The New Era of Programmable Systems

The next breakthrough in processing and system design methodology comes from the merger of the most advanced technologies from Xilinx and IBM.

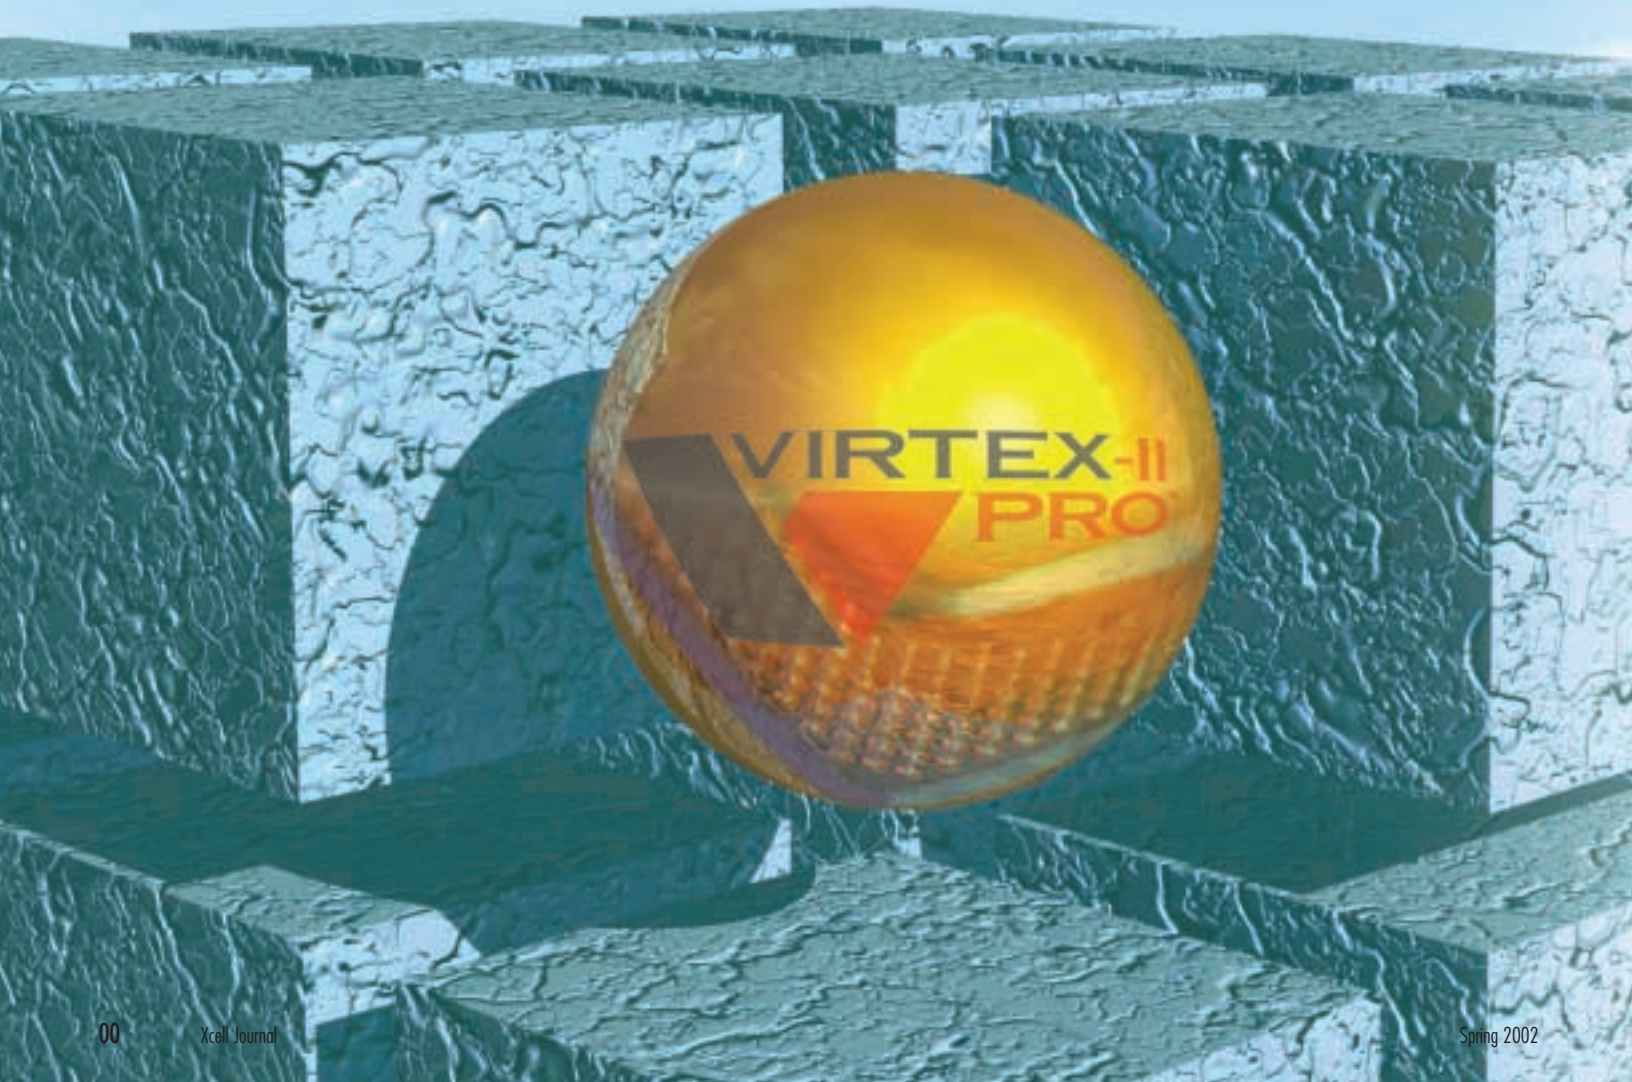

by Babak Hedayati
Sr. Director of Product Solutions and Partnerships
Xilinx, Inc.
babak.hedayati@xilinx.com

Over the course of the semiconductor revolution, with the help of Moore's Law, FPGAs have grown to densities of 10 million gates. They have consumed key system-level functions such as block memory, clock management, digitally controlled impedance matching, embedded multipliers, 844 Mbps LVDS I/Os, and many other functions. As the densities increased so did the insatiable hunger for soft IP cores for simple functions, complex DSP algorithms, networking protocols, interfaces, and so on.

The advantage gained by increased FPGA capabilities – sometimes unbelievable to some – has been adopted by thousands of design teams looking to improve their time to market by targeting segments of their system to Virtex™-II FPGAs whenever possible. Yet few could initially imagine the possibilities of a "programmable system" when Xilinx talked of immersing 300 MHz IBM® PowerPC™ 405 processors into the Virtex-II FPGA fabric and embedding high-speed multi-gigabit serial I/Os around it.

Many immediately saw the value of integration for cost reduction, increased performance, and reliability. Others saw the potential for its incredible flexibility and scalablity for implementing specialized and high-speed interfaces. The idea of extreme hardware parallel processing, and multiple processing on the same device, enticed many system designers to consider the great possibilities of such a solution.

