

FINAL REPORT FOR AWARD # 9875177

Kenneth J Christensen ; *U of South Florida*
CAREER: Performance Evaluation of Gigabit Ethernet Networks, A Systems and
Experimental Approach

Participant Individuals:

Graduate student(s) : Kenji Yoshigoe; Zornitza Genova; Aamir Shaikh
Research Experience for Undergraduates(s) : Zane Reynolds
Graduate student(s) : Ahmed Aslam
Pre-college teacher(s) : Tahvia Shaw
Research Experience for Undergraduates(s) : Ashley Hopkins

Partner Organizations:

Cisco Systems, Inc.: Financial Support

We have submitted a Cisco University Research proposal to continue the work that Ahmed started during his one semester of funding from this CAREER program. The Cisco proposal addresses modeling of Parallel Packet Switches (PPS). The requested funding is \$26K to support Ahmed for one year.

The proposal is currently under review by Cisco.

Other collaborators:

I have collaborated with Alan George (faculty) at the University of Florida, Neil Gunther (consultant) at Performance Dynamics, Aju Jacob (student) at the University of Florida, Joe Rogers (graduate student and now staff) at the University of South Florida, Allen Roginsky (mathematician) at IBM Corporation, and Marcel Waldvogel (researcher) at IBM Corporation. Neil, Aju, Joe, Allen, and I have worked on collaborative papers. Involved in this collaboration have been the students funded by this CAREER grant. Alan and I established a new workshop on high-speed local networks. Marcel and I co-edited a special issue of Computer Communications on Terabit switching.

Activities and findings:

Research and Education Activities:

This 1999 to 2003 NSF CAREER program addressed the 'network of the future' where both fixed-size cell and variable-length packet networks will co-exist. The educational emphasis was on developing undergraduate research involvement and the development of two new courses (one in capacity planning and performance evaluation and the other a projects course for a department proposed MS in Information Systems degree program). The three stated research objectives of this program were to 1) characterize traffic from representative gigabit applications, 2) advance new algorithms for achieving high-throughput in input-buffered only Ethernet switches and full-duplex repeaters,

and 3) study backwards compatible CSMA methods for large-span, half-duplex Ethernets. Based on referee comments and the recommendation of the NSF program director at the time of the grant, the third objective was refocused to investigate gigabit-speed switching at higher layers.

Research activities:

Activities completed for the three objectives are:

Objective 1 - Final results have been obtained on traffic characterization of traces taken on the Gigabit Ethernet test bed funded by this grant. Work was also completed on comparing Internet1 and Internet2 traffic on the USF gigabit backbone LAN. This work entailed tracing traffic from the USF Internet1 and Internet2 access points. The trace files are available on the project website (or by direct request) for use by other researchers.

Objective 2 - Simulation models have been built for Input Queued (IQ), Combined Input and Crossbar Queued (CICQ), and Virtual Input Queued Parallel Packet Switch (VIQ PPS) switches. The project web site contains all of the switch models and also includes models of other switch architectures (e.g., output queued switch) used for comparison.

Objective 3 - The concept of a URL router was investigated. A method of using CRC32 for compressing URL lists into digests was investigated and evaluated. New aggressive hashing methods were studied in the context of application-layer switching. Trace files and simulation models are available directly from Zornitza Genova (funded by this grant as a PhD student) who is now an Assistant Professor at the University of North Florida.

Education activities:

For the educational component, two REU students and one RET teacher were supported. In addition, the Internet-accessible tools page was developed and continues to grow with the addition of new tools. The tools page (<http://www.csee.usf.edu/~christen/tools/toolpage.html>) continues to get about 800 hits per month.

Two new courses were developed at the undergraduate level, they are:

* CIS 4930 - Capacity Planning
(<http://www.csee.usf.edu/~christen/class9/class9.html>)

* CIS 4910 - Senior Project
(<http://www.csee.usf.edu/~christen/class6/class6.html>)

The stated goal in the proposal was to develop the project course at a graduate level as part of a planned MS in Information Systems. This graduate track did not materialize, so the projects course was aimed at undergraduate student. The Capacity Planning course introduces undergraduate students to an analytical aspect of computer science. Several students commented that this was the only class in their curriculum where they used their math training. The Senior Project class was retooled from a independent-study style format to one of student teams working on industry-contributed projects. Relationships

with local industry have been fostered (we have worked with over 12 central Florida companies in the second year of this course). This course is very popular with students due to its 'real world' focus that helps students transition into industry. Formal evaluation of this new projects course is ongoing as part of the exit interviews of graduating students.

