Platform FPGA- The Future of Logic Design

Xilinx and its partners are building the high-performance technology platform on which the designs of the future will emerge.

By: Wim Roelandts, CEO, Xilinx

The revolution in logic design continues, bringing dramatic performance improvements and new capabilities that help you create the systems of the future, and get them to market faster than ever before. FPGAs were once just interconnect routing and logic gates; then we added dedicated hard cores for memory, clock management, and I/O. Now, FPGAs are becoming the platform on which a combination of complex hard cores and flexible soft cores combine with an abundance of programmable logic gates to give you the best possible performance, along with the ease-of-use and time-to-market advantages for which FPGAs are well known. Plus, we can bring you these advantages at a lower cost than ever before.

Xilinx has entered a number of strategic partnerships and has acquired key technologies for creating the new programmable logic platform. Here is an overview of our recent activities.

The IBM Partnership

Our recent partnership with IBM brings us two immediate and dramatic benefits: the power of the PowerPC™ hard core, and the advanced CMOS manufacturing capability of IBM's state-of-the-art facilities. IBM gets intellectual property (IP) from Xilinx to help reduce defect densities and improve manufacturing productivity. This partnership has far-reaching implications, and gives both companies a significant competitive advantage.

The PowerPC Core

IBM's PowerPC module and CoreConnect™ bus will soon be integrated into Virtex-II FPGAs. With this powerful combination, you can achieve performance that was never possible before, and you can quickly develop unique system-level applications with greater ease. We found that the PowerPC was the most-used processor in high-end designs; our communication and computing customers use the PowerPC because it has good performance, and it has a lot of peripherals and other functions that make it easy to use.

Processors like the PowerPC are often used as logic engines for low speed, very complex logic; they allow you to write detailed programs that perform intricate condition checking and control functions. However, because a processor basically executes one instruction at a time, it's slower than actual gates which can operate in parallel.

Now, in one Platform FPGA, you will have the best of both worlds; you have the dedicated PowerPC processor for complex control applications, and you have programmable logic gates for very-high-speed data paths. The big advantage of having this all on one chip is that you can very quickly move data from the PowerPC processor to on-chip peripherals or custom logic, which may be hard cores, soft cores, or unique designs created with programmable gates. This will give you much higher performance than you get using separate chips, which must pass signals through their slower I/O interfaces.

We are implementing the PowerPC processor and other dedicated functions (such as memory, clock management, multipliers, and I/O interfaces) as hard cores, to give you the best possible performance. We will compliment these hard cores with over 50 soft core peripheral functions. By keeping most of the peripherals as soft cores, you can choose only those functions that you need, and create custom designs with ease.

Advanced CMOS Manufacturing Capability

IBM is one of the most advanced CMOS semiconductor companies in the world, with device manufacturing technologies that are typically a year ahead of most other companies. Our partnership with IBM gives Xilinx access to this manufacturing technology, and a tremendous competitive advantage. To be competitive in our marketplace we have to push manufacturing technology to the limits. By using the most advanced manufacturing processes, we can reduce the size and cost of transistors, which enables us to continue building bigger and bigger FPGAs and reduce the costs of existing devices.

For IBM, FPGAs are the ideal "process drivers" to test and refine their advanced manufacturing processes. Because FPGAs have very regular structures and they allow us to address almost every square micron of space on a chip, it makes them ideal for troubleshooting problem process areas. So, Xilinx gets advanced manufacturing technology and IBM gets devices that help them drive their manufacturing process to maturity. Thus, we can achieve better yields, faster, and that means lower costs.

Xilinx has been the leader in developing programmable logic technology, and we have expanded the market dramatically–today, the PLD business is growing 40% faster than the regular IC business. Our current technologies, bigger densities, higher speeds, and lower costs, are expanding the market much faster than in the past; with IBM, we are pushing it even further.

Gigabit Serial I/O Capability

Many new systems today are requiring much faster data transfer between systems, boards, and devices, due primarily to the ever increasing demand for faster networks. Very high speed (gigabit per second) serial I/O capability promises to solve this difficult problem.

Traditionally, data has been shared by using parallel busses such as PCI (Peripheral Component Interconnect). However, there are inherent limitations with shared busses. To increase the speed of a shared bus, you can either increase the speed of each wire (which is very difficult to do because there are many of them), or you can increase the number of wires (which takes more and more I/O pins). For example, PCI was once just 32 bits wide; now you also have 64-bit PCI–and that's not enough. The problem with this approach to increasing bandwidth is that at some point you reach a level of decreasing return; the extra pins and the need for shared bus protocols limits the performance and makes it prohibitively expensive.

The best solution we have for this bandwidth bottleneck is to use point-to-point connections, over a single pair of wires, operating at very high speeds. Currently, with this technology, you can achieve a data rate of two to three gigabits per second. The big advantage of this method is that you use less wires and less power, and the total amount of data you can move is in fact higher than with a typical parallel bus.

To create a gigabit serial I/O channel, a hard core is needed; you cannot achieve these speeds with soft cores in an FPGA. The hard core does several functions; it receives and transmits the data, and it also recovers the clock (because you can recover the clock from the data, a single pair of wires is all you need for data transfer). The hard core must also serialize and de-serialize the data. By de-serializing the data to a 16- or 32-bit internal bus, the data speed is then reduced by 16 or 32, which an FPGA can easily handle.

With gigabit serial I/O, all the de-serialization is done within the chip. When the work is done, you can serialize the data again and send it out to other devices, over a single pair of pins. This is a very efficient and low cost way to transfer data.

We recently made several announcements regarding our commitment to gigabit serial I/O.