Many engineers expected the next breakthrough in processing and system design methodology to come solely from the masters of the microprocessor world and leading ASIC vendors – but the breakthrough has come from the merger of the most advanced technologies from Xilinx and IBM.

The need for high-speed communication and increased bandwidth has driven the rapid evolution of technology throughout multiple industries. Design challenges associated with integration, high speed interfacing, higher performance processing, and new design methodologies must be solved, and the rapid rate of change in technology demands hardware programmability – this time at the system level.

Digital Convergence

The convergence of voice, video, data processing, and packet processing both on the infrastructure equipment and consumer products is putting immense pressure on corporations and their engineers. They must now incorporate computing, networking, wireless, and video imaging technologies that previously existed stand-alone in their respective markets.

This digital convergence has been an inevitable reality since the early days of the electronics industry, where specialized equipment and devices demonstrated their true potential as soon as they were connected with other devices. First came the telegraph, then the telephone, computers, video, the Internet, storage, wireless, and the infrastructure behind it all. Now the world is incredibly crowded with new consumer technologies that seem to pop up on a daily basis incorporating new features from the digital revolution. They target individual niche areas, compete, and often become extinct, just as suddenly as they were brought to market. This extinction is often caused by rapidly changing standards and requirements, or by competitive products putting tremendous pressure on corporations and their engineering teams.