Findings:

The findings of this CAREER program have made a significant contribution to the direction of future single-stage switch designs. These contributions are described in the Contributions within Discipline section of this report. The findings are documented in 20 publications presented in the list of references in this section. The key findings by objective were:

Findings for objective 1:

We discovered that the traffic characteristics of bulk data transfer are dependent on operating system (Linux and Windows were compared), whether memory-to-memory or disk-to-disk, and data rate of the Ethernet link [18]. In general, disk-to-disk and higher data rate resulted in the greatest burstiness of traffic. Furthermore, it was found that shaping only 1% of the shortest interdeparture times could significantly reduce packet loss and queueing delay in the network. From these findings we designed and implemented a sockets shim layer where shaping of traffic is done within the system and does not require leaky bucket or other shaping hardware on the adapter Ethernet adapters typically do not have shaping hardware in the same way that ATM adapters do. The traffic shaping was shown to reduce queueing delay. To study Internet1 versus Internet2 traffic characteristics we collected Cisco NetFlow traces from each network [14]. These NetFlow traces were then used to drive a fluid-flow model to study queueing behavior. We believe that this was the first use of NetFlow traces to drive a fluid-flow model. The findings showed that traffic on the Internet1 and Internet2, despite a roughly similar application mix, had very different characteristics. For the Internet2 traffic as input (but not for the Internet1 traffic), increasing buffer size in a modeled single server queue driven with trace traffic as the input had little effect on packet loss. This is indicative of heavy-tailed traffic. We also found that most of the bit loss (caused by packet loss at the finite-buffer queue) for a full set of flows was caused by the largest 1% of flows by size. Thus, again a 'top 1%' effect was seen, but this time in flows and not interarrival times. A first order shaping of these 1% flows in our model resulted in a significant reduction of bit loss. This has significant implications to traffic shaping in large networks [14]. We specifically addressed traffic shaping for bandwidth allocated services for gigabit Ethernet in [15]. We proposed an adapter-implemented, byte-based leaky bucket approach - called RATE as a counterpart to existing 802.3 PASUE - suitable for emerging Ethernet-in-the-first-mile applications. The students involved with these findings were Zane Reynolds (undergraduate REU student), Aamir Shaikh (MS student), and Kenji Yoshigoe (PhD student). Zane is now in graduate school at USF, Aamir is employed in industry, and Kenji expects to complete his PhD in summer 2004.

Findings for objective 2:

Buffered crossbar switches are not a new architecture - they go back to the mid-1980s. However, it is only recently with continued increases in VLSI density and the use of Virtual Output Queues (VOQs) at the input that such switches have become feasible to implement. Our findings were the first to show the potential of the Combined Input and Crossbar Queued (CICQ) switch as the next generation single stage switch [7, 16, 17, 19]. In existing Input Queued (IQ) switches the input and output ports are tightly coupled so that switch matrix scheduling (e.g., iSLIP) can be achieved. In the CICQ switch this tight coupling is not needed thus reducing complexity with the possibility of scaling-up to much higher speeds. A very low-cost and simple 24-port, 10-Gbps CICQ switch was designed and evaluated in [8]. Also in [8] is the design and evaluation of a very fast Masked Priority Encoder (MPE) round-robin poller. This MPE design was shown to be faster than any existing poller design when implemented in FPGA technology. A key problem with both existing IQ and the CICQ switch architectures is instability for unbalanced inputs. This instability is typically remedied by using internal speed-up. We proposed and evaluated a simple threshold and bursting method which achieved stability and did not require speed-up [5]. This method was analyzed for the CICQ switch. The CICQ switch exhibits fewer stability problems than a purely IQ switch. This is very significant. We also further investigated round robin poller designs resulting in the Overlapped Round Robin (ORR) design that is completely speed scalable [6]. The CICQ switch architecture is amenable to natively switching variable length packets. This was a key motivation in exploring the CICQ architecture in the first place. The variable length packet properties of this switch were explored in [1, 2]. Fairness is an issue when dealing with variable length packets. A new method of block forwarding (where an internally transferred block can cross packet boundaries) was proposed to ensure fairness [2, 7]. The student involved with these findings was Kenji Yoshigoe (PhD student). Kenji is continuing his research in CICQ switches to study removing all feedback from the switch crossbar to the input ports. This employs rate control internal to the switch (in the input ports) and we hope will allow for another order of magnitude speed scaling of the CICQ switch architecture.

