

Xilinx Unveils New FPGA Architecture to Enable High-Performance, 10 Million System Gate Designs

New Virtex-II Architecture Delivers Twice the Performance of the Virtex Family

Press Backgrounder

Xilinx has unveiled the first details of the revolutionary Virtex™-II architecture, which has up to 10 times the system gate density and twice the performance of the original Virtex family. The new density and performance milestones are made possible by combining a new logic, memory, and routing fabric with leading-edge process technology.

New architectural innovations include more powerful, higher capacity configurable logic blocks (CLB), new Active Interconnect™ technology, higher bandwidth SelectI/O™ input/output (I/O) interfaces, a larger and more flexible allocation of distributed and block RAM, and enhanced arithmetic and embedded 18-bit multiplier capability. Together, these architectural features provide significant improvements over the original Virtex family, including:

- Ten times the system gate density, with up to 10 million system gates
- Twice the system frequency performance, with internal system clocks up to 200 MHz
- Quadruple the I/O performance, with 800+ Mbps bandwidth per I/O signal
- Quadruple block RAM capacity, with True Dual-Port™ 18-kbit block RAM
- Quadruple the distributed RAM capacity, with 128-bit single-port or 64-bit dual-port distributed RAM

The Virtex-II architecture builds on the highly successful Virtex platform, which is used in place of custom application specific integrated circuits (ASICs) in many leading-edge applications, including wireless base stations, telecom central switches, complex networking systems, video servers, and medical imaging systems. Virtex series devices were the first to enable system manufacturers to incorporate more than one million system gates into a single programmable device. This capability has fundamentally changed the dynamics of the system industry by allowing development cycles and incremental upgrades for complex systems to become dramatically shorter and faster, respectively.

The Virtex-II architecture further extends this capability to 10 million system gates and a 200 MHz system frequency. This will enable further system performance increases by incorporating large IP cores, digital signal processing (DSP) subsystems, and complex switching capabilities into a single programmable device. Businesses and consumers alike will benefit more quickly from improved wireless communications, faster Internet access, higher quality video services, and faster information services.

Virtex-II Architecture Addresses Challenges of 10 Million Gate Designs

- *Routability*
- *Software and IP efficiency*
- *System performance*

The Virtex-II architecture is optimized to address the complex challenges of supporting 10 million system gates, at very high data bandwidths, and for real-world applications. The challenges include routability, software and IP efficiency, and system performance.

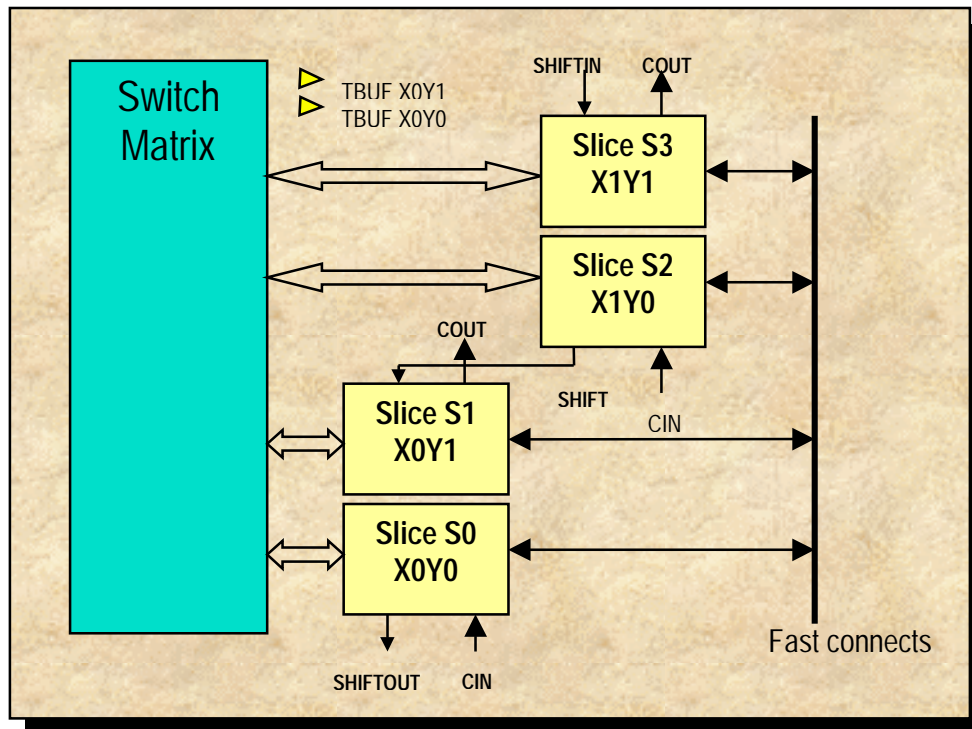
In ultra-high density field-programmable gate arrays (FPGAs), the routing structure is especially important for determining device utilization, routing delays within the design, and system development time. The productivity achieved using high-level design language (HDL) and intellectual property (IP) methodologies becomes extremely important for designs with very high densities, as these techniques become the only practical ways to design very complex applications. System performance becomes a greater challenge at higher densities, as subsystem capabilities such as arithmetic and cache functions become integrated into a single device. While the original Virtex family has proven successful at the one million-system gates level, the Virtex-II architecture is developed to handle designs with 10 times the density and twice the internal system performance.

