Xilinx Virtex®-II Product Backgrounder

January 2001

The Virtex-II Platform FPGA solution is the result of the largest silicon and software R&D effort in the history of programmable logic, the goal of which was to revolutionize complex single-chip design. Improvement in engineering productivity, silicon efficiency, and system flexibility are all premier aspects of this new platform. To ensure success in the market, complex systems must achieve rapid time-to-market while incorporating state-of-the-art memory and interconnect fabric, general purpose embedded hard-macro, application-specific soft IP library and global design management capabilities for clocking, system noise, and design protection. Virtex-II devices provide all of these features, pre-engineered to keep design cycles short.

The Virtex-II family provides the Xilinx IP-Immersion™ technology, which incorporates abundant on-chip memory options and advanced routing resources for supporting complex designs that use intellectual property (IP). Key to the Virtex-II platform is hard-macro building blocks and a rapidly growing library of soft IP blocks. For the first time in the programmable logic industry, innovative Virtex-II features enable designers to create true complete systems on a single platform, with all the speed, performance and advanced features required for today's cutting edge applications.

The Virtex-II family of Platform FPGAs is a complete programmable solution**,** with densities up to ten million system gates. The inherent flexibility of Xilinx FPGA devices allows unlimited design changes throughout the development and production phases of the system—cutting design cycles to months versus years. Today's complex designs in fast moving markets require improved productivity, reduced design risk, and higher system flexibility. The Virtex-II platform will accelerate the industry-wide shift away from inflexible custom ASICs toward FPGAs in applications such as optical networking systems, gigabit routers, wireless cellular base-stations, modem arrays, and professional video broadcast systems.

The Virtex-II architecture was developed to seamlessly integrate a wide variety of new hard-macro building blocks now in development, including the PowerPC™ processor, CoreConnect™ high-speed internal bus, and channel-bonded 3.125 Gbps serial interfaces. These new hard-macros will dramatically increase the signal processing and data transmission capabilities available in a single-chip solution.

Virtex-II "at a glance" Features

- IP-Immersion™ architecture enables Platform FPGA
- High logic capacity, up to 10 million system gates
- High memory capacity, up to 4.5 Mbits
- Fourth generation segmented routing technology optimized for fast, wide busses
- Embedded 250MHz multiplier hard macro, with up to 192 multipliers in single device
- Flexible SelectI/O-Ultra™ technology supporting 840 Mbps I/Os, with up to 1108 user I/O pins (554 differential I/O pairs)

- 18-kbit True Dual-Port™ block RAM/ROM
- 18-bit by 18-bit multipliers
- "Glitch-free" clock multiplexing bitstream
- Hard- and soft IP support for RapidI/O™, PCI-X, LDT, POS PHY L3/L4, Flexbus 4, and OIF SPI-4 interfaces and Lightning Data Transport™ (LDT)
- Xilinx Controlled Impedance Technology (XCITE™) technology, providing built-in impedance matching on all single-ended I/Os for signal integrity
- Digital clock management hard-macro supporting de-skew and precise frequency/phase manipulation
- Bitstream encryption for design protection
- Future support for integrated processors and high-speed serial interfaces

Virtex-II Target Applications

- *Optical networking and high-end DSP systems*
- *Complex high-speed data transmission and manipulation*
- *IP-based systems using hard-macro and soft IP blocks*

The Virtex-II solution was developed to enable rapid development of the two most technically challenging system applications: data communications and digital signal processing systems. These systems are characterized by the need for high logic integration, fast and complex routing of wide busses and extensive pipeline and FIFO memory requirements. The Virtex-II family incorporates the highest system gate capacity in the industry with up to ten million system gates. A new Active Interconnect™ architecture is optimized for predictable routing delays, an advanced memory array architecture with up to 4.5 Mbits of on-chip memory is built in to every device. On-chip support for high-speed I/O standards with up to 1108 user I/O pins is included. Applications incorporating DSP functionality, such as echo cancellation, forward error-correction, and image compression/decompression all benefit from the abundance of embedded high-speed 18 bit x 18 bit multiplier blocks within Virtex-II Platform FPGAs. The unique features of the revolutionary Virtex-II architecture make it ideal for optical networking products, storage area networks (SANs), Voice-over-Internet-Protocol (VoIP) systems, video broadcast systems, medical imaging systems, wireless base-stations, and Internet infrastructure products.

