.xilinx.com/xcell/

Support Across The Board.™

The logo consists of the lowercase letters 'e', 'x', and 'p' in white, each enclosed within a black circle. The circles are arranged horizontally and overlap slightly.

Add Functionality to Your Prototype Platform



Available EXP Modules

- EXP to P160 Adapter
- EXP Prototype
- Analog Devices EXP Adapter Module
- Video EXP Module
- High-Speed ADC EXP Module
- High-Speed DAC EXP Module
- Interface EXP Module

EXP Baseboards

- Xilinx Spartan-3A DSP Starter Kit
- Xilinx Spartan-3 PCI Express Starter Kit
- Xilinx Virtex-4 FX60 PCI Express Board
- Xilinx Virtex-5 LX Development Kit
- Xilinx Virtex-5 LXT PCI Express Board
- Xilinx Virtex-5 SXT PCI Express Board

Avnet Electronics Marketing introduces the EXP specification – a versatile expansion interface for FPGA development boards that allows designers to add application specific daughter cards and easily connect to FPGA I/Os.

Key EXP Features

- Royalty free, public domain specification
- Full and half card options
- 168 (full module) and 84 (half module) user I/O
- Single-ended and differential pair signaling
- High-performance: > 100 MHz single-ended and > 300 MHz differential

Learn more about EXP at www.em.avnet.com/exp



Enabling success from the center of technology™

800-332-8638
www.em.avnet.com

IS YOUR CURRENT FPGA DESIGN SOLUTION HOLDING YOU BACK?




FPGA Design | Ever feel tied down because your tools didn't support the FPGAs you needed? Ever spend your weekend learning yet another design tool? Maybe it's time you switch to a truly vendor independent FPGA design flow. One that enables you to create the best designs in any FPGA. Mentor's full-featured solution combines design creation, verification, and synthesis into a vendor-neutral, front-to-back FPGA design environment. Only Mentor can offer a comprehensive flow that improves productivity, reduces cost and allows for complete flexibility, enabling you to always choose the right technology for your design. To learn more go to mentor.com/techpapers or call us at 800.547.3000.

DESIGN FOR MANUFACTURING + INTEGRATED SYSTEM DESIGN
ELECTRONIC SYSTEM LEVEL DESIGN + FUNCTIONAL VERIFICATION

**Mentor
Graphics**[®]
THE EDA TECHNOLOGY LEADER

EMBEDDED

Unlock your future



Enter the New Era of Configurable Embedded Processing

Adapt to changing algorithms, protocols and interfaces, by creating your next embedded design on the world's most flexible system platform. With the latest processing breakthroughs at your fingertips, you can readily meet the demands of applications in automotive, industrial, medical, communications, or defense markets.

Architect your embedded vision

- Choose MicroBlaze™, the only 32-bit soft processor with a configurable MMU, or the industry-standard 32-bit PowerPC® architecture
- Select the exact mix of peripherals that meet your I/O needs, and stitch them together with the new optimized CoreConnect™ PLB bus

Build, program, debug . . . your way

- Port the OS of your choice including Linux 2.6 for PowerPC or MicroBlaze
- Reduce hardware/software debug time using Eclipse-based IDEs together with integrated ChipScope™ analyzer

Eliminate risk & reduce cost

- No worry of processor obsolescence with Xilinx Embedded Processing technology and a range of programmable devices
- Reconfigure your design even after deployment, reducing support cost and increasing product life

Order your complete development kit today, and unlock the future of embedded design.

Get the Complete Embedded Solution



www.xilinx.com/processor



www.xilinx.com/processor

At the Heart of Innovation

Xcell journal

PUBLISHER	Forrest Couch forrest.couch@xilinx.com 408-879-5270
EDITOR	Charmaine Cooper Hussain
ART DIRECTOR	Scott Blair
DESIGN/PRODUCTION	Teie, Gelwicks & Associates 1-800-493-5551
ADVERTISING SALES	Dan Teie 1-800-493-5551 xccladsales@aol.com
TECHNICAL COORDINATORS	Kevin Kitagawa Kevin Tanaka Tamara Snowden Silvia Gianelli
INTERNATIONAL	Piera Or, Asia Pacific piera.or@xilinx.com Christelle Moraga, Europe/ Middle East/Africa christelle.moraga@xilinx.com Yumi Homura, Japan yumi.homura@xilinx.com
SUBSCRIPTIONS	All Inquiries www.xcellpublications.com
REPRINT ORDERS	1-800-493-5551



www.xilinx.com/xcell/

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

© 2008 Xilinx, Inc. All rights reserved. XILINX, the Xilinx Logo, and other designated brands included herein are trademarks of Xilinx, Inc. All other trademarks are the property of their respective owners.

The articles, information, and other materials included in this issue are provided solely for the convenience of our readers. Xilinx makes no warranties, express, implied, statutory, or otherwise, and accepts no liability with respect to any such articles, information, or other materials or their use, and any use thereof is solely at the risk of the user. Any person or entity using such information in any way releases and waives any claim it might have against Xilinx for any loss, damage, or expense caused thereby.

Behind the Wheel

One of the benefits of growing older is that you gain perspective. You have a longer time span from which to observe change and gauge its relative positive and negative consequences.

For example, I can remember that my parents' new car came luxuriously equipped with an AM radio that had chrome metal push-button tuning, electric windows, a driver's-side power seat that went forward and back, and a really cool telescoping radio antenna that magically raised and lowered into the fender like some kind of futuristic space vehicle.



Cars back then were basically heavy electro-mechanical devices with a lot of chrome.

Now fast-forward through a myriad of technologies ranging from personal computers, wireless networks, flat-panel displays, mobile phones, global positioning systems, and the Internet, and we arrive at today's automobile – a highly sophisticated, computerized transportation system.

What a trip!

The Xilinx® Automotive (XA) product family is ideal for many advanced automotive electronics modules and systems, ranging from the latest driver assistance and infotainment systems to reconfigurable instrument clusters and electronic control unit gateways. Essentially, we help make the automobile environment more entertaining, informative, and productive.

Xilinx is also responding to rapidly changing consumer product requirements in the low-cost, high-volume markets with domain-optimized PLDs, custom IP, and development boards, creating solutions for two key consumer product segments: displays and handsets.

Today we are witnessing a convergence of high technologies that are changing forever the way we communicate and do business with each other. Limited only by our imagination, there seems no end in sight.



Forrest Couch

Forrest Couch
Publisher

AUTOMOTIVE



15

Marriage Made in Heaven?

Creating standard automotive platforms with maximum differentiation.

PORTABLE HANDSETS



40

Designing Portable Handsets Using CoolRunner-II CPLDs

CPLDs can improve processor-based handsets.

AUTOMOTIVE



25

Scalable and Flexible In-Vehicle Networking

FPGAs and full IP solutions give automotive engineers options in optimizing their electrical architectures.

AUDIO/VIDEO



53

A High-Speed Broadcast Video Connectivity Solution


Xilinx Spartan-3E and Spartan-3A FPGAs, National Semiconductor PHY, and the Xilinx protocol stack provide a cost-effective and flexible approach to the challenges of multi-rate broadcast.

12 **Cover**

At the Heart of Consumer and Automotive Innovation

Xilinx offers silicon, development tools, IP, middleware, and design services for the automotive electronics market while meeting the industry's temperature, quality, and reliability requirements.





VIEWPOINTS

Letter from the Publisher: Behind the Wheel.....	5
Wim Roelands: Looking Ahead in 2008.....	8
Kevin Morris: Short Stack with Syrup.....	10
Ivo Bolsens: Enabling Efficient Packet Processing.....	66

FEATURES

Cover

At the Heart of Consumer and Automotive Innovation.....	12
---	----

Automotive

Marriage Made in Heaven?.....	15
Block Matching for Automotive Applications on Spartan-3A DSP Devices.....	16
A Compact Multimedia Display Development Platform for Automotive and Industrial Markets.....	20
Scalable and Flexible In-Vehicle Networking.....	25

Portable Handsets

Designing GPS Systems Using CoolRunner-II CPLDs.....	29
Decrease Processor Power Consumption Using a CPLD.....	32
Supporting Multiple SD Devices with CPLDs.....	37
Designing Portable Handsets Using CoolRunner-II CPLDs.....	40
Easing Design Challenges with CoolRunner-II CPLDs.....	43

Audio / Video

Designing Digital Displays with Spartan-3 Generation FPGAs.....	46
Taking Device DNA Technology to the Next Level.....	49
A High-Speed Broadcast Video Connectivity Solution.....	53

RESOURCES

Technical Papers and Literature.....	58
Automotive ECU Development Kit.....	59
Spartan-3AN FPGA Starter Kit.....	60
CoolRunner-II CPLD Starter Kit.....	61
Protect Your Brand with Spartan-3 Generation FPGAs.....	62
Titanium Dedicated Engineering.....	63

WEB XCLUSIVES

These articles are available at: www.xilinx.com/publications/xcellonline/xcell_63/	
Implementing Industrial Ethernet.....	Web
Designing Efficient, Synchronized, and Reliable Motion Networks.....	Web
Speeding Up Control Automation with EtherCAT.....	Web

Looking Ahead in 2008

Xilinx is poised for growth in high-volume consumer applications.



by Wim Roelandts
CEO and Chairman of the Board
Xilinx, Inc.

Last year was yet another successful year for Xilinx, with record sales of our 90-nm and 65-nm devices in consumer and communications applications. Nowhere was this growth more evident than Asia, where revenues increased nearly 500 percent.* Our customer base continues to expand throughout Asia as more and more designers adopt Xilinx solutions to meet technical and time-to-market demands.

The year 2007 also brought the largest single order in our history: 8 million units for a single socket. As the million-unit orders continue, we envision a time where anyone who does logic design will consider Xilinx. We are perhaps halfway there. Industry analysts project further growth for PLDs as new requirements throughout the electronics industry force the need for flexible architectures that can cope with not only current applications but future and possibly unknown features. As systems continue to increase in complexity, designers will naturally make the transition to a programmable fabric.

The beauty of programmable logic is its flexibility to meet the demand of any specific “hot” application.

Looking Ahead

If economic conditions remain favorable, we anticipate yet another year of growth in 2008, particularly in consumer, surveillance, automotive, and communications applications. The flat-panel display market is one of the fastest growing segments in the electronics industry, thanks to declining prices and government mandates promoting digital programming. FPGAs are playing a key role in enabling many of the key technologies behind the latest displays.

In fact, our Spartan™ Generation devices can be found in many of the latest flat-panel displays. This is a high growth market for Xilinx as panel makers expand production capacities to meet consumer demand. According to iSuppli, worldwide shipments of LCD and PDP TVs are expected to reach nearly 80 million units in 2009.

With the ongoing threats to global security, demand for high-performance surveillance video products continues to grow. Today's surveillance video systems must not only record – they must analyze data in real time. This requires higher quality video and DSP processing capabilities, something our devices are very good at. This is a big opportunity for Xilinx and a major reason we've focused on constantly pushing the DSP performance capabilities of our FPGAs.

In the automotive space, we're seeing tremendous growth in driver assistance applications. Previously only available in high-end luxury cars, electronic driver assistance systems are making their way into the average vehicle. Market analysts expect driver assistance system demand to grow by 5x between 2007 and 2012.

Mobile devices also offer an excellent opportunity for PLDs as portable electronic manufacturers continue to rush new products to market. Today's designers are now looking beyond the fixed architecture of ASICs and ASSPs to discover the innate design flexibility and time-to-

market benefits of programmable logic. Rather than demanding higher performance, consumers now demand ease of use – the iPhone is an excellent example. PLDs allow designers to differentiate their products with ease-of-use features while providing faster time to market.

Finally, in the communications sector, we believe that growth will accelerate as wireless service providers continue their quest to upgrade the infrastructure and triple play continues to be deployed. Next-generation base station deployments must conquer the challenge of continually reducing cost (as measured by cost per channel) while at the same time adding increased functionality to support new services, protocols, and changing subscriber usage patterns. Programmable logic is the ideal solution for wireless base stations, providing not only a flexible design platform but allowing service providers to make future upgrades from a remote location. This can literally save millions in the long run.

Conclusion

Xilinx pioneered the industry's transformation toward a focus on vertical markets and engaging with customers at a system architecture level. Today, this transformation lets us address our customer's complex design challenges by providing them with innovative, flexible, and compelling solutions that help them achieve their objectives of cost management, time to market, and leadership.

And while the beauty of programmable logic is its flexibility – PLDs can be used in virtually any electronic system so there is no immediate need to develop new products to meet the demand of any specific “hot” application – providing just the right mix of features for a given application is a great way to reduce system cost. By optimizing our devices for a given domain, we offer the broadest range of solutions in the industry, providing a unique mix of core capabilities, such as logic, memory, parallel and serial I/Os, embedded processors, DSP function-

ality, and other functions suited to specific application requirements.

Xilinx will continue to expand its efforts in delivering optimized solutions for a broad set of markets and applications. In addition to providing just the right amount of features in our silicon, we're working hard to provide just the right amount of IP, design tools, and peripherals required to deliver an “out-of-the-box” solution. ●●●

**Source: Xilinx Financials, September quarter FY08 versus September quarter FY06.*

Xilinx Chief Willem Roelandts Elected SIA Chairman

On November 14, 2007, the board of directors of the Semiconductor Industry Association (SIA) elected Willem (“Wim”) P. Roelandts, chairman, president and CEO of Xilinx, as its 2008 board chairman. Roelandts succeeds Richard Templeton, president and chief executive officer of Texas Instruments. Hector de J. Ruiz, chairman and CEO of AMD, was elected vice chairman.

“Wim Roelandts and Hector Ruiz bring many years of leadership in the micro-electronics industry to the positions of chairman and vice chairman of the SIA board,” said SIA President George Scalise. “That experience will be invaluable as the semiconductor industry continues its efforts to secure additional funding for university research, immigration, and education reform to ensure access to a highly skilled workforce, and tax reform to level the global playing field.

“Leadership in innovation is the key to leadership in technology,” said Roelandts. “The SIA public policy agenda is focused on ensuring that the U.S. has a business climate that encourages and supports innovation. I welcome the opportunity to be a spokesman for our innovation agenda.”

Short Stack with Syrup

Looking between the layers of the non-volatile Spartan-3AN family.



by Kevin Morris
Editor – FPGA and
Structured ASIC Journal
kevin@techfocusmedia.com

Non-volatile FPGA is one of today's oddest market segments.

Bound by the ill-defined characteristic of “non-volatility,” the field of available devices is diverse from the ground up, with strikingly different architectures, approaches, and benefits.

Actel, long the leader in non-volatile FPGAs, began with antifuse technology (which is one-time programmable) and followed later with flash-based FPGAs. QuickLogic fielded an antifuse-type approach from the beginning, and Lattice Semiconductor joined the fray over a year ago with their LatticeXP hybrid devices, embedding a flash boot PROM on the FPGA for rapid, on-chip configuration.

With the new, non-volatile Spartan™-3AN family, it is clear that Xilinx approached the non-volatile problem differently. Often, on your board, you'd pair a low-cost SRAM FPGA like a Spartan-3 device with a commodity flash memory for storing the bitstream and possibly other user data. Xilinx just took the die for that flash memory and the die for that FPGA, stacked them using packaging techniques now common in cell phones and other space-constrained devices, and dropped them into a single package. Voila! A non-volatile FPGA.

Hey, wait, that's totally cheating, isn't it? Does stacking the two die vertically instead of horizontally on your board somehow change a regular-old volatile SRAM FPGA into a newfangled non-volatile one? Aren't



there rules about these things? Somebody bring in a referee. What's the catch?

In reality, there is no “catch.” Like any engineering solution to a problem, there are trade-offs. To understand the ones involved in Spartan-3AN FPGAs, we need to examine our possible reasons for wanting a non-volatile FPGA in the first place, and then see how the various options stack up.

As long as we don't pop open the package and start asking too many questions (we'll discuss opening the package later, so keep those Dremel tools powered down for the moment, lil' hacker dudes), we can assume that a non-volatile FPGA is one that doesn't require an external boot PROM. You apply power to the pins and (after some time passes) the FPGA is ready to go, independent of any other devices on the board.

Why do we care? Well, we might want the smaller footprint of a single chip. We might want simpler board design. We might hope for lower power consumption. We might want additional design security. We might be trying to reduce total BOM cost. We might want instant (or near-instant) power on. Compared with Spartan-3A devices, how does the new Spartan-3AN family do on these issues? Yes, yes, no, sorta, and probably not.

Xilinx Spartan-3AN FPGAs are available in a subset of the package options of the Spartan-3A family. Probably because the stacked-die packaging process is more expensive, Xilinx isn't rolling out the entire fleet of packages for Spartan-3AN devices. The

company says they will probably be adding additional packaging options based on customer demand, so if the package you want is not on the list – vote early and often. The good news is that the package you choose will be the only FPGA-related thing you have to buy, inventory, place, and connect to your board. You won't need a separate boot PROM as you would with an SRAM-based FPGA. Board real estate and complexity is definitely reduced.

The FPGA portion of a Spartan-3AN device has the same power consumption as a corresponding 90-nm Spartan-3A device because – hey, it's the same die. As a review – Spartan-3A devices have some nice power-saving features for an SRAM FPGA. Its 90-nm technology gives it pretty good dynamic power consumption. In addi-

tion, Spartan-3A devices (and thus also Spartan-3AN devices) have two additional power modes called “Suspend” and “Hibernate” that allow you to drop static (standby) power by 40% with a fast wake-up time (suspend mode) or by 99% with a “normal” wake-up time (hibernate mode). This means that the device is still more power-hungry than other super-low-power non-volatile FPGAs like QuickLogic’s PolarPro or Actel’s Igloo, but it maintains other advantages. (Remember what we said about trade-offs?)

The flash portion of the Spartan-3AN FPGA will, of course, have the same power profile as the commodity flash memory device that it is. If you dropped one of those on your board next to a Spartan-3A device, your total power characteristics should be almost identical. One could argue that having the flash separate gives you more options on the flash side than defaulting to the built-in flash in Spartan-3AN FPGAs. If one argued that, one would be correct.

Security Risks

Design security is one of the issues most often cited for wanting a non-volatile FPGA over a conventional SRAM design. The important thing to remember here is that security is relative. Every option available on the market provides differing amounts and types of security. Almost nothing can protect against the best-funded security risks such as those funded by governments. As you move down the scale on attackers’ budgets and time constraints, however, security measures start to shake out at different levels.

The biggest design security risk with SRAM-based FPGAs is the availability of the configuration bitstream going between the boot PROM and the FPGA during startup. Since FPGAs are standard parts, if the contents of the boot PROM can be cloned, the design is stolen. The normal way to deal with this is to encrypt the bitstream in the boot PROM, and to have an encryption key stored in the FPGA that will allow the bitstream to be decrypted upon load. The problem with SRAM (and non-volatile devices in general) is that any encryption key

stored in SRAM is lost when the device is powered down. This leaves us with a problem in moving encrypted bitstreams that has been addressed two ways: 1) by attaching an external battery to the device to maintain the encryption key and 2) by adding permanent fuses to each device for storing the encryption key.

The Spartan-3AN device does neither of these. Not exactly, anyway. There is a permanent serial number stored in each device, and that serial number can be used to implement various security schemes as you see fit. It is not, however, used to decrypt the incoming bitstream. Bitstream protection is done by obscuring the connection between the FPGA and the boot PROM (now stacked on top of the FPGA inside a single package). While more secure than a separate boot PROM, the scheme is still far from impermeable as far as configuration stream protection goes. (OK, you can fire up those Dremel tools now.)