Configurable Logic Block (CLB)

- *Twice the logic versus the Virtex series*
- *Enhanced wide multiplexer capability*
- *Deeper distributed RAM*
- *Arbitrary length shift registers*

The Virtex-II architecture is comprised of CLB tiles that include logic cells and routing resources. Virtex-II CLBs are significantly enhanced to address a higher level of system integration. Compared to the Virtex series CLB, the Virtex-II CLB contains twice the logic capacity, with four logic slices (each containing two logic cells with individual lookup tables and dedicated registers). In addition, specialized logic structures are added to allow fast complex functions to be implemented using adjacent CLBs. These capabilities increase the amount of logic that can be combined to form fast complex functions, such as a 16:1 multiplexer within one CLB, or a 32:1 multiplexer within two adjacent CLBs. This type of capability is important for networking and mass storage applications that route large numbers of data busses at high switching frequencies.

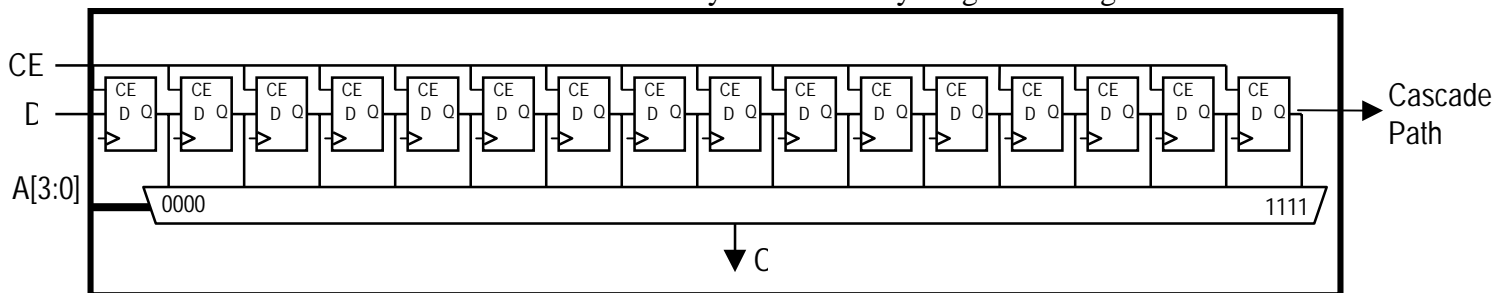
Figure 1: Virtex-II CLB Showing Four Logic Slices



The Virtex-II CLB also allows each of its eight lookup tables (LUTs) to be configured as 16-bit distributed RAM or as 16-bit shift registers. Distributed RAM within a single CLB can be used independently or combined for up to 128x1 single-port or 64x1 dual-port SRAM. This capability can be efficiently used for content addressable memories (CAMs), register banks, and data caches. In shift-register mode (SRL16), the LUT can be ‘unraveled’ into a fast cascadable 16-bit shift register with a variable output tap. The Virtex-II architecture allows these shift registers within each CLB, as well as in adjacent CLBs, to be cascaded to form arbitrary lengths. The SRL16 capability allows 16 registers to be implemented with a single logic cell, which is a 16x-density improvement over other competing FPGA architectures.