IP-Immersion Capability

- *The Xilinx Active Interconnect technology for predictable routing performance*
- *Hard-macro building blocks*
- *The Xilinx Smart-IP™ software technology*
- *Bitstream encryption for design protection*

The Virtex-II solution uses an innovative memory-based logic and interconnect fabric for implementation of high-speed signal routing. The interconnect methodology uses a fourth generation segmented routing structure, called Active Interconnect technology, to ensure predictable routing performance during design iterations. Each routing switch connection is buffered at the output, which provides a constant routing delay independent of the signal fanout. This reduces development time compared with older FPGA architectures and ASICs, by minimizing the number of routing delay changes required during each design iteration.

Hard-macro building blocks within the Virtex-II FPGA include popular, high-performance sub-functions such as memory, multiplier modules, and I/O macros. These custom implementations require minimal die size when compared to soft implementations, and they provide significant performance and logic efficiency benefits.

The hardware architecture supports Smart-IP technology, which allows subsystem IP blocks to have constant timing within the IP block without requiring fixed physical placement or a fixed target device. IP blocks delivered using Smart-IP technology are available from Xilinx as well as third party vendors, and they allow re-targeting into any comparable Virtex-II device with the same timing performance. This allows designs to be re-used across the entire Virtex-II family.

An important feature of the IP-Immersion architecture is the protection of internal IP. The bitstream encryption feature allows the bitstream to be encoded with a secret key, thereby protecting the internal design against theft when the bitstream is intercepted. The key, stored inside the device, prevents the disassembly of internal user logic.

Interconnect Engine for Fast, Wide Busses in Networking Applications

- *High-capacity RAM for wide FIFOs*
- *Active Interconnect technology for abundant, time-invariant routing*
- *Soft IP support for wide LVDS- and ULVDS-based I/O interface*
- *16 global clock nets supporting multiple high-frequency clock domains*

The Virtex-II architecture incorporates a number of sought-after features that specifically support wide data widths in complex networking and transmission systems. Modern complex systems operate at multiple clock domains, with large IP-based subsystems operating independently. Large, wide FIFOs and buffer memories are needed for handling fast and wide inter-subsystem data transfer. These wide busses are required both internally for intra-chip communications and externally for switched fabric communications. For example, wide 32+ bit data busses may drive multiple ULVDS high-speed interface standards for data transfer across a backplane, for point-to-point communications, or to implement high-speed multi-cast bus standards.

These requirements challenge and exceed the capabilities of current programmable logic devices, which lack the gate capacity, memory and routing resources, performance, and architecture flexibility to fully support these designs. The Virtex-II solution is the first embodiment of the Platform FPGA specifically targeted to improve “ease of speed” in the development and production of these complex systems.

Fully Digital Clock Management

- *16 low-skew pre-engineered clock domains*
- *Glitch-free clock multiplexing capability*
- *Up to 12 digital clock managers supporting phase-lock, frequency synthesis, and dynamic phase shift*

Each Virtex-II FPGA contains 16 global clock buffer-multiplexers, allowing up to 16 low-skew clock domains within each device. Additional clock signals can also be routed using the abundant low-delay routing resources, enabling a high degree of design flexibility. Each global clock buffer-multiplexer provides a 2 input/1 output, glitch-free," clock multiplexing capability, which allows advanced control switching in shared resource systems. The clocking circuitry automatically eliminates clock spikes that are less than the output clock periods, ensuring a smooth transition in swapping clock signals.

The Virtex-II family also supports up to 12 Digital Clock Managers (DCMs) within a single device, each of which is capable of phase-locking onto a reference clock, synthesizing precise frequency using user-specified multiplication factor and divisor, and static or dynamic phase shift control. For deep submicron process technologies, the digital implementation of these systems provides a higher level of stability compared with analog implementations and also eliminates the need for specialized analog power supply requirements.