The Conexant™ Partnership

Xilinx recently entered into a strategic development and licensing partnership with Conexant Systems, to integrate their SkyRail™ 3.125 Gbps serial transceiver technology into our next generation Virtex-II FPGAs. This hard core is the fastest core available in CMOS today, and will be available in the second half of 2001 in a select offering of several different Virtex-II devices. In our Virtex-II architecture you will get more than 20 different I/O standards, plus several of these gigabit serial I/O channels.

The high-speed SkyRail transceiver is compliant with industry standards such as Gigabit Ethernet and Fibre Channel in addition to the emerging 10-Gigabit Ethernet (IEEE 802.3ae) standard. By integrating quad transceivers, which are used to create 10-gigabit attachment unit (XAUI) interfaces, a single FPGA can interface to both 10-Gigabit Ethernet and OC-192c. The high-speed transceiver is also compliant with the 2.5 Gbps InfiniBand™ architecture standard being created by the InfiniBand Trade Association.

The RocketChips Acquisition

High-speed serial I/O capability is so important, we decided not to stop at the 3.125 Gbps speed offered by the Conexant core–we are developing the technology further. That's why we recently acquired a company called RocketChips, which is very active in creating high speed serial I/O cores. RocketChips already has a product that is very similar to the Conexant core, and they plan to develop even higher speed cores operating at 5 to 10 Gbps.

RocketChips' gigabit and multi-gigabit serial CMOS transceiver technologies provide solutions for a wide range of serial system architectures in networking, telecommunications, and enterprise storage markets. Their products include serial backplane transceivers (Single and Quad 3.125 Gbps transceivers), telecom transceivers (SONET OC-48 and OC-192), enterprise storage transceivers (Fibre Channel, Ethernet), and networking transceivers (Gigabit Ethernet, 10 Gbps Ethernet, and InfiniBand).

PMC-Sierra Partnership

Xilinx recently announced the availability of POS-PHY™ Level 3 Link Layer and Physical Layer cores. These cores provide solutions for the emerging Packet Over SONET (POS-PHY) applications, and both cores are compatible with the POS-PHY Level 3 interface specified by the SATURN® Development Group. With these cores, broadband system designers can rapidly develop highly functional, scalable, and standards-based equipment to increase the speed of networks up to 2.5 Gbps, and support the exploding growth of IP traffic over SONET/SDH backbones.

Xilinx has also been active in the Optical Internetworking Forum (OIF) and the ATM Forum to drive POS-PHY Level 4 acceptance. And, we are the only FPGA company to demonstrate over 800 Mbps operation, confirming that we can provide the full speed capability to support the 10 Gbps OC-192 draft standard at the OIF (OIF2000.088.2).

Serial Protocol Standards

To use these high speed serial I/O channels effectively, you need well defined protocols and networking standards. Xilinx actively supports all of the emerging standards, including:

- Lightning Data Transport™ (LDT) - A chip-to-chip interconnect that provides much greater bandwidth per I/O channel. It can achieve a bandwidth of 6.4 Gbps per eight-wire link width, and can support up to 32 links.
- InfiniBand™ - This newly designed interconnect system utilizes a 2.5 Gbps wire speed connection with one, four, or twelve wire link widths. Promoted by an association comprising industry leaders such as, Compaq, Dell, HP, IBM, Intel, Microsoft, and Sun Microsystems, InfiniBand intends to deliver a channel based, switched fabric technology.
- XAUI - A quad transceiver utilizing 3.125 Gbps serial links to create a 10 gigabit attachment unit interface (XAUI). Multiple XAUI interfaces can be implemented to allow a single chip to interface to both 10 Gigabit Ethernet and OC-192c.
- Fibre Channel - A high-bandwidth serial standard offering 1.06 Gbps data rates scalable to 2.12 or 4.24 Gbps. It is capable of carrying multiple existing interface command sets, including Internet Protocol (IP), SCSI, IPI, HIPPI-FP, and audio/video.
- Gigabit Ethernet + 10 Gbit Ethernet - This includes devices compliant with the IEEE 802.3 alliance.
- ATM (OC-12, OC-48, OC-192) - This includes support for OC-12 (622 Mbps), OC-48 (2.4 Gbps), and OC-192 (10 Gbps).
- RapidIO™ - A next-generation switched-fabric interconnect architecture for embedded systems that is optimized for both high bandwidth and low latency. Initial implementations are expected to exceed 1.0 Gbps throughput based on clock rates of 250 MHz and higher.

These standards all use the same physical interface, so you can use our hard I/O cores for all of them. Then, we implement the level-2 protocols in programmable logic (soft cores), so you can quickly create designs using any of these standards. This gives you a lot of flexibility and it helps you interface directly to on-site networks.

What Does It Mean?

An FPGA is no longer just gates and routing. Over the years we have added more and more hard cores, such as memory, clock management, and arithmetic functions. Now we are driving the technology a major step further by adding hard CPU cores and high speed serial I/O cores. Combine these dramatic technology advances with our high performance development tools, our unique Internet Reconfigurable Logic capability, our extensive training and support services, our state-of-the-art manufacturing capabilities, and our ongoing partnerships with other industry leaders, and you get a logic design solution that can breathe life into your new designs.

The role of the FPGA is changing; it is becoming a platform on which the combination of soft cores, programmable logic, and hard cores gives you the best possible design solution–the speed of a custom ASIC and the time-to-market advantages of a flexible FPGA.

THE ROLE OF THE FPGA IS CHANGING; IT IS BECOMING A PLATFORM ON WHICH THE COMBINATION OF SOFT CORES, PROGRAMMABLE LOGIC, AND HARD CORES GIVES YOU THE BEST POSSIBLE DESIGN SOLUTION–THE SPEED OF A CUSTOM ASIC AND THE TIME-TO-MARKET ADVANTAGES OF A FLEXIBLE FPGA.