QUESTIONS AND ANSWERS
FOR XILINX VIRTEX SERIES

Q. Why do you say, "Xilinx is redefining the FPGA"?

Until Virtex series, the measuring criteria for an FPGA has focused on density and performance. Virtex series both significantly exceeds these current standards and offers more. In developing a device capable of implementing one million gates, devices need more than just gates and routing. System features such as advanced clock management; direct interface to existing and future voltage and signal standards; multi-granularity high bandwidth memory integration; and power management are critical features for what we call the Xilinx Virtex technology employed in Virtex devices. Offering this technology across from 50,000 to one million gates redefines what is possible with an FPGA.

Q. What do you mean by "system level integration capabilities" or "system level features"?

Terms such as system on a chip (SOC) and system level integration (SLI) have been used to define a device or design methodology that allows multiple functions of a system to be integrated into a single device. Such devices must offer not only the density required to implement the system, but also the capability of integrating these multiple functions easily with high performance. The system level features of Virtex devices allow this ease of integration. Delay locked loops (DLLs) allows high performance communication between the different functions. Select IO allows direct connection to external devices, high-speed busses and backplanes. SelectRAM+ memory allows efficient integration of high bandwidth memory functions. Thermal management capabilities provide constant monitoring of environmental conditions to allow systems to run at peak performance.

Q. How can an FPGA address system-level requirements?

For any device to address system-level requirements, they must provide high density, high performance and an ease of integrating multiple functions in a single device. Virtex FPGAs address these requirements by providing up to one million gates with a routing architecture that allows true 100+ MHz performance. In addition to providing the necessary resources and routing, Virtex series offers key system features in clock management, memory hierarchy, highly flexible IO, and power management allowing Virtex to be the center of any high performance system.

Q. How are these requirements being addressed now?

Until the Virtex series, ASIC devices, specifically standard cell devices, have been used to address the requirements for system-level integration. Virtex devices will offer the time to market and flexibility advantages of an FPGA to the system-level design community.

Q. Who are the decision-makers for designing in a Virtex device? Are these different people than other FPGA decision-makers?

The Virtex series redefines what is possible for an FPGA. Current FPGA designers will like the increased density, increased performance, and increased feature set of Virtex devices. In addition, Virtex will get the attention of decision-makers that have not had the option to consider FPGAs in the past because of its limited functionality. At many companies, a System Architect will separate key system components from glue logic functions. Historically, FPGAs get their first consideration for implementation in the glue logic functions. Because of performance, density or feature requirements, the key system components are commonly earmarked for ASIC design. With Virtex FPGAs, System Architects can consider the time to market and flexibility advantages of FPGAs for all components in a system.

Q. What are some of the new markets and applications that Virtex will enable for Xilinx?

Virtex devices will find homes in the traditional strong markets for FPGAs: networking, telecommunication, high-end computing, industrial and wireless markets. These are key markets for Virtex devices since the time to market and flexibility advantages of FPGAs have always been valued here. Within these strong markets, Virtex devices offer new application possibilities for FPGAs. The Virtex family is already being considered for applications that previously required ASICs. Some examples are SONET data stream processing for OC-3, OC-

12 and OC-48 standards, main processing engine for massively parallel compute network, variety of protocol controllers and switch matrix for next generation network equipment, and backplane controllers for ultra high speed board to board communication.

Q. What has allowed you to bring down the cost for such a robust feature set?
Virtex devices are designed for optimal efficiency using advanced 0.22 micron, 5 layer metal process technology. This delivers the highest number of gates for any given silicon area. The new features introduced with the Virtex architecture also were designed for efficient implementation using the state of the art process technology. Combining the highest density process with an architecture optimized for this process results in the highest number of gates per die area even with all the powerful new features.

Q. What is the advanced 0.22-micron process and is it different than the 0.25-micron process?
The 0.22-micron process is the second-generation 0.25-micron process from Xilinx representing a 10 percent shrink over the first generation. Xilinx is still the only programmable logic vendor shipping on 0.25 micron and has extended its leadership by delivering devices on its second generation before other vendors deliver their first. This aggressive process results in the most efficient use of die area, fastest devices, and the lowest power consumption.

Q. Which features are unique to Xilinx?
The system level features in Virtex devices include multiple delay locked loops (DLL), multi-standard I/O, hierarchy of RAM capabilities, and thermal management. Each feature by itself, although similar to features in other programmable devices, is unique to the industry in the specific implementation.