Figure 2: Single logic cell configured as SRL16 shift register (with up to 8 per CLB).

Note: These can be cascaded to efficiently form arbitrary-length shift registers.



The SRL16 capability is useful for very fast and efficient DSP pipeline operations, linear feedback shift registers, and small fast FIFOs. For example, wireless base stations can implement

digital filters with scheduled multiplication operations by using 16 SRL16 units for storing coefficients. This function utilizes only 16 logic cells, as compared to 256 logic cells in competing architectures that lack the SRL16 capability. This results in a 16-fold increase in logic efficiency, along with simplified routing requirements and faster pipeline performance.

Table 1: CLB Logic Capability

	Virtex-II	Virtex Series
Number of Logic Slices	4 (8 logic cells)	2 (4 logic cells)
Widest multiplexer without additional logic	32:1	8:1
Distributed RAM within CLB	128x1 single-port 64x1 dual-port	32x1 single-port 16x1 dual-port
Shift register using SRL16	128-bit within CLB, any length with inter-CLB cascade	Four independent 16-bit within CLB

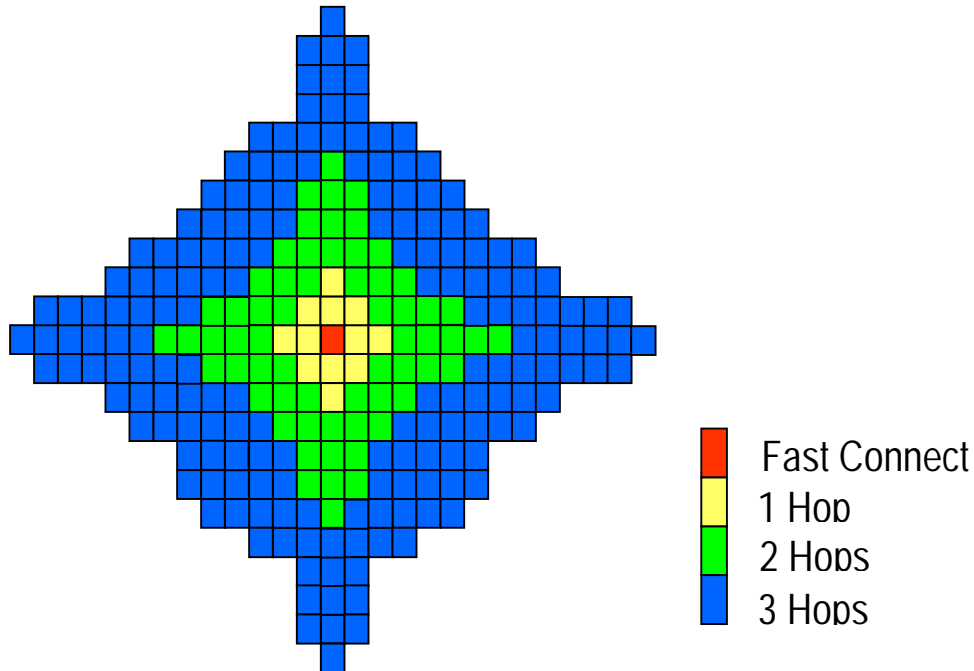
Routing Architecture

- *Active Interconnect™ technology: fourth generation segmented routing*
- *Vector-based routing with support for Smart-IP™ technology*
- *Fully buffered routing switches for fast, predictable routing delay*

The new Active Interconnect technology in Virtex-II, is developed for superior device utilization and fast software compile times, crucial for designs with up to 10 million system gates. Active Interconnect technology employs a fourth generation segmented routing architecture, which offers superior scalability at higher densities. Unlike CPLD-based fixed-length routing FPGA architectures, segmented routing architectures can be optimized in terms of routing length distribution to match the theoretical routing requirements with excellent silicon efficiency. In particular, routing studies as well as theoretical relationships, like Rent's Rule, predict that the number of different routing lines vary more slowly than the square of the density. Furthermore, routing length requirements have a distribution with an abundance of short single- and double-CLB lengths, with a gradually decreasing requirement for longer routing signals. The Virtex-II routing structure offers an abundance of single-, double-, and longer-length routing resources that allow for fast place-and-route software performance with full device utilization and consistently high performance. In contrast, nonsegmented routing architectures scale inefficiently as the square of density increases, which results in large silicon areas and poor routability, leading to slower performance, especially for high-density designs.

