

Introducing the New Virtex-4 FPGA Family

The latest FPGAs from Xilinx set new records in capacity, capability, performance, power efficiency, and value.



by Eric Goetting
Vice President &
General Manager,
Advanced Products Division
Xilinx, Inc.
erich.goetting@xilinx.com

Welcome to the Xilinx® Virtex-4™ edition of the *Xcell Journal*. We've created this special issue to show you the new Virtex-4 FPGA family, and how its innovations enable the creation of next-generation systems that do more than ever thought possible only a few years ago.

In this article, I'll take you behind the scenes for a guided tour of some of the new technologies, as well as a bit of the inspiration and rationale behind them.

With more than 100 innovations, the Virtex-4 family represents a new milestone in the evolution of FPGA technology. After conducting extensive interviews with leading design engineers worldwide, we knew that they wanted the following things in an advanced next-generation FPGA family:

- Higher performance
- Higher logic density
- Lower power
- Lower cost
- More advanced capabilities

At the heart of the Virtex-4 FPGA is our next-generation 90 nm triple-oxide 10-layer copper CMOS process technology.

It's relatively easy to deliver on one or two of these items – our challenge was to deliver all of them at the same time. We did this through a combination of innovative process and circuit design, process development, the ASMBL architectural approach, and the use of advanced embedded functions.

Development work on the Virtex-4 family (code-named “Whitney” after the highest mountain in the continental United States) began more than two years ago. It represents the creativity and dedication of hundreds of engineers, spanning integrated circuit design and layout, software and IP development, process development, testing and characterization, systems and applications engineering, technical documentation, and product marketing.

One of the most remarkable developments embodied in the new Virtex-4 FPGA family is the ASMBL architecture, which represents a fundamentally new way of constructing the FPGA floor plan and its interconnect to the package. First of all, ASMBL enables I/O pins, clock pins, and power and ground pins to be located anywhere on the silicon chip, not just along the periphery as with previous approaches. This in turn allows power and ground pins to be brought directly into the center of the silicon die, thereby significantly reducing on-chip IR drops that can occur with the largest FPGAs running at the highest frequencies.

Clock input pins are also located in the center of the die, which reduces clock latency. This is because clock networks need to have equal delay to all endpoints (that is, minimum skew), and thus the clock must emanate from the center. In periphery-connected clock input pins, the signal first traverses from the edge of the die to the center, and is then distributed to all regions. The Virtex-4 ASMBL design eliminates this traversal completely, and thus directly reduces the clock network propagation delay.

In addition to its electrical advantages, ASMBL provides another significant benefit in that it allows a more flexible – and thus more precise – allocation of on-chip resources.

That in turn has enabled us to offer Virtex-4 devices in three unique platforms, each with a different mix of on-chip resources:

- The LX platform, optimized for logic applications
- The SX platform, optimized for high-end DSP applications
- The FX platform, optimized for embedded processing and high-speed serial applications

A Look Inside the Virtex-4 FPGA

At the heart of the Virtex-4 FPGA is our next-generation 90 nm triple-oxide 10-layer copper CMOS process technology. While that's quite a lot of adjectives, every one of them is incredibly important. The first, 90 nm, refers to the “drawn” gate length of the smallest transistors. As transistors get smaller, they get faster, use less dynamic power, and enable higher complexity at lower price points. Chip designers think in terms of “transistor budgets,” which are now in the billion transistor range.

Triple-Oxide 90 nm CMOS Technology

Triple-oxide technology refers to the number of transistor oxide thicknesses available in the process. More oxide thicknesses allow more tuning of performance and power in the device circuitry, and enable Virtex-4 devices to deliver industry-leading performance while dramatically lowering power consumption.

One of our key inputs from many engineers was that performance and power were very important constraints in their systems designs, and that they needed both high performance and low power. With a dual-oxide 90 nm process, we would have had to choose performance or power. This wasn't good enough. By employing a triple-oxide 90 nm process, we achieved high performance and low power.