Moving Toward Total Cost Management

Rapid product extinction makes executives question why their companies are spending so much capital on multi-million dollar ASIC NREs (Non-Recurring Engineering charges) and design automation contracts – because today's economic and technological conditions often require design changes midway in the development process. Success in today's marketplace is accomplished by getting to market first, not by designing for high volumes and the lowest unit cost – spending immense amounts in advance on creating custom ASICS, without a guaranteed future, is a recipe for failure.

Many companies are faced with huge inventories of ASICs and ASSPs that cannot be

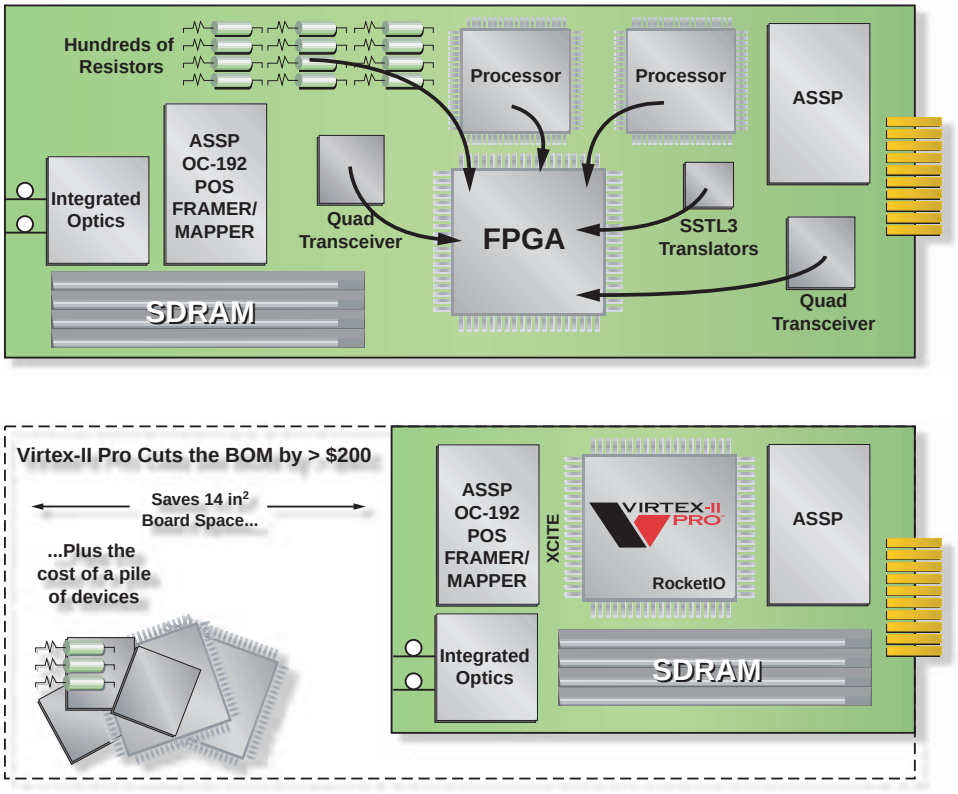

Figure 1 - Leading edge price/performance through integration and reduction of board size

re-targeted to multiple products to reduce the risk. Though ASICs are often less expensive for high volume designs, they incur much higher overall costs, higher risk, longer development cycles, and less time in market. It's time for corporations to evaluate the total cost of their development strategy as a whole and avoid falling into the pitfalls of separating capital investments, development costs, production costs, obsolescence costs, and inventory management.

The coming generation of computers and telecommunications equipment is very different from prior evolutions of information technology because it is dramatically reversing the age-old wisdom of creating specific devices for each application. Current systems are integrating computers, cell phones, game systems, cameras, appliances, automobiles, offices, and homes. Eventually, we will likely have only a few types of super systems remaining that synthesize and extend the capabilities of all current systems.

One of the key trends to reduce cost, increase performance, and increase the reliability of systems has been through integration. However when designers integrated their systems into custom ASICs, they increase inventory risks and require large initial up front investments. Hence, companies find it difficult to stop midway to change their designs. In addition, smaller companies or start ups find that they must commit the majority of their funding just to develop their platform – and sometimes they have a difficult time getting a large ASIC supplier to entertain their development.