Looking beyond the CICQ switch, we believe that a time will come where link speeds will greatly exceed memory speeds. This will make IQ switches of any kind infeasible to implement and motivates the need for Parallel Packet Switches (PPS). The Virtual Input Queued (VIQ) PPQ was proposed and evaluated [3, 11]. This switch architecture was shown to have both better performance (lower packet delay) and better implementation feasibility than existing PPS designs. The student involved with these findings was Ahmed Aslam (PhD student). Ahmed continues his work in PPS architectures. Ahmed is exploring the use of redundant memories to support overlapped read and writes from a switch buffer.

Findings for objective 3:

For objective 3 we investigated methods for high-speed switching at the application layer. In particular, we focused on URL switching for Content Distribution Networks (CDNs). At the application layer the

number of addresses (URLs in our case) and their dynamic nature is significantly greater than at the IP layer. Thus, methods for efficient sharing, storing, and look-up of application layer addresses (URLs) needed to be investigated [20]. The idea of the URL router as a key component of a CDN was proposed [4]. The use CRC32 as a signature for a URL was investigated and found to be better than existing Bloom filter approaches for updating of digests, storage efficiency, and look-up time [10, 13]. New hashing methods to exploit the CRC32 URL signatures were explored in [4]. The new aggressive hashing algorithm, a type of move-to-front hashing method, was proposed and found to perform very well in the look-up of CRC32 URL signatures. The student involved with this objective was Zornita Genova Prodanoff (who finished her PhD in August 2003). Zornitza continues work in this area as an Assistant Professor at the University of North Florida.

Other findings:

Simulation was the major tool used for this CAREER program. In the proposal we planned for tool development. A remote execution tool was developed [12] and enhanced [9]. The student involved with this objective was Ashley Hopkins (undergraduate REU student).

References:

- [1] K. Christensen, K. Yoshigoe, A. Roginsky, and N. Gunther, 'Performance Evaluation of Packet-to-Cell Segmentation Schemes in Input Buffered Packet Switches,' submitted to the IEEE International Conference on Communications in September 2003.
- [2] K. Yoshigoe, K. Christensen, and A. Roginsky, 'Performance Evaluation of New Scheduling Methods for the RR/RR CICQ Switch,' submitted to Computer Communications in July 2003.
- [3] A. Aslam and K. Christensen, 'A Parallel Packet Switch with Multiplexors containing Virtual Input Queues,' submitted to Computer Communications in April 2003.
- [4] Z. Genova and K. Christensen, 'Managing Routing Tables for URL Routers in Content Distribution Networks,' to appear in the International Journal of Network Management.
- [5] N. Gunther, K. Christensen, and K. Yoshigoe, 'Characterization of the Burst Stabilization Protocol for the RR/RR CICQ Switch,' Proceedings of the IEEE Conference on Local Computer Networks, pp. 260-269, October 2003.
- [6] K. Yoshigoe, K. Christensen, and A. Roginsky, 'Design of a High-Speed Overlapped Round Robin (ORR) Arbiter,' Proceedings of the IEEE Conference on Local Computer Networks, pp. 638-639, October 2003.
- [7] K. Yoshigoe and K. Christensen, 'An Evolution to Crossbar Switches with Buffered Cross Points,' IEEE Network, Vol. 17, No. 5, pp. 48-56, September-October 2003.

