

NetFPGA—An Open Platform for Teaching How to Build Gigabit-Rate Network Switches and Routers

Glen Gibb, *Member, IEEE*, John W. Lockwood, *Member, IEEE*, Jad Naous, Paul Hartke, and Nick McKeown, *Fellow, IEEE*

Abstract—The NetFPGA platform enables students and researchers to build high-performance networking systems using field-programmable gate array (FPGA) hardware. A new version of the NetFPGA platform has been developed and is available for use by the academic community. The NetFPGA platform has modular interfaces that enable development of complex hardware designs by integration of simple building blocks. FPGA logic is used to implement the core data processing functions while software running on an attached host computer or embedded cores within the device implement control functions. Reference designs and component libraries have been developed for the CS344 course at Stanford University, Stanford, CA, and taught at a series of tutorials held in the United States, United Kingdom, India, China, Australia, and Europe. The open-source Verilog, C, Perl, and Java reference design is available for download from the project website.

Index Terms—Field-programmable gate arrays (FPGAs), Internet, networks, protocols, routing, switches.

I. INTRODUCTION

HIGH performance network switches and routers enabled the rapid growth of the Internet. Gigabit Ethernet switches are widely deployed to interconnect computers in local area networks (LANs). Multigigabit/second links are used to transport Internet protocol (IP) packets across wide area networks (WANs). At most universities, students only learn to build networking systems with software. The students that do take a hands-on course in computer networking typically write software programs that send and receive packets through user-space sockets. Students who take advanced courses write software for the kernel that interfaces directly with a Linux operating system. While systems implemented with software may be able to send and receive some of the packets to and from a gigabit/second Ethernet line card, software alone is not suitable for switching, routing, and processing most of the traffic that appears on high-speed networks.

Manuscript received July 6, 2007; revised November 26, 2007. Published August 6, 2008 (projected). This work was supported by Xilinx, Cisco, Huawei, Google, Agilent, Juniper, Micron, Cypress, Broadcom, Digilent, Synplcity, and Grants (NSF EIA-0305729; CNS-0551703) from the National Science Foundation. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or other sponsors of this project.

The authors are with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305-9030 USA (e-mail: grg@stanford.edu; jwlockwd@stanford.edu; jnaous@stanford.edu; phartke@stanford.edu; nickm@stanford.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TE.2008.919664



Fig. 1. Photograph of a NetFPGA installed in a Desktop PC.

Commercial vendors of high-speed networking equipment use application specific integrated circuits (ASICs) and/or field-programmable gate arrays (FPGAs) to accelerate the switching, routing, and processing of packet data. To be competitive, students need to understand how these hardware-accelerated systems operate. Using the NetFPGA platform, students can build and prototype their own hardware-accelerated networking systems and run live traffic through the hardware. Using the NetFPGA, traffic on four gigabit Ethernet links can be processed at the full gigabit/second rate. By processing data with hardware, back-to-back packets can be scanned at gigabit rates [1], [2].

For teachers of hardware or networking classes, a set of reusable infrastructure and courseware has been developed. By starting with this infrastructure, students need not build systems from scratch. Gateware and software provided with the NetFPGA enables the platform to receive and transmit packets, buffer data in memory, and communicate with the host processor. A reference router enables the NetFPGA to operate as a full line-rate Internet router. The courseware, consisting of Web pages, homework assignments, machine problems, and presentation slides, provides instructors with the resources to teach students how to build an Ethernet switch and an Internet router.

The NetFPGA platform fits into a standard desktop or rack-mount computer. The hardware augments a PC with an expansion card that has four ports of gigabit Ethernet attached to the peripheral communication interconnect (PCI) bus. The FPGA on the platform directly handles all data-path switching, routing, and processing of Ethernet and Internet packets, leaving software to handle only control-path functions. A photograph of a PC with the NetFPGA installed is shown in Fig. 1.