- Virtex devices are the first to provide high performance, working clock management circuitry with its DLL. Others have introduced devices with phase locked loops (PLL) that have not worked or only worked to lower frequencies.
- Having four DLLs on a single device is also unique and allows for complete system clock management as compared to single chip clock management.
- The multi-standard I/O is implemented such that even standards not defined today are likely to be supported with the Virtex device. Others have offered a finite number of options limited to different voltage standards.
- With Virtex devices, both voltage and swing characteristics are independently programmable.
- Virtex devices are the first to embed both distributed and block memory that will allow the right type of high bandwidth memory. Distributed memory has proven very valuable in high-speed data processing systems where a small amount of wide data is needed. The block memory in Virtex provides extremely efficient data buffering on and off chip and temporary storage for processing kbytes of information found in graphics and networking applications. Having full dual port, two independent read/write ports, is also unique to the Virtex block memory. This will allow for bidirectional buffers that do not suffer performance degradation seen with dual port memories from others.
- Thermal management control is a first in the industry.

Q. Why is temperature monitoring important?
Virtex devices are inherently low power (milliwatts of quiescent power) because of the 2.5V operation. The largest device has 75 million transistors—10 times the Pentium II—and the architecture supports performance over 100 MHz which can lead to tens of watts. By offering temperature monitoring, Virtex devices can be designed to perform right up to the limits of the packaging technology. No performance is compromised to account for worst case conditions that may never happen.

Q. What is the DLL useful for?
Clock synchronization is critical in high performance applications. The four Virtex delay locked loops (DLLs) provide a mechanism to remove all clock skew in a system as well as minimize clock to output delays. In ultra

high performance applications, clock skew and clock to output delays can easily be more than 50 percent of the critical path. With Virtex DLLs these values can be reduced to nearly zero effectively doubling possible system frequency.

Q. What are the advantages over PLL?

Compared to a phase locked loop (PLL), a DLL has a number of advantages. The Virtex DLL is a complete digital implementation that simplifies the processing and the future scalability of the DLL. The Virtex DLL does not accumulate jitter because each clock edge in essence restarts the circuitry. Finally, the Virtex DLL has been proven to operate at frequencies over 200 MHz while competing PLL solutions in programmable logic devices either don't work or are limited to frequencies below 100 MHz.

Q. What is the difference between all these I/O standards?

The different I/O standards supported by Virtex come in two flavors. The common standards such as LVTTL, LVCMOS2, PCI33 and PCI66 operate at either 3.3V or 2.5V and swing about a point that is common. The more exotic standards like GTL, GTL+, SSTL, HSTL, AGP and CCT are specifically designed for low power and high performance. The voltage levels and the swing points differ and are designed for specific applications. GTL and GTL+ standards are for low power high-speed backplane interconnect. SSTL, HSTL and CTT standards are specific to certain high-speed memory devices. AGP is specific to high performance graphics. To support all of these with single device, Virtex SelectIO offers programmability to the output voltage waveforms as well as programmability to the swing point.

Standard		**Application**	*Xilinx is lone supplier*
LVTTL	Low Voltage Transistor-Transistor Logic	General Purpose 3.3V Logic	
LVCMOS2	Low voltage CMOS	General Purpose 2.5V Logic	
PCI 33MHz 3.3V	Peripheral Component Interface	33 MHz PCI performance for 3.3V systems	
PCI 33 MHz 5.0V	Peripheral Component Interface	33 MHz PCI performance for 5V systems	
PCI 66 MHz 3.3V	Peripheral Component Interface	66 MHz PCI performance for 3.3V systems	
GTL	Gunning Transceiver Logic	High-speed processor interface	✓
GTL+	Gunning Transceiver Logic	High-speed backplane driver	✓
HSTL-I	High-Speed Transceiver Logic	High-speed SRAM	✓
HSTL-III	High Speed Transceiver Logic	High-speed SRAM	✓
SSTL3-I	Stub-Series Terminated Logic	High-speed SDRAM	✓
SSTL3-II	Stub-Series Terminated Logic	High-speed SDRAM	✓
SSTL2-I	Stub-Series Terminated Logic	High-speed SDRAM	✓
CTT	Center-Tap-Terminated	High-speed memory	✓
AGP	Advanced Graphics Protocol	High-speed processor interface	✓

Q. Do I need to order different Virtex devices for the different standards?

No, every I/O on every Virtex device is capable of supporting every standard. To interface to a specific standard, a designer drops in the type of I/O buffer required. For example, to implement a GTL input buffer, a designer simply specifies IBUF_GTL. The configuration of the I/O buffer is done during device configuration and the specific I/O detail is contained in the bitstream generated by the Xilinx software solutions.

Q. Does the Virtex family support applications for reconfigurable computing? Is it partially reconfigurable?
Virtex devices do support reconfigurable computing. Each device can be reconfigured to provide different functions at different times. A specific high-speed configuration port allows bitstreams to be downloaded into Virtex devices at a rate up to 400 Mbit/second, which allows even the largest device to be completely reconfigured in milliseconds. Virtex devices will also provide for partial reconfiguration, meaning a portion of the device can be changed while the rest of the device continues to operate.