- [8] K. Yoshigoe, K. Christensen, and A. Jacob, 'The RR/RR CICQ Switch: Hardware Design for 10-Gbps Link Speed,' Proceedings of the IEEE International Performance, Computing, and Communications Conference, pp. 481-485, April 2003
- [9] A. Hopkins, 'Remote++: A Script for Automatic Remote Distribution of Programs on Windows Computers,' Proceedings of the 41st ACM Southeast Regional Conference, March 2003.
- [10] Z. Genova and K. Christensen, 'Efficient Summarization of URLs using CRC32 for Implementing URL Switching,' Proceedings of the IEEE Conference on Local Computer Networks, pp. 343-344, November 2002.
- [11] A. Aslam and K. Christensen, 'Parallel Packet Switching using Multiplexors with Virtual Input Queues,' Proceedings of the IEEE Conference on Local Computer Networks (LCN), pp. 270-277, November 2002.
- [12] K. Christensen, 'REMOTE: A Tool for Automatic Remote Execution of CSIM18 Simulation Models,' Proceedings of the Annual Simulation Symposium, pp. 134-142, April 2002.
- [13] Z. Genova and K. Christensen, 'Using Signatures to Improve URL Routing,' Proceedings of the IEEE 2002 International Performance, Computing, and Communications Conference, pp. 45-52, April 2002.
- [14] J. Rogers and K. Christensen, 'A Fluid-Flow Characterization of Internet1 and Internet2 Traffic,' Proceedings of the IEEE Conference on Local Computer Networks, pp. 509-513, November 2001.
- [15] K. Yoshigoe and K. Christensen, 'RATE Control for Bandwidth Allocated Services in IEEE 802.3 Ethernet,' Proceedings of the IEEE Conference on Local Computer Networks, pp. 446-453, November 2001.
- [16] K. Christensen, 'Design and Evaluation of a Parallel-Polled Virtual Output Queued Switch,' Proceedings of the IEEE International Conference on Communications, pp. 112-116, June 2001.
- [17] K. Yoshigoe and K. Christensen, 'A Parallel-Polled Virtual Output Queued Switch with a Buffered Crossbar,' Proceedings of the IEEE Workshop on High Performance Switching and Routing, pp. 271-275, May 2001.
- [18] A. Shaikh and K. Christensen, 'Traffic Characteristics of Bulk Data Transfer using TCP/IP over Gigabit Ethernet,' Proceedings of the IEEE International Performance, Computing, and Communications Conference, pp. 103-111, April 2001.
- [19] K. Christensen, 'A Parallel-Polled Virtual Output Queued (PP-VOQ) Switch,' IEE Electronics Letters, Vol. 36, No. 22, pp. 1902-1903, October 26, 2000.
- [20] Z. Genova and K. Christensen, 'Challenges in URL Switching for Implementing Globally Distributed Web Sites,' Proceedings of the Workshop on Scalable Web Services, pp. 89-94, August 2000.

Training and Development:

Graduate student mentoring was a key part of this CAREER program. However, undergraduate and K-12 teacher mentoring also played a key role. Three graduate students (Aamir, Ahmed, Kenji, and Zornitza) have been funded through this grant for varying amounts of time. Two undergraduate REU students were funded through supplements (Ashley and Zane) and one K-12 teacher was funded through a supplement (Tahvia Shaw).

* Aamir completed his MS in May 2001 and is currently employed in industry.

* Ahmed is expecting to complete his PhD in fall 2004 and will be seeking a position in industry. Ahmed presented his work at one conference (IEEE LCN).

* Kenji is expecting to complete his PhD in summer 2004 and is applying for an academic position in the USA. He has already had a campus interview with West Georgia State University and a telephone interview with University of Hawaii at Hilo. Kenji presented his work at two conferences (IEEE LCN and IEEE IPCCC).

* Zornitza has completed her PhD and is now an Assistant Professor at the University of North Florida. Zornitza has presented her work at three conferences (SWS workshop, IEEE LCN, and IEEE IPCCC).

* Ashley completed her BS in 2003 and is currently employed by Harris Corporation in Melbourne, Florida. Ashley presented her work at one conference (ACM SE regional conference).

* Zane completed his BS in 2002 and is currently a PhD student at the University of South Florida.

* Tahvia is middle school teacher at Terrace Community School, a public charter school. Tahvia's RET experience focused on ways of bringing computer science into her fifth grade classroom and using modern computer software tools (StudyWorks) for teaching mathematics. Tahvia developed several lesson plans which are in active use at her school and are generally available to others via her website. The middle students have enjoyed the new material and accompanying hands-on lessons.

Outreach Activities:

Outreach focused on two areas: 1) participating in the yearly College of Engineering Expo where labs are opened-up to the general public, and 2) having K-12 involvement via an RET supplement.