II. THE NETFPGA PLATFORM

A photograph of the NetFPGA (version 2.1) platform is shown in Fig. 2. At the center of the NetFPGA platform is a large Xilinx FPGA device. Surrounding the FPGA are




Fig. 2. Photograph of the NetFPGA platform.

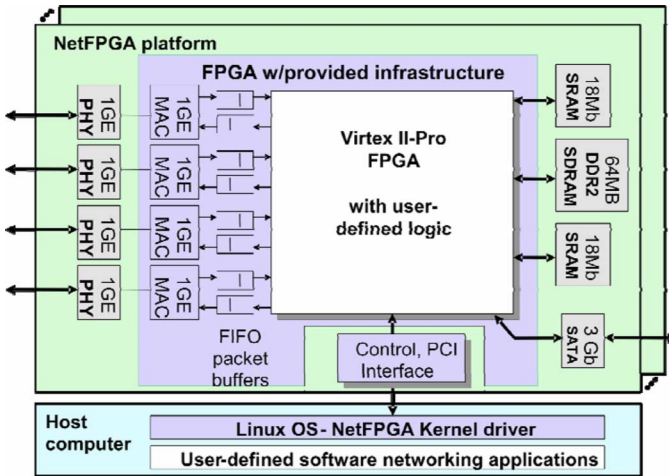


Fig. 3. Block diagram of the NetFPGA.

four memory devices—two static RAMs (SRAMs) and two second-generation double data rate (DDR2) SDRAM devices. On the left side of the platform, a quad-port physical-layer transceiver (PHY) is provided that enables the platform to send and receive packets over four standard twisted-pair Ethernet cables. On the right side of the board, two serial advanced technology attachment (SATA) connectors on the platform allow multiple NetFPGAs within a system to exchange data at high speeds without using the PCI bus.

A. Components of NetFPGA

A detailed block diagram that shows the major components of the NetFPGA platform is shown in Fig. 3. The NetFPGA 2.1 platform contains one large Xilinx VirtexII-Pro 50 FPGA, which is programmed with user-defined logic. The core clock on this device can be configured by the user to run at either 125 or 62.5 MHz. The NetFPGA platform also contains one small Xilinx Spartan II FPGA that implements the control logic for the PCI interface to the host processor.

Two external Cypress SRAMs are arranged in a configuration of 512 k words by 36 b (18 Mb/4.5 MB total) and operate synchronously with the FPGA core logic at either 62.5 or 125 MHz. One bank of external Micron DDR2 SDRAM is arranged in a configuration of 16 M words by 32 b (64 MB total). Using both edges of a separate 200 MHz clock, the memory has a bandwidth of 400 MWords/s (1600 MB/s = 12 800 Mb/s).

The Broadcom quad-port gigabit PHY sends packets over standard category 5, 5e, or 6 twisted-pair cables. The PHY

interfaces with four gigabit Ethernet media access controllers (MACs) instantiated as soft cores within the FPGA. The NetFPGA also includes two multigigabit interfaces with SATA connectors that enable multiple NetFPGA boards in a system to exchange traffic directly without use of the PCI bus.

B. Computer-Aided Design Tool Flow

A number of computer-aided design (CAD) tools are required to simulate and implement designs for the NetFPGA. Designs are simulated using Mentor Graphics ModelSim to verify correct logical operation. Verilog source code is synthesized using either Xilinx ISE or Synplicity's Synplify Pro. Logic is mapped, placed, and routed using Xilinx ISE backend tools. Finally the resultant bitfile is programmed into the FPGA using a command-line program that runs on the host PC. Debugging of hardware circuits can be performed via the JTAG port using Synplicity Identify or Xilinx ChipScope on-chip logic analyzers.

C. Interface to Software

A Linux kernel device driver allows user space software on the host PC to send packets and receive packets via the Ethernet ports on the NetFPGA board. The driver also allows application programs to read or write the memory-mapped registers that correspond to a design's on-chip registers, SRAM, and DRAM. In total, 128 MBytes of host memory address space is reserved for each NetFPGA board installed in a host computer.