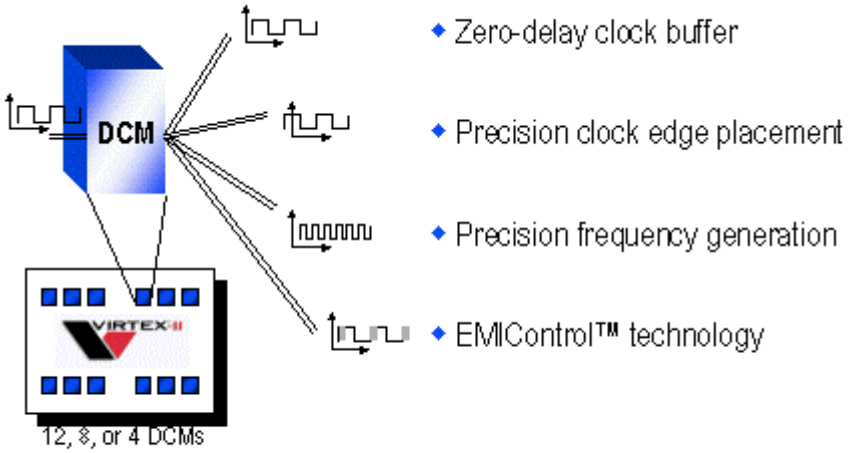

Figure 1: Digital Clock Manager Capabilities

There are three key features supported by the DCM including zero-delay clock buffering, precise clock frequency generation, and precise clock edge placement. The DCM uses a digital delay line with internal control feedback to match the phase and period of a reference clock, thereby producing an effective zero-skew version of the reference clock. This enables the DCM to generate 50/50 duty outputs ideal for double data rate application.

The DCM creates a synthesized frequency based on a multiplication and divisor factors up to 4096. This provides immense flexibility for users to generate the precise clock frequency customized for their applications. This enables simplified system clock generation for applications such as E1/T1 translation, video clock generation, and crossbar data switching.

The DCM allows the output clock signal to be phase-shifted a specified amount compared to the reference clock. The input reference clock period is divided internally into 256 time units, with

the output clock able to align to any of these times. The user can control the phase shift with a static value of –255 to +255 in the software, or the phase shift may be dynamically controlled using a synchronous increase/decrement control signal generated within the system. This unique capability provides precise clock edge adjustments to be made during the end of the development phase, or it may be incorporated into a synchronization sequence for optimizing high-speed data transfer.

Memory-Based Data-Path Fabric

- *Versatile lookup tables (LUTs) supporting 16-bit distributed memory*
- *Block RAM/ROM functions supporting multiple FIFOs*

The heart of the IP-Immersion architecture is the memory-based data path fabric that supports synchronous and asynchronous data transfer among different IP-based sub-systems. The proprietary architecture provides an array of configurable logic blocks (CLBs) that encompass both distributed memory elements and routing resources. Each CLB comprises eight 16-bit units of memory, configurable as a 16-bit RAM, a 16-bit variable tap shift-register, or a 4-input look-up table (LUT), with internal hardware logic to fully interconnect these resources into larger memory sections or wide logic functions. This allows up to 128-bit RAM, 128-bit shift-register, and 32 input/1 output multiplexers to be implemented without requiring routing resources or extra routing delays.

The IP-Immersion architecture also provides for up to 192 Block RAM (BRAM) modules interspersed within each Virtex-II device, each with 18 Kbits (over 2 Kbytes) memory. Each BRAM supports high-speed, dual-port synchronous read/write operations, and allows asymmetric port width, three write modes, and bitstream programmable contents after power-up. The BRAM hard-macro block forms the basis for efficient implementation of asymmetric-port FIFOs, ROM-based lookup tables, finite-state machines, and CAMs.

Additional memory requirements, for higher capacity or specialized memories, form the last stage of the IP-Immersion memory hierarchy. External memories such as DDR SDRAM, QDR™ SRAM, and specialized CAM devices are supported via hard-macro blocks within the SelectI/O circuitry, including dedicated DDR registers and built-in support for HSTL, SSTL, and other high-speed I/O standards.

The full memory hierarchy of the IP-Immersion architecture, consisting of distributed memory, block memory, and high-speed memory interfaces, provide a powerful memory-based data-path fabric to support IP-based systems.

Powerful SystemIO™ Functionality

- *SelectIO-Ultra™ technology supporting physical interfaces for new high-speed communication standards*
- *Xilinx XCITE digitally controlled impedance technology*

The Virtex-II Platform FPGA showcases the industry's premier I/O interface technology, SystemIO technology, to fully address all aspects of system connectivity in high-performance

designs. System connectivity consists of the physical interface and the protocols required to offer higher bandwidth. Virtex-II SystemIO technology uses SelectI/O-Ultra block to provide the fastest and most flexible electrical interfaces. All Virtex-II I/O pins supports full SelectIO-Ultra functionality, to increase design flexibility. Each user I/O pin can support any of over 20 interface electrical standards. These standards form the soft IP building blocks for important new interface protocols such as PCI-X 133 MHz, RapidI/O, POS PHY Level 4 (16 data/clock at 832 Mbps), Flexbus 4 (16 data/clock at 832 Mbps), SPI-4 (16 data/clock at 832 Mbps), and LDT.