The Information Systems Laboratory (that houses the gigabit test bed funded by this grant) was opened to the general public for the annual Engineering Expo for all four years of the grant. Approximately 200 middle and high school students and about 30 non-students (adults) toured the lab each year. We set-up a PC with an online quiz on the topic of 'What is a gigabit?'. On completion of the quiz visitors got a picture of themselves, taken by a webcam, printed on the labs

homepage. This was very popular with the middle and high school students.

The RET supplement supported a school teacher (Tahvia Shaw) from the Terrace Community School middle school. Included in the supplement funding was money to purchase 25 copies StudyWorks software (StudyWorks is a scaled-down Mathcad) for the school. The school is now using StudyWorks and the lesson plans that Tahvia developed while on her RET assignment at the university. This serves to encourage the interest of middle schools in science, math, and technology.

The principal investigator has been a Florida state science fair judge for three consecutive years during this CAREER program.

Journal Publications:

A. Shaikh and K. Christensen, "Traffic Characteristics of TCP/IP Bulk Data Transfer in Gigabit Ethernet", *Proceedings of the IEEE 2001 International Performance, Computing, and Communications Conference*, vol. 1, (2001), p. 103. Published

K. Christensen, "Design and Evaluation of a Parallel-Polled Virtual Output Queued Switch", *IEEE 2001 International Conference on Communications*, vol. 1, (2001), p. 112. Published

Z. Genova and K. Christensen, "Challenges in URL Switching for Implementing Globally Distributed Web Sites", *Proceedings of the Workshop on Scalable Web Services*, vol. August, (2000), p. 89. Published

K. Christensen, "A Parallel-Polled Virtual Output Queued (PP-VOQ) Switch", *IEE Electronics Letters*, vol. 36, (2000), p. 1902. Published

K. Yoshigoe and K. Christensen, "A Parallel-Polled Virtual Output Queued Switch with a Buffered Crossbar", *Proceedings of the IEEE Workshop on High Performance Switching and Routing*, vol. 1, (2001), p. 271. Published

K. Christensen, "REMOTE: A Tool for Automatic Remote Execution of CSIM18 Simulation Models", *Proceedings of the 35th Annual Simulation Symposium*, vol. , (2002), p. 134. Published

Z. Genova and K. Christensen, "Using Signatures to Improve URL Routing", *Proceedings of the IEEE 2002 International Performance, Computing, and Communications Conference*, vol. , (2002), p. 45. Published

K. Yoshigoe and K. Christensen, "RATE Control for Bandwidth Allocated Services in IEEE 802.3 Ethernet", *Proceedings of the IEEE 26th Conference on Local Computer Networks*, vol. , (2001), p. 446. Published

J. Rogers and K. Christensen, "A Fluid-Flow Characterization of Internet1 and Internet2 Traffic", *Proceedings of the IEEE 26th Conference on Local Computer Networks*, vol. , (2001), p. 509. Published

K. Yoshigoe and K. Christensen, "An Evolution to Crossbar Switches with Buffered Cross Points", *IEEE Network Magazine*, vol. 17, (2003), p. 48. Published

A. Aslam and K. Christensen, "Parallel Packet Switching using Multiplexors with Virtual Input Queues", *Proceedings of the 27th IEEE Conference on Local Computer Networks (LCN)*, vol. , (2002), p. 270. Published

Z. Genova and K. Christensen, "Efficient Summarization of URLs using CRC32 for Implementing URL Switching", *Proceedings of the 27th IEEE Conference on Local Computer Networks (LCN)*, vol. , (2002), p. 343. Published

K. Yoshigoe, K. Christensen, and A. Jacob, "The RR/RR CICQ Switch: Hardware Design for 10-Gbps Link Speed", *Proceedings of the IEEE 2003 International Performance, Computing, and Communications Conference*, vol. , (2003), p. 481. Published

A. Aslam and K. Christensen, "A Parallel Packet Switch with Multiplexors containing Virtual Input Queues", *Computer Communications*, vol. , (2003), p. . Submitted

K. Yoshigoe, K. Christensen, and A. Roginsky, "Design of A High-Speed Overlapped Round Robin (ORR) Arbiter", *High Speed Local Networks Workshop at the 28th IEEE Conference on Local Computer Networks (LCN)*, vol. , (2003), p. 638. Published

N. Gunther, K. Christensen, and K. Yoshigoe, "Characterization of the Burst Stabilization Protocol for the RR/RR CICQ Switch", *IEEE Conference on Local Computer Networks (LCN)*, vol. , (2003), p. 260. Published