The NetFPGA device driver is a loadable kernel module (LKM) that should be compatible with any Linux 2.6 kernel. Loadable kernel modules allow drivers to be installed in the kernel without requiring a full kernel recompilation. The driver provides four network interfaces, named nf2c0 to nf2c3, to user space.

User-space programs send and receive network packets to the NetFPGA via the Linux networking stack using a standard socket interface. Programs initiate communication with the NetFPGA using standard `socket()` and `bind()` functions and then use the `connect()` or `listen()` functions to open a channel of communication with the NetFPGA. Programs then use standard functions such as `read()`, `write()`, `send()`, and `recv()` to send and receive data to/from the NetFPGA.

Direct memory access (DMA) is used to transfer complete Ethernet frames between the host and NetFPGA—these DMA transfers are initiated by the driver but performed by logic in the NetFPGA. This allows the host to perform concurrent operations while the I/O is in progress.

User-space programs can perform memory-mapped reads and writes via I/O control commands (`ioctl`). An `ioctl` command performs a single word transfer to or from the NetFPGA. These calls provide a convenient method to access registers and memory on the board.

D. Availability of Hardware

NetFPGAs are available to the community directly from a third-party vendor who manufactures the hardware. This company, Digilent, Inc., has a long history of making low-cost eval-

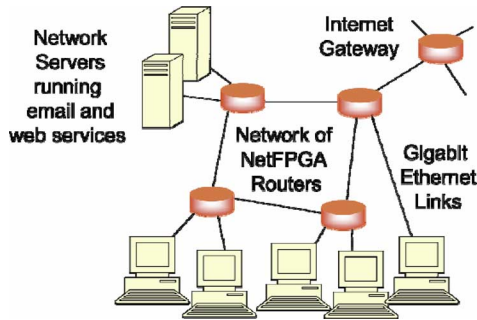


Fig. 4. Sample topology of a network with five NetFPGA routers.

uation boards through their partnerships with the Xilinx university program (XUP).

III. STANFORD'S CS344 ROUTER DESIGN COURSE

The CS344 course at Stanford University, Stanford, CA, is a graduate-level, project-based course that teaches Ph.D., Master's, and advanced undergraduate students how to build an Internet router in just eight weeks [3]. Students who take the class learn how to develop both the hardware and software components for a high-speed router. Each team consists of one student to work on hardware and one or two students to write router control software.

Using the NetFPGA, student teams build a fully functional, gigabit/second Ethernet switch and router. Together, teams interconnect multiple NetFPGA platforms in a network that routes packets over several gigabit Ethernet links. The topology of one such network implemented with five NetFPGA routers is shown in Fig. 4. Forwarding tables are implemented in hardware, and use values populated by the routing software to forward packets along the shortest path between machines.

A. Student Background

Students undertaking the hardware design task are expected to enter the class with basic hardware design experience. The assumption is made that these students already understand the basic concepts of synchronous logic and hardware pipelining. Most computer engineering students entering CS344 have already coded at least a small project using Verilog.

Students undertaking the software design task are expected to have taken a networking class. They are assumed to have already been exposed to the concepts of network stacks and socket programming. Students should also be proficient in programming in C.

Hardware students are also encouraged to have taken a networking course prior to enrolling in this course, so as to have a basic understanding of addressing and routing. Students without this background are expected to do background reading during the first few weeks of the course.

B. Course Structure

This course is conducted as a quarter-long (10 weeks) course at Stanford University. Students taking the course typically take two other classes during the quarter, or are working on a part-time research project. Students are expected to allocate

10–20 hours per week to complete the class assignments. The course is structured with periodic assignments that have regular deliverables. A short lecture is given before each assignment to describe the objectives and milestones. Teams meet with the teaching staff on a weekly basis. The teaching staff provides guidance to each team, and monitors progress to ensure that students do not fall behind on the assignments.