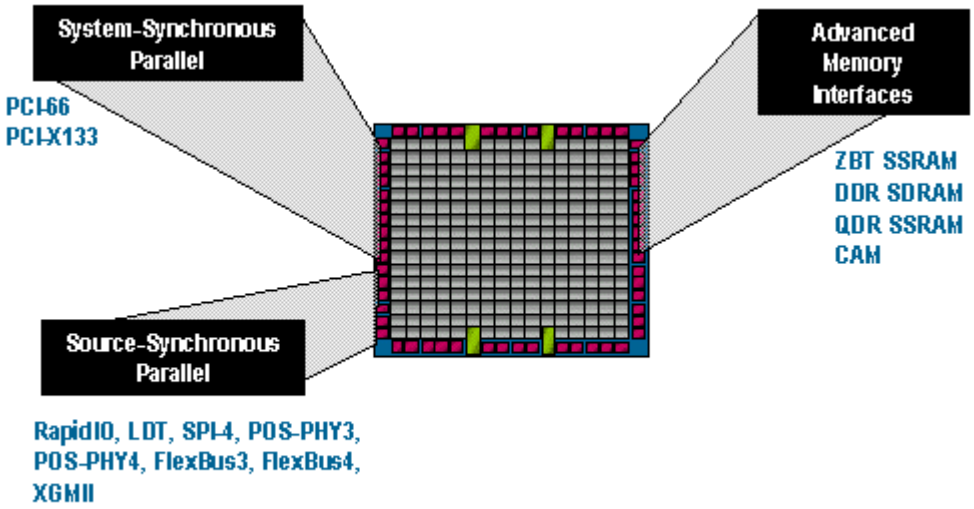

Figure 2: Virtex-II SystemIO™ Technology

The Virtex-II Platform FPGA also supports the Xilinx XCITE on all single-ended I/Os—the first such capability available in the semiconductor industry. The XCITE option is available on each user output, whereby the output impedance is matched to an external reference impedance dedicated to one of 8 I/O banks within each device. This capability eliminates the need for most external termination resistors and allows high-precision impedance matching required for high-speed data transfer.

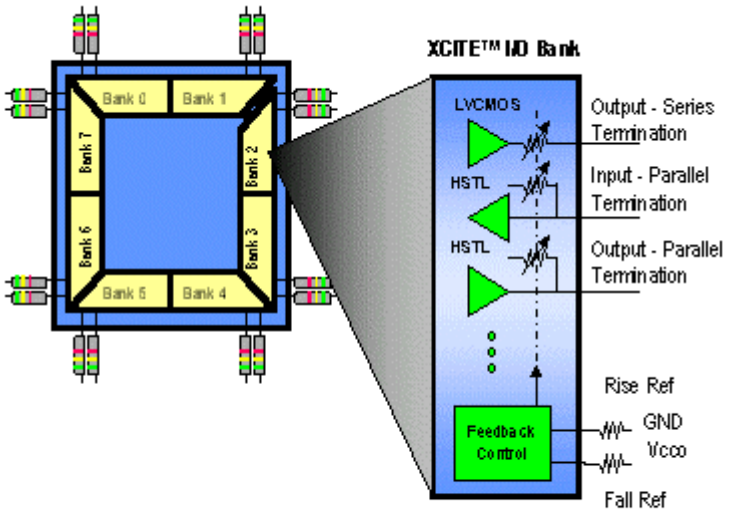

Figure 3: XCITE™ Digital Controlled Impedance

This provides a unique capability within Platform FPGAs that is unavailable in ASICs, whereby Virtex-II designs may be used in different electrical environments by matching local impedance requirements. In contrast, ASICs designed for one particular board impedance spec may not be able to work in a different board environment and cannot be tuned for adjustment.

Bitstream Encryption Strengthens Design Security

- *Triple DES*
- *Multiple key support*

The bitstream encryption hard-macro function provides a high level of security against the unauthorized copying of protected designs. The ability to encrypt bitstream is available for the first time in FPGAs to support the Platform FPGA methodology. During the design compilation process, a user selectable key is provided to the design software in order to generate an encrypted bitstream. During the manufacturing process, the key is loaded into the Virtex-II device via the serial JTAG port. Once loaded, the key will remain indefinitely within the Virtex-II device by supplying power to the special decryption power pins using a battery or other power source.