Because the Spartan-3AN FPGA has room for two complete configurations in the flash, it’s obvious that Xilinx thinks you’ll be updating the bitstream at some point in the product life cycle. The Spartan-3AN device even has a nice feature that’ll let it boot on the original “golden” bitstream if a bad download or other corruption causes the new, updated one to malfunction. The problem is, any mechanism you have for delivering an updated bitstream to the device will be snooperable by prying eyes (or prying logic analyzers, really). The bitstream obscuring method doesn’t hold up as well when the bitstream has to travel from your home base to the device in the field. Be sure to design your security scheme accordingly.

As far as the lower cost question goes, you’ll probably pay a bit of a premium for Spartan-3AN devices compared with the same Spartan-3A device with the same flash memory. When you factor in the other costs of an additional device such as board area and inventory and such, it may be a wash. On the positive side, you certainly won’t be taking a big BOM hit for using the new family. Xilinx projects prices in the \$5 range in volume.

It’s also important to look at Spartan-3AN devices in comparison with other non-volatile options. After all, from a marketing perspective, the Spartan-3AN FPGA is probably most important because it gets Xilinx into the non-volatile FPGA race. Before, a list of non-volatile FPGA vendors would have included Actel (for their antifuse and flash-based devices), QuickLogic (for their antifuse devices), Lattice (for their LatticeXP hybrid devices) and arguably even Altera for their Max II (which is really an FPGA marketed as a CPLD). Now, that list includes Xilinx with the Spartan-3AN family.

With a 90-nm SRAM-based FPGA at its core, Spartan-3AN devices will be faster and have higher density than the other entrants. Because it uses commodity flash bonded to the FPGA, it will offer more non-volatile storage than the other entrants. On the down-side, it will take longer to configure (in the range of 100 milliseconds because of streaming the configuration over a serial flash versus microseconds or even near-zero for truly non-volatile devices like antifuse and Actel’s flash). It will likely be less secure than the other options, and although its power consumption will be very good, it will consume more power in both active and standby modes than the other non-SRAM alternatives.

The Spartan-3AN family offers devices ranging from 50K to 1.4M “system gates” (corresponding with Spartan-3A devices XC3S50A to XC3S1400A) with flash memory ranging from 1M to 16M. Three of the devices, the XC3S200AN, XC3S700AN, and XC3S1400AN are available now, with the remaining devices (XC3S50AN and XC3S400AN) slated for second quarter.

Xilinx Spartan-3AN FPGAs are likely to be popular in space-, cost-, and power-constrained high-volume applications that need some of the features of non-volatile FPGAs with the performance and density of leading-edge SRAM devices. In addition to bringing Xilinx officially into the non-volatile game, the unique characteristics of the Spartan-3AN family gives us another valuable trade-off point in our FPGA selection. ●●●

At the Heart of Consumer and Automotive Innovation

Xilinx offers silicon, development tools, IP, middleware, and design services for the automotive electronics market while meeting the industry's temperature, quality, and reliability requirements.

by Kevin M. Kitagawa
 Director, Worldwide Marketing, High-Volume Products
 Xilinx, Inc.
kevin.kitagawa@xilinx.com

As we move toward an increasingly digital world, consumers want to experience content such as audio, video, and data anytime and anywhere. Better and inexpensive wired and wireless infrastructure is available to deliver content for a better consumer experience, and product manufacturers are consistently challenged to come up with new innovations to improve that experience.

As an example, digital music players, HDTV digital displays, automotive electronics, cell phones, and broadband Internet access have all changed consumer expectations and behavior on how to access and experience content at home, in the car, or on the move. Predicting which innovations will be successful often requires potentially risky product decisions, and manufacturers are trying to incorporate feedback directly from consumers into product development cycles early to ensure success of their products. This often translates to the need for quick product development cycles to get innovative products to market quickly.

The initial costs of developing ASICs have been exponentially rising with each process geometry shrink, making this path less attractive for most applications. Also,



ASSPs typically cannot meet the needs of every customer, and silicon manufacturers must make compromises to implement the features customers request the most.

Xilinx® FPGAs and CPLDs are at the heart of innovation. These products provide flexible yet low-cost solutions to meet the application-specific requirements of high-volume products such as digital displays, set-top boxes, automotive rear-seat entertainment systems, smartphones, and video equipment.

Xilinx in Consumer Electronics

Consumer electronics manufacturers need to differentiate their products from their competitors. They are often driven to deliver the latest technologies and innovations as quickly as possible to the consumer to stay ahead in a highly competitive market-

place. Later in the product life cycle, manufacturers focus on driving product costs down to deliver more and more value at lower prices to the consumer.

As an example, HDTV digital display manufacturers are challenged with the evolution of new flat-panel technologies and interface standards while trying to improve image quality to address any shortcomings



Xilinx in Automotive Electronics

New and future vehicles are increasingly reliant on electronics in order to allow OEM automakers to differentiate their products. These electronics systems range from driver information and rear-seat entertainment systems to newer driver-assistance technologies such as night-view and lane-departure warning systems.

The Xilinx Automotive (XA) family of products brings the scalability and flexibility of FPGAs and CPLDs to the automotive electronics market, allowing automotive tier-ones to bring out innovative features while also giving them platform development capabilities (even in a fragmented application sector), thus lowering costs.

With the increasing use of LCD/TFT displays, multiple networking protocols, video, and graphics, Xilinx XA solutions allow for quick development cycles and flexibility of customization through to production. By using Xilinx XA programmable logic, automotive tier-ones can quickly add new features, improve existing features, or change interfaces simply by making changes within the FPGA fabric without needing to wait years until the next time the electronics are redesigned. FPGAs also allow designers to look again at their overall application architecture and build systems based on their needs instead of what is available in semiconductor hardware today.

Conclusion

With Xilinx high-volume Spartan, CoolRunner™-II CPLD, and XA product families, manufacturers can afford to deliver more innovative products quickly to market, knowing that they have the flexibility of programmable logic minimizing their risk.

Decreasing costs from our suppliers, along with leading-edge technology, have allowed us to extend lower prices to our customers, accelerating our success in high-volume design wins moving into production. With cost points starting in the low single digits, Xilinx FPGAs and CPLDs can eliminate the need to move to ASICs or ASSPs in production in high-volume applications such as consumer or automotive electronics, while still allowing endless customization to meet your customer requirements. ●●●

in the quest to deliver the perfect TV picture. Variations in panels used in different models often have different specifications or even different suppliers.

FPGA programmability enables manufacturers to easily accommodate these variations using image-enhancement algorithms implemented in a Spartan™ FPGA and allowing the same hardware design to be utilized in a whole family of products. Native support for most differential I/O interfaces in Spartan FPGAs makes it easy to directly interface with the new panel interfaces. These advantages make Spartan FPGAs a compelling option in many consumer electronics products.

Xilinx is Driving Automotive Solutions



As the automotive electronics market grows in the areas of infotainment, driver assistance, and driver information systems, Xilinx® devices are at the heart of each new innovation.

- **Image Processing and Recognition** – Offering faster development and product differentiation for driver assistance applications.
- **Video and Graphics** – Allowing for highly scalable and cost-efficient implementations of infotainment and driver information systems.
- **Application Development Platforms** – Providing developers a quick start and serving as a comprehensive hardware base to satisfy their expectations.
- **Xilinx Technology Leadership** – Leading the way in programmable logic devices, one of the fastest growing segments of the semiconductor industry.
- **Vehicle Networking** – Designing and supporting key automotive-specific network interfaces.

For more information on Xilinx automotive products, visit

www.xilinx.com/automotive

Marriage Made in Heaven?

Creating standard automotive platforms with maximum differentiation.

by Nick DiFiore

General Manager, Automotive Division

Xilinx, Inc.

nick.difiore@xilinx.com

Within the Automotive Business Division of Xilinx, we are very excited by the changes we see coming. Programmable logic is developing rapidly, as are the requirements of the automotive market. These changes are providing some great opportunities to create solutions to complex problems currently plaguing automotive design.

The advantages of programmable logic have always been flexibility, scalability, and fast time to market. Today, these are advantages of key value in the automotive electronics market, as consumer products with highly complex and varying networking and interface standards are introduced into vehicles at a rate never seen before.

Vehicle differentiation and high value-add is often provided by the OEMs through the electronics, but at the same time, engineering resource and time constraints mean that tier ones need to supply common platforms and reusable designs. PLDs in automotive electronics make for a great partnership, but like all partnerships, there are also many challenges.

As a new technology in this market, the first challenge is in providing not just cutting-edge silicon, but full solutions to demonstrate and prove the technology out to our customers. To achieve this goal, we have made a large investment at Xilinx to recruit both highly experienced automotive system architecture engineers and veteran automotive semiconductor personnel to work directly with customers and determine what silicon and solutions would be a true value-add in this market. We have also invested in engineers to develop IP both



internally and externally through our LogiCORE™ IP and automotive AllianceCORE programs.

To continue our own learning and also drive the market forward, Xilinx has joined a number of consortiums, such as AUTOSAR, JASPAR, the FlexRay Consortium, and the MOST Cooperation. By doing this, we have been able to put together an arsenal of fully scalable in-vehicle networking and video/graphics solutions.

These critical investments have resulted in complete solutions for applications, such as video displays (“A Compact Multimedia Display Development Platform for Automotive and Industrial

Markets,” also in this issue of *Xcell Journal*), driver assistance systems (“Block Matching for Automotive Applications on Spartan-3A DSP Devices” in this issue), and a flexible Xilinx MOST solution (see our demonstration at CES).

The second challenge is in meeting the very high-quality and reliability standards that the automotive market requires of semiconductor suppliers. The advantages of working with automotive customers were recognized at the very highest levels of management within Xilinx. By achieving ISO-TS16949 certification, implementing the AEC-Q100 qualification standard, and creating a separate XA line of products in order to control bill-of-materials (BOM) and production sites,

we can improve the quality and logistic processes for all of our customers and ensure that we are on the path of continuous improvement, which is a key part of our corporate philosophy.

The last four years of the automotive XA program at Xilinx have been both challenging and exciting, but our success in becoming the number-one programmable logic supplier (based on Semicast data) in the automotive electronics market shows that we are on the right track.

Our next challenge is to ensure that even with all of the possible opportunities there are for programmable logic, we continue to work closely with our automotive customers to ensure that we focus on the right solutions to meet their future needs in architecture and design. ●●●



Block Matching for Automotive Applications on Spartan-3A DSP Devices

The Spartan-3A DSP FPGA outshines VLIW DSP-CPU's in automotive driver assistance applications.

by Daniele Bagni
DSP Specialist
Xilinx, Inc.
daniele.bagni@xilinx.com

Paul Zoratti
Automotive Senior System Architect
Xilinx, Inc.
paul.zoratti@xilinx.com

Automotive engineers are adapting numerous intelligent technologies to assist humans in operating vehicles safely. Prominent technologies in vehicle systems include radar-, ultrasonic-, and camera/vision-based sensing. Collectively referred to as driver assistance (DA) systems, these technologies are designed to facilitate safe driving during adverse conditions and potentially dangerous roadway situations.

The first generation of camera-based DA systems is now available on a variety of production vehicle models. The majority of such systems provide drivers with a video image of the environment around the vehicle. The most prevalent systems are parking/backup aids that use a rear-facing camera to capture the scene behind the host vehicle and display the image on a radio/navigation system screen, or on a small display in the instrument cluster.

A second generation of camera-based systems is in development and testing, with some limited deployment. Rather than merely providing an image to drivers, these second-generation systems apply image processing and analytics to extract information from the video stream and characterize and evaluate the vehicle environment. If warranted, the driver receives an appropriate warning.



As engineers gain real-world experience in characterizing the vehicle environment, future DA techniques will increase in complexity, offer greater utility to consumers, and enhance the performance of other vehicle subsystems. Figure 1 summarizes the variety of current and future DA features.

Advanced Processing Requirements

Processing requirements for DA systems can exceed the capabilities of current automotive-grade serial DSP processors. In addition, a growing need exists to bundle multiple DA features together, based on a single suite of vision sensors to drive consumer value.

For example, a forward-looking vision module may need to simultaneously support lane departure warning, intelligent headlamp control, and sign recognition functionality – all of which require different processing algorithms. Therefore, the DA market offers a real opportunity for FPGAs to provide system value through raw processing performance, configuration flexibility, and device scalability.

Image processing and analytics functionality for vision-based DA schemes can include spatial/temporal filtering, lens-distortion correction, image sharpening, contrast enhancement, edge detection, pattern matching, object recognition, object tracking, and, in some cases, graphical overlay. Of particular interest is a form of pattern-matching functionality that supports motion estimation or stereo disparity calculations.

To illustrate the performance value of FPGA processing, consider the following vision-based system: a wide-VGA resolution imaging device (752 x 480 pixels) generating video at a 30-Hz frame rate (fps), and the need to estimate the motion (or flow) of objects from one frame to another. One algorithmic approach – also suitable for stereo ranging disparity calculations – is to partition the image into blocks (say, 4 x 4 pixels in size) and evaluate a match criteria for each block in the first frame to a location in the second frame over a specified search area (say, 20 x 20 pixels).

A common match criteria is to find the minimum absolute error (MAE) of the pixel intensities between the 4 x 4




Figure 1 – Driver assistance features




Figure 2 – An example of SIMD: SAD operation on a quadruple of 8-bit samples

block in the first image and the pixels within the search area on the second image by using an operator called SAD (sum of absolute differences).

Our 4 x 4 block matching example requires more than 250 MMAE/s (millions of MAE calculations per second), since (752 pixels) x (480 lines) x (20 x 20 pixel search area) x (30 fps) / (4 x 4 pixel block size) = 270,720,000 MAE/s. MAE indicates the final matching error of a 4 x 4 pixel block, whereas SAD refers to the sum of absolute differences calculation on four individual pairs of elements. Therefore, every MAE requires four SAD operations.

Processing Options

Processing options at the automotive design engineer's disposal include very long instruction word (VLIW) DSP-CPU and FPGAs. FPGAs offer dramatically higher

processing capability than any existing VLIW DSP-CPU. This is because of the architecture of the FPGA: the large amount of functional units in parallel (including programmable MACs) allows the FPGA to achieve 10-30 times better performance than that offered by any DSP, depending on the application implemented, even if the clock frequency of the FPGA is much lower than the clock frequency of the DSP-CPU. Using the block-matching operation example, we will demonstrate that Xilinx® FPGAs have superior performance to any VLIW DSP-CPU processor.

SAD and MAE Calculations in VLIW DSP-CPU Processors

A SAD operation on four elements of 8-bit pixel video data could be implemented in a single-instruction-on-multiple-data (SIMD) within a 32-bit architecture DSP-

CPU, thus efficiently executing the equivalent of 11 elementary instructions in only one cycle, as shown in Figure 2.

For example, the Nexperia PNx1500 media processor equipped with the TriMedia 32-bit VLIW-CPU can implement two quadruple SAD instructions on a single clock cycle for 8-bit pixels with two cycles of latency. Associated with every long instruction word are as many as five elementary RISC/SIMD instructions

per clock cycle, of which only two can be SAD (named "8meii" in the TriMedia databook).

Therefore, the MAE on a 4 x 4 block size would require five clock cycles, as shown in Table 1: two cycles for pipelining two quadruple SADs (cycle 1 for sad1/sad2, cycle 2 for sad3/sad4) and three cycles for accumulating the partial results (cycles 3, 4, and 5). Hence, a 300-MHz Nexperia PNx1500 processor could compute a peak of 60 MMAE/s if processing a single block.

By processing more than one 4 x 4 block at a time, the peak performance can improve a little. For example, the MAE of two 4 x 4 blocks in parallel could be computed in seven cycles, thus achieving 85.71 MMAE/s, and three blocks could be processed in nine cycles, or 100 MMAE/s.

The maximum amount of blocks that can be processed in parallel is limited first by the number of SIMD SAD operations allowable in any long instruction word; second by the number of general-purpose registers of the VLIW-CPU; and third by the scheduling algorithms of the optimizing compiler. Overall performance goes into saturation when adding more blocks, which is why we do not consider more than three MAEs processed in parallel.

The Texas Instruments (TI) TMSD320DM6437 digital media processor has a long instruction of eight elementary RISC operations per cycle through two separated data paths: each one of four slots per cycle. Its VLIW-CPU can execute as many as two SAD instructions per cycle (named "subabs4" in the TI DM6437



databook), each with a latency of one cycle. However, to accumulate the partial results, a three-cycle latency SIMD MAC operation must be performed (named “dotpsu4”), with a constant 0x01010101.

Therefore, a 600-MHz TI DM6437 DSP-CPU could compute an MAE in seven cycles (as shown in Table 2), thus achieving a peak performance of 85.71 MMAE/s for 4 x 4 pixel blocks. If two blocks are processed in parallel, we get nine cycles and 133.33 MMAE/s, whereas for three blocks we get 11 cycles and 163.64 MMAE/s – again lower than our 250 MSAD/s requirement.

De-Rating VLIW DSP-CPU Performance

Thus far, we have assumed 8 bits per pixel, which is very suitable for 32-bit architecture DSP-CPU processors. However, new CMOS image sensors have a higher resolution range: 12 to 14 bits per pixel. For these data types, the classic quadruple 8-bit sub-words SIMD of 32-bit architectures are less effective and must be replaced by dual 16-bit half-words SIMD, in which the sub-

word parallelism is only two. Therefore, peak performance degrades significantly, because more clock cycles are necessary to compute an MAE.

Table 3 shows what could be the pseudo assembly code of the SAD computation on the TI VLIW DSP-CPU when using 16-bit sub-word instructions, considering the correct latency and the functional issue slot able to deliver such instructions. As a result, eight cycles are necessary for one 4 x 4 block while 10 and 12 cycles are required, respectively, for two and three blocks processed in parallel. Corresponding peak performance is then 75 MMAE/s, 120 MMAE/s, and 150 MMAE/s. All of these numbers are smaller than the ones achieved with 8-bit sub-word instructions.

Spartan-3A DSP FPGA SAD and MAE Performance

To fill the processing performance gap between Spartan™-3 and Virtex™-4 devices, Xilinx introduced the Spartan 3A-DSP 1800A and 3400A FPGAs. These

devices incorporate a modified version of the DSP48 slices found in Virtex-4 devices. In addition, 3A-DSP parts include a larger number of on-chip memory (block RAMs). Both of these enhancements, along with a price point suited to high-volume applications, make 3A-DSP devices a good fit for automotive vision-based DA systems.

Figure 3 shows the scheme of SAD computation for a quadruple of 12-bit pixels on the Spartan-3A DSP 1800 (XC3SD1800A-4FG676) device. This implementation was done with the System Generator for DSP design flow (a bit-true, cycle-accurate, synthesizable library provided by Xilinx in the Simulink tool). The amount of resources required is 121 slices (236 LUTs and 140 flip-flops). By replicating this structure four times and adding the partial results, we get the scheme for a whole 4 x 4 block, which requires 508 slices (990 flip-flops and 606 LUTs) with a throughput of one (which means that at any clock cycle we can start the computation of a new MAE) and a latency of seven cycles.