At the start of the course, students are provided with the NetFPGA hardware platform, Verilog HDL code that implements an Ethernet repeater, and skeleton software code that provides a framework for implementing the router's control path. The hardware and software members of each team work independently during the early stages of the course except for meetings where they discuss the interfaces between hardware and software.

In the first assignment, the hardware members of each team extend the design of the given two-port Ethernet repeater to implement a four-port learning Ethernet switch. This assignment provides the hardware students with an opportunity to familiarize themselves with the CAD tools, the NetFPGA platform, and the format of fields within an Ethernet packet.

In the next assignment, hardware students replace the switching logic with routing logic that forwards packets using the IP address in the header. The students compare the destination IP address in the packet to entries stored in the software-generated routing table to find a longest matching prefix. According to which of the entries matches, they then look up the next-hop using the address resolution protocol (ARP) and update the destination address field in the Ethernet header.

While some team members build hardware, the rest write software to implement the control path. Students use the virtual network system (VNS) in order to test control programs independently of the hardware [4]. The control path software includes a routing algorithm based upon OSPF, a command-line interface, and the logic to route packets in software (necessary when the hardware forwards packets to software, such as in the case of an ARP table miss). Also included in the implementation is support for ARP and ICMP protocols.

At the midpoint of the course, the hardware and software members integrate their designs. The hardware and software members test their designs together to ensure that the hardware and software components properly communicate. Teams test the interoperability of their routers with one another by interconnecting their routers and verifying that they exchange shortest path first (PW-OSPF) properly exchanges routes between the machines.

In the final few weeks of the course, student groups each implement an advanced feature for their router. Students choose their project and discuss their plans with the instructor. Examples of features implemented by students during the spring 2007 quarter include a network address translation (NAT) module, a multicast router, access control lists (ACLs), and DES encryption.

C. Hardware Design

A core feature of the NetFPGA platform is a modular design that is easy to understand and modify. The class and reference




Fig. 5. NetFPGA's router data path.

material encourage students to carry that modularity through to their resultant designs.

The starter code is structured around the following set of components: the user data path, the Ethernet I/O blocks (encapsulating the MAC and clock domain crossing logic), and the host interface logic (enabling communication and packet transfer between the host and NetFPGA). Students add logic to perform routing or switching to the user data path.

The user data path is implemented as a pipeline, using multiple modules. An arbiter reads packets from the Ethernet or from the host interface then feeds them through the pipeline. Because packets are multiplexed together, no crossbar is needed within the data path. Data is demultiplexed as packets are written to output queues. Peak bandwidth through the user data path is 8 Gb/s in the FPGA fabric. A block diagram of the data path is shown in Fig. 5.

The main components implemented by the students within the user data path were the input arbiter and an output port lookup module. The output port lookup module was built as a set of multiple submodules that holds the ARP and routing tables, performs longest prefix match (LPM) on the routing table, performs next-hop lookups on the ARP table, verifies the IP header, and updates the necessary fields in the packet. The output queues and the modules, except for the user data path, were provided to the students as part of the starter code.

D. Hardware Testing

Students first test their designs in simulation. The NetFPGA test environment uses Perl scripts to instruct the simulator to inject packets into the PHY, perform DMA transfers, read packets from the PHY, monitor a DMA transfer, and read/write registers through the PCI interface. Cycle-accurate simulation allows basic design errors to be identified and corrected prior to synthesis. Once the hardware design appeared to be functional, students synthesized their circuits using the Xilinx ISE CAD tools and downloaded the resulting bitfile into the Virtex FPGA on the NetFPGA. Simulation worked well for testing small volumes of traffic, but was found to be far too slow for testing large volumes of traffic.

Students were given access to PCs that allowed them to test their designs in real hardware at full gigabit line rates. Each

PC had both a NetFPGA card and a quad-port gigabit network interface controller (NIC) installed. Initially, each port of the NIC was connected to a port on the NetFPGA. Students ran utilities such as ping, traceroute, iperf, and Wireshark to inject and monitor real traffic to and from the NetFPGA.

