Swiss Engineers Use FPGAs to Link

Designers at Supercomputing Systems AG (Zürich, Switzerland) took a novel approach in designing the GigaBooster parallel computer system — combining multiple processing elements built from standard components with a fast, low-latency communication scheme implemented in hardware. To implement the critical communications control logic, they chose the world's leading FPGA family — the XC4000 Series.

> *"To implement the critical communications control logic, they chose the world's leading FPGA family — the XC4000 Series."*

Based on the "alpha7," a prototype system designed at the electronics laboratory of the ETH (the Swiss Federal Institute of Technology), the GigaBooster is a parallel computer containing seven individual processing elements (PEs) on a single board. Each PE is constructed from a DECchip 21066 Alpha processor, one Mbyte of cache memory, up to 128 Mbytes of DRAM memory (256 Mbytes in the "root" PE), a SCSI controller, and two special registers that control clock, reset, interrupts and similar functions. Each PE is accompanied by two banks of FIFO buffers dedicated to interprocessor communication and connected to a common 72-bit bus. All the special registers and FIFO buffers are controlled by a central communications controller realized within three XC4013 FPGAs. All the PEs run the Digital UNIX operating system, providing access to over 3000 applications.

A new communications protocol called Intelligent Communication was developed to provide fast communications and communication control, enabling the system to take full advantage of the processing power in each PE. This protocol, implemented directly in hardware using the FPGAs, allows fast, low-latency data exchange among the processors, and a programming model with simple and efficient code. The in-system programmable nature of the XC4013 FPGAs was key to the development of this protocol; during development, various approaches were tested and compared simply by reconfiguring the FPGA devices.

In the original alpha7 system design, the communications controller was squeezed into just two XC4013 devices.

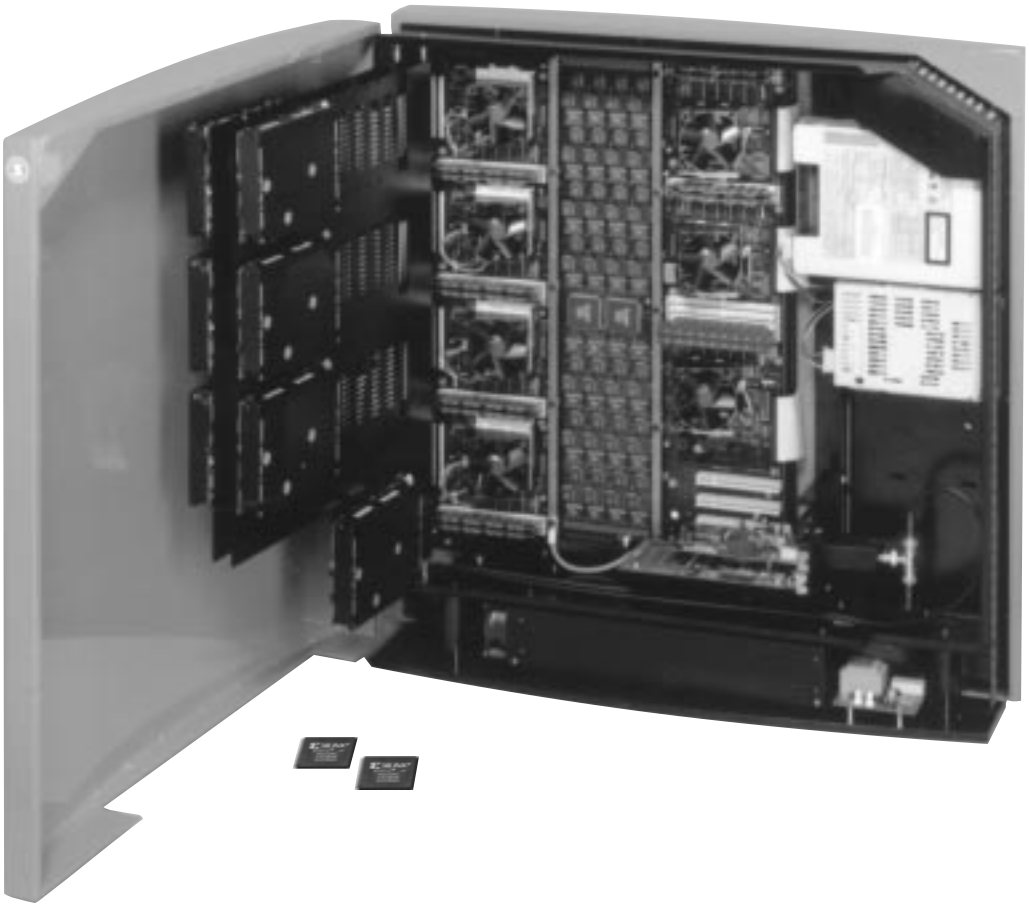

Supercomputing Processors

Supercomputing Systems

In the GigaBooster system, the controller was redesigned into three XC4013 FPGAs to allow room for expansion. One FPGA holds several small state machines, an abundance of control registers, and other glue and interface logic; this design uses about half of the available logic blocks, but all of the I/O pins. The other two FPGAs implement the logic directly involved in the gathering and redistribution of data from the processing elements, including a 42-bit counter and a large register/comparator file for each PE. The first of these FPGAs is more than 90% utilized, and connects to five of the processing elements.

About 40% of the second FPGA contains the identical logic for the remaining two PEs. Additional logic is dedicated to monitoring the communications behavior of applications running on the system, and the remainder of this FPGA can be used to support a module slot for an additional processing element or an interface to an optical high-speed network. The XC4000 architecture's built-in carry logic was critical to attaining acceptable performance from the large counters and comparators in this implementation. The design runs at 20 MHz.

The FPGA designs were developed on a Sun workstation using Viewlogic's Powerview tools and the XACT® development system. This combination "forms a very comfortable development environment," according to Hansruedi Vonder Müehll, the design engineer at the Swiss Federal Institute of Technology who was responsible for the original design of the communications controller. Some floorplanning was required, and both functional and timing simulation were used to debug and verify system operation.

The FPGA's readback capability also was exploited during system debug, and is now used to "dump" the state of the communications controller in the event of a communication failure. Using readback, the values of internal counters, registers, and state machines can be extracted and analyzed.

In summary, the use of reconfigurable XC4013 FPGAs was key to the implementation of interprocessor communication protocols directly in hardware, as opposed to the more-traditional software approaches. The resulting high-performance communications management allows the system to tap the full processing power of each of its Alpha processors, delivering 1.6 Gigaflops of peak performance in an affordable and compact system. ◆

"The use of reconfigurable XC4013 FPGAs was key to the implementation of interprocessor communication protocols direclty in the hardware."