Utilizing a 150-MHz clock frequency (of the maximum 250 MHz the device could achieve), we would need only two parallel structures occupying approximately 6% of the device area to achieve 300 MMAE/s and meet the 250 MMAE/s requirement of our example application. This leaves ample resources available to implement other image processing func-

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5
Cycle 1	sad1=8meii(A1,B1)	nop	sad2=8meii(A2,B2)	nop	nop
Cycle 2	sad3=8meii(A3,B3)	nop	sad4=8meii(A4,B4)	nop	nop
Cycle 3	sad12=sad1+sad2	nop	nop	nop	nop
Cycle 4	sad34=sad3+sad4	nop	nop	nop	nop
Cycle 5	tot=sad12+sad34	nop	nop	nop	nop

Table 1 – Pseudo assembly code for the MAE computation on the Nexperia/TriMedia VLIW DSP-CPU, with quadruple 8-bit sub-word parallelism

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8
	L1	S1	M1	D1	L2	S2	M2	D2
Cycle 1	d1=subabs4(A1,B1)	nop	nop	nop	d2=subabs4(A2,B2)	nop	nop	nop
Cycle 2	d3=subabs4(A3,B3)	nop	sad1=dotpsu4(d1, 0x01010101)	nop	d4=subabs4(A4,B4)	nop	sad2=dotpsu4(d2, 0x01010101)	nop
Cycle 3	nop	nop	sad3=dotpsu4(d3, 0x01010101)	nop	nop	nop	sad4=dotpsu4(d4, 0x01010101)	nop
Cycle 4	nop	nop	nop	nop	nop	nop	nop	nop
Cycle 5	nop	nop	nop	nop	nop	nop	nop	nop
Cycle 6	sad13=sad1+sad3	nop	nop	nop	sad24=sad2+sad4	nop	nop	nop
Cycle 7	tot = sad13+sad24	nop	nop	nop	nop	nop	nop	nop

Table 2 – Pseudo assembly code for the MAE computation on the TI VLIW DSP-CPU, with quadruple 8-bit sub-word parallelism

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8
	L1	S1	M1	D1	L2	S2	M2	D2
Cycle 1	d1=sub2(A1,B1)	d3=sub2(A3,B3)	nop	nop	d2=sub2(A2,B2)	d4=sub2(A4,B4)	nop	nop
Cycle 2	ad1=abs2(d1)	nop	nop	nop	ad2=abs2(d2)	nop	nop	nop
Cycle 3	ad3=abs2(d3)	nop	nop	nop	ad4=abs2(d4)	nop	nop	nop
Cycle 4	z13=add2(ad1, ad3)	nop	nop	nop	z24=add2(ad2, ad4)	nop	nop	nop
Cycle 5	nop	nop	s13=dotp2(z13, 0x00010001)	nop	nop	nop	s24=dotp2(z24, 0x00010001)	nop
Cycle 6	nop	nop	nop	nop	nop	nop	nop	nop
Cycle 7	nop	nop	nop	nop	nop	nop	nop	nop
Cycle 8	tot = s13 + s24	nop	nop	nop	nop	nop	nop	nop

Table 3 – Pseudo assembly code for the MAE computation on the TI VLIW DSP CPU, with dual 16-bit sub-word parallelism

tions, data routing pipes, memory interface controllers, and a 32-bit MicroBlaze™ embedded processor for serial processing and external communications.

For reference, by utilizing only 70% of the whole FPGA device again at 150 MHz, the Spartan 3A-DSP 1800A device could process up to 23 blocks in parallel (70% x 16,640 slices/508 slices/block = 23 blocks). This corresponds to a peak performance of 3,529 MMAE/s, which is at least 25 times greater than the 600-MHz TI DSP-CPU's peak performance.

Conclusion

Using an automotive vision-based application example, we've shown how to leverage the programmable parallelism of a medium-sized, low-cost Xilinx FPGA to provide required processing performance over VLIW DSP-CPU's. Table 4 summarizes the results of our analysis.

Note that for MAE calculations on 4 x 4 blocks of 12-bit pixel data, the Spartan-3A DSP outperforms the TI TMS320DM6437 by 2 times at only one-fourth the clock speed. Furthermore, resource utilization on the FPGA is only 6%, thus allowing for the implementation of other image processing functions (in parallel if necessary) on the same device.

On the other hand, the VLIW DSP-CPU is totally busy during the SAD calculations, leaving little opportunity to conduct other simultaneous functions by consuming the available slots of the serial processor long instructions.

We were quite conservative about the FPGA estimated clock frequency (150 vs. 250 MHz), as well as the motion estimation search area (the larger the search area, greater the number of MAEs to be computed). A 30 x 30 search area, for example, would require 609 MMAE/s, far beyond the VLIW DSP-CPU capability but using only 12% of the slices on the 1800A device.

Finally, in the MAE implementation we did not use at all of the DSP48 MAC units: we estimate that a 4 x 4 block of 12-bit input data MAE could take 400 slices (782 flip-flops and 400 LUTs) and four DSP48s by replacing the 100-slice adder tree with four DSP48 units.

Thus, the Spartan-3A DSP 1800A device is well suited for vision-based applications requiring significant processing horsepower,

flexibility, and scalability such as those found on future generations of automotive driver assistance applications. ●●




Figure 3 – System Generator for DSP scheme of the SAD computation for a quadruple of four 12-bit samples in the Spartan-3A DSP 1800 device

Device	Configuration	Clock Freq.	Performance
Example Application Requirement	752 x 480 image 4 x 4 pixel block 20 x 20 pixel search area	N/A	> 250 MMAE/s
Philips Nexperia PNX 1500 VLIW DSP-CPU	8-bit pixel depth Parallelism: Three 4 x 4 blocks	300 MHz	100 MMAE/s
TI TMSD320DM6437 VLIW DSP-CPU	8-bit pixel depth Parallelism: Three 4 x 4 blocks	600 MHz	163.64 MMAE/s
TI TMSD320DM6437 VLIW DSP-CPU	12-bit pixel depth Parallelism: Three 4 x 4 blocks	600 MHz	150 MMAE/s
Xilinx Spartan-3A DSP 1800A FPGA	12-bit pixel depth Parallelism: Two 4 x 4 blocks (approximately 6% of device resources)	150 MHz	300 MMAE/s
Xilinx Spartan-3A DSP 1800A FPGA	12-bit pixel depth Parallelism: 23 4 x 4 blocks (approximately 70% of device resources)	150 MHz	3,450 MMAE/s

Table 4 – Summary of results

TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Evaluate the System Generator for DSP design tool.
- Learn more about and purchase the Spartan-3A DSP 1800A device.
- Buy the XtremeDSP Development Platform – Spartan-3A DSP 3400A Edition.



A Compact Multimedia Display Development Platform for Automotive and Industrial Markets

Xilinx FPGAs and Xylon IP cores shorten design cycles and lower production costs for multimedia applications.

by Davor Kovacec
CEO
Xylon d.o.o.
dkova@xylon.hr

Market leadership in the fast-moving electronics market continuously demands innovative and cost-effective products. With an ever-shrinking time-to-market window, design tools and pre-designed IP blocks are important elements of success.

The right development platform can make all the difference between success and failure. It is a balancing act between the contradictory requirements of having a standard but still highly configurable solution.

Combining Xilinx® FPGAs with the Xylon logicBRICKS IP cores library, Xylon's feature-rich logiCRAFT 3 compact multimedia display development platform is an ideal solution, addressing the time-to-market and flexibility needs of the high-volume customer base. You can quickly turn system designs running on this generic FPGA development platform into specialized products. Such a design approach enables a large portion of design reuse through different hardware (IP cores) and software modules. You can reuse these same modules in many system designs for different applications.

The logiCRAFT Platform

The logiCRAFT 3 compact multimedia display development platform (Figure 1) is an FPGA-based board using hardware modules in the form of IP cores to develop FPGA functionality. The platform is designed to support a wide variety of audio and video sources and handle a





Figure 1 – LogiCRAFT 3 compact multimedia display development platform

variety of display types. With several standard communication interfaces, you can easily integrate the logiCRAFT 3 platform into a larger system. Figure 2 shows logiCRAFT 3's functional blocks.

A wide variety of system interfaces such as audio/video (A/V) I/Os, line drivers, gener-

al-purpose I/Os, and COG power supplies makes this platform extremely versatile.

The logiCRAFT 3 platform is primarily aimed at the automotive market, including applications such as navigation, infotainment, rear-seat entertainment (RSE), and driver assistance. Other multimedia applications are equally applicable, such as consumer, medical, and measurement instrumentation or factory automation applications.

Figure 3 shows block schematics of an automotive RSE system. In this particular example, the logiCRAFT 3 is used in the car headrests and serves to deliver A/V content.

The platform's flexibility can help reduce overall solution costs as well as time to market. For example, the platform's capability to fully support a number of different displays enables you to make selections when required for cost or availability. Implementing high-quality wireless audio

through two embedded IR headphone audio outputs for the RSE is also an opportunity to reduce system cost. The FM modulation is implemented using the logiAIR IP core, enabling the use of a low-cost field effect transistor (FET) for IR LED driving. There is also the flexibility to enable Bluetooth with the addition of a Bluetooth module.

Audio and video streams, as well as control data, are transferred to the RSE headrest unit by using an embedded high-speed Gbps digital link. The low data-rate content, such as push-button status, remote controller data, touch-screen data, and others are transferred using the same link.

The link itself requires only two twisted pair lines. One line may also be used for powering the platform, resulting in a significant reduction in infotainment wiring costs. You can have the A/V devices connected to the logiCRAFT 3 through the composite video blanking and sync




Figure 2 – LogiCRAFT 3 platform's functional blocks

(CVBS) or S-video inputs (video) and stereo line inputs (audio). An iPod can also be connected as required.

The logiCRAFT 3 platform can also be used in stand-alone human machine interface (HMI) applications such as:

- Instrument cluster applications (Figure 4)
- Car radio and navigation
- Industrial and medical instrumentation (Figure 5)

In these applications, you can use either the CAN network connection for all data transfer or the standard RS232 channel for communication and debugging. For applications requiring HMI control, an embedded IR remote control sensor, touch-screen interfaces, and GPIO for buttons and LEDs are provided.

You can also add features by using Xilinx LogiCORE™ IP, other third-party IP cores, or by designing a custom circuit.

The list of key IP cores that can be utilized on the logiCRAFT 3 includes:

- logiCVC-ML – compact multilayer video controller
- logi2D – 2-D graphics accelerator
- logiCAN – CAN 2.0B-compatible network controller
- logiUART – universal asynchronous receiver/transmitter
- logiWIN – versatile video input
- logiMEM – flexible SDRAM/DDRAM/flash memory controller
- logiI2S – I²S transceiver
- logiRC – IR remote controller receiver
- logiAIR – digital FM modulator for IR headphones

Spartan™-3E FPGA hardware resources, dedicated multipliers, digital clock managers (DCMs), and block RAMs ensure that IP cores operate at higher speeds and consume less area on the FPGA. In addition, FPGA configuration from an external byte-width memory provides fast board startup time, which is of a great importance in automotive applications. All Xylon IP cores are carefully designed for




Figure 3 – Rear-seat entertainment system




Figure 4 – Platform/FPGA structure for rear-seat entertainment application




Figure 5 – Platform/FPGA structure for instrument cluster application

the lowest FPGA resource utilization, making them particularly suitable for high-volume applications.

Platform Features

- Small form factor 155 mm x 115 mm (4" x 3")
- Supports Spartan-3E 3S250E/500E/1200E 256-pin FPGAs
- 8 MB of NOR flash
- 16 MB/32-bit-wide SDRAM running up to 133 MHz
- Up to two simultaneous video inputs selectable from:
 - Four CVBS or two S-video PAL/NTSC/SECAM inputs
 - One LVDS channel (clock + three data pairs) directly connected to the FPGA
 - High-speed LVDS Gbps digital transceiver
- 8-bit general-purpose I/O
- Touch-screen controller
- CAN interface and single-wire CAN interface
- RS232 interface
- Interface to external Bluetooth module
- iPod control interface
- One stereo audio input selectable from four stereo line inputs
- Microphone input
- Audio line output
- Headphone output
- Two stereo IR headphone audio outputs
- Video output configurable as:
 - CVBS/S-video PAL/NTSC/SECAM output
 - One LVDS channel (clock + three data pairs) directly from FPGA
 - Digital RGB interface
- Power supply for COG displays, including GAMMA correction and VCOM circuits
- Power and control output for CCFL backlight inverters and LCD backlights

RSE Example

Figure 4 shows the platform/FPGA structure for an RSE application where the FPGA design is based on Xylon's IP cores. In such applications, A/V and control data is usually transferred between an infotainment master board and the electronics, which are embedded into car headrests through a number of cables.

The LogiCRAFT 3 platform, with its support for a high-speed Gbps digital link,

dramatically reduces the complexity and cost of the wiring required. You can transfer all digital video, audio, and bi-directional control data using only two twisted pair lines. You can also discard the power cable, because the platform can be powered through one of the two twisted pair lines.

Besides the reduction in wiring complexity, the logiCRAFT 3 platform also integrates many functions otherwise supported by dedicated integrated circuits into just one low-cost Spartan-3 XC3S250E FPGA. You will not need additional electronics for IR audio FM modulation, IR remote protocol decoding, or I²C bus control, among others. The logiCRAFT 3 represents an ideal solution for saving your budget without having to sacrifice anything from the set of common RSE features.

Instrument Cluster Example

Figure 5 shows the platform/FPGA structure for an instrument cluster, where the powerful yet compact logiCVC multilayer display controller will help you create a state-of-the-art instrument panel. With a 60-fps refresh rate, you will be able to animate on-screen gauges and give them the motion smoothness of their analog counterparts.

In addition, the logi2D graphic accelerator, with its bitmap image rotating capability, offers additional flexibility for instrument animation. You can easily




Figure 6 – Platform/FPGA structure for industrial application

design an instrument panel with GUI builder or graphic libraries to efficiently use Xylon's graphics acceleration IP cores.

Communication between the vehicle engine control unit data and the platform is easy by using the onboard CAN interface with our logiCAN IP core. You can also connect a suitable camera using the Gbps digital link, thus adding a parking assist capability to the list of features. You can use a standard PAL/NTSC/SECAM as well.

With the help of our logiWIN IP core, you can de-interlace scale and position incoming video to optimize it for your particular display. All FPGA functionality required to support the instrument cluster configuration will fit into one Spartan-3 XC3S1200E FPGA. This way, you can achieve an outstanding cost-to-performance ratio and provide a very competitive solution for today's demanding and fast-changing markets.

Industrial Application Example

Figure 6 shows the platform/FPGA structure where a LogiCRAFT 3 compact multimedia display development platform serves as a central unit for monitoring and controlling industrial processes, with the FPGA design based on Xylon IP cores. As in the previous example, the logiCVC/logi2D

combination, supported with dedicated graphic software libraries, gives you everything you need for making and combining screen elements required for successful process monitoring and control. You can also use an embedded touch-screen controller to make human-machine interaction intuitive and logical.

For handling classical buttons and indicators, an I²C-controlled GPIO interface is at your disposal. If your application calls for real-time environment surveillance or real-time video process monitoring, the platform's PAL/NTSC/SECAM video decoders provide an excellent solution.

By introducing the logiWIN IP core into your FPGA configuration, you can easily handle decoded video signals in terms of de-interlacing, scaling, and positioning. The platform also provides the means to create audio alarms in an emergency situation. To implement this, you can use an embedded audio codec in combination with the logiI2S IP core. The platform-process communication can be established by using the on-board CAN interface and logiCAN IP core.

If you wish to connect the platform to a PC, you can do that through the platform's RS 232 interface. To host all of this functionality, logiCRAFT 3 utilizes a Spartan-3

XC3S1200E FPGA, which gives you enough room to develop the optimal configuration.

However, your actual final production configuration may require less FPGA resources, so a smaller and less-expensive FPGA would suffice. The logiCRAFT 3 addresses this issue by providing pin-compatible scalability between Spartan-3 XC3S500E and XC3S250E FPGAs, thus delivering an additional feature for tailoring your final product.

Conclusion

The value of the LogiCRAFT 3 compact multimedia display development platform from Xylon lies in the high number of supported display types; the innovative system architecture, small form factor, and broad feature set; and the ability to be configured for performance, price, low development, and production cost with virtually no obsolescence.

Based on the Xilinx Spartan-3E family, with its low cost-per-gate and feature-rich architecture, the logiCRAFT 3 provides an excellent base on which to build.


For more information about the LogiCRAFT 3 compact multimedia display development platform and logicBRICKS Xylon IP library, please e-mail sales@logicbricks.com or visit www.logicbricks.com.



Scalable and Flexible In-Vehicle Networking

FPGAs and full IP solutions give automotive engineers options in optimizing their electrical architectures.

by Kevin Tanaka
Worldwide Automotive Marketing
and Product Planning Manager
Xilinx, Inc.
kevin.tanaka@xilinx.com



Over the past 10 years, in-vehicle networking architectures have become much more complex. Although the number of in-vehicle networking protocols has been reduced, the number of networks actually being deployed has increased dramatically. This leads to network architecture scalability issues and semiconductor device optimization to match each application and network's actual needs.

FPGAs, once thought of as a solution for development purposes only, have come down in price to the point where many issues can be resolved and put into production at a lower overall system cost than the traditional ASIC or ASSP solution. All of the major FPGA suppliers to the automotive market are now ISO-TS16949 certified, making programmable logic devices a mainstream technology in the automotive market.

In-Vehicle Network Electrical Architecture

During the last ten years, many proprietary OEM automaker networking protocols have made way for more standardized global protocols like CAN, MOST, and FlexRay. This lets semiconductor vendors concentrate on building devices with those specific protocols built in, bringing more competition and lower prices to tier-ones and more module interoperability at the OEM level. However, in today's vehicle electrical architectures, both OEMs and tier-ones still struggle with many issues.

Engineers can partition and build networking strategies in several different ways. High-end vehicles can have anywhere up to seven different network buses running simultaneously. For instance, a single vehicle could have a LIN loop for mirrors, a low-speed CAN loop at 500 Kbps for low-end functions like seat or door control, a

high-speed CAN loop at 1 Mbps for body control, another high-speed CAN loop for driver information systems, a FlexRay loop at 10 Mbps for real-time driver assistance data, and a MOST loop at 25 Mbps for control and media streaming within or across various infotainment systems like navigation or rear-seat entertainment.

On the other hand, low-end vehicles may have no more than a single LIN or CAN loop, with all of the other modules working on their own with almost no interaction. Each OEM automaker deals with inter-module communication and vehicle network topology differently, and each vehicle platform is different, making it difficult for tier-ones to develop reusable module architectures with the correct interfaces. Uncertainty of the final architecture into which a module will go is an area where FPGAs excel.



Because of their fixed hardware architecture, ASICs, ASSPs, and microcontrollers are usually either under- or over-resourced with no flexibility. The programmability (and re-programmability) of the FPGA allows for the simple addition or subtraction of on-chip channels (for example, channels of CAN), along with the re-use of IP. With this flexibility, an optimized solution for the number and type of networking interfaces can be built into a module quickly.

Semiconductor Implementation of Network Protocols

The scalability of the FPGA for the number and types of interfaces is not its only merit. In the case of ASSPs, ASICs, and microcontrollers, the peripheral macros are implemented in hardware, making them inherently inflexible. In an FPGA environment, the networking interface IP itself can be optimized depending on the IP being used.