Z. Genova and K. Christensen, "Managing Routing Tables for URL Routers in Content Distribution Networks", *International Journal of Network Management*, vol. , (), p. . Accepted

K. Yoshigoe, K. Christensen, and A. Roginsky, "Performance Evaluation of New Scheduling Methods for the RR/RR CICQ Switch", *Computer Communications*, vol. , (), p. . Submitted

K. Christensen, K. Yoshigoe, A. Roginsky, and N. Gunther, "Performance Evaluation of Packet-to-Cell Segmentation Scheme in Input Buffered Packet Switches", *International Conference on Communications*, vol. , (), p. . Submitted

Other Specific Products:

Special issue of a journal

Guest editor for special issue on Issues and Trends in Terabit Switching in *Computer Communications* (with Marcel Waldvogel from IBM-Zurich), Vol. 25, No. 6, April 2002. The journal issue appeared in 2002.

Software (or netware)

The REMOTE tool was developed to allow for automatic remote execution of console mode programs (such as CSIM18 simulation models) on network-connected Windows PC's. The REMOTE tool has been updated to the REMOTE++ tool. This update was completed by an REU student (Ashley Hopkins).

The relationship of the remote simulation tool to the proposal is as follows. The development of this tool was described in the proposal for laboratory infrastructure as, "Another tool that will be developed will be a means of parallel independent execution of discrete-event simulation programs building on the work in [41]. This tool will enable network-attached computers to be used to complete simulation studies faster than could a single computer. Peer researcher and student contributions will be solicited for these research tools." We have developed (and now improved) this tool and used it to improve the efficiency (i.e., execution time) of simulation studies on new switch architectures. This tool is generally available. No statistics have been compiled on who is using this tool, but several email inquiries

have been received and responded to.

This software is available on a website (<http://www.csee.usf.edu/~christen/career/remote.html>). A short description of this tool has appeared in the software tools section of the September issue of IEEE Network Magazine.

Software (or netware)

The Christensen "tools page" is a collection of C source code tools or performance modeling and evaluation. Included are tools to solve probability matrices, characterize time series, generate time series with given distributions, simulation models of queueing system, TCP/IP sockets programs, and others. The unique feature of these tools is their source code nature - they are intended to be used to teach concepts (e.g., how to iteratively solve a probability matrix, how to generate exponentially distributed random variates, how to estimate the Hurst parameter using R/S analysis, and so on). The tools page generates over 800 hits per month and results in several email discussions every month. It appears that graduate students and instructors around the world are using (and, in some cases, finding bugs in) these tools. The use is both in the classroom and for graduate-level research (e.g., as components in larger tools). The relationship to this project is that many of the tools were used in the research and teaching components (of this project).

The software is available on a website (<http://www.csee.usf.edu/~christen/tools/toolpage.html>).

Internet Dissemination:

<http://www.csee.usf.edu/~christen/career/main.html>
<http://www.csee.usf.edu/~christen/tools/toolpage.html>

These two URLs are for 1) the project homepage, and 2) tools page. All models, trace files, etc. are available on the project homepage. The tools page is described elsewhere in this report.

Contributions:

Contributions within Discipline:

The contributions of this CAREER program have made an impact to the discipline. This impact can be measured, in part, by how other researchers reference and use our work.

Contributions from objective 1:

The contributions from objective 1 are to show that higher data rates result in greater burstiness of the carried traffic irrespective of the applications generating the traffic. We also showed that a '1% effect' exists where either 1% of the interarrival times between packets or 1% of flows cause the majority of packet losses from queue overflow. This has implications to traffic shaping (i.e., that only a very small amount of traffic needs to be shaped to achieve significant

reductions in packet loss).

