Fraser Research

A not for profit company

Annual Report 2004

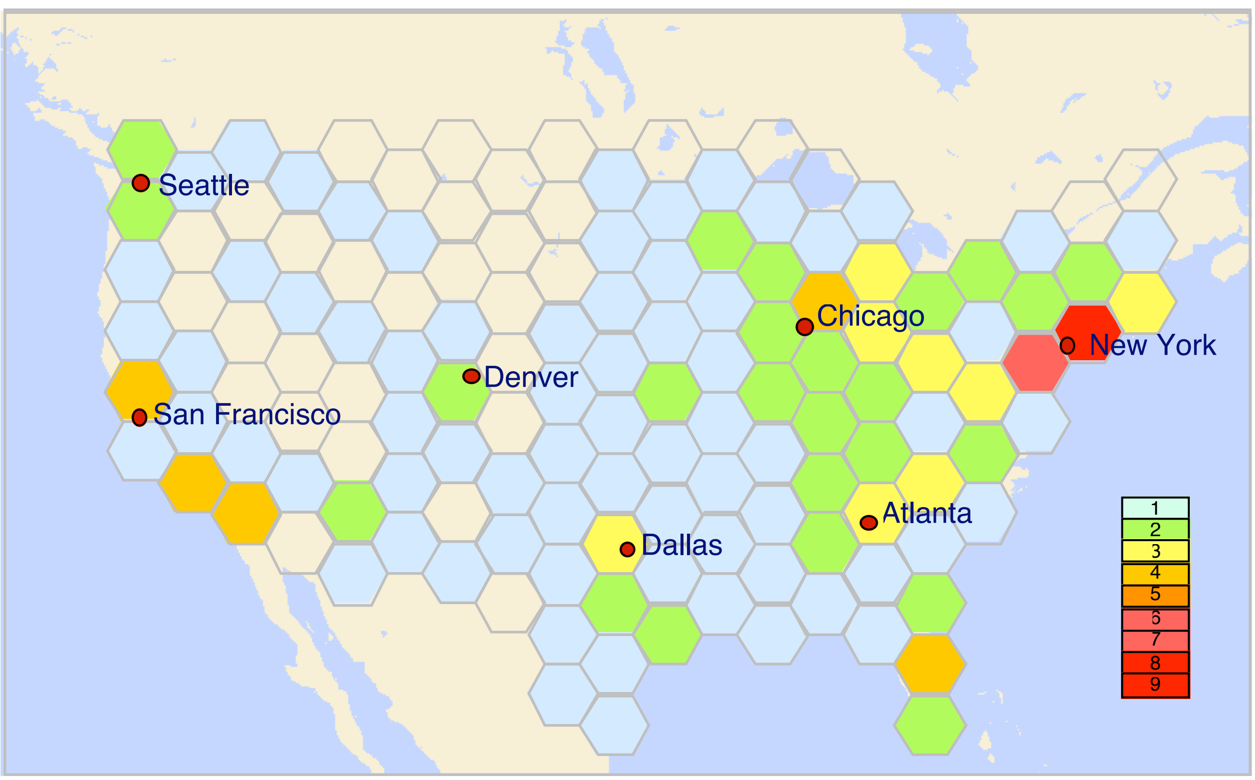

Major support provided by The National Science Foundation

100 x 100 *Award # ANI 0331588: 100 Mb/sec for 100 Million Households*

FAST Copper *Award # CNS-0427687 'FAST Copper': Dynamic Optimization of Resources in Frequency, Amplitude, Space, and Time for Broadband Access Networks*

Fraser Research
182 Nassau Street, Suite 301
Princeton, NJ 08542

Voice: 609 497 7337
Fax: 609 497 7335
www.FraserResearch.org

Contents

Letter from the President

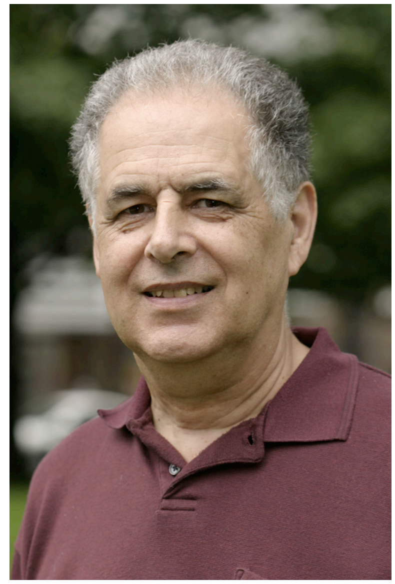

One hundred and forty years ago, the industrial revolution followed a course which has its echo in events that are taking place today: a proliferation of new technology with broad implications, reinvention of established processes, creation of new infrastructure, innovations in retailing ... From this comparison it is possible to see how revolutionary change founded in technology makes its way through society. In the early days progress may be paced by scientific development but soon it is the insights of a new generation that stimulate the introduction of new applications, the development of new markets, and the formulation of new challenges which in turn drive technical progress.

Fraser Research aspires to be part of that cycle - sensitive to the needs of society yet possessing the wit to identify and address fundamental problems. The critical challenge on which Fraser Research is presently focused lies within the communications infrastructure. Today's Internet is an outstanding achievement, but the principles of its design are showing their age. Technical progress in communications, electronics and computing have made extraordinary advances in the past three decades, and much has been learned about how a global integrated services network will be used. Therefore, even as massive infrastructure deployments are now being made, it is important to look forward towards the next generation network that will emerge in response to user needs and technical opportunities that are now apparent.

The most evident need is for a robust communications service with solid protection against abuse of privacy and challenges to security. Robust, reliable and always available communications will become critical to the prosperity and welfare of nations. It will be used by billions of people and it must be useable by everyone regardless of their level of education. Network operations and maintenance will require hundreds of thousands of people. So even though the network will become the world's largest machine where all parts must work together, it must not be so complex that only a select few can maintain it. As is usual in engineering, simplicity will be among the most difficult of goals to achieve.

In the pages of this, our first annual report, the reader will find a sampling of the problems that we have begun to address. While the four research activities described here relate to the lowest layers of network architecture, our work includes a broad perspective on networking and related aspects of computing. We believe that it is important to revisit in detail the fundamental tenets of packet switched communications because of weaknesses that we perceive in the present deployments and the rapidly diminishing opportunity for change at