For example, with Xilinx® LogiCORE™ CAN or FlexRay networking IPs, users can flexibly program the number of transmit and receive buffers along with the number of filters. In traditional hardware solutions, an engineer working with a CAN controller would typically only have three configuration choices: 16, 32, or 64 message buffers. The Xilinx scalable MOST network interfacing solution includes network controller IP that can be configured for either master or slave operation and a host of IP, such as asynchronous sample rate converters (ASRC), data routers, or encryption engines for copy protection, depending on the level of system functionality and available processing outside of the FPGA.

The IP allows for optimization and the ability to push into lower density devices for low-end solutions and higher density

devices for high-end solutions, often using the same package footprint on the module target board. Also, for each main protocol, middleware stacks and drivers have been developed to round out the solution. This type of scalability as well as the versatility of an FPGA-based solution is just not possible in a traditional automotive hardware solution.

All of the major FPGA vendors have soft microprocessors that can be efficiently implemented in the fabric for control functions and can run at speeds that rival some of those embedded in hardware. Another major advantage of the FPGA architecture is the ability to offload processing functions from the microprocessor and partition by using the parallel DSP processing capabilities found in either multipliers or hard MACs on-chip, increasing overall performance and throughput.

We've Come a Long Way

Programmable logic devices have come a long way in becoming a mainstream technology in the automotive market. The field has evened on the reliability side, and FPGA technologies are allowing scalable and flexible integration that has never before been possible in traditional ASIC, ASSP, or microcontroller architectures. Overall production system costs are reduced because of shortened development cycles, advanced process technologies used by programmable logic device vendors, and the economies of scale that a programmable device inherently brings with it.

With the key IP and solutions for in-vehicle networking coming to fruition and the added performance capabilities of FPGA architectures, programmable logic devices will become a major player in helping alleviate some of the engineering difficulties inherent in developing in-vehicle electrical architectures. ●●



**FPGA Powered –
That's it!**

FPGA-Module: TQM hydraXC –


**smallest, most universal
Hardware Platform for
Reconfigurable Computing**

- Based on XILINX Spartan 3, Virtex 4 and Virtex 5 technology
- Ethernet 10/100, USB 2.0, RTC
- SPI-, NAND-Flash, DDR2 / SDRAM
- Size 2.13 Inch x 1.73 Inch (54 mm x 44 mm)
- Programmable VCC IOs

**Embedded solution with
TQM hydraXC for**

- Fast time to market
- Economical series production
- Highest flexibility
- Hardware reduction

:: Starter kits available ::



Starter kit
with module

TQ components


Email: info@tqc.de
www.tq-group.com

Infoline: (+49) 81 53 / 93 08 - 333

TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Download more information about the Xilinx Automotive Networking LogiCORE solution.
- Buy a Xilinx Automotive ECU Development Kit.

GET ON TARGET



LET XCELL PUBLICATIONS HELP YOU GET YOUR MESSAGE OUT TO THOUSANDS OF PROGRAMMABLE LOGIC USERS.

Hit your target by advertising your product or service in the Xilinx *Xcell Journal*, you'll reach thousands of qualified engineers, designers, and engineering managers worldwide.

The Xilinx *Xcell Journal* is an award-winning publication, dedicated specifically to helping programmable logic users – and it works.

We offer affordable advertising rates and a variety of advertisement sizes to meet any budget!

Call today :
(800) 493-5551
or e-mail us at
xcelladsales@aol.com

Join the other leaders in our industry
and advertise in the *Xcell Journal*!



www.xilinx.com/xcell/



Spartan™-3 Generation Display Solutions

inrevium
by Tokyo Electron Device Ltd.



Complete development solution with pre-verified reference designs and IPs for digital image processing applications.

Reference Designs (for Free):

- 7:1 LVDS Rx/Tx display link
- RSDS/Mini LVDS Tx
- 266Mbps DDR SDRAM Controller
- Color Temperature Correction
- Precise Gamma Correction(PGC)
- Image Dithering Engine(IDE) etc.

IPs (from Xilinx Partners):

- Apical Limited
- Dynamic Range Compression(iridix)
- Noise Reduction(NR)

Spartan-3A Display Solutions Board TB-3S-1400A-IMG

RoHS
Compliant

Features

Xilinx Device: XC3S1400A-FGG484, XCF08P VO48
Memory: DDR SDRAM 512Mb x 16bit x 4(32bit x 2ch)
Interfaces: RS-232C x1
High Speed I/O Connector x4
Option Board: DVI Tx and Rx

Spartan-3E Display Development Kit HW-SPAR3E-DISP-DK-UNI-G-PROMO

RoHS
Compliant

Features

Xilinx Device: XC3S1600E-FG484C, XCF08P-FS48
Memory: DDR SDRAM 512Mb x 16bit x 4 (32bit x 2ch)
Flash Memory: 16Mbit SPI
Interfaces: RS-232C x1
High Speed I/O Connector x4
Option Board: DVI Tx and Rx, Pin header adapter cards x2

TOKYO ELECTRON DEVICE LIMITED



World Headquarters

1 Higashikata-cho, Tsuzuki-ku, Yokohama, Kanagawa, 224-0045, Japan
Phone: +81-45-474-7028
E-mail: psd-sales@teldevice.co.jp

US office

2953 Bunker Hill Lane, Suite 300 Santa Clara, CA 95054, USA
Phone: +1-408-919-4772
URL <http://www.inrevium.jp/eng/>

Worldwide Distributors:



Nu Horizons Electronics Corp.

Phone: +1-888-747-NUHO (6846)
URL <http://www.nuhorizons.com>



HiTech Global Design & Distribution, LLC

Phone: +1-408-781-8043
URL <http://www.hitechglobal.com>



Designing GPS Systems Using CoolRunner-II CPLDs

Expanding functionality using low-power CoolRunner-II CPLDs.

by Arthur Yang

Technical Marketing Manager, Micro Scale Division
Xilinx, Inc.

arthur.yang@xilinx.com

Asking for directions is becoming a foreign concept. It's fast dying out, much like the art of folding a map or a newspaper. Its impending extinction can be attributed to the increasing popularity of the Global Positioning System (GPS).

GPS is present in a growing number of products: automobiles, cell phones, personal data assistants, even wristwatches. With each GPS vendor showcasing a product line of several dozen GPS products, consumers are overwhelmed with the selection.

Therefore, success lies in product differentiation and specialization. Xilinx® CoolRunner™-II CPLDs are the ideal chips to add features or interface with another device to make your GPS product stand out from the rest – without breaking the power budget.

The Power Advantage

Both portable and in-dash GPS devices must adhere to a strict power budget. CoolRunner-II CPLDs have the advantage of minimal power consumption without the need for sleep mode states. The smallest CoolRunner-II device has a quiescent power of 13 μ A. In some portable applications, a sleep mode is sufficient and a wake-up time in the hundreds of milliseconds is acceptable.

In some cases, however, the current design states are lost. When using sleep mode, your entire device shuts down. You must rely on a secondary device to poll for interrupts and initiate a wake-up sequence. CoolRunner-II devices offer a unique sleep mode without the associated design headaches through its DataGATE feature, a self-contained and user-configurable circuit that allows you to disable as much of the device as you desire.

By enabling DataGATE, you can turn off whichever inputs you choose in several nanoseconds, thereby shifting the power consumed by the CPLD closer to a quiescent state. This could be done periodically, such as when polling for an interrupt, or it could be dependant on a



particular state of operation. Thus, portions of the CoolRunner-II device can remain active while others are in standby.

Let's look at an example to explain further. The CPLD sits on an address and data bus between a microprocessor, mobile SDRAM, and SD flash card. Data moves between each of the devices. The CPLD monitors the data activity and blocks traffic based on the particular function. If the current task is to move data from the microprocessor to the SD card, then the CPLD blocks any data to and from the SDRAM using DataGATE.

Within the CPLD, you can write code to simply decode the function being performed (either from watching the address lines or by parsing frame data) and then enable/disable whichever data path is required. This is a trivial example that is possible to a much lesser degree with other programmable logic devices (PLDs) by using output 3-states.

The advantage of DataGATE is that it effectively 3-states at the input of the CPLD. Power savings increase because the CPLD I/O as well as core circuitry are placed in a quiescent state, whereas alternate solutions require the entire PLD to stay in active mode.

The Security Advantage

GPS system pricing makes it an attractive target for product cloning. To help prevent this, you can utilize CoolRunner-II CPLD's read/write protect security to implement a security system that prevents cloning by overbuilding. Overbuilding occurs when a contract manufacturer orders extra components for a given production run and then builds extra products that are identical to the authentic versions.

The concept is that the system requires interaction with the CPLD to perform any desired function. You can implement this security in any number of ways. The most straightforward solution would be to have a CPLD act as a data-traffic cop, directing data between the individual devices on the board.

A more complex solution involves using the CPLD to implement a block cipher. This would have the microprocessor submit a random stream of data to be encrypt-

ed by the CoolRunner-II device and returned. This data would then be decrypted and verified against the original data.

The CoolRunner-II CPLD's security is implemented through multiple programming bits, so someone attempting to determine which bits are relevant for security must find and disable several bits out of tens of thousands in the non-volatile array. The security in CoolRunner-II CPLD read/write security prevents readback and programming on top of the existing pattern. So the device

blank CPLDs is not useful without access to the programming file.

Achieving Product Differentiation in GPS Units

Original consumer GPS units were straightforward. They simply gave location information in the form of latitude and longitude. Today's GPS units not only offer real-time maps and directions, they offer MP3 playback or integration with cell phones via Bluetooth. There are also market-specific GPS systems with features such as traffic updates for in-car navigation,




Figure 1 – GPS block diagram

will not permit extraction of the programmed JEDEC file, nor will it allow someone to overlay a modified version of the code on top of an existing pattern. To reprogram the device, the entire device must first be erased, in which case the design information is lost.

A critical factor in any of these flows is that pre-programmed CPLDs must be provided to the contract manufacturer from a trusted source (Xilinx or an authorized Xilinx distributor). This prevents overbuilding – because simply ordering more

ultra-small units for runners or cyclists to measure their pace, units with sonar for fishermen – even GPS-enabled collars for keeping tabs on the family pet.

Each product requires an interface to something different – and that's where the CoolRunner-II CPLD is a perfect fit. The following topics demonstrate some uses for CoolRunner-II CPLDs. I've also listed associated Xilinx application notes at the end of this article. Figure 1 demonstrates functions where Xilinx CPLDs can benefit a GPS application.



Xilinx CoolRunner-II CPLDs demonstrate their usefulness in lowering power, increasing security, and providing connectivity solutions for today's GPS systems. With our portfolio of small form-factor packages, these devices can fit into the smallest portable packages such as cell phones, dog collars, and wristwatches.

SD Card Interface

SD memory (in its various physical formats such as mini-SD and micro-SD) has separated from the pack of memory interfaces that competed a few years back. Although other interfaces such as MMC and Compact Flash are still around, they represent only a small percentage of the market. Flash cards are essential for any product with picture or MP3 playback; both features are becoming more prevalent in GPS handheld units.

Mobile SDRAM Interface

You can simplify microprocessor code by allowing the CPLD to act as the memory interface. If your high-end model requires multiple memory modules but your low-end product does not, let the CPLD code change, not your microprocessor.

Level Translation

As GPS moves into newer consumer areas, they connect to components that are not yet optimized for low-voltage operation. CoolRunner-II CPLDs have a minimum of two I/O banks (increasing to four I/O banks in the largest device), allowing multi-voltage interfaces to be addressed easily. Supported voltage standards include 1.5 V, 1.8 V, 2.5 V, 3.3 V, SSTL 2-1, SSTL 3-1, and HSTL-1. 5-V interfaces are supported with some external circuitry.

Keypad Scanner

Many of the lower end GPS models cannot use a touch-screen interface because of cost constraints or size limitations. Even in the expensive models, a few buttons are designated for certain features such as power or volume control. Some form of keypad or button interface is used in most models. This is an ideal use for CoolRunner-II CPLDs because when there is inaction from users, the CPLD remains in a quies-

cent state and can immediately respond to a user key press without having to wake up from sleep mode. Furthermore, it can be designed to verify user data before waking up the rest of the system.

For example, many cell phones require that you press two keys in sequence before waking up to ensure that the buttons weren't accidentally pressed.

Microprocessor Interface

A common role for CoolRunner-II CPLDs is port expansion. Many microprocessors simply lack sufficient I/O for the multitude of devices with which they must communicate. CoolRunner-II CPLDs allow a product platform design to add and change modules without changing the core processor.

Serial Peripheral Interface (SPI)

SPI is a common interface used by many peripherals, including flash storage chips, LCDs, touch screens, and temperature sensors. Its popularity can be attributed to a simple four-wire interface and improved throughput over I²C or SMBus.

Conclusion

Xilinx CoolRunner-II CPLDs demonstrate their usefulness in lowering power, increasing security, and providing connectivity solutions for today's GPS systems. With our portfolio of small form-factor packages, these devices can fit into the smallest portable packages such as cell phones, dog collars, and wristwatches. Personally, I can't wait until they embed a GPS into car keys.

TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Start evaluating CoolRunner-II CPLDs for your GPS application and purchase the CoolRunner-II Design Kit.
- Learn more from these Xilinx application notes:
 - XAPP347, "Decreasing Processor Power Consumption using a CoolRunner-II CPLD"
 - XAPP395, "Using DataGATE in CoolRunner-II CPLDs"
 - XAPP906, "Interfacing with Multiple SD Devices with CoolRunner-II CPLDs"
 - XAPP398, "CompactFlash Card Interface for CoolRunner-II CPLDs"
 - XAPP394, "Interfacing to Mobile SDRAM with CoolRunner-II CPLDs"
 - XAPP785, "Level Translation Using Xilinx CPLDs"
 - XAPP512, "Implementing Keypad Scanners with CoolRunner-II"
 - XAPP799, "An SMBus/I²C Compatible Port Expander"
 - XAPP341, "UARTs in Xilinx CPLDs"
 - XAPP386, "CoolRunner-II Serial Peripheral Master"



Decrease Processor Power Consumption Using a CPLD

Utilizing a CPLD to offload operations from the system microprocessor keeps the processor in a power saving mode longer and contributes to significant power savings.

by Mark Ng
Applications Engineer
Xilinx, Inc.
mark.ng@xilinx.com

One of the most critical factors in designing handheld, portable electronics today is reducing overall system power consumption. With increased consumer expectations, portable devices require longer battery life and higher performance. Even power reductions on the order of 10 mW are crucial to portable system designers and manufacturers.

Several design techniques are used by designers today to significantly reduce overall system power consumption, such as:

- Reducing operating voltage
- Optimizing system and CPU clock frequency
- Eliminating spikes of large current consumption during the power up sequence
- Efficiently managing system battery operation
- Efficiently managing operating mode of system devices
- Minimizing bus activity
- Reducing bus capacitance
- Reducing switching noise

These are just a few examples of design techniques for reducing the power consumption in any end application.

One of the most important power saving techniques mentioned in this list is the ability to manage the operating mode of devices in the system. Many manufacturers today offer devices with power saving modes that temporarily suspend the device from its normal operation. These devices have the option to power down or transition to a non-functioning state if the device is not active for a specific amount of time. This feature is available on many of today's microprocessors and microcontrollers. By taking advantage and managing the operating mode of large power consumers on a PCB, such as the processor, the overall power consumption of the system can be reduced significantly.

Reducing power consumption not only involves correct management of the operating mode of a device, but designing a system to take advantage of the modes a device can operate within. Offloading operations of the microprocessor allows it to stay in its low-power state for a longer amount of time. One way to reduce system power is to allow a low-power programmable logic device, such as a CPLD, to manage these offloaded operations. This article will describe this possibility, along with types of operations that allow a processor to remain in a low power state longer, thereby reducing system power consumption.

Microprocessor Operating Modes

In some portable applications, the CPU can consume 30% of the overall system power. Figure 1 illustrates the typical power consumption of system components in a Web Pad application.

Microprocessor power consumption can range from 720 uW to 1 W during normal operation. Microprocessor operating modes vary by part and manufacturer and include modes such as Normal, Run, Sleep, Suspend, Standby, Stop, and Idle operation. Operating modes can vary in power consumption as much as 230 mW between states. Normal operation of some low power microprocessors can be as little as 250 mW.




Figure 1 – Web Pad power consumption

Example

Note: The microprocessor reference provided is an example to illustrate the power consumption in different operating modes. Since no standard method exists to determine power consumption, the data provided in this document is based on data provided by the manufacturer and is to be used for reference only. Please see “References” for more information.

To illustrate the difference in power consumption of operating modes in a microprocessor, an example is provided. Figure 2 illustrates the power consumption of the Intel StrongARM SA-1110 microprocessor operating modes. The power dissipation numbers shown in Figure 2 are determined by operating at 206 MHz with a nominal external voltage supply of 3.3 V and internal voltage supply of 1.8 V.

Operating modes of the StrongARM processor include Normal, Idle, and

Sleep. In Normal operation, the CPU is full-on, with the device fully powered and receiving active clocks. In Idle mode, even though power is applied to the CPU and other components, all clocks to the CPU are stopped, with only clocks to peripheral devices active. In Sleep mode, power to the CPU and other peripheral components is disabled. Sleep mode disables all functions except the real-time clock, interrupt controller, power manager, and general purpose I/O.

Operating Mode Control

Microprocessors with power saving modes have an on-board power management controller. Operating modes allow the operating system or software application to temporarily suspend the CPU. The microprocessor executes a series of instructions to be placed into a power saving state. Once in a power down mode, several components of the microprocessor can still respond to system interrupts.

For example, the Idle mode of the StrongARM SA-1110 processor saves significant power, but certain modules remain powered, such as the LCD, memory, and I/O controllers. Even though the clock to the CPU is stopped, peripheral modules are still active. The Idle mode can still consume a significant amount of power, on the order of 100 mW. By placing the processor into the Sleep mode, only active modules are powered to respond to interrupts and wake up signal requests. Sleep mode consumes even less power than Idle mode; current consumption can be less than 100 uA.




Figure 2 – Intel StrongARM SA-1110 power consumption

Inserting an external device to respond and handle system interrupts can reduce the operations required of the processor. By allowing the microprocessor to stay in its power down mode as long as possible, significant power savings can be realized.

For a microprocessor to return to normal operation from a power down mode, an event must occur. The following events can wake up the processor, but vary based on manufacturer, part, and current operating mode:

- Hardware reset
- System interrupt
- General-purpose I/O interrupt
- Real-time clock interrupt
- OS timer interrupt
- Peripheral interrupt
- External wake-up signal

Upon recognition of an enabled wake-up event, the microprocessor will begin a series of steps to wake up from a power down state. Figure 3 illustrates the general flow for a processor waking up from a power down mode.




Figure 3 – Wake-up sequence

CPLD Design

Operating modes are utilized when the microprocessor is idle for a specific amount of time. When a microprocessor receives an enabled interrupt, the processor will respond to the interrupt request. When the processor is responding to the interrupt, it will operate in its run or normal mode. Reducing the number of interrupts to the processor will increase the time the processor is in a power saving state. Ideally, if the microprocessor does not have any instructions to execute, it will remain in a power saving mode forever. Inserting an external device to respond and handle system interrupts can reduce the operations required of the processor. By allowing the microprocessor to stay in its power down mode as long as possible, significant power savings can be realized.