A fully integrated system-level solution such as the Virtex-II Pro™ family solves all of these problems. Offering multiple gigabit serial I/Os, the fastest FPGA solution in the world, up to four Power PC 405s, XCITE™ controlled impedance technology, and other system-level features, you get a smaller board size, lower overall costs, and faster time to market, as illustrated in Figure 1.

The Ultimate Connectivity Solution

For a long time, the original PCI bus was the industry standard. To increase the bandwidth, many designers began to use bridges, continuing with the parallel, shared bus strategy. Then the standard moved to 64-bit 66 MHz versions, and later to PCI-X running at 133 MHz.

The problems with continuing this strategy are obvious. Wider busses require more pins and higher cost, and moving to higher frequencies causes signal integrity issues. Plus, the shared bus created more overhead and less bandwidth predictability. Although PCI will be used for years to come, today's performance-hungry applications have already moved toward packet switched LVDS-based parallel methodologies such as POS PHY Level 4, Flexbus 4, RapidIO, Hyper Transport, and others. This requires smarter protocols and point-to-point interfacing between devices and boards. The move was welcomed because it increased bandwidth. However, it often resulted in increased clock skew and signal integrity issues.

In the last two years companies like Xilinx, in partnership with Conexant (and the later acquisitions of companies such as RocketChips), have discovered how to implement mutli-gigabit serial I/Os in CMOS technology, making it a cost effective method of delivering point-to-point serial switched interconnections. This means much higher performance without any side effects. Other companies in the silicon industry, such as ASSP companies and standards committees, are now quickly adopting this strategy to reduce cost, increase reliability, and increase bandwidth, as shown in Figure 2 and Figure 3.

The Processing Revolution

The quest for higher performance processing is evident in many applications. Companies have traditionally turned to "farms" of expensive high-performance processors to achieve the performance they need. In doing so, they have usually faced prohibitive costs along with the massive efforts of managing and partitioning their tasks throughout the processor farm.

Hardware oriented companies such as networking, telecom, and wireless infrastructure developers have taken the lead by implementing parallel processing in hardware. For example, using Xilinx XtremeDSP™ solu-