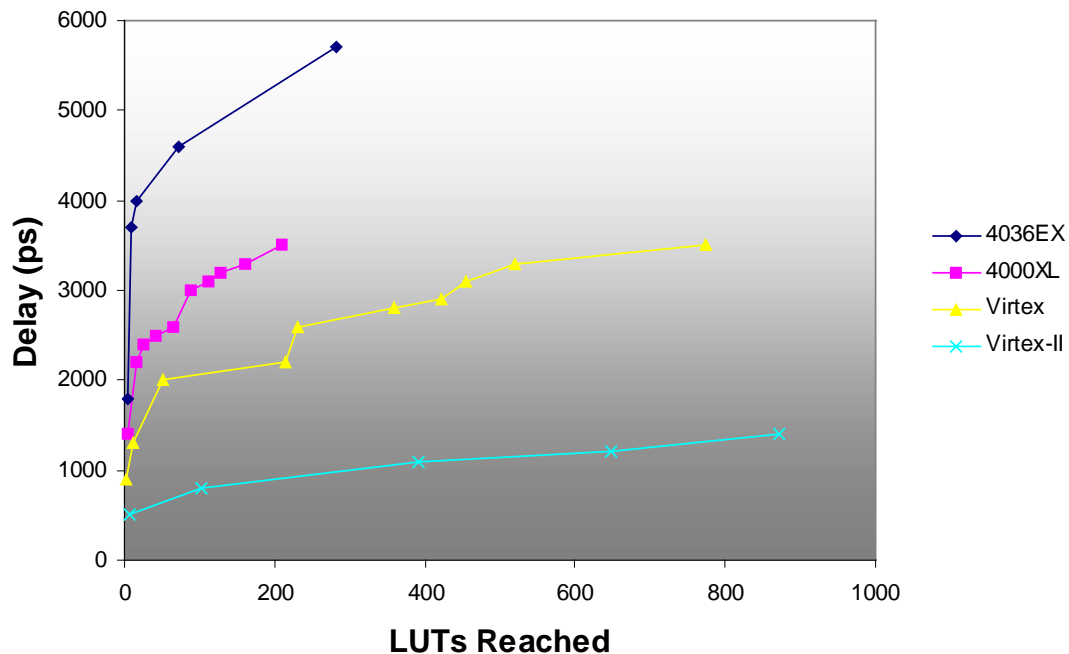
Active Interconnect technology leverages the strength of vector-based routing in the Virtex series. Vector-based routing enables regional routing delays to be quantified by the vector distance. This allows IP core timing to be preserved by embedding relative placement information—a key component of the Xilinx Smart-IP technology. The Virtex-II architecture further improves on routing capability by expanding the distance that a single routing connection can address. This, in conjunction with the high-capacity CLB structure, provides superior routability with minimal routing delays versus current generation FPGAs.

Figure 3: Number of CLBs reached using Virtex-II Active Interconnect technology



Active Interconnect technology provides active drivers for all routing connections, as opposed to other FPGA routing architectures that rely on passive transistor pass-gates for connecting routing lines. This innovative technology allows predictable, fanout-independent routing delays that support HDL- and IP-based design methodologies, which cannot tolerate pervasive routing delay changes during design iterations. In contrast, passive routing structures can suffer unpredictable delays due to additional capacitive loading, for example when additional fanout is added during design iterations. Active Interconnect technology achieves predictability within IP cores by supporting Smart-IP technology and provides predictable routing delays for their interconnects. The stability of routing delays allows greater engineering productivity and faster time-to-market, by reducing system simulation requirements and re-engineering efforts during design iterations.

Figure 4: Routing Delay vs. Number of Accessible Lookup Table.