Utilizing a low power programmable logic device to supplement the microprocessor will save system power and increase system battery life. The industry's latest CPLD offerings simultaneously deliver high performance and low power

consumption. Standby current of a typical low power CPLD is less than 100 uA. Figure 4 illustrates using a reprogrammable CPLD to interface to incoming system interrupts. Utilizing an external data acquisition device to offload interrupt requests required of the microprocessor will reduce overall system power.

System Interrupts

Depending on the end application for the processor, a variety of external devices may interrupt the processor. These interrupts include both data acquisition and data processing requests. By separating data processing interrupts to the microprocessor, data acquisition interrupts can now be serviced by the external CPLD. Utilizing a CPLD to handle data acquisition interrupts will offload interrupt requests to the microprocessor and save power.

Categorization of the type of data acquisition interrupts to the CPLD will depend on the end application. Peripheral devices or incoming data demanding a response to incoming data can be classified as data




Figure 4 – System block diagram

acquisition interrupt requests. Data acquisition interrupts include:

- Memory access interrupts
- Communication interfaces such as I²C, UART, SPI, or ISA
- General-purpose I/O interrupts
- LCD interface interrupts

This is not a complete list of interrupts that can be processed by the CPLD, but provides a starting point for the system design.

Operational Flow

Figure 5 illustrates the main operational flow for the design of a CPLD. Once a valid external interrupt is recognized by the CPLD, it will determine if it contains the functionality to process the interrupt. Once the CPLD has processed the interrupt, it can assert an interrupt to the processor for any data processing requests needed. If the CPLD is unable to process the interrupt, the interrupt is passed to the processor. The CPLD also monitors the operating state of the processor.




Figure 5 – CPLD flow diagram

Functionality

The low-power CPLD design consists of an interrupt interface and controller to handle interrupt requests, the functionality to process the interrupt, and a processor interface. The main functions of the CPLD are described in more detail next and separated as follows:

- Interrupt interface for system devices
- Interrupt controller
- Interface with peripheral devices for interrupt processing
- Microprocessor interrupt interface
- Microprocessor operating mode interface

Interrupt Interface

The interrupt interface of the CPLD receives all external device interrupt requests previously recognized by the microprocessor. The interrupt interface determines if the CPLD is capable of processing the interrupt request. The CPLD handles data acquisition interrupts that

request data receiving and storage capabilities. If the CPLD is unable to process the interrupt, the interrupt is passed to the microprocessor.

The CPLD interrupt interface provides the masking capability for all interrupt sources and the ability to determine the interrupt source. Programmable logic provides flexibility to change the trigger mode, which includes a high or low level and falling or rising edge sensitivity. The CPLD interrupt control registers are similar to the registers in the microprocessor.

Interrupt Controller

The CPLD interrupt controller emulates the functionality that exists in the system microprocessor. The interrupt controller interprets from which device the data acquisition interrupt was received and initiates the processing of the interrupt. The CPLD processes the data acquisition interrupt request that would have otherwise interrupted the microprocessor.

The interrupt controller initiates the action to process the request. An example of this is an application where the CPLD is receiving data from a remote device. The device is requesting to write the data being sent into memory. The CPLD interrupt controller recognizes a valid interrupt and initiates the memory interface to interpret the data.

Peripheral Device Interfaces

The CPLD provides the interface to system devices that are needed in processing interrupt requests. Device interfaces that are needed are dependent on the end application. When an external device interrupts the CPLD to read or write data into a memory component, that particular memory interface is needed in the CPLD design. The types of interfaces needed can range from memories to LCD interfaces to communication interfaces such as PCI, UART, SPI, and ISA.

Microprocessor Interrupt Interface

The CPLD, like any external device requesting services of the processor, has the capability to interrupt the microprocessor. The CPLD must be able to interrupt the microprocessor once a data acquisition operation is complete. The designer has the option to set the priority level of interrupt requests from the CPLD and whether or not interrupts received from the CPLD will wake the processor from a power down state.

Microprocessor Operating Mode Interface

Depending on the system microprocessor, the CPLD will be able to recognize the operation state of the processor. Some microprocessors provide external pins that represent the current operating mode. Depending on the CPLD and microprocessor design, the CPLD could recognize the current operating state of the processor and determine whether to assert an interrupt to the processor to execute a waiting interrupt. For example, if a low priority interrupt is received by the CPLD and the processor does not need to transition from its low power state, the CPLD can create a register indicating pending interrupts. Then when the processor wakes, the interrupt pending register can be read by the microprocessor.

Benefits

Figures 6 and 7 illustrate the power savings that may be realized in a typical battery-operated device using a leading-edge, low-power CPLD (Figure 7) versus a stand alone microprocessor design (Figure 6). The power requirements of the CPLD are minimal compared to the power savings realized by keeping the microprocessor in its low power modes for a longer amount of time. Standby current of a typical low power CPLD is on the order of 100 uA. The operating power consumption depends on the application and clock frequency. For a 64-macrocell CPLD fully populated with 16-




Figure 6 – Stand-alone processor power consumption




Figure 7 – Power consumption of processor and CoolRunner-II CPLD

bit counters and a 50-MHz clock, ICC is around 10 mA. Note that the actual power savings realized will depend on the system design, including the type of microprocessor and the CPLD design.

Along with power savings attained using a CPLD, interrupt response time is reduced. The peripheral device no longer has to wait the delay time for the microprocessor to wake from a power saving state. Additional design savings can be realized and include:

- Reducing the number of interruptions to the processor
- Reducing the number of processor wake-up cycles over a length of time
- Reduction of clock frequency without impact on throughput
- Running the processor at a lower frequency for data processing operations
- Running the CPLD at a higher frequency for data acquisition operations

Conclusion

Designing a power-sensitive application involves not only using software for power management, but utilization of hardware design techniques. Designing a low-power CPLD to keep a microprocessor in a low power operating state longer can significantly reduce system power consumption. The latest CPLDs on the market today offer a flexible combination of low power and high speed for any end application.

References

Intel StrongARM SA-1110 Microprocessor Developer's Manual. June 2000.

DragonBall Power Management. Motorola Semiconductor Application Note. Motorola, Inc. 1998.

Geode GX1 Processor Series Low Power Integrated x86 Solution. National Semiconductor. October 2000.

Turley, Jim. Microprocessors for Consumer Electronics, PDAs, and Communications. Microprocessor Report. September 1999.

EDN Access. 27th Annual Microprocessor/Controller Report. September 2000.

Xilinx Global Trade Show Calendar

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

For more information and the most current schedule, visit www.xilinx.com/events/.

North America

January 7-10, 2008
Consumer Electronics Show
Las Vegas, Nevada

April 14-17, 2008
National Association
of Broadcasters
Las Vegas, Nevada

April 15-17, 2008
Embedded Systems Conference
San Jose, California

Europe

February 11-14, 2008
3GSM
Barcelona, Spain

February 26-28, 2008
Embedded World
Munich, Germany



Supporting Multiple SD Devices with CPLDs

Creating an SD multiplexer using CPLDs.



by Mark Ng
Staff Applications Engineer, CPLD Applications Engineering
Xilinx, Inc.
mark.ng@xilinx.com

There has been an increasing demand to add multiple Secure Digital (SD) devices in a single system. The problem, however, is that most host devices/processors – for example Intel PXA270, TI OMAP, or Qualcomm MSM processors – only provide a single SD interface. Fortunately, Complex Programmable Logic Devices, otherwise known as CPLDs, can be used to allow host devices to support any number of SD devices. This article details a scalable, auto-sensing bi-directional multiplexer-based design.

Figure 1 shows a generalized CPLD usage model to incorporate any number of SD ports for a given host device that only

has a single native SD interface. The CPLD is placed between the host controller and the SD devices. As such, the CPLD part performs a bi-directional multiplexing function, allowing the host to communicate with any selected SD device. More importantly, this design has no directional control pins, which means that the CPLD automatically detects the direction of data flow.




Figure 1 – Using CoolRunner-II CPLDs to provide additional SD ports

This implementation is extremely flexible and scalable, meaning that the number of SD ports can be increased or decreased as desired. The design also supports any of the defined SD card modes – SPI, 1-bit, or 4-bit data modes.

While the primary purpose of using a CPLD device in this type of application is to provide additional SD ports to the host controller, secondary benefits include level translation and logic isolation between the host and the SD card. Figure 1 shows the case where the host is 1.8 V but the SD Devices are 3.3 V. The industry's latest CPLDs provide negligible standby current and ultra-low dynamic power consumption. Hence, incorporating a complex programmable logic device in your system will not significantly impact your power budget.

Compliance with the SDA Specification

The SDA (Secure Digital Association) specification states that one SD bus can only support one SD device. The clock pin can be shared, but the DAT[3:0] and CMD lines must be unique for every SD device. See Figure 2 for additional details.

This reference design is fully compliant with the SDA Specification. The following section will show you how to satisfy the above requirements while supporting any number of SD devices using a controller with a single bus.

CPLD Design

A block diagram showing typical use of this design for two SD devices sharing the same SD host interface can be seen in Figure 3. Conceptually, the design can be viewed and used as a bi-directional multiplexer. The host device controls the




Figure 2 – SD system bus topology

CPLD via the ‘Select’ signals, thereby dictating which SD device to communicate with. Once an SD device has been selected, the logic in the CPLD device automatically detects the direction of data flow and allows data to stream accordingly (either from the host to the SD card or from the SD card to the host). A directional control pin is not required, thereby making this design easy to use.

The host can access each SD device individually without affecting the state of the other when the multiplexer is switched accordingly. If neither the host nor the SD is driving data, the CPLD allows the system to be in the default high impedance with weak pull-up state. The primary purpose of this circuit is to provide additional SD capability to the host, but this circuit can also be used to provide level translation and/or logic isolation.

Implementation Details

Figure 4 shows the actual logic circuit for a 1:2 bi-directional multiplexer design, which can be described using VHDL. In the initial condition or idle state, the Host and SD cards should be high impedance with a weak pull-up. Hence, the circuit in Figure 4 is designed to 3-state the CPLD’s output buffers, thereby allowing the external pull-up resistors to take effect. Register A (A_REG) and Register B (B_REG) are both designed to be initialized to logic ‘0’ upon power-up.

The SD cards are selected via the ‘Select’ inputs to the CPLD. When ‘Select’ is logic ‘0,’ SD1 is chosen and when ‘Select’ is logic ‘1,’ the SD1 device is chosen. For simplicity




Figure 3 – Block level diagram: A bidirectional multiplexer

ty while describing this circuit, let’s assume in the following discussion that the Host is only choosing to communicate with SD1.

The auto-directional control aspect of this design is implemented in the following manner – a transaction is initiated when either the host or SD1 drives Low. For example, if the host wants to send data to the SD1 device, the host would begin by driving the A side Low. Upon driving Low, the logic in the circuit detects the Low going edge and responds by enabling the ‘B’ output buffer, but continues to keep the ‘A’ output buffer disabled. Specifically, when A is driven Low, a rising edge is delivered to

the clock input of A_REG. After clocking, A_REG’s Q output becomes logic ‘1’ and therefore prevents B_REG from receiving a clocking event. In parallel with the A_REG clocking and triggering, gate B1 outputs a logic ‘1’ when A goes Low. This enables the ‘B’ Output Buffer and, ultimately, B will follow A and drive Low.

Conversely, when it is driven from Low to High, gate B1 outputs a Low and 3-states the B output buffer. This forces B to go High via the external pull-up resistor. Once the A and B sides are both High, A_REG and B_REG are reset to 0. This process is repeated indefinitely. The reverse




Figure 4 – SD multiplexer circuit for two SD devices

happens when SD1 attempts to drive data toward the host. Additionally, if the host wishes to communicate with the SD2 device, the 'Select' inputs to the circuit are set to a logic '1' and the sequence of events are similar to the above.

Design Verification

Simulation Results

Functional and timing simulations have been extensively performed on this circuit using ModelSim, and test stimuli have been included with the reference design. Figure 5 shows some simulation results.

In the first part of Figure 5, the Select input is held Low. A dotted white line denotes a "Weak 1" condition, or in other words, represents a pulled-up state. In the first transaction, the host attempts to drive data toward SD1, and SD1 follows accordingly. Immediately after, the SD1 device attempts to drive data toward the host, and the host follows. Similar events happen when the Select input is driven Low. The host drives data toward the SD2 device, then the SD2 device drives data toward the host.

Hardware Results

Xilinx created an SD Multiplexer demo board, and has used this board to verify this bi-directional multiplexing design. Figure 6 shows the demo board, which features a CoolRunner™-II XC2C32A CPLD located in the center. Two SD card sockets are located along the top edge of the board. The bottom-most portion of the board is designed to mimic the physical dimensions of an SD card. Figure 7 shows the demo board plugged into a USB SD card reader. As expected, the XC2C32A allows a PC to communicate with either the first or second SD card flawlessly.

Device Utilization

Table 1 shows device utilization statistics for various implementations. As stated in the SDA specification, there are three signaling modes defined for SD cards: SPI Mode, 1-bit SD Data Transfer Mode, and 4-bit SD Data Transfer Mode. This design can be easily adapted for any chosen mode. The design can also accommodate any number of SD expansion ports, with default VHDL code set to two ports.



Figure 5 – Simulation results

Voltage and Current Considerations

The SDA specification contains stringent voltage and current requirements for SD cards. Programmable logic devices are ideal for this application because they are extremely low power and have features such as I/O Banking. The I/Os can be configured as 1.5 V, 1.8 V, 2.5 V, or 3.3 V, allowing them to interface to any SD device. CPLDs also contain I/O Banks, which allow for voltage translation capabilities between the processor and SD card.

The extremely low power nature of modern complex programmable logic devices allow for standby operation as low as 15 µA. The addition of a low power CPLD in a system will minimally impact the current budget.

VHDL Download

The VHDL files to compile and simulate these designs are located at: www.xilinx.com/products/silicon_solutions/cplds/resources/coolvhdlq.htm.

Conclusion

As SD devices gain in popularity, the need will increase for ways to support more than one SD device with host controllers. This article provides a verified solution to the problem at hand. This solution will give designers the flexibility to implement two or more SD devices into a system.



Figure 6 – Xilinx SD multiplexer demo board



Figure 7 – Xilinx demo board plugged into a USB SD card reader

Number of SD Expansion Ports	Device	Macrocell Utilization (SPI or 1-Bit Data Mode)	Macrocell Utilization (4-Bit Data Transfer Mode)
1	XC2C32A	13 Macrocells	19 Macrocells
2	XC2C32A	21 Macrocells	30 Macrocells
3	XC2C64A	27 Macrocells	39 Macrocells

Table 1 – Device utilization statistics for various implementations

Designing Portable Handsets Using CoolRunner-II CPLDs

CPLDs can improve processor-based handsets.

by Mark Ng
Staff Applications Engineer
Xilinx, Inc.
mark.ng@xilinx.com

Portable consumer electronic designs such as cell phone handsets, PDAs, and MP3 players are typically very high-volume products. Because of this, product designers look first to ASIC or ASSP methodologies to pack the greatest functionality into tiny, portable packages.

This solution “hits the mark” for dense functionality and usually has acceptable power consumption. But the consumer world is rapidly changing – features envisioned at one point in time can become obsolete within a matter of months, as competitors react to ever-changing technologies and market dynamics to deliver differentiated solutions.

The term “cutthroat” is frequently used to describe this level of competition. Mistakes are not tolerated, and mistakes are expensive. Choosing the correct ASSP or designing an ASIC correctly every time is nearly impossible. Mitigating these circumstances is crucial to sustaining market share.

Today’s designers are now looking beyond the fixed architecture of ASICs and ASSPs to discover the innate design flexibility and time-to-market benefits of programmable logic. Xilinx® CPLDs offer portable device designers a viable alternative to stan-

dard cell technology, providing one of the lowest cost, lowest power CPLDs in the industry with CoolRunner™-II devices.

Since 2001, the Xilinx CoolRunner-II CPLD family has provided designers with pricing low enough to beat that of discrete logic devices, allowing designers to easily implement a wide variety of logic functions in a single package. In this article, I’ll demonstrate ways to expand beyond the limitations of today’s ASIC/ASSP portable handset solutions, with simple, cost-effective, low-power programmable logic using CoolRunner-II CPLDs. As most handsets are OMAP-, XScale-, or i.MX-based designs, I’ll describe solutions to several specific problems, with links to application notes that provide in-depth details.

Level Translation

Interfacing two chips of different voltage standards is a common problem. Every type of memory is not made at every volt-

age standard, and microprocessors are offered at many voltages. Matching standards can be as simple as introducing level translators, but they are expensive and take more area than might be desired. Using a CPLD is a better solution and offers substantially greater flexibility. All CoolRunner-II CPLDs are capable of translating between two voltages, and some can handle as many as four.

CoolRunner-II CPLD I/O banks easily translate between voltages ranging from 1.5 V to 3.6 V in a single chip, as shown in Figure 1. But this totally disregards the programmability of the devices. You get the translation as part of the whole package, which means you get a bundle of logic, flip-flops, power reduction resources, and I/O buffers frequently priced below level translator chips. XAPP785 explains the details on taking advantage of this powerful feature to expand the capabilities of your OMAP, XScale, or i.MX designs.



Pin Expansion

High pin-count ASICs are more expensive than low pin-count ASICs, in general. If your logic needs dictate a low capacity but your I/O requirements dictate a high capacity, you may be paying for logic you will never use to gain the pins. One solution to this is adding a CoolRunner-II CPLD to operate as a “pin expander,” as shown in Figure 2.

The basic idea is to identify GPIO pins that typically operate at a slow speed. Then, rather than assign ASIC pins to them, attach CoolRunner-II CPLD pins to the slow-moving GPIO signals, serialize the signals, and import them to the ASIC on fewer net pins. Serializing/deserializing is done through simple, efficient shifting, and can drop the pin counts dramatically on expensive ASICs. Xilinx application note XAPP799 shows how to do this through an I²C port, but you can use other methods.

As an alternate viewpoint, OMAP, XScale, and i.MX processors provide specific pin mixes to support the applications their vendors deem appropriate. This doesn't mean that you must agree. CoolRunner-II CPLD pin expansion permits you to create your own GPIO pins of assorted voltages and additional capabilities (pulsing, PWM, individually 3-stated).

Pin Swizzling

CPLDs offer the ability to rearrange your pinouts when PCB layout errors occur. This valuable quality is key to keeping you on schedule and within financial and power budgets. Correcting misconnections on a board without having to re-spin the PCB can shave weeks to months off of product schedules.

CoolRunner-II CPLDs are built from powerful logic blocks using programmable logic arrays that can reassign pin logic at will. You will be amazed at how well these devices retain pinouts through multiple edits yet permit re-assigning a design

onto different pins as needed. The CoolRunner-II family data sheet explains the architecture and points you to application notes that give all the detail you will need to understand the value of PLAs.

Power Control

Quick power up is one of the strengths of CPLDs. Containing their own configuration cells permits CoolRunner-II CPLDs to power up and direct the activities of other chips as they subsequently arise. This includes some power regulators, which may

be sequenced by CoolRunner-II CPLDs, as well as other controlling signals that need to be well defined early in board operation. Xilinx application note XAPP436 describes some of these capabilities.

Power Reduction

XScale-, OMAP-, and i.MX-based chipsets all include some version of the ARM microprocessor. This is not a surprise. Advanced RISC machines started early with developing low-power methods to operate microprocessors. Subsequently, the licensing vendors have all added their own methods to further reduce processor power. Typical power reduction operations include clock gating, voltage throttling, and on-board memory management to reduce transfers within the device. These are sometimes referred to as run, wait, doze, sleep, and hibernate.