At a larger volume of traffic, students found problems with their hardware that they had not observed in simulation. These problems were generally caused by corner cases not covered in simulation. To debug the hardware, signals on the FPGA were analyzed using Xilinx Chipscope, an on-chip logic analyzer. Chipscope allowed the students to monitor and capture the sequence of events that caused a packet to be dropped or corrupted.

Later in the course, students interconnected different machines to create network topologies with multiple routers. In these networks, students performed experiments to ensure that multihop routes were created correctly. The students unplugged Ethernet ports to verify that the routers responded properly to link failures and were able to reroute packets along the next-best path. Multiple groups tested their routers together to verify that the routers were interoperable.

E. Student Feedback

Students were asked to provide informal feedback at the conclusion of the class. Common feedback received indicated that the class was interesting, personally satisfying and quite challenging. Students also indicated that the workload was relatively high. One student wrote that *“Overall it was awesome. I came in without knowing much about networks or routers, and I left with a wealth of information. It must sound weird but I truly liked everything since I have never seen this stuff in such detail beforehand.”*

Another student wrote: *“It was one of the most satisfying, interesting, and time consuming courses I’ve taken at Stanford.”*

The two most common issues reported by the hardware students were: 1) problems meeting timing and 2) problems identifying elusive bugs. The resource requirements of the starter code were reduced for the 2008 version of the class, making it easier for the place and routing tools to meet timing requirements. The timing issues can also be significantly eased by targeting to the 62.5 MHz clock. Additional steps have also been undertaken to further simplify the process of writing simulation tests to aid in identifying more bugs before testing in hardware.

IV. NETFPGA INFRASTRUCTURE

The NetFPGA is an open platform that can be used at other universities to teach courses like CS344. The components that implement the functionality to switch and route packets can be reused to teach courses like CS344 at other institutions.

Most NetFPGA designs include both hardware components and software components. For example, with the four-port router, all of the forwarding functions are implemented as hardware components using gateware developed in Verilog, while the software components that implement and control the PW-OSPF is written in C, C++, and Java.

Multiple circuits have been developed for the NetFPGA that include: a four-port network interface card, a four-port learning Ethernet switch, and a four-port IP Version 4 (IPv4) router. The router not only forwards packets, but also negotiates

routes using a simplified version of the OSPF protocol called PW-OSPF.

NetFPGA infrastructure was designed to be capable of handling packets at gigabit/second line rates. The modules provide a high degree of visibility into the operation of the device through use of memory-mapped registers that can be accessed by the CPU on the host. These registers are used to configure the mode of the hardware, set the sizes of buffers, and monitor the progress of packets through the pipeline.

A. The Reference Router Gateway

The reference hardware design gateway consists of a modular pipeline with multiple components. Components are modular and can be reused in other projects that build upon the function of the base router. By starting with a working reference design, the barrier to entry for implementing complex network designs is lowered for students and researchers, who are mainly interested in adding a feature rather than rebuilding a complete network system from scratch. The modular design allows the components to be easily shared, thus facilitating the development of a community-supported project.

B. Student Design Projects

Once the students completed assignments to implement the Ethernet switch and Internet router, they were given the opportunity to implement their own advanced feature that built upon the reference router during the last two weeks of the course. The NetFPGA allows packets to be created, modified, or dropped. Packets can be sent to and from the Ethernet ports and the host processor. Packets can also be forwarded between any of the four physical ports on the platform.

During the spring 2007 quarter of CS344, four advanced features from the 2007 class were NAT, time synchronization, multicast and packet capture. All four features were implemented within a two-week period.

The first team implemented NAT to allow multiple hosts to share a single IP address. The NAT device translated packet headers between a publicly-accessible IP address and a private address space for hosts behind the NAT device. The team added an extra lookup to identify the IP address/port translation to apply. Packets without a match in the hardware table were forwarded to software for further processing. Frequently-used translations were stored in a small table in hardware, while a much larger table of infrequently-used translations was stored in software.