Process Technology Migration

The Virtex-II architecture is the next Xilinx® FPGA platform for supporting additional system features. The architecture allows efficient silicon implementations in advanced deep submicron process technologies down to 100-nanometer (0.10-micron) feature sizes, with over a half billion transistors on a single device. Increasing system demands for very high density FPGAs and the emerging role of FPGAs as process technology drivers have together created the following trend—that FPGAs have become among the most complex digital semiconductors in the world, with higher transistor counts than state-of-the-art microprocessors.

The innovations in Virtex-II architecture provide for rapid migration to future process generations, which will have transistors with very short channel lengths and lower threshold voltages. These pose significant challenges in the current generation of competing FPGA architectures. These architectures typically use the transistor as a controlled resistor, whereby the transistor source and drain have a very high resistance when the transistor is "off" or a relatively low resistance when the transistor is "on". In the "on" state, the lines are connected together and driven by the same driver, with part of the routing being driven through the resistance of the transistor. However, this basic structure does not scale well with further process shrinks, where the resulting routing delay can increase greatly due to reduced transistor widths. In contrast, the patented Active Interconnect technology allows the elimination of transistor pass-gate only structures, thereby minimizing routing delays.

IP and Software for 10 Million System Gate Designs

To achieve high productivity at multi-million gate densities, next-generation system designs require full compatibility with an IP-based design methodology. Xilinx pioneered Smart-IP technology with the original Virtex series to enable predictable performance of IP cores. This proprietary technology allows relational placement information to be embedded in the cores. The Xilinx Active Interconnect technology in the Virtex-II architecture extends the predictability between the IP blocks to facilitate easy integration of multiple cores. The latest Alliance Series version 3.1i software incorporates hierarchical floorplanning, integrated modular design to facilitate engineering team work, dramatic run-time improvements for timing closure, and incremental design flows to further optimize productivity for designs with 10 million system gates. Furthermore, the Virtex-II architecture is fully supported in the version 3.1i software release, as well as the latest synthesis tools from the leading EDA vendors, enabling designers to start their multi-million gate designs today.

System Bandwidth

System bandwidth is increasingly limited by chip-to-chip communications in addition to the bandwidth of chips themselves. This trend has given rise to many specialized I/O standards for memory, graphics controllers, and other subsystem functions. Continuing the history of supporting leading-edge I/O standards, the Virtex-II architecture is designed to enable over 800 Megabit per second (Mbps) I/O performance while supporting single-ended and differential I/O standards currently supported in the Virtex series.

RapidI/O™ Support

In addition, enhanced SelectI/O technology in the Virtex-II architecture integrates RapidI/O support. RapidI/O technology is a high-speed, packet-based interconnect technology used in leading communications processors, host processors, and networking DSPs for 10Gbps bandwidth. This new I/O technology addresses the need for further bandwidth increases in the networking market. The Virtex-II architecture supports this new technology with full electrical compliance with signaling standards, in conjunction with an enhanced SelectRAM memory hierarchy to support packet storage and manipulation options required by those applications.

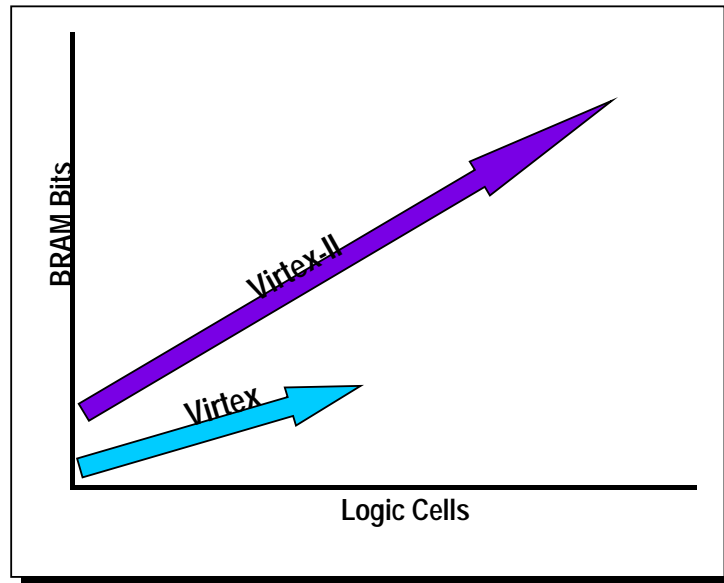
Memory Bandwidth

Xilinx pioneered the SelectRAM memory hierarchy to offer flexible on-chip and off-chip memory resources within a single FPGA device. The SelectRAM memory hierarchy enables Virtex FPGAs to be the first programmable devices to offer three different memory resources: distributed RAM, block RAM, and high-performance interfaces to external RAM and CAM devices.