Also, operating systems like Symbian have added “power awareness” to the mix, so that unused resources can be parked in the lowest power mode possible for the current tasks being executed. This all works well and lowers processor power. However, lowering power in the rest of the system exceeds the scope of these methods.

Enter CoolRunner-II CPLDs. CoolRunner-II CPLDs are designed to be inherently low-power parts. That is important, but alone is not enough. CoolRunner-II special features also can be used to lower the power in other devices. Using clock dividers and Xilinx DataGATE technology can reduce power in many (if not all) of the chips on your design, as shown in Figure 3. Blocking power to other chips can also reduce electromagnetic fields being propagated on your board and emanating from your system. This powerful signal blocking technique can pay off in many ways.

Logic Consolidation

Having three two-input AND gates, two three-input OR gates, and a




Figure 1 – CoolRunner-II CPLD-level translation of TI OMAP signals




Figure 2 – CoolRunner-II CPLD pin expansion of XScale processor




Figure 3 – DataGATE blocking extraneous switching to various devices

Schmitt buffer package on your board can burden your bill of materials (BOM), eat away at your power and cost budgets, and lower your reliability. Collecting that stray logic into a consolidated, low-power CoolRunner-II not only solves these problems but stores additional unused logic right there on your board – ready to use with future improvements/edits. WP214 shows what you can expect from collecting logic gates/flip-flops into CoolRunner-II CPLDs. Table 1 summarizes the “burn rate” for logic.

Conclusion

CoolRunner-II CPLDs are quickly becoming the standard for low-power, low-cost, high-volume, portable consumer products. This article has focused on how these powerful products can make life easier when building systems with OMAP, XScale, and i.MX processors, but CoolRunner-II CPLDs work just as well with many other processors to add functionality, save power, and get products to market fast.

Function	Macrocells	P-Terms	Flip-Flops
Shift Register (Simple)	1 per bit	1 per bit	1 per bit
Counter (Simple)	1 per bit	1 per bit	1 per bit
2:1 Mux	1	2	0
4:1 Mux	1	4	0
8:1 Mux	1	8	0
8-bit Loadable Shifter	8	16	8
8-bit Loadable/SL/SR Shifter	8	24	8
8-bit Loadable Counter	8	16	8
8-bit Load/Up/Dn Counter	8	24	8
Full Adder / Bit	2	7	0/1 (optional)
2:4 Decoder	4	4	0
3:8 Decoder	8	8	0
4:16 Decoder	16	16	0
8-bit Equality Comparator	1	16	0
And/Nand Gate (1-40 Inputs)	1	1	0
Or/Nor Gate (1-40 Inputs)	1	11	0
Ex-Or/Ex-Nor (2-3 Inputs)	1	2-3	0
Level Translator (Per Bit)	1	1	0

Table 1 – Macrocell “burn rate” for common TTL functions

TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Read the application note, “Level Translation Using Xilinx CoolRunner-II CPLDs.”
- Read the application note, “An SMBus/I2C-Compatible Port Expander.”
- Read the application note, “Managing Power with CoolRunner-II CPLDs.”
- Read the application note, “Using CoolRunner-II Advanced Features” to learn how to do signal blocking.
- Read the white paper, “The Real Value of DataGATE” to learn how much power you can save with DataGATE.

Supporting Your Future
HUNT ENGINEERING
www.hunt-rtg.com

USB connected Programmable FPGA systems

V-II Pro PowerPC

- Virtex-II Pro XC2VP7
- 256 Mbytes DDR Memory
- Configurable digital I/Os
- PowerPC boot FLASH
- USB 2 or Standalone

Software Defined Radio

- Virtex-II FPGA 1M gates
- 2 ch 125Mps A/D and D/A
- TI C6203 DSP
- 32Mbytes SDRAM
- Configurable Digital I/O
- USB 2 or Standalone

Imaging with Virtex-4FX

- Virtex-4 FPGA FX12
- 128Mbytes DDR Memory
- CameraLink connection
- VHDL and PowerPC Imaging Libs
- USB 2 or Standalone

Programmable hardware with cables, device drivers, loading tools, examples and Power Supply. Systems can be used connected to a PC using USB, or can function standalone (without USB) using the initialisation PROMs.

sales@hunteng.co.uk
+44 (0)1278 760188
www.hunt-rtg.com

Easing Design Challenges with CoolRunner-II CPLDs

The CoolRunner-II CPLD Starter Kit promotes ease of use.

by Arthur Yang

Technical Marketing Manager, Micro Scale Division
Xilinx, Inc.

arthur.yang@xilinx.com



You shake your head in frustration as you stare at the flat line on the oscilloscope. For the umpteenth time, you shuffle through the data sheets on the table and check the input waveforms and pin connections. The clock ticks past six; with a sigh, you reach for the phone. Another late evening with soldering irons, neon-colored scope traces, and takeout pizza. What a way to spend a Friday night.

Prototyping with a new device doesn't have to be such a headache. Imagine a full-featured design kit that is easy to use and inexpensive. It would have a number of different visual outputs, LEDs for status values, numeric displays for counters and state machines, and character displays for text messages. It would have several I/O ports to test interface standards and peripherals, and user interface buttons and switches. It would be easy to get power readings. It would interface to a PC so the software guys could interact with it. And it would be cheap enough to not even cause a blip on the budget.

Working together, Xilinx and Digilent Inc. have partnered to create the new CoolRunner™-II Starter Kit, making what you once imagined a reality.

One feature built into the CoolRunner-II CPLD Starter Kit is real-time power measurement and reporting. XMeter is a program that reads power measurements for the 1.8-V power supply and charts the data in real time.

CoolRunner-II CPLD Low-Power Features

At the heart of the CoolRunner-II CPLD Starter Kit is a 256-macrocell CoolRunner-II CPLD, an ultra-low power programmable logic device. The CoolRunner-II family has quiescent power as low as 13 μA and a number of design features that help reduce dynamic power as well. Let's briefly discuss two of these features, the clock divider and DataGATE.

CoolRunner-II devices, starting from the 128-macrocell device, have a built-in clock divider circuit. This component allows you to not only save user logic but also lower power by reducing the frequency of a clock that may not need to run at full speed. The clock divider provides divisions of 2, 4, 6, 8, 10, 12, 14, and 16 and outputs directly onto a global clock net. If used in conjunction with dual-edge triggered registers, you can obtain odd values of division.

DataGATE is effectively an input 3-state. When signals toggle at an I/O pin, current flows at the I/O transistors as well as in all of the traces within the CPLD. There are many situations where you can ignore specific I/Os, such as in idle states or don't-care conditions. In these situations, you can disable unnecessary inputs to save current.

DataGATE is also useful for signals that have slow rise or fall times. The largest amount of power consumed in the I/O transistors occurs when the voltage at the pin is roughly one-half of V_{CCIO} . You can use DataGATE to periodically enable an input to sample its state and then disable it to save power. Without DataGATE, very slow switching signals will waste multiple milliamps of current.

CoolRunner-II CPLD Kit Features

The CoolRunner-II CPLD Starter Kit has eight peripheral module expansion ports, each with four I/O pins. Each of

these ports is capable of accepting any of 29 peripheral modules (PMods). The board itself has two push buttons, and the kit comes with three PMods to get you started quickly: quad switches, a PS/2 connector, and a dual seven-segment display.

PMods offer total flexibility without the expense of time and money. This means no more time wasted soldering devices onto a breadboard, no more paying for an on-board Ethernet PHY that you'll never use, and no need to purchase an expensive programming cable or design software. Simply buy the modules that you need for your project and plug them into any of the I/O expansion ports. Update the pin constraints in your design file and you're ready to go.

A number of PMods cover the user interface. There are simple switches (such as quad switches, buttons, or a rotary switch) and more complex PMods (such as the PS/2 adapter that facilitates connection to a standard PS/2 keyboard). An RS232 adapter even allows data transfer using a UART.

For display purposes, there are several LEDs on the board itself, along with the provided dual seven-segment display. For more detailed display options, there is a 16 x 2-character display PMod.

Here is a short list of available PMods:

- Analog to digital: two 12-bit 1 MSPS A/D
- Digital to analog: two 12-bit 1 MSPS D/A
- 16-Mb SPI flash
- Infrared light detector
- H-bridge for DC motor drive
- Speaker/headphone amplifier
- Servo-motor connector

XMeter: Built-In Power Measurement

One of the greatest features of a CoolRunner-II CPLD is its low power operation. One of our more popular demonstrations is the "fruit-powered" demo, where we run the device from the electrical current generated by a few pieces of grapefruit.

But fruit won't replace the AAA battery, so you need a way to exactly determine the operating lifetime for your design. Power estimators have their place, but measuring the real thing is critical.

One feature built into the CoolRunner-II CPLD Starter Kit is real-time power measurement and reporting. XMeter is a program that reads power measurements for the 1.8-V power supply and charts the data in real time. Figure 1 shows the XMeter program in operation. The sample design was consuming 200 μA at full frequency; when DataGATE was enabled, the current dropped down to ~86 μA .




Figure 1 – XMeter showing DataGATE power savings

Note that DataGATE power savings are entirely design-dependant. You can implement DataGATE in a number of different ways based on your design requirements, so expected power savings will vary from this example.

XMeter makes it easy to load up different versions of your code and determine which one has the optimum power consumption profile, taking the guesswork out of power optimization.




Figure 2 – PC application communicating with CoolRunner-II CPLD

PC-to-Kit Communication

The programming cable provided with the kit is a standard USB to mini-USB cable available at any electronics store. Chances are one was included with your digital camera or cell phone. JTAG communication between the PC and the


board is handled through a Cypress EZ-USB chip, so to program the board you would use Digilent’s Export programming software. (The board supports programming from iMPACT, but this requires a separate Xilinx programming cable not included in the kit.)

The advantage of this communication method is that an application programming interface (API) allows you to create your own PC-based programs and communicate directly with the CPLD. Figure 2 shows a program that allows you to read from and write to registers within the CoolRunner-II CPLD.

Conclusion

The CoolRunner-II CPLD Starter Kit is optimized for low power operation and removes many roadblocks for prototyping. Ease of use is covered by pre-built modules that can be plugged into the board without warming up the solder gun. The suggested resale price of this kit is

less than \$50. Every chip on the board was selected for low power and power measurement is built-in, so easy power benchmarking and design evaluation is covered. This kit offers low power, low cost, high functionality, ease of use, and no compromises. Get your evenings back. ●●



XCELL JOURNAL GOES DIGITAL.

The digital edition of *Xcell Journal* combines the award-winning look of our printed magazines with all of the multimedia capabilities the Web has to offer.

An enhanced viewing technology, viewable on any Internet browser, lets you browse, search, and make notes, as well as e-mail authors and click through to our advertisers’ highlighted websites.

We invite you to sign up for your free subscription today – visit:

www.xilinx.com/xcell/



TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Order the CoolRunner-II CPLD Starter Kit.
- Read about DataGATE power consumption savings in the white paper, “The Real Value of CoolRunner-II DataGATE.”
- See a complete list of available PMods.
- Download the Digilent Adept SDK API and transport design.



Designing Digital Displays with Spartan-3 Generation FPGAs


Programmable solutions improve your time to market.



by Glenn Crow
Strategic Marketing Manager, General Products
Division
Xilinx, Inc.
glenn.crow@xilinx.com

The dynamics of the digital display market to constantly strive for improvement while reducing total system cost has caused models to change between two to four times per year. These changes can be challenging as displays evolve; incorporating the right feature set in a very short time to market (TTM) becomes the standard business model to lead the pack or simply to stay competitive. Choose the wrong features and the product will be slow to gain market share; more importantly, it will fail to reach its full potential in a market that turns several times a year, affecting revenue.

Extending the life of an existing ASIC is imperative to meet cost goals. The TTM of an ASIC cannot keep pace with the dynamics of this rapidly paced market without revenue loss caused by a missed feature inclusion or design modification. When an ASIC requires a re-spin because of a design error or newly required feature, the market window of opportunity can be missed.



Functionality Supported by Xilinx

Figure 1 – This diagram shows five different blocks marked with a yellow dotted line on two separate boards that can be implemented in a Spartan-3 generation device. Depending on the design requirements and cost target, one or all of the Xilinx-supported blocks on the video/tuner or panel display board could be implemented in a single device.

By incorporating a Xilinx® Spartan™-3 Generation FPGA alongside an ASIC, the life of the ASIC can be increased, allowing the development and engineering costs to be spread out over a larger number of devices. This enables fast TTM and allows many new features to be incorporated at any point throughout the design process.

Coupling the ASIC with an FPGA provides optimum design flexibility, allowing the design to be quickly upgraded and keep pace with the market while also getting the maximum benefits from portions of the design that will be re-used in future products.

For more than five years, Spartan-3 Generation FPGAs have helped designers of digital displays by reducing system cost and complexity while integrating new and unique features. These FPGAs provide the industry's lowest cost solution by eliminating the need for external components.

In this article, I'll explain how Xilinx consumer display solutions incorporate complex video image enhancement algorithms effortlessly, implement hassle-free

commonly used high-speed panel interfaces, and provide instant access to the industry's widest support of I/O interfaces for changing standards and protocols. In addition, they keep your design and IP protected with design-level security technology and support green requirements with dual-power management modes.

Video Image Enhancement

Xilinx Spartan-3 Generation FPGAs provide pre- and post-image processing functions to enhance or extend the life of existing ASICs/ASSPs. As display technology and sizes change, adjustments to the image may be required that the ASIC/ASSP cannot perform. A custom image processing algorithm can also be used for product enhancement and differentiation.

As shown in Figure 1, blocks 1 and 2 can easily implement functions such as motion-adaptive temporal noise reduction, de-interlacing artifact filters, intra-frame noise reduction, color space converters, and dynamic range compression to improve image quality. These functions and more

can be easily implemented with embedded multipliers, DSP blocks, and block RAM to create product differentiation.

Hassle-Free Panel Interfacing

Each new generation of displays integrates a wide variety of panel sizes and types containing distinct interface, timing, and image characteristics. Newer LCD panels incorporate even more complexity, with features such as dynamic backlight and ambient light control. Spartan-3 Generation FPGAs support the latest interface I/O standards, as well as flexible timing controllers to adapt to different panel timing variations.

In addition, you can easily implement image processing to improve the image quality of each panel and enhance features like backlight control for truer colors, deeper blacks, and more detail in dark scenes. This represents some of the functionality that can be accomplished in blocks 2 and 3 (Figure 1).

The best way to keep pace with new panel technology as well as panel interface standards is with the only FPGAs that have



native support (without external components) for the latest differential standards such as Transition Minimized Differential Signaling (TMDS) and Point-to-Point Differential Signaling (PPDS).

Instant Access to Changing Standards and Protocols

Today's displays must support a variety of standards and protocols that range from video input such as DVI and HDMI to internal memory, ASSP/ASIC, and panel interfaces. Spartan-3 Generation FPGAs can help you stay connected with existing interfaces like DDR, DDR2, USB, and Firewire (IEEE1394) standards. Blocks 4 and 5 in Figure 1 show some of the many interfaces/features that may change from model to model.

As consumers demand more features on digital displays, the only way to quickly react is through a programmable solution. Spartan-3 Generation FPGAs offer the ability to support network, hard drive, and multiple memory interfaces as well as the latest digital video interfaces.

The support of different standards or different combinations of standards is as easy as reconfiguration with MultiBoot. MultiBoot allows you to change the functionality by simply reconfiguring the Spartan-3 Generation FPGA, which supports low-cost serial, parallel, and platform flash configuration devices.

Protection with Built-In Design-Level Security

Protecting valuable, proprietary image enhancement algorithms and timing control systems from cloning, overbuilding, and reverse engineering has become increasingly critical. Low-cost, robust security that can be designed quickly to protect

product improvements without impeding quick TTM is a key feature.

Xilinx offers many levels of security that range from hidden bitstreams to advance data manipulation. Spartan-3 Generation FPGAs offer the right security solution to help protect key IP, data, and the complete design in high-volume, low-cost products. With the ease and flexibility to change the security, each model or generation of displays could have a different scheme, making tampering or theft of key circuitry or IP even more difficult.

Go Green

Increased global regulations are requiring more consumer products to reduce standby power consumption. The Spartan-3 Generation offers multiple power-saving modes to help meet your power budget goals. All Spartan-3A, 3AN, and 3A DSP device platforms offer suspend and hibernate modes. The hibernate mode can reduce static power by as much as 99%, while the suspend mode will reduce static power by at least 40%.

Conclusion

Spartan-3 Generation FPGAs offer today's digital display designers the optimal low-cost programmable solution. No other solution combines the flexibility of programmable logic with the native support of cutting-edge I/O and interface support, a wide array of IP, flexible design-level security, and multiple development boards and starter kits to quickly design and test your video applications.

No extra components needed for the built-in security and access to standards that keep you ahead of the game give you the ability to re-use your workhorse ASIC or ASSP and extend its usefulness. ●●●

New Xilinx User Community Works for You 24/7

Launched on August 13, 2007, the Xilinx® User Community (<http://forums.xilinx.com>) is an interactive message board designed for sharing knowledge and information about Xilinx products and technology. Users can start threads, share expertise, exchange ideas, peruse a wealth of valuable information, and submit questions to other users or Xilinx employees.

The community has the following categories, each comprising several boards:

- General discussion
- Silicon devices
- Design tools
- Intellectual property
- Boards and kits

So far, the international Xilinx User Community has more than 1,000 accounts registered, with more than 800 threads posted from around the world.

Don't hesitate. Join the community and start sharing now. Together, let's make the Xilinx User Community an extremely useful resource for all Xilinx users.

Please e-mail forumadmin@xilinx.com if you have any feedback about this community.



TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Learn more about Xilinx digital display solutions.
- Learn more about Spartan FPGAs.
- Download your free ISE™ WebPACK™ software design tools and begin designing immediately.



Taking Device DNA Technology to the Next Level

Xilinx and Helion work together to make low-cost FPGA designs more secure.

by Maureen Smerdon
Strategic Marketing Manager
Xilinx, Inc.
maureen.smerdon@xilinx.com

Graeme Durant
CEO
Helion Technology Limited
graeme@heliontech.com

In today's world, device cloning and unauthorized overbuilding have driven design security to a higher level of importance. In consumer electronics applications, design security has become as important as the media encryption flowing across the device. Example applications include broadband access, wireless networking, routers, high-definition DVD, and DVRs.

In 2007, Xilinx introduced its Device DNA technology, offering a range of security features to safeguard against reverse engineering, cloning, and unauthorized overbuilding. Xilinx offers this revolutionary capability in its latest low-cost Spartan™-3 FPGAs, specifically Spartan-3A, Spartan-3AN, and Spartan-3A DSP devices.

To expand the capabilities of this technology, Xilinx worked closely with Helion, a Xilinx Alliance IP partner specializing in data encryption and security solutions, to develop Device DNA Checker, an IP block specifically designed to work in conjunction with Spartan-3 devices to provide designers with an even more solid and robust security solution.

Device DNA

Before understanding how the Helion Device DNA Checker works together with Spartan-3 devices, let's look at the core of Xilinx Device DNA technology.

Xilinx Device DNA technology is a permanent, factory-set ID code that is different in every device. Starting at 57 bits and programmable up to any bit length to meet user requirements, the Device DNA enables advanced 64-bit, 128-bit, or even more complex 256-bit algorithms. This unique ID can be used to tie a design to a specific FPGA.