The second team implemented a simplification of the IEEE 1588 protocol. This protocol provides a mechanism for precision time synchronization, on the order of tens of nanoseconds, of devices over Ethernet. Implementation required deterministic time stamping of packets at arrival/departure. No elastic buffers could be used between the wire and the point of time stamping. The team successfully demonstrated synchronization between multiple routers by showing submicrosecond time synchronization.

The third team implemented a multicast circuit that allowed a single packet to be copied to multiple destinations using hard-



Fig. 6. Photo from the CS344 final presentations in spring 2007.

ware. The team processed join and leave messages to add and remove the host from a multicast group table in hardware.

The fourth team built a circuit to capture packets transiting through the router and forward those packets to the host. The implementation compiled Berkeley packet filters (BPF) into a form suitable for the hardware. The hardware contained a set of circuits that matched parts of a packet simultaneously.

C. Presentation of Results of Advance Features

A photo of the CS344 final project presentations, which took place on June 5, 2007, is shown in Fig. 6. Students presented their projects to the class and to a panel of judges from industry, who were invited to evaluate the results of the student projects.

V. RELATED WORK

Other platforms have been used to implement some of the functionality that is provided by the NetFPGA platform. Similar to the NetFPGA, the RiceNIC utilized a VirtexII-Pro FPGA to implement the function of a network interface card [5]. Modules on the field-programmable port extender (FPX) platform have also been used in the classroom to implement Internet routing functions [6], [7].

Several research projects have used the NetFPGA platform to implement customized packet switches with advanced features. With Ethane, the NetFPGA was used to provide fine-grained control on how packets were switched in a LAN [8]. The NetFPGA has also been integrated into the DETER testbed where it was used to implement intrusion prevention functions [9].

ACKNOWLEDGMENT

The authors would like to thank the contributions of G. Watson, J. Weaver, and R. Raghuraman, who contributed to earlier versions of the NetFPGA platform. J. Luo contributed to the DMA.

REFERENCES

- [1] J. W. Lockwood, N. McKeown, G. Watson, G. Gibb, P. Hartke, J. Naous, R. Raghuraman, and J. Luo, "An open platform for gigabit-rate network switching and routing," in *Proc. Int. Conf. Microelectronic System Education*, San Diego, CA, Jun. 2007, pp. 160–161.

- [2] G. Watson, N. McKeown, and M. Casado, "NetFPGA: A tool for network research and education," in *Proc. Workshop Architecture Research Using FPGA Platforms*, Austin, TX, Feb. 2006, CD-ROM.
- [3] Build an Internet Router, Stanford University Computer Science 344, 2007 [Online]. Available: <http://cs344.stanford.edu>
- [4] M. Casado and N. McKeown, "The virtual network system," in *Proc. 36th SIGCSE Tech. Symp. Computer Science Education*, St. Louis, MO, 2005, pp. 76–80.
- [5] J. Shafer and S. Rixner, "A reconfigurable and programmable gigabit Ethernet network interface card," Dept. Electrical and Computer Engineering, Rice Univ., Houston, TX, 2006, Tech. Rep. TREE0611.
- [6] J. W. Lockwood, "Platform and methodology for teaching design of hardware modules in Internet routers and firewalls," in *Proc. Int. Conf. Microelectronic System Education*, Las Vegas, NV, Jul. 2001, pp. 56–57.
- [7] J. W. Lockwood, C. Neely, C. Zuver, and D. Lim, "Automated tools to implement and test Internet systems in reconfigurable hardware," *SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 3, pp. 103–110, Jul. 2003.
- [8] J. Luo, J. P. M. Casado, N. McKeown, and J. Lockwood, "Prototyping fast, simple, secure switches for Ethane," in *Proc. Hot Interconnects 15 Symp.*, Stanford, CA, Aug. 2007, pp. 73–82.
- [9] N. Weaver and V. Paxson, "Stress-testing a Gbps intrusion prevention device on DETER," presented at the DETER Symp., Arlington, VA, Jun. 2006 [Online]. Available: <http://www.isi.edu/deter>

Glen Gibb (S'03–M'04) received the M.S. degree from Stanford University, Stanford, CA, in 2005. He is working towards the Ph.D. degree in electrical engineering at Stanford.