The 10-layer copper refers to the number of metal interconnect layers and their

material, which is copper rather than aluminum (the traditional material). More layers provide more routing in less space and shorter connection distances. Copper reduces resistance compared to aluminum, and thus speeds signal interconnect and reduces on-chip power-distribution IR drop. As clock rates go up and voltages go down, these considerations have become increasingly important, and have driven the industry-wide shift to copper interconnect.

The Virtex-4 logic fabric was completely re-engineered to fully take advantage of the 90 nm triple-oxide CMOS process, resulting in the highest performance fabric ever, with system clock rates in excess of 500 MHz (at three LUT levels). At the same time, static power was cut in half compared to 130 nm Virtex-II Pro™ devices, as was dynamic power.

Thus, while some industry pundits were proclaiming that the future of deep submicron CMOS devices was getting hotter and hotter, with chip temperatures destined to reach that of rocket nozzles and the surface of the sun, the Virtex-4 design's creative approach has turned that conventional wisdom on its head, resulting in overall power reductions of 50% compared to our previous 130 nm generation. In many applications, such as DSP functions, power levels are reduced even more – as much as 90%. No wonder design engineers say that Virtex-4 FPGAs are cool – they literally are.

High-Performance Clocking

Clocks were rated as one of the most important and critical FPGA resources in our surveys of design engineers. Quantity, quality, connectivity, frequency, duty cycle, jitter, and skew all made a big difference.

To take clocking to the next level in Virtex-4 devices, all global clock resources were made fully differential, thereby reducing skew, jitter, and duty-cycle distortion. This marks the first implementation of differential clocking in a programmable logic device. Not only that, but the number of global clocks was increased to 32, for every

device, and internal connectivity options enhanced to allow any region to use any 8 clocks simultaneously.

500 MHz Synchronous Memories and FIFOs

On-chip synchronous block RAM was enhanced to run at 500 MHz. Built-in support for first-in first-out (FIFO) memories was included directly in the block RAM unit, enabling the same 500 MHz operation for FIFOs (approximately a 2X speedup over fabric-based FIFOs), while eliminating the need for any additional logic cells or complex FIFO designs.

If you're designing systems requiring ECC (error checking and correcting) memory, Virtex-4 devices have built-in ECC support, with single-bit correct and double-bit detect. ECC is common in infrastructure equipment in networking, telecom, storage, servers, instrumentation, and aerospace applications, and provides the highest levels of data integrity. Like the integrated FIFO support, the integrated ECC eliminates the cost and delay of fabric-based solutions.

Speaking of on-chip memory, Virtex-4 devices continue to offer SelectRAM™ memory, whereby each LUT is transformed into a 16 x 1 RAM, ideally suited for building high-speed register files and local buffers.

At the other end of the spectrum, interfaces to external memory devices such as DDR, DDR2, QDR-II, and RLDRAM-II are dramatically enhanced through our new ChipSync™ technology, which offers memory interface speeds at rates limited only by the speed of the external memory devices.

The new Virtex-4 ML461 Advanced Memory Development System contains fully functional and hardware-proven reference designs for all of today's most popular memory technologies. If you plan to use external memory, I highly recommend that you check this out.

DSP Performance of 256 TeraMAC/s

In the DSP domain, we incorporated some of the world's fastest multiply accumulate (MAC) technology. The XtremeDSP™ slice can perform an 18 x 18 signed multiply and 48-bit accumulate every 2 ns.

The Virtex-4 LX, FX, and SX platforms include the breakthrough XtremeDSP technology. With the new SX platform we did something completely new – we dramatically increased the ratio of DSP units to logic cells. Given the highly integrated nature of XtremeDSP slices, they need only small amounts of logic fabric to implement most common DSP functions, and thus increasing the ratio provides a significant increase in DSP compute power per unit silicon area. In fact, SX devices provide a 10X performance increase per unit cost over previous solutions.

Power is dramatically reduced as well, with more than a 10X reduction for multiply/add functions from previous FPGA solutions. The Virtex-4 SX55 contains 512 XtremeDSP slices, providing an aggregate DSP compute performance of 256 TeraMAC/s, making it one of the most powerful DSP devices ever manufactured.