Every design can hold a unique authentication algorithm to further increase design security and reduce reverse engineering and potential threats. This allows you to precisely select the complexity of the algorithm and

directly control the amount of security needed. With this flexibility, you can spend as much or as little on security as necessary, depending on the security needs of your specific application.

The authentication algorithm you choose is implemented in the FPGA and takes in the ID value and generates a unique result for that ID. The result is then stored wherever you choose, such as in external memory or internal flash (for Spartan-3AN FPGA devices only). Each time the FPGA is powered, the generated and stored results must match in order to authorize that device. The algorithm is the secret to the security because it is known only to you, and is almost impossible to determine without access to the original design source.

Designers have full flexibility in customizing algorithms for both authentication as well as responses to authentication failures. With its embedded flash, the Spartan-3AN device further enhances security by hiding any configuration communication from the outside, making it extremely difficult to understand the design contained within the FPGA.




Figure 1 – Device DNA design-level security




Figure 2 – ATM transactions and Device DNA security flow

This security or licensing process is designed with complete flexibility in mind. You can easily change the security or licensing process from model to model, thus increasing design security. The read-only Device DNA is accessible through either the external JTAG port or the internal DNA port for easy connection to the security algorithm.

If a cloner or overbuilder copies the bitstream and places it into another FPGA, the Device DNA of the new FPGA will be different. After using the algorithm to check the Device DNA, the design will return an unauthorized or a failed result, allowing the user or designer to determine how to respond to the security breach.

Figure 1 shows a graphical representation of the security flow.

The Device DNA security process is like an ATM transaction. To withdraw money from an ATM, you insert your ATM card and enter your PIN on the touch pad. If your card and associated PIN match the ID stored at the bank, your transaction is approved and your money is available. If the match fails, your transaction is rejected and you do not get your money. Figure 2 shows the how the ATM transaction compares to the security flow.

Unauthorized Operation

Users have more flexibility in the determination of what happens when a design

is non-authorized or illegitimate. The simplest way is to disable functionality. This can easily be accomplished by utilizing global control signals such as resets, 3-states, or clock gating.

During normal operation, the device is powered up and the bitstream is loaded for configuration of the FPGA. The security algorithm reads the Device DNA and generates an active value. It then compares the active value with the check value stored during the initial setup. If the check value is equal to the active value, the device will operate normally.

You can design your product to respond in one of the following ways when the two values do not match:

- No functionality. The design completely ceases to function. This can be easily implemented in a Spartan-3 FPGA by using global control signals like 3-state, gated clocks, or flip-flop clock enable.
- Limited functionality. The design has partial or basic operation but key functionality is disabled or bypassed. This response allows a third-party test house or contract manufacturer to build and test while preventing overbuilding. It also allows the system to be run in evaluation or demo mode.
- Time bomb. The design operates with full functionality for a predetermined amount of time before shutting off. This response allows a third-party test house or contract manufacturer to build and test. It also allows the system to operate in demo mode or for IP evaluation.
- Self-destruction (Spartan-3AN devices only). Uses flash sector erase and lock-down protection to erase all sectors and permanently lock flash memory to all zeros. This response prevents repeated unauthorized access attempts.

Authentication Algorithms

The techniques necessary to create a solution using Device DNA and protect a product against overbuilding and cloning are described both in this article and in Xilinx white papers such as WP266, "Security




Figure 3 – Device DNA Checker block diagram

Solutions Using Spartan-3 Generation FPGAs” (www.xilinx.com/documentation/white_papers/wp266.pdf). However, actually implementing such a solution quickly and securely and ensuring that it does not use excessive logic resources in your device can be time-consuming. Thus, there is a strong case for using a ready-made and tested IP block specifically designed for the purpose.

Helion, using its well-established cryptographic IP and expertise, developed a ready-made solution comprising a pair of IP cores, allowing designers to easily implement Device DNA technology in their designs.

The Helion approach uses a selection of strong cryptographic functions to implement the security algorithm and various obfuscation techniques to ensure that any reverse-engineering attacks are made suitably difficult. Most importantly, the design is highly parameterized so that each user can make their implementation unique to them. Other customers using these IP blocks – even Helion itself – cannot work around the scheme.

The Helion Device DNA Checker

The Helion Device DNA Checker comprises a ready-made IP building block that is dropped into the user design to sit alongside the main FPGA application. It has relatively low resource requirements and is intended to coexist easily with the main design. Its interface is simple; it requires an input from the Device DNA block in the FPGA, and after processing generates a



Figure 4 – Using Device DNA in a software application

“pass/fail” flag. This flag can be used to degrade normal operation of the rest of the design if the check does not pass. Simple “go” and “done” signals start the check process. Figure 3 shows a block diagram of the Helion Device DNA Checker.

Helion’s Device DNA Checker uses both the Device DNA code plus additional check bits stored in non-volatile system memory to enhance security. These additional bits are a mixed-up combination of randomly generated bits that form part of the check algorithm; randomly generated bits that are ignored (sometimes called “salt” bits); and the final stored check code itself. This additional data is simply streamed into the Checker after the Device DNA bits via the cascade input to the Device DNA block.


So that each customer using this solution can be confident that their protection is unique to them, the exact operation of the Checker is highly parameterized and not fixed by Helion. You set your own secret “configuration value” when the

Checker block is synthesized. This is effectively a long string of bits that define the low-level operation of each part of the check algorithm used in the block. Because this is set at synthesis time, the configuration value can result in significantly different logic for each user.

Neither Helion nor any other user of this technology can replicate the algorithm without access to the secret configuration, adding another layer of protection into the system. In the diagram, the functions of the blocks shown in pink are all affected by this secret compile-time configuration value.

Internal Operation

Within the Checker, the input stream goes one of three ways. It is either used in the cryptographic hash, used as the stored check code (to be compared to the generated active code after the hash is complete), or discarded. The choice of which bits go where is determined by the user configuration previously outlined and is different for each user.



Note: Function of blocks in pink are controlled by user configuration at synthesis time

Figure 5 – Using Device DNA as a hardware download

Any bits destined for hashing or for the check code go through “obfuscation” functions before being used. These are intended to further obscure the processing being performed. Their exact definition is again driven by the user configuration outlined above, so each user’s implementation will be different. Even the choice of cryptographic hash function used is determined by your secret configuration. Once the appropriate bits have been hashed to form an active code and the stored check code has been regenerated, the two are compared and the final pass/fail result is indicated.

Getting Authorized

For each legitimate product that has a unique Device DNA and contains the DNA checker, the product must be activated by generating the additional Device DNA bits to be stored in non-volatile memory. This can be achieved in two ways:

1. Software application. In this case, a software application supplied by Helion models the algorithms used in the Checker to calculate the necessary additional Device DNA data. It takes in a list of authorized Device DNA values and a copy of the customer’s secret compile-time configuration and generates data sets to be placed into the systems’ non-volatile memories (see Figure 4). This application could be used remotely to generate data sets

on demand, perhaps automatically over the Internet. Clearly this software and the user-configuration information must only be distributed to highly trusted parties; otherwise the whole scheme is compromised. However, remotely enabling manufactured units may be a great way to achieve separation between trusted and non-trusted parties in the manufacturing process.

2. Hardware download. In this case, an optional second IP block called the Helion Device DNA Generator implements the same algorithms as the Checker to generate the necessary additional Device DNA data (Figure 5). You would put together a special programming bitstream incorporating this block using the same user configuration as the Checker at synthesis time. This does not include any of your normal design, as it is only required for configuring non-volatile user memory with the appropriate additional

Device DNA data. This FPGA image could be self-contained, taking the data generated by the Helion block and burning it into the appropriate non-volatile memory, thereby enabling that unit.

Again, this special bitstream must only be distributed to highly trusted parties. It would need to be loaded once into the manufactured hardware to authorize the unit, so it cannot be used remotely. This may be more appropriate in some manufacturing processes than the alternative software approach, however.

Conclusion

Helion’s solution for accessing the benefits of Xilinx Spartan-3 FPGA Device DNA technology is both off-the-shelf yet unique to each user. It is easy to use and requires minimal device resources. It can be deployed in an existing design very quickly, offering the benefits of protection against overbuilding and cloning during outsourced manufacture. 🌈

TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Learn more about the Spartan-3A device family or the Spartan-3AN device family.
- View a security demo on the Spartan-3AN device family.
- Learn more about Helion, www.heliontech.com.

A High-Speed Broadcast Video Connectivity Solution

Xilinx Spartan-3E and Spartan-3A FPGAs, National Semiconductor PHY, and the Xilinx protocol stack provide a cost-effective and flexible approach to the challenges of multi-rate broadcast.

by Bob Feng
Spartan Applications Engineer
Xilinx, Inc.
bob.feng@xilinx.com

Mark Sauerwald
SDI Applications Engineer
National Semiconductor
mark.sauerwald@nsc.com

Today's high-speed video application designers have significant challenges in addressing both the digital IP and analog physical interface requirements of their products. Because digital and analog components often have very different requirements, trying to support both in one ASSP chip often compromises the quality or cost-effectiveness of the solution. It can also be difficult to find a solution that has exactly the right IP and physical interface without waste in either area, or with the flexibility to meet the requirements of multiple standards.

A new chip set featured by Xilinx and National Semiconductor combines the best of the digital and analog worlds into one

highly integrated solution. This solution, including the protocol IP stack, is handled by the Spartan™-3E or Spartan-3A FPGA silicon. The analog section is handled by National Semiconductor SDI PHY products for the greatest signal quality with the lowest jitter. It allows professional audio/video broadcast (AVB) system developers to concentrate more on their own specific video content processing functionality and IP instead of the front-end interface connectivity.

SDI Video Standards

Serial digital interface, or SDI [SMPTE-259M], is a broadcast industry standard widely adopted today to transport uncompressed standard-definition (SD) video signals over a single coaxial cable. By definition, SDI typically supports data rates of 270 Mbps to cover screen formats of 480i at 60 Hz (480i60).

High-definition (HD) SDI, or HD-SDI [SMPTE-292M], boosts the bit rate up to 1.485 Gbps to support high-definition formats like 720p60 and 1080i60.

Three-gigabit SDI, or 3G-SDI [SMPTE-424M], further extends the serial digital throughput up to 2.97 Gbps in order to carry the highest screen resolution: 1080p60.

National Semiconductor PHYs

National Semiconductor offers a complete portfolio supporting the physical layer transmission for SDI applications. National's new family of SDI serializers and deserializers have speed grade options supporting SD SMPTE 259M at 270 Mbps, HD SMPTE 292M at 1.485 Gbps, and the new 3-Gbps standard (3G-SDI) SMPTE 424M at 2.97 Gbps (Table 1).

	Soft SERDES	Pixel Processing
SD-SDI	27 MHz	27 MHz
HD-SDI	148.5 MHz	74.25 MHz
3G-SDI	297 MHz	148.5 MHz

Table 1 – FPGA design frequency domains



National's LMH034x family highlights superior analog performance:


- Ultra-low output jitter: 50 ps typical at HD and 3-Gbps rates (Figure 1)
- Exceptional input jitter tolerance: 0.6 UI minimum (Figure 2)
- Integrated, high-precision PLL for serial clock reference and data recovery
- Integrated cable driver in LMH0340 transmitter

- Integrated serial re-clocked loop through and driver
- Low power consumption
- TX: 435 mW
- RX: 590 mW
- No external VCOs or clock required

In addition to leading-edge analog performance, National's LMH family reduces the traditional parallel bus between the PHY device and the host

FPGA from a 20-bit, single-ended interface to a five-channel low-voltage differential signaling (LVDS) interface. This innovative narrow differential bus reduces EMI and simplifies board layout by reducing the number of traces on the interface and using fewer pins on the host FPGA. Additionally, National's discrete PHYs do not require any external VCOs or jitter-reducing PLLs (Figure 3).


The combined National/Xilinx Spartan solution enables low-cost FPGAs into the high-end AVB market supporting SD, HD, and 3-Gbps data rates for professional video applications.



Equipment: Tektronix CSA8000 sampling scope with 20-GHz sampling heads.

Input Signal: PRBS 2¹⁵-1
Data Rate: 2.97 Gbps

Figure 1 – LMH0340 3-Gbps output jitter: 30 ps at HD and 3G rates



Data Rate: 2.97 Gbps
Equipment: Agilent J-BERT

Figure 2 – LMH0341 minimum input jitter tolerance: 0.6 UI

Spartan Features for Video Applications

The Spartan-3E and Spartan-3A FPGA families suit many aspects of video applications by offering high performance, high density (logic and I/O), great flexibility, and scalability with unique, cost-effective features such as:

- 50,000 to 1.6 million system gates
- True LVDS differential I/O drivers at more than 666 Mbps, with internal termination on receiver for direct chip-to-chip communication
- Double data rate (DDR) I/O registers at more than 300 MHz to increase effective bandwidth beyond 600 Mbps
- 18-Kb dual-port block RAMs at more than 200 MHz for FIFOs and data buffering
- Dedicated 18 x 18 multipliers at more than 200 MHz for high-speed digital signal processing
- Digital clock managers (DCMs)
- Clock deskew
- Frequency synthesis
- High-resolution phase shifting
- Wide frequency range (5 MHz to more than 300 MHz)
- Full programmability to easily modify the design during development or in the field, or to support multiple standards in a single solution




Figure 3 – SDI bill of materials reduction

- 20:5/5:20 LVDS soft serialization and de-serialization (SERDES)
- SMPTE scrambling/descrambling
- Video framer/de-framer
- CRC and line number insertion
- Rasterization
- ANC insertion
- Video standard detection and flywheel

The FPGA design is effectively divided into two frequency domains: “soft SERDES” and “pixel processing,” as illustrated in Table 2. The clock frequency used in the soft SERDES is typically only half of the serialization bit rate, by leveraging the DDR technique. On the other hand, the pixel processing clock frequency is determined by the relevant video transmission format: 74.25 MHz for 720p60 and 148.5 MHz for 1080p60.

The timing closure challenge is mainly on the soft SERDES side, as 297-MHz operation is required to achieve 594 Mbps across all of the differential channels. The Xilinx® Spartan applications team has been offering this soft SERDES reference design in a beta version since May 2007. Since then, Xilinx and National Semiconductor have conducted extensive testing. All three data rates have passed BERT test suites developed by Xilinx. Figure 4 illustrates the basic SERDES construct.

Xilinx has a long history of supporting SDI interfaces in the Virtex™ family of FPGAs. XAPP514, “Audio/Video Connectivity Solutions for the Broadcast Industry,” is a video connectivity IP and reference design book that details all aspects of the protocol stack: SDI, HD-SDI, DVB-ASI, SDTV/HDTV test pattern generation, and even embedded audio. Xilinx and National Semiconductor are actively working to port these highly valuable reference designs into Spartan-3E and Spartan-3A FPGAs. Figure 5 illustrates a list of successfully ported reference blocks used for demonstration purposes based on an internal evaluation board. Figure 6 illustrates an SMPTE 75% color bar display generated by the board.




Figure 4 – Basic soft SERDES construct in Spartan-3E FPGAs

- Software and IP to quickly implement key features of video applications
- Design examples and reference boards to get started quickly

By using FPGAs, you can be compliant to industry standards while differentiating yourself from your competitors. Such differentiation may be too difficult to find using an ASSP solution and too expensive to address with an ASIC. The flexibility of a programmable solution provides faster time

to market, while field updates provide longer time in market. Numerous standards (and versions) cause uncertainty, so designs need flexibility in transmission schemes, MPEG profiles, display formats, and color correction.

Interconnect Soft SERDES and Protocol IP Stack

While the National Semiconductor PHY takes care of the SDI physical interface, the FPGA plays an essential role in supporting all digital functions in the protocol IP stack, including:



Target Applications

Xilinx low-cost Spartan-3 generation FPGAs have been used successfully in a wide range of consumer and professional video applications. These include a JVC professional broadcast HDV camera/recorder using the Spartan-3E FPGA. The combination of the Spartan FPGA for the digital logic and National Semiconductor PHY for the analog interface opens up new possibilities in high-end applications in professional video, broadcasting, and digital cinema. Applicable products include high-definition video cameras, digital video recorders, video editors, and display monitors.

Conclusion

The power of Xilinx Spartan-3E and Spartan-3A FPGAs, combined with a proven National Semiconductor SD/HD/3G-SDI transceiver and the XAPP514 protocol IP, delivers a truly cost-effective solution to the ever-increasing data throughput requirements of broadcast video applications. Although the complete hardware solution is available today, a complete SDI evaluation kit will be offered by Xilinx distribution partner Avnet in the first quarter of 2008. ●●

Product ID	Description	Max Data Rate	Data Rates Supported	SMPTE Standards Supported
LMH0340	Serializer and Driver	3G	2.97G 1.485G 270M	424M 292M 259M
LMH0341	Reclocking Deserializer	3G	2.97G 1.485G 270M	424M 292M 259M
LMH0040	Serializer and Driver	HD	1.485G 270M	292M 259M
LMH0041	Reclocking Deserializer	HD	1.485G 270M	292M 259M
LMH0050	Serializer	HD	1.485G 270M	292M 259M
LMH0051	Deserializer	HD	1.485G 270M	292M 259M
LMH0070	Serializer and Driver	SD	270M	259M
LMH0071	Reclocking Deserializer	SD	270M	259M

Table 2 – National Semiconductor PHY families

TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Learn more and select the best Spartan FPGAs for your applications.
- Download XAPP514, "Audio/Video Connectivity Solutions for the Broadcast Industry."

Topics Relevant to You Delivered in One Monthly Newsletter!

Thanks for your feedback! We have now created a new Xilinx newsletter that will provide you all of the information you need in one location. This newsletter will replace the multiple topic newsletters you have currently been receiving.

It's simple! Choose only the topics you want to receive and your customized newsletter will be delivered to your inbox every month.

Register now for your chance to win free software, design kits, or one of 50 gift cards.

<http://newsletter.xilinx.com>



GET PUBLISHED



WOULD YOU LIKE TO WRITE FOR XCELL PUBLICATIONS? IT'S EASIER THAN YOU THINK.

Interested in adding “published author” to your resume and achieving a greater level of credibility and recognition in your peer community? Consider submitting an article for global publication in the highly respected, award-winning *Xcell Journal*.

The *Xcell* team regularly guides new and experienced authors from idea development to published article with an editorial process that includes planning, copy editing, graphics development, and page layout. Our author guidelines and article template provide ample direction to keep you on track and focused on how best to present your chosen topic, be it a new product, research breakthrough, or inventive solution to a common design challenge.

Our design staff can even help you turn artistic concepts into effective graphics or redraw graphics that need a professional polish.

We ensure the highest standards of technical accuracy by communicating with you throughout the editorial process and allowing you to review your article to make any necessary tweaks.

Submit final draft articles for publication in our Web-based *Xcell Online* or our digital and print *Xcell Journal*. We will assign an editor and a graphics artist to work with you to make your papers clear, professional, and effective.

For more information on this exciting and highly rewarding opportunity, please contact:

Forrest Couch
Publisher, Xcell Publications
xcell@xilinx.com



www.xilinx.com/xcell/



Featured White Papers

From high-volume consumer to automotive to global security, Xilinx white papers help you stay on top of technology applications.