With the Virtex series, Xilinx provides an unprecedented memory-to-logic ratio that enables products such as the memory-rich Virtex-EM family to satisfy today's data-intensive applications. The Virtex-II architecture takes this ratio to the next level. As illustrated in Figure 5, the Virtex-II architecture is well positioned to deliver next-generation FPGA products with

abundant memory resources, ideally suited for ever more data-intensive applications running on Internet infrastructure products.

Figure 5: Unprecedented Memory to Logic Ratio



The Virtex-II architecture further optimizes the SelectRAM™ memory with built-in parity bit storage and options for 1x to 36x data width configurations. This allows better medium-density data storage for internal FIFOs, DSP coefficient storage, and local cache memory. Each True Dual-Port™ block RAM can be configured for different data widths, with full read/write capability on each port. Building on the strength of True Dual-Port memory, new user-selectable "read-before-write" and "no-output-change-write" modes are added for the block RAM, which further improves the efficiency of block RAM operations. These two new modes enable the output port to optionally read previous contents or remain unchanged during a single write cycle. This eliminates separate read and write cycles, resulting in better efficiency and higher bandwidth for pipeline and DSP operations.

Distributed RAM in the Virtex-II architecture can be chained together, up to 128 bits deep. This capability is ideally suited for CAM, register banks, and data cache blocks, commonly used in DSP pipelining applications.

In addition, with built-in DDR I/O support, the Virtex-II architecture further improves memory bandwidth by enabling seamless interfaces to leading-edge DDR (double data rate) and QDR (quad data rate) memory at over 300 Mbps at each of the read and write busses.

Processing Bandwidth

The Virtex-II architecture includes embedded multiplier and enhanced arithmetic capability for high-bandwidth operations. The multiple high-speed 18-bit multiplier blocks available in this architecture enable theoretical performances exceeding 0.6 Tera MAC (trillion multiply-accumulates per second). The Virtex-II architecture is ideally suited for high-end DSP applications, such as wireless base stations and high-end video and image processing systems. In these applications, the embedded multiplier capability in the Virtex-II architecture can be used in conjunction with standard digital signal processors to increase overall design throughput. The IP leveraging the improved arithmetic capability performs hardware-intensive FIR (finite-impulse response) filtering and FFT (fast Fourier transform) operations, while the digital signal processors execute the decision-intensive operations.

The overall logic performance of the Virtex-II architecture is also enhanced two-fold over original Virtex levels. This is achieved with a combination of process technology and architectural enhancements, and it further increases the overall processing bandwidth of system applications.

System Timing Bandwidth

With the constant demand for higher bandwidth, complex system clock management techniques are required to address this challenge. The Virtex-II architecture further enhances the integrated system clock management solution, including improved global clock distribution and increased number of digital delay locked loops (DLLs). Combined with the copper interconnect technology employed, this feature substantially improves the clock skew characteristics, both for global and local clock networks. By minimizing clock skew, overall system bandwidth is increased by reducing the needed margin on required setup and clock-to-output parameters for all components. Furthermore, the DLLs operate in excess of 400 MHz to support high system frequencies required in next-generation complex systems.

Conclusions

Xilinx redefined the fabric and function of FPGAs with the Virtex series, enabling designers to architect systems with an FPGA at the core of their design. By significantly enhancing logic and routing structures, the new Virtex-II architecture further enriches the fabric for designs with up to 10 million system gates. The new features addressing high-performance and low-power requirements of tomorrow's complex system designs are unmatched in any other programmable architecture. The Virtex-II architecture offers revolutionary enhancements in all aspects of Virtex system-level capabilities and pushes logic density and bandwidth performance to the next level. Yielding twice the performance with only half the power consumption of the original Virtex family, the new Virtex-II architecture is slated to provide the next generation of advanced FPGAs capable of solving the system-level and bandwidth challenges of tomorrow's leading-edge communications, networking, and multimedia systems.