The state-of-the-art XtremeDSP slice employs new “silicon algorithms” developed by a company called Arithmatica™. Many different architectures exist for implementing multiplication, and the Arithmatica architecture is truly a breakthrough. We are excited to see it available for the first time to FPGA users. For more information, visit Arithmatica's website at www.arithmatica.com.

The Evolution of Advanced I/O Technology

I/O continues to be a critical success factor for today's systems designers. During the last decade, we have seen four major changes in I/O. First was the shift away from 5V, the result of the need to scale voltages as we scaled the transistor. This in turn led to the plethora of I/O standards that we are all familiar with today: SSTL, HSTL, LVDS, and LVCMOS 1.5. The Virtex-4 SelectIO™ resource continues to lead the industry, supporting virtually every I/O standard in use today on every pin.

XCITE On-Chip Termination

The second major change was the transition from lumped loads to transmission line loads – again the direct result of Moore's Law. As transistors got faster and clock rates increased, I/O edge rates

increased as well. But because the propagation speed of signals is a constant, dictated by the speed of light, we entered the realm in which a signal on one end of a wire was no longer the same as the signal on the other end of the same wire. This is what transmission lines are all about, and their appearance during the last few years has driven a sea change in all aspects of signal interconnect and I/O design.

To make sure that these signal “waves” don't start “splashing” uncontrollably, transmission lines need to be driven, built, and received using proper signal integrity approaches, the most critical of which is termination. Traditionally implemented with discrete resistors on the PCB, termination layouts can become exceedingly difficult around high-density pinouts like those used in FPGAs. This often dictates more PCB layers and thus more system cost.

Virtex-4 FPGAs include our third-generation of XCITE™ integrated digitally controlled termination technology. Offering a precisely controlled source impedance at the output drive pin, it is designed to enable the driving of transmission lines without external components, with maximum speed and signal integrity, and with straightforward PCB layout and layer stack-ups.

Likewise, on inputs, XCITE offers parallel termination for single-ended inputs and true differential termination for differential inputs. Termination occurs on the end of the transmission line at the die, not on the way there on the PCB, offering maximum signal integrity. Many customers report that the XCITE technology has saved them many PCB layers, increased PCB packing density, and saved them substantial dollars in their bill of materials.

Source-Synchronous Interfaces

The third major change was the shift from system-synchronous to source-synchronous interfaces. Traditional system-synchronous interfaces work by distributing a single clock to all transmitters and receivers in the system, and transmitting data between source and destination within a single clock cycle. This makes the data rate inversely proportional to the sum of clock-

to-out, transmission line delay, and input setup time.

Typically, system synchronous interfaces top out at speeds in the range of 100 MHz. To go faster, source-synchronous interfaces transmit a clock along with the data, and the receiver uses this clock to capture the data. Using this technique, along with double-data-rate transmissions, enables parallel I/O data rates in excess of 1 Gbps.

The challenge of source-synchronous interfaces is that each interface generates a new clock domain at the receiver. On top of this, to operate at high speeds, the precise alignment of clock and data at the receiver is paramount. To address this new world of source-synchronous interfaces, Virtex-4 devices include the breakthrough ChipSync technology. ChipSync units lie between the SelectIO technology and the core FPGA fabric, are available on every I/O pin on the device, and serve to transmit and receive high-speed source-synchronous data and clocks, achieving speeds of 1 Gbps per pin pair.

On the receiver, precise digital delay lines work internally to align data signals to each other, and then to align these to the received clock. The captured data is synchronized and transferred to the selected FPGA core clock domain.

To operate at maximum data rates, the transmit and receive units include parallel-to-serial and serial-to-parallel conversion units, respectively. Using ChipSync technology is virtually automatic for most designs, as it is utilized automatically in the various Xilinx IP cores and reference designs.