Q. What is the utilization percentage for the Virtex devices?
The five layer metal process allows for massive amounts of routing. The design for the routing of Virtex is to allow full utilization of the million-gate device. Because of the segmented routing nature of the Virtex device, all family members will have a routing structure that is designed to support high performance for the million-gate device. Designers will directly benefit from this routing with very fast compile times, high performance from push button design flows, and full utilization of all device resources.

Q. Does the gate-counting method differ from the traditional Xilinx method? What percentage of RAM use are you assuming in these gate counts? How do you account for the Block RAM?
The gate counting used on Virtex devices is consistent with the system gate counting used for Spartan and XC4000 devices. The basics of the system gate count equation are each logic cell used as logic provides 12 system gates; each logic cell used as distributed memory provides 64 system gates (16 bits x 4 gates/bit). Since larger devices typically require more memory the percentage of logic cells used as distributed memory ranges from 5—30 percent across the family. With Virtex additional capabilities add to the system gate count. Each DLL used provides 7000 system gates and each logic cell used as a programmable delay element provides 112 system gates. Finally, block memory adds to the gate capacity of Virtex and each bit of block RAM memory provides four system gates. Up to 40 percent of block memory usage is factored into the system gate count.

Q. Is there JTAG support in Virtex devices? Yes

Q. Will all future FPGA families from Xilinx be derived from the Virtex architecture?
Xilinx invented the FPGA and continues to drive the standard platforms. Virtex represents the fourth generation of FPGA standards, preceded by the XC2000, XC3000 and XC4000 architectures. New families will use the Virtex architecture as their basis.

Q. When will software support be available for Virtex?
Virtex software support is available today. Xilinx and other EDA vendors including Exemplar, Synopsys and Synplicity offer breakthrough technology in performance via VHDL and Verilog synthesis. The latest software releases, the Xilinx Alliance Series and Foundation Series 1.5 software, offer full support for the Virtex family of FPGAs. New timing driven compile time benchmarks have been set with these tools by offering typical rates of 200,000 gates/hour for timing driven place and route.

Q. What is "vector-based interconnect"?
"Vector-based interconnect" conveys the result of the massive routing capacity offered by Virtex devices and represents the next generation of the Xilinx segmented-routing structure. With five layers of metal, the routing architecture of Virtex has been designed such that the minimum amount of metal is required to implement a given interconnect. This means that delays are closely estimated as a function of the straight line distance (vector length) of the interconnect. The designer benefits of vector-based interconnect are: 1) predictably high performance due to routing efficiency; 2) consistent performance across devices due to consistent routing architecture; 3) fast compile times due to the massive amount of routing tracks; 4) and a cores-friendly architecture. Add period

QUESTIONS AND ANSWERS
FOR XILINX VIRTEX SERIES

Q. What make Virtex series so "cores friendly"?

The suitability of Xilinx FPGAs to a high performance, highly predictability and easy to use core methodology is further extended by the features of the Virtex architecture. While the vector-based interconnect allows more cores with higher performance requirements to be implemented in HDLs than in older architectures, the fact that the routing scheme is a next generation segmented routing structure means that the benefits of Smart-IP are still applicable. Smart-IP allows the core designer to leverage the absolute maximum performance for the fastest cores without giving up predictability or ease of customization.

The Virtex memory capabilities further extend the range of cores that can be easily implemented in an FPGA. SelectRAM distributed memory is leveraged for small FIFO's, data buffers and memories at the places that they are needed in the design. When larger memory blocks are required, the new SelectRAM+ memory blocks provide a wealth of size and bus width options to support single-port or multi-port memory for larger storage needs. Add to this the ability to interface to high speed SDRAMs using the SelectI/O and you have a memory hierarchy that gives Virtex devices the capability of powering even the most demanding memory-based applications.

This wealth of enhancements ranging from the improvements that allow faster multipliers to the system level features such as the DLLs and clock management capabilities has resulted in Virtex being the most core friendly architecture ever. Its feature set is unparalleled and is especially important at the one million gates density level. It will enable cores implementation that is not possible in any other device.

Q. When will the 64 bit/66 MHz PCI core be available?

Virtex is capable of delivering a 64 bit/66 MHz PCI compliant design. We are in the process of productization of this core. Its availability will be announced later.

Q. What are the devices in the family?

	Logic cells	System gates	Block RAM bits
XCV50	1,728	57,906	32,768
XCV100	2,700	108,904	40,960
XCV150	3,888	164,674	49,152
XCV200	5,292	236,666	57,344
XCV300	6,912	322,970	65,536
XCV400	10,800	468,252	81,920
XCV600	15,552	661,111	98,304
XCV800	21,168	888,439	114,688
XCV1000	27,648	1,124,022	131,072