The Core Story: *A Breakthrough*

This year, Xilinx will ship products with unprecedented logic density. Our FPGAs already offer you 250K system gates, an order of magnitude greater than what was available just a few years ago. Devices twice as large are expected around mid-year, and our first million-gate FPGA will be sampling by the end of 1998.

How will you create the logic to fill these huge devices?

It's obvious that designing one gate at a time is not going to work. The answer lies with intellectual property (IP), or cores, which are predefined system functions, used for nearly a decade by designers of traditional mask-programmed custom ASICs.

It's only recently that cores have started making inroads into FPGA designs. Three main reasons account for this new migration. First, FPGAs are now large enough to accommodate cores, and a surprisingly large and diverse library of compatible cores is coming into existence. These cores include functions such as PCI and PCMCIA bus interfaces, digital signal processing algorithms, RISC microprocessors, standard peripheral controllers, and asynchronous transfer mode (ATM) functions. The expansive logic resources of our FPGAs, coupled with cores, allow you to create system-level designs on a single chip.

"The great attraction of cores is that they allow you to quickly and reliably create the most difficult sections of your designs. For small designs, cores are a welcome convenience. But for larger designs, they are becoming a necessity."

The great attraction of cores is that they allow you to quickly and reliably create the most difficult sections of your designs. For small designs, cores are a welcome convenience. But for larger designs, they are becoming a necessity.

For example, one of our customers, a forward-looking company, created an embedded application that combines a RISC processor core with several DSP core functions, communicating over a PCI bus whose interface logic was also created by a core. This would have been a monumental undertaking if designed from scratch.

New tools are already on the horizon that will make the task of grouping multiple cores on a single FPGA even easier. These tools will allow you to pull cores from a common library and place them at predetermined locations to make the most efficient use of resources and achieve your performance requirements.

A second reason cores are coming to the FPGA marketplace is performance. The latest generation of FPGAs operate at system speeds in the 80—100MHz range, and they will soon exceed 150MHz, fast enough to handle 66 MHz PCI 64 or communications protocols such as a 155Mbps synchronous optical network (Sonet). These functions are significant design challenges in themselves, and it's a great benefit to you if the functionality is available as a core.

Third, intellectual property developers have traditionally built their cores around standard high-level description languages such as VHDL and Verilog, the mainstream tools of ASIC designers that provide a large degree of flexibility. These languages are now becoming the basis for more and more PLD tools, and that is one of the many things attracting IP developers to the programmable logic market. In fact, developers are discovering that FPGAs are excellent prototyping vehicles for cores.

Developers can silicon-test their designs directly on programmable logic devices and polish the code much quicker and with less expense than they could by going through an ASIC vendor

in Time to Market

by Rich Sevcik, Senior Vice President of Software, Xilinx

and lining up a customer as a development partner. Our SRAM-based PLDs permit core designers to "rewire" the devices immediately by reprogramming them with new designs.

Moreover, the growing use of FPGAs presents independent IP developers, and programmable logic vendors themselves, with a "mass market" for their products. The worldwide universe of programmable logic customers numbers in the tens of thousands, compared to several hundred very large companies that buy IP for their high-volume mask-programmed ASICs.

Cores are also helping you answer the classic "make or buy" question. For example, as ubiquitous as it is today, the PCI interface remains a complex standard, rife with timing-critical specifications. Making a PCI interface from scratch can add from six to nine engineering months to a design. Buying it, on the other hand, can mean a substantial saving of time and money, especially for PCI designs where engineering time, cost, and volume may not justify going to a traditional ASIC solution. Additionally, buying cores frees you to concentrate on the intellectual value they add – beyond the basic PCI bus interface – to the product that's on the drawing board.

The Xilinx LogiCORE PCI interface illustrates one model of how cores are being delivered in the FPGA market today. Our PCI offering consists of pre-defined functions (target and initiator) that allow you to create a complete PCI interface on a single FPGA, and still have ample logic remaining to create the unique back-end interface required for your application.

Two points illustrate why the market has quickly accepted a product like the LogiCORE PCI core (more than 300 electronic equipment manufacturers have licensed it to date). First, the design is pre-verified and tested, ensuring that it will comply with the rigorous PCI specification. Second, it has a pre-defined layout, and it's optimized for our FPGA architecture. Therefore, timing for critical paths is fixed, ensuring predictable and consistent performance.

Cores are flowing from a growing community of independent IP developers who are coming to realize that their cores must be tuned for a particular device architecture using the tools designed to program that device. Power consumption, performance, and core predictability vary considerably based on differences in the FPGA vendors' place-and-route tools, device interconnect structure, and on-chip memory resources.

"Cores are flowing from a growing community of independent IP developers who are coming to realize that their cores must be tuned for a particular device architecture using the tools designed to program that device."

Design verification and device optimization are critical elements for the success of PLD cores, whether they are sold and supported by device suppliers or by IP developers. In fact, independent IP developers are beginning to align themselves closely with programmable logic vendors in order to accomplish this. Xilinx, for example, has partnered with nearly two dozen IP providers worldwide through its AllianceCORE program, and expects to expand the number of partners in the program this year. Such partnerships help to ensure that cores will reach you only after they are verified and optimized, and only when a strong support system is in place.

The momentum is clearly behind PLD cores, and during 1998 you can expect to see significant new developments in this segment of the market. Larger and faster devices, new FPGA architectures, powerful tools, and targeted IP offerings are shaping up to combine cores and FPGAs into true system-level solutions. ◆