Networking interfaces such as SPI-4.2 and HyperTransport™, and memory interfaces such as DDR, DDR2 SDRAM, and QDR II SRAM, all employ the Virtex-4 ChipSync technology. And if you're designing your own source-synchronous interface, the ChipSync wizard gives you complete control and an easy-to-use GUI that lets you dial in exactly what you want to build.

Multi-Gigabit Serial Interfaces

The fourth major change in I/O has been the rapid adoption of high-speed serial interfaces. For years, serial interfaces were limited to long-distance communications,

such as those used in fiber-optic links in the SONET/SDH world and the Ethernet links like 100BASE-T.

A key breakthrough occurred in the late 1990s, in which high-speed serial transceivers (which traditionally had been designed using complex process technology such as GaAs [Gallium-Arsenide]) were for the first time created using advanced design techniques using standard CMOS. Once implemented in CMOS, these transceivers had lower cost and much lower power, and could even be integrated into complex CMOS chips.

Virtually overnight, gigabit serial technology changed from a rare, expensive, and power-hungry technology to a common, low-cost, and very power-efficient technology. This has been the economic and technical impetus behind the industry's "Serial Tsunami," in which interface after interface has shifted from parallel to gigabit serial links. Two common examples are visible in today's computer architectures, with the shift from parallel PCI to 2.5 Gbps serial PCI-Express™, and the shift from the parallel ATA drive interface to the Serial ATA interface.

There are more than a dozen multi-gigabit serial interfaces in widespread use today, with more being introduced every year. The Virtex-4 FX family provides our third-generation RocketIO™ multi-gigabit serial transceiver technology. Spanning speeds from 622 Mbps to more than 10 Gbps, each Virtex-4 RocketIO transceiver is programmable and can implement a myriad of speeds and serial standards. Link-layer IP is available for such standards as PCI Express, Serial-ATA, FibreChannel, Gigabit Ethernet, and Aurora, to name a few.

In addition, Virtex-4 FX devices each include multiple embedded tri-mode (or 10/100/1000) Ethernet MACs, making implementation of compliant Ethernet devices simpler and faster than ever.

Application-Specific Embedded Processing

Virtex-4 embedded processing solutions include full support for both MicroBlaze™ 32-bit soft CPUs on all devices, and embedded PowerPC™ 32-bit RISC CPUs on all Virtex-4 FX devices.

The number of CPUs in one device is limited only by your imagination, and of course by the available logic cells. The high-performance PowerPC CPU runs at clock rates up to 450 MHz and delivers up to 702 DMIPS each. The first PowerPC processor available by any manufacturer on 90 nm, the PowerPC processor is incredibly power-efficient, using only 0.29 mW/DMIPS. This makes it among the lowest power micro-processors available from any manufacturer worldwide.

New Auxiliary Processor Unit (APU) technology connects the CPU to the FPGA fabric, enabling implementation of acceleration hardware for virtually any application. Once only the domain of high-budget ASIC and ASSP design teams, the Virtex-4 FPGA's architectural ability to combine application-specific hardware acceleration with high-performance RISC CPUs shatters traditional barriers of cost, time-to-market, and risk.

During the next few years, I expect to see more and more instances of application-specific acceleration, as it truly offers the ability to deliver very high performance at low cost and low power. A recent research program completed within Xilinx Research Labs, led by Dr. Kees Vissers, demonstrated a 20-fold speedup for an encryption/decryption (IPSEC) application over the base PowerPC processor. Using only 135 mW, it outperforms a 3.2 GHz Pentium™-4, while at the same time reducing power by 99%. That, in my opinion, is what state-of-the-art embedded processing is all about.

Conclusion

I hope that you've enjoyed reading a bit about the Virtex-4 Platform FPGA and the factors that drove its design. From the breakthrough ASMBL architecture and the triple-oxide 90 nm CMOS process technology, to the world's most capable embedded processing and multi-gigabit serial solutions, Virtex-4 devices offer an unparalleled set of enabling technologies for your next-generation systems designs. I look forward to seeing the creativity of the world's designers in tomorrow's products. ●●●