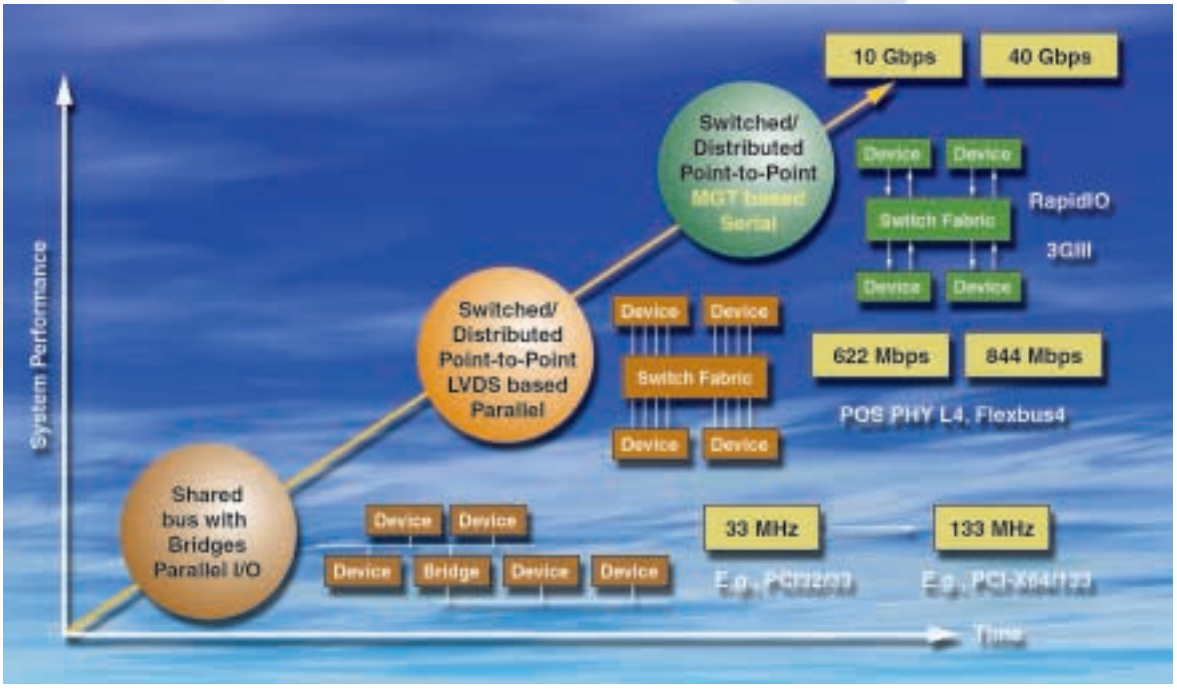

Figure 2 - The trend towards switched and distributed serial I/O

Figure 3 - Virtex-II Pro FPGAs support all connectivity standards

tions, with a single clock cycle they can process massive amounts of multiply accumulate functions (over 600 billion MACs per second). On the other hand, when companies have had to deal with multiple smaller tasks, they have traditionally turned to multiple processors, and optimized the code for processing each smaller task; a technique often used in network processors, as shown in Figure 4.

Now for the first time, through the Virtex-II Pro programmable system, designers get both high-performance processing and distributed processing through multiple PowerPC processors immersed in the FPGA fabric. With Virtex-II Pro FPGAs, the whole is much more than the sum of the parts.

Enabling a New System Development Paradigm

Design teams can now make system-level tradeoffs and optimization throughout the design cycle. Traditionally, architecture-level teams have had to make such tradeoffs early in system definition phase – deciding what goes into hardware what gets implemented as software code. This restriction has been the cause of many delayed products and products that have been unsuccessful, because of the inability to make adjustments for new features or performance optimizations during the design phase. Electronic design automation (EDA) companies have partially addressed the problem by developing system-level tools (such as behavioral partitioning tools and so on) to make the system-level tradeoffs easier.

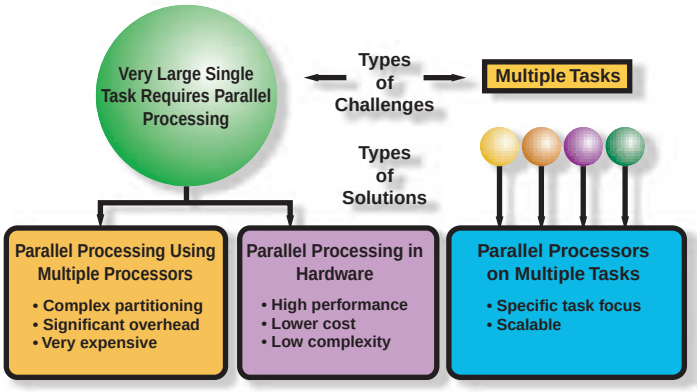

Figure 4 - Virtex-II Pro FPGAs provide higher performance processing

Now, with Xilinx technology, design teams can make tradeoffs and optimizations throughout the system design; creating a more integrated system with higher performance and faster time to market. They can even make changes to their hardware and software in the field, after the product is in the customers' hands, to fix bugs or implement new features. New business models can now be developed for programmable system design.

The Virtex-II Pro solution provides a standard programmable system platform, fully supported by embedded development tool vendors such as Wind River Systems; EDA companies such as Cadence, Mentor Graphics, and Synopsys; and system-level tools companies such as Celoxica and The Mathworks – all industry leaders and strategic partners to Xilinx. These partnerships enable a new development paradigm for programmable systems design. This overall solution of devices, software, cores, and partnerships means that companies like yours can now rest a little easier.

Welcome to the Programmable World

Xilinx and its partners are taking the next step in the evolution of programmable logic by creating a new event – Programmable World 2002. Here, you will hear industry leaders and visionaries discussing the latest Platform FPGA solutions, including detailed technical training for the PowerPC, Wind River tools, and others. From implementation techniques for multi-gigabit serial I/Os to digital signal processing, you will hear experts from more than 50 companies discussing this revolution in logic design.

Programmable World 2002 will be held simultaneously in multiple locations throughout North America and Europe. You won't need to travel, but if you do, it's all free with breakfast and lunch included.

April 17th you can see it all. To get the full details, or to register for the technical sessions, go to: *www.xilinx.com/pw2002*. Be sure to reserve your seat now; attendance is limited.