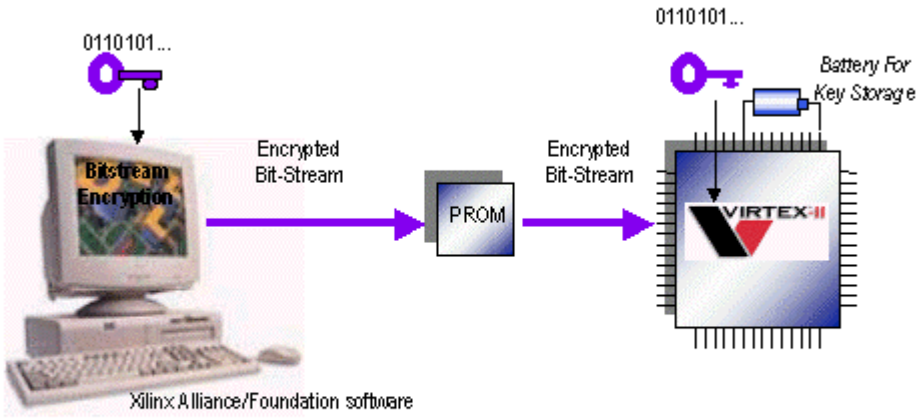

Figure 4: Bitstream Encryption for Design Security

With the key loaded, each time the design bitstream is loaded (such as during system power-up or reset), the Virtex-II FPGA is able to decrypt the design bitstream internally, and enable the device initialization. Even if the bitstream is read by a competitor, it is impossible to decode or disassemble the bitstream without the decryption key.

The encryption/decryption scheme utilizes the well-known DES format, using up to 168-bits of key length with triple-DES. This methodology is highly secure against many different types of attacks, and offers superior protection against design copying.

Each Virtex-II FPGA allows up to two sets of triple-DES keys, each of which may be used to encrypt a design bitstream that is accepted into the device. This offers a protected way for different companies to develop designs for a single protected FPGA without disclosing the same key to multiple parties.

The bitstream encryption feature also allows the Virtex-II Platform FPGAs to become a “virtual foundry” for semi-custom chip set development. It will allow third party providers of IP and design services to develop protected chip sets for specific customers, thereby improving overall design productivity and time-to-market, as well as simplifying licensing agreements. The design can also be modified and improved over time by using Internet Reconfigurable Logic (IRL™) technology or other means, by reprogramming the device with design improvements or feature enhancements.

For example, a echo cancellation chip-set may be developed for VoIP applications using Virtex-II Platform FPGAs by a third party IP developer. The design may be published with bill-of-materials, encoded bitstream files, and demonstration boards. Prospective end-users may analyze the cost and evaluate the design. Once finalized, the end-user can contract directly with the developer for additional services, customization requirements, and licensing terms.

The unique bitstream encryption feature of the Virtex-II solution enables full protection for complex designs, which is necessary in the increasingly competitive marketplace.

Platform FPGAs are the solution for today's system level design needs

- *System level density and speed*
- *Dynamic impedance matching to increase the design lifetime*
- *Bitstream encryption to protect whole-chip IP*
- *Field reprogrammability for maximum flexibility*

The Virtex-II Platform FPGA solution is targeted at implementing many high-performance, complex functions that previously required custom ASICs. For the first time, ASIC oriented designs with several million-gate requirements and up to 420 MHz system frequencies may be implemented within new Platform FPGAs. This brings to bear the inherent advantages of FPGAs to ASIC designs, including a shortened development cycle, increased engineering productivity, lower development costs, and fast design reprogrammability.

In addition, Virtex-II Platform FPGA solution provide additional capabilities that are not possible with ASICs, including XCITE technology and IRL capability. Using the XCITE technology, Virtex-II designs may be customized to different boards with different impedance requirements, and prolongs the product life cycle of designs. Unlike ASICs, the outputs of Virtex-II devices can be customized after the design is finalized to precisely match board impedance. The bitstream encryption feature of the Virtex-II devices, along with IRL technology, allow**s** simple and efficient field updates while eliminating design theft.

The new Virtex-II Platform FPGA solution will allow many leading-edge system companies to avoid the increasing challenges of designing and maintaining custom ASICs, while enjoying increased time-to-market competitive advantages, higher engineering productivity, more powerful and flexible design methodology, and reduced development costs.

—30—