His research interests include switch and router design and implementation, FPGA architectures and digital logic implementation.

John W. Lockwood (S'88–M'96) received the B.S., M.S., and Ph.D. degrees from the University of Illinois, Urbana-Champaign, in 1991, 1993, and 1995, respectively.

He is currently a Consulting Associate Professor in the Department of Electrical Engineering, Stanford University, Stanford, CA, where he leads the alpha and beta release program and organizes the worldwide NetFPGA tutorial program. Prior to joining Stanford, Lockwood worked at Washington University, St. Louis, MO, where he was granted tenure in the Department of Computer Science in 2006. At Washington University, Lockwood led the Reconfigurable Network Group (RNG) and developed the field-programmable port extender (FPX) to enable rapid prototype of extensible network modules in field-programmable gate array (FPGA) hardware. He served as the Principal Investigator on Grants for the National Science Foundation, Xilinx, Altera, Nortel Networks, Rockwell Collins, and Boeing. He has worked in industry for AT&T Bell Laboratories, IBM, Science Applications International Corporation (SAIC), the National

Center for Supercomputing Applications (NCSA), and served as Cofounder of Global Velocity, a networking company focused on high-speed data security. His research interests include prototype of reconfigurable hardware, implementation of systems that provide Internet security, and content processing technologies to enable development of new types of networks.

Dr. Lockwood is a member of the Association for Computing Machinery, Tau Beta Pi, and Eta Kappa Nu.

Jad Naous received the B.Eng. degree in computer engineering from McGill University, Montreal, QC, Canada, in 2005, and the M.S.E.E. degree from Stanford University, Stanford, CA, in 2007. He is working towards the Ph.D. degree in electrical engineering at Stanford.

In 2006, he was a Graduate Intern for Sun Microsystems Labs, where he worked on the next generation switch project. In 2007, he joined Agilent Technologies Labs as a Graduate Intern where he helped implement special devices for the IEEE1588 Precision Time Protocol using NetFPGA.

Paul Hartke is working towards the Ph.D. degree at Stanford University, Stanford, CA.

He served as a Teaching Fellow Instructor for the courses EE183: Advanced Logic Design Laboratory and EE108A: Digital Systems I at Stanford. He is currently at Xilinx, San Jose, CA, working on networking solutions based on the Virtex II Pro FPGA family. Previously, he worked as an Intern for Archway Digital Solutions, Abrizio, and Hewlett-Packard Laboratories.

Nick McKeown (S'91–M'95–F'05) received the B.E. degree from the University of Leeds, Leeds, U.K., in 1986, and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 1992 and 1995, respectively.

He is a Professor of Electrical Engineering and Computer Science at Stanford University, Stanford, CA, where he is Faculty Director of the Clean Slate Program. From 1986 to 1989, he worked for Hewlett-Packard Labs, Bristol, U.K. In 1995, he helped architect Cisco's GSR 12000 router. In 1997, he cofounded Abrizio, Inc. (later acquired by PMC-Sierra), where he was CTO. He was Co-founder and CEO of Nemo Systems, which is now part of Cisco Systems. His research interests include the architecture of the future Internet, and tools and platforms for networking teaching and research.

Dr. McKeown is the STMicroelectronics faculty scholar, the Robert Noyce faculty fellow, a fellow of the Powell Foundation and the Alfred P. Sloan Foundation, and recipient of a career award from the National Science Foundation. He is also a fellow of the Royal Academy of Engineering (U.K.) and the Association for Computing Machinery. In 2005, he was awarded the British Computer Society Lovelace Medal. In 2000, he received the IEEE Rice Award for the best paper in communications theory.