Contributions from objective 2:

The CICQ switch overcomes many limitations in IQ switches that use iterated maximal-match scheduling methods. Our work has been cited in a recent IBM Technical Report (F. Abel et al. 'A Four-Terabit Single-Stage Packet Switch...') as 'The renewed interest in CICQ is demonstrated in [4]-[9]...' where [9] is our May 2001 HPSR paper. This May 2001 HPSR paper was the first paper to show that CICQ switches can have lower delay than well known iSLIP scheduled IQ switches. Our September 2003 IEEE Network magazine article titled 'An Evolution to Crossbar Switches with Virtual Output Queueing and Buffer Cross Points' will be, we expect, a seminal paper defining the CICQ switch as the future architecture for single-stage switches. Of significant interest is to see that Nick McKeown of iSLIP fame is now also focusing on CICQ architectures (Stanford HPNG Technical Report TR03-HPNG-061501 - <http://tiny-tera.stanford.edu/~nickm/papers/index.html>). Our contribution was to be among the first to realize the potential of CICQ switches and to begin the investigations into understanding the performance and feasibility of this architecture. We were also the first to understand the potential for the CICQ switch to be able to natively forward variable length packets. The contribution of this are to reduce the needed internal speed-up by two-fold and to eliminate the hardware needed for segmentation and reassembly of packets. Manolis Katevenis at the University of Crete is now exploring variable length packet switching by CICQ switches (and cites our work as one of only two others studying variable length packets in CICQ switches). Our work has been referenced in at least seven publications already. Direct contributions also include designs for faster round robin polling circuits. We hope and expect to see these designs implemented in future switches that use round robin polling.

The PPS VIQ switch architecture was only recently published (in the IEEE LCN conference in November 2002) and has not yet been referenced in other publications. This work has introduced a new use of Virtual Input Queues (VIQs) as a re-sequencing buffer. Although VIQs have already been used in PMC-Sierra Enhanced TT1 chip set product (to implement flow-control mechanism between Line-cards and core switch), no prior work has considered VIQs as a re-sequencing buffer. We showed that re-sequencing can be done simpler and faster using VIQs in the re-sequencing buffer. This will be helpful to industry to design a multi-terabit packet switch using off-the-shelf existing switches cost-effectively and quickly. The VIQ PPS is both feasible to implement (we believe this to be the case -- to test this we have applied for a Cisco University Research Grant to model the VIQ PPS at the hardware level using FPGA technology) and has better performance than any existing implementable PPS architecture.

Contributions from objective 3:

The contributions for the work in objective 3 are less obvious in the published literature. To date, there have been no citations of the only very recently published work from this objective. The key contribution from the work in objective 3 is to define an efficient way to route in CDNs. It is remarkable that the simple CRC32 can be

used as a very effective signature for variable length URLs. The new aggressive hashing algorithm is a contribution to the discipline.

Contributions to Other Disciplines:

Our work in round robin poller designs has application well beyond packet switch architectures. Other areas - including large-scale industrial engineering applications - use round robin polling for scheduling of tasks.

Contributions to Education and Human Resources:

The participants in this CAREER program were very diverse. This includes US women students, and international students from India, Pakistan, and Japan. One significant education and human resource contribution of this program has been the completion of a woman PhD (Zornita) who has taken an academic position at a regional university (the University of North Florida). This serves to strengthen the position of women within Computer Science and Engineering where women are an under-represented group and serve as a role model for future women students both at the University of South Florida and the University of North Florida. This program has also provided research experiences to two undergraduate students (Ashley and Zane) and one K-12 teacher (Tahvia). Zane has continued his studies at the graduate level. The work that Tahvia completed during her RET assignment at the university lives on in the form of lessons and software tools (StudyWorks) at her middle school.

Contributions to Resources for Science and Technology:

The test bed (which consists of six PCs attached to a Gigabit switch, which in turn is directly connected to the USF gigabit backbone) funded from this CAREER program is now a resource for students in the Department of Computer Science and Engineering. The test bed is used by students for class projects and research. Several teams of undergraduate senior project students have already used the test bed.

The High-Speed Local Networks (HSLN) workshop was initiated in the third year of this CAREER program in partnership with Alan George at the University of Florida. This workshop was started to bring the high-speed local area networking community together. The workshop runs under the umbrella of the IEEE Local Computer Networks (LCN) conference. Papers at the HSLN workshop have focused on cluster interconnects and switch architectures. This workshop attracted about 25 papers in both 2002 and 2003 for a single track of 12 accepted full papers. A poster session was also part of the workshop. For 2004 this workshop was handed-off to a new group of general and program chairs. This new yearly workshop is a resource for science and technology.

Contributions Beyond Science and Engineering:

The benefits from this research address the scalability of the Internet at its key bottleneck - the core routers. A better Internet is a contribution to resources beyond science and engineering.

Categories for which nothing is reported:

Products: Book or other one-time publication