“Solutions Without Compromise: How to Match Consumer Timelines with Automotive and Industrial Requirements and Reduce Risks of Obsolescence Using FPGAs and Intellectual Property”

www.logicbricks.com/pdf/Xilinx%20Electronic%20Displays%20Conference%202007.pdf

FPGAs are enabling a new automotive world, where compromises in (and between) flexibility, scalability, time to market, and system cost are no longer required. In a number of different automotive applications, these value propositions are key.

Multimedia and infotainment systems are always adapting to different types of displays. Graphic controllers have short lifetimes that do not fit automotive and industrial embedded requirements. Why not think about a solution in programmable logic that can drive any display, instead

of using different graphic controllers and changing them for different displays?

“Automotive FPGAs Offer Innovative Solutions to Automotive Displays Challenges”

www.techonline.com/showArticle.jhtml?articleID=193103236&queryText=Automotive+FPGAs+Offer+Innovative+Solutions+to+Automotive+Displays+Challenges

There is a continuing trend in the automotive industry to improve the comfort, safety, convenience, productivity, and entertainment experience for the driver. This trend is spurring increased use of a myriad of display technologies including TFT, LCD, and OLED.

Along with the benefits of these new technologies are unique challenges that traverse a wide spectrum of issues, including high-performance, low-EMI signaling;

complex timing control; high-performance image processing; and integration with a wide spectrum of automotive buses – all influenced by continuously evolving standards and degrees of adoption across the automotive industry. Although these challenges are significant, low-cost automotive FPGAs are taking a lead by providing comprehensive solutions and driving the use of a wide range of display technologies.

This white paper explores the various uses of automotive display technologies and associated challenges, highlighting the benefits and cost-effectiveness of FPGAs in addressing these challenges within the context of a rear-seat entertainment system.

“Security Solutions Using Spartan-3 Generation FPGAs”

www.xilinx.com/bvdocs/whitepapers/wp266.pdf

Whether boarding a plane, closing the front door, or beginning your next-generation circuit design, security has become a significant issue. In our homes, we try to build in the right amount of security to protect ourselves against theft. Security is rapidly becoming a necessity in the electronics industry as well.

It is important to understand why security issues have escalated to the forefront in electronics design. One reason is the alarming amount of counterfeited goods. These goods threaten the U.S. economy and have a significant effect worldwide in the consumer market, according to the Anti-Counterfeiting Coalition.

This white paper identifies the top design security threats, explores basic levels of security, and describes how new, low-cost Spartan™-3A, Spartan-3AN, and Spartan-3A DSP FPGAs from Xilinx can help protect your products and profits. ●●

Automotive ECU Development Kit

Xilinx offers a complete development platform ideal for designing intelligent automotive subsystems.

Xilinx has teamed up with Si-Gate GmbH (www.si-gate.com), a leading provider of automotive electronics solutions, to develop an automotive electronic control unit (ECU) development kit for in-vehicle networking systems, infotainment, driver assistance, and driver information systems.

What's Included

- XA1600E development board equipped with a Spartan™-3E XC3S1600EFG484 FPGA, with 1.6 million system gates and interfaces for expandability
- EDK and ISE™ software evaluation tools
- Resource CD with reference designs for easy evaluation
- Power supply with universal adaptor
- Programming cable
- Custom serial cable
- QuickStart guide

Key Features

- Xilinx device: XC3S1600EFG484 FPGA
- Complete out-of-the box development system
- On-board hardware interfaces such as CAN 2.0C, Ethernet 10/100, USB 2.0, SPI, and SCI
- Support for pre-verified Xilinx® and partner IP for high- and low-speed CAN, FlexRay, and MOST network interfaces (the latter requires a daughtercard)



Price: \$1,495

Part Number: HW-XA3S1600E-UNI-G

Targeted Applications

- Automotive subsystems
- In-vehicle networking systems
- Infotainment
- Driver assistance and information systems

TAKE THE NEXT STEP

(Digital Edition: www.xcellpublications.com/subscribe/)

- Learn more about the Automotive ECU Development Kit.
- Buy online from Avnet, Nu Horizons, or your local distributor.
- E-mail designkits@xilinx.com for more information.



Spartan-3AN FPGA Starter Kit

With its out-of-the-box guarantee, you'll be testing and designing FPGA solutions in minutes with the new Spartan-3AN Starter Kit.

The general-purpose non-volatile Spartan™-3 Generation FPGA platform evaluation board includes DDR2 memory, 10/100 Ethernet PHY, A/D circuitry, VGA and RS232 connectors, and an expansion mechanism. Robust subsystems allow you to:

- Multi-boot pre-programmed reference designs in a single session
- Prototype countless applications with support for 26 single-ended and differential I/O standards
- Connect to DDR2 memory for rapid data transfer using a proven memory interface controller
- Shorten development time with pre-verified design files and schematics
- Safeguard your designs with exclusive Spartan-3AN FPGA features

What's Included

- Evaluation board
- Xilinx® ISE™ software, WebPACK™ software, and ISE Foundation software evaluation
- 100-240 V, 50/60 Hz power supply with universal plug adaptors
- Quick-start guide
- Programming cable
- Product collateral

Key Features

- Xilinx Devices: Spartan-3A FPGA (XC3S700A-FG484) and Platform Flash PROM (XCF04S-VOG20C)
- Clocks: 50-MHz crystal oscillator on-board, open slot for optional user-installed clock



Price: \$225

Part Number: HW-SPAR3AN-SK-UNI-G-PROMO

- Memory: 4-Mb Platform Flash PROM, 32 M x 16 DDR2 SDRAM, 32-Mb parallel flash, two 16-Mb SPI flash devices
- Analog Interface Devices: Four-channel D/A converter, two-channel A/D converter, signal amplifier
- Connectors and Interfaces: Ethernet 10/100 PHY, JTAG USB download port, two 9-pin RS-232 serial port, PS/2-style mouse/keyboard port, 15-pin VGA connector capable of 4,096 colors, one FX2 100-pin and two 6-pin expansion connectors, 20 user I/O available on stan-

dard header pins, stereo mini-jack for PWM audio, rotary/push-button function switch, eight individual LED outputs, four slider switches, four push-button switches

- Display: 16-character, two-line LCD

Targeted Applications

- Markets: Consumer, telecom/datacom, servers, storage
- Applications: General prototyping

TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Learn more about the Spartan-3AN FPGA Starter Kit.
- Buy online from Avnet, Nu Horizons, or your local distributor.
- E-mail designkits@xilinx.com for more information.

CoolRunner-II CPLD Starter Kit

The ideal CPLD low-power design platform.

This design kit is an ideal platform to evaluate and implement your CPLD designs quickly. A unique evaluation board with eight peripheral connectors allows you to select from a variety of peripheral expansion modules (sold separately) to expand the capabilities of your design.

What's Included

- X-Board with XC2C256-TQ144
- Resource CD, including reference manual, software, application notes, web seminars, and data sheets
- USB cable
- ISE™ software WebPACK™ CD
- PMod-PS2 (PS2 module to connect a keyboard or other PS2 device to the X-Board)
- PMod-SSD (seven-segment display module)
- PMod-SWITCH (switch module for interfacing to four slide switches)

Key Features

- XC2C256 CoolRunner™-II CPLD in TQ144 package
- USB interface
- Eight 6-pin connectors for standard peripheral modules
- One 26-pin and one 16-pin connector for daughterboards
- Built-in temperature and power monitor
- Free reference designs



Price: \$49.95

Part Number: HW-CRII-SK-G

Targeted Applications

- Low-power design
- Logic consolidation
- Memory control and interfacing
- Display control
- Microprocessor interface
- State machines
- ADC interfacing
- Stepper motor control
- RFID telemetry
- Bluetooth or Wi-Fi interface
- Power or heat monitor
- MP3 player
- UART
- Smart Card reader

TAKE THE NEXT STEP

(Digital Edition: www.xcellpublications.com/subscribe/)

- Learn more about the CoolRunner-II CPLD Starter Kit.
- Buy online from Avnet, Nu Horizons, or your local distributor.
- Buy the Peripheral Module Bundle (HW-CRII-PM-ACC-G). Extend your CPLD functionality by adding peripheral expansion modules to your system, such as ADC, motor control, serial flash, and more for \$99.
- View the CoolRunner-II Starter Kit Demo.
- E-mail designkits@xilinx.com for more information.

Protect Your Brand with Spartan-3 Generation FPGAs

Build a low-cost security solution for high-volume applications.

In today's world, security is a huge concern for our global society. Whether boarding a plane, closing the front door, or beginning your next-generation circuit design, security has become a significant issue.

In our homes, we try to build in the right amount of security to protect ourselves against theft. Security is rapidly becoming a necessity in the electronics industry as well. It is important to understand why security issues have escalated to the forefront in the electronics design field.

One reason is the alarming amount of counterfeited goods that are the result of theft. These goods threaten the economy and have a significant effect worldwide in the consumer markets, according to the Anti-Counterfeiting Coalition. The World Customs Organization estimates that counterfeiting accounts for 5% to 7% of global merchandise trade, equivalent to lost sales of as much as \$512 billion in 2004. This threat grows by more than 12% per year, tarnishing the reputation and long-term credibility of genuine brands.

High-volume solutions, served by Xilinx® Spartan™-3 Generation FPGAs and CoolRunner™-II CPLDs, include technologies such as configuration data protection, hidden bitstream, active defense, and Device DNA design-level security for low-cost hardware and software IP protection. Taken together, these features enable designers to implement a low-cost, highly robust security solution to deter reverse-engineering, cloning, and overbuilding.

How Device DNA Works

The Device DNA security mechanism is similar to an ATM transaction, in which the active value generated from the card-plus-pin combination is compared to a number stored in a bank computer that authorizes or rejects the transaction.

Similarly, the unique 57-bit Device DNA number is used with a customer-defined security algorithm to generate an

active value. The active value is compared to a pre-stored check value to determine whether design functionality can proceed.

Our Device DNA demo highlights how easy it is to implement a low-cost security solution for high-volume applications with Device DNA technology. It requires either the Spartan-3A FPGA Starter Kit or the Spartan-3AN FPGA Starter Kit. Documentation is provided. 




Figure 1 – Similarities between Device DNA security and ATM transactions

TAKE THE NEXT STEP (Digital Edition: www.xcellpublications.com/subscribe/)

- Download the Device DNA demo.
- Read the white papers, "Security Solutions Using Spartan-3 Generation FPGAs" and "Advanced Security Schemes for Spartan-3A/3AN/3A DSP FPGAs."
- Read the brochure, "Design Security for High-Volume Applications."

Titanium Dedicated Engineering

This Xilinx program provides skilled experts when you need them.

by Jannis McReynolds
Senior Manager, Services Marketing
Xilinx, Inc.
jannis.mcreynolds@xilinx.com

In today's environment, the pressure to produce more complex products faster and at a lower cost is only increasing. Making these challenges even more complex are such issues as compressing design cycles, implementing new technologies, overcoming design obstacles, and augmenting your knowledge base.

In the automotive electronics market, technical challenges are ever-present in the areas of in-vehicle infotainment, comfort, and convenience, as well as in gateway and driver-assistance systems.

The consumer marketplace continues to have time-to-market pressures, along with the continued advancement of digital consumer applications for home networking and information appliances.

Xilinx® Titanium Dedicated Engineering helps you meet your specific design challenges by producing the exact results you need when you need them. Our technical team provides faster issue resolution, expert design advice, and risk mitigation. Your result is faster time to market, lower production costs, and a competitive edge.

What is Titanium Dedicated Engineering?

Titanium Dedicated Engineering sends you a dedicated application engineer on a contract basis. The engineer can work remotely or at your site. Whether you are starting a new project or are nearing your deadlines and need that extra expertise to meet design performance, Titanium Dedicated Engineering is the answer.

Titanium engineers bring expertise and experience from multiple, focused engagements in areas such as embedded support,



general debugging, general tool usage, performance optimization, timing closure, triple modular redundancy support, and PlanAhead™ software floorplanning.

Key Features

- Skilled at providing solutions and experts in Xilinx knowledge
- Trusted advisors proven to perform under time pressures
- Easily integrates into team environments
- Available for quick deployment
- Can work on- or off-site

Benefits

- Quick and direct access to a dedicated engineer
- Leading-edge Xilinx expert design process knowledge
- Leverage in-factory support of product experts

What Our Customers Have to Say


“The Xilinx team responded with extremely short notice and solved a major bug for us in an incredibly short time. The engineer was extremely knowledgeable. He not only solved the bug, but also advised our team how to get far better performance on the design. He paid for himself by allowing a transcontinental team to immediately become productive after days of inactivity tracking this bug.”

– Large multinational customer

Conclusion

Regardless of where you are in your design cycle, turn to Titanium Dedicated Engineering. To learn more, visit www.xilinx.com/titanium.

To purchase Titanium Dedicated Engineering in North America, contact your local Xilinx sales representative, or e-mail fpga@xilinx.com. In Europe, contact your local Xilinx sales representative, or e-mail eurotitanium@xilinx.com.



NO TOYS

Real Tools for ASIC Prototyping

Prove your design with high speed FPGA hardware logic emulation, that plugs directly into your PC. This 12+ million gate, 8-lane PCIe board, is ready to go: fixed PCIe core, extensive on board memory and eight independent clock networks will get your design running — *fast!* The DN9000K10PCIe-8T features six Virtex-5, LX330 chips with 100% of FPGA resources user-available.

There are no unreliable cables or stacks of overheating boards. The price, with the biggest and baddest FPGAs, is less than \$70,000. Why pay more for a “some assembly required” toy? Visit www.dinigroup.com to see our complete selection of real high-speed prototyping solutions.

The
DiNI
Group



X-ray vision for your designs

Agilent 16998A High-performance FPGA Analyzer

Agilent's New 16998A FPGA Analyzer



Agilent portable and modular logic analyzers

- Increased visibility with FPGA dynamic probe
- Customized protocol analysis with Agilent's exclusive packet viewer software
- Accurate and reliable probing
- Pricing starts at \$9,970



Now you can see inside your FPGA designs in a way that will save weeks of development time.

Agilent's 16998A high-performance FPGA analyzer allows you to access different groups of signals inside your FPGA for debug—without requiring design changes. You'll increase visibility into internal FPGA activity by gaining access up to 128 internal signals with each debug pin.

Deep memory is essential to capture enough system activity to get to the root cause of complex problems. With the FPGA analyzer you receive an upgrade to 4 M acquisition memory and the FPGA dynamic probe software for FREE.

Whether your signals are single-ended or differential, you can debug your FPGAs and surrounding system faster and more effectively with Agilent's 16998A high-performance FPGA analyzer.

www.agilent.com/find/fpga

© Agilent Technologies, Inc. 2007
Windows is a U.S. registered trademark of Microsoft Corporation



Agilent Technologies

Enabling Efficient Packet Processing



by Ivo Bolsens
CTO
Xilinx, Inc.
ivo.bolsens@xilinx.com

I recently attended an intriguing Clean Slate workshop at Stanford University. The topic focused on the idea of reinventing the Internet infrastructure in order to solve some of its shortcomings and enable new applications and services (<http://cleanslate.stanford.edu>).

Driven by an increase in rich media applications such as video, service convergence, and broadband access deployment, global IP traffic is showing explosive growth. Furthermore, IP data network traffic is highly dynamic, with constantly changing demands and substantially different characteristics when compared to telephony voice.

Today, IP backbones are designed to have core routers directly connected through static circuits. This assumption of static links means that links must be grossly over-designed to deal with peak traffic uncertainties and link failures. As IP traffic grows, such a network design implies bigger links and bigger routers – neither of which can be scaled cost-effectively. Moreover, the reliability and security of the networking infrastructure to support a growing demand of dependable applications becomes a serious challenge.

Network Virtualization

The Clean Slate researchers agreed that “network virtualization” would become an essential technology to provide future value-added services and robustness.

Network virtualization is a capability that allows you to create multiple yet very different concurrent networks on top of the same physical infrastructure of diverse resources, and would imply a functionality

A clean-slate approach to reinventing the Internet.



that extends beyond current packet routing and forwarding.

Today, traditional networking features nodes with fixed architectures. Flexibility in networks comes from packet switching itself – guided by look-up tables – plus the relatively slow evolution of protocol development and node software updates. To support desirable concepts like virtualization and technological convergence, it will be necessary to be much more flexible in terms of node architectures and protocol definition. Therefore, data processing (whether packetized or something else) may be carried out using very different styles, depending on widely differing applications. The protocols providing the rules for such processing will also vary widely and require much more dynamic definition.

From a Xilinx® perspective, FPGA technology is experiencing a parallel shift that complements the future direction of networking. Many people see FPGAs as ASIC substitutes; once programmed, the devices have the characteristics of fixed hardware. Today, however, FPGAs are providing a basis for delivering flexible, “soft” processing architectures that can be matched to particular problem instances.

For example, you can essentially build your own network processing unit

(NPU) every time rather than forcing problems onto fixed architectural solutions such as an ASSP or NPU (be it a connection of multi-threaded engines or deeply pipelined parallel data paths). Moreover, the soft architecture can evolve over time in response to changes in demand and usage.

Today’s FPGA architectures enable the creation of soft packet processing architectures that, depending on application functionality, can be configured as highly pipelined parallel data paths or networks of multi-threaded micro-engines.

These soft processing architectures allow the memory architecture, the connection of the different packet processing functions, and the packet operational units to be tailored to the application at hand. This results in the implementation of packet processing functions in FPGAs that are outperforming NPUs and ASSPs in throughput, power consumption, and cost by at least one order of magnitude.

An important component is the availability of programming tools that harness the highly concurrent capabilities of the FPGA and move away from traditional hardware design paradigms toward new, flexible protocol implementation paradigms.

ACCELERATE VERIFICATION



With The Confirma™ ASIC/ASSP Verification Platform From Synplicity®

Reduce verification time from months to days

The Confirma platform is a tightly-integrated, easy to use, and comprehensive at-speed ASIC/ASSP verification solution that dramatically accelerates functional verification of ASICs, ASSPs, and SoC designs.

The three major components of the Confirma Platform are already considered best-in-class tools: the Certify® ASIC verification software implements an existing ASIC design in multiple FPGAs; the HAPS™ High-performance ASIC Prototyping System™ is the execution engine; and the Identify® Pro software, with its revolutionary TotalRecall™ technology, provides



HAPS-54 Virtex-5 Board Offers Prototypers 8 Million ASIC Gates

full visibility and a seamless interface into industry-standard software simulators for easy debugging. When used together, the tools provide the highest-performance ASIC and ASSP verification platform available today.

Visit <http://www.synplicity.com/confirma> to learn more about reducing verification time.



Synplicity
600 West California Avenue
Sunnyvale, CA 94086
Phone: (408) 215-6000
Email: info@synplicity.com

At the heart of innovation



Enabling the next generation of digital consumer applications

Xilinx programmable logic devices (PLDs) are the technology of choice for today's consumer products. Already pervasive in flat-panel displays, smart handsets, automotive infotainment, home networking and much more, Xilinx PLDs are also enabling tomorrow's consumer technology innovations.

Low-Power, Low-Cost, High-Security Solutions

Known for their design flexibility, low-cost, enhanced security and time-to-market advantages, Xilinx solutions are optimized for a variety of end markets and applications. Combine this with custom IP, starter kits and award-winning design tools, and you can see why Xilinx is the world's leading PLD provider.

Visit our website to learn more about Xilinx solutions—at the heart of innovation in the consumer marketplace.

Visit us at
CES

Meeting Place
#35768
South Hall

 **XILINX**[®]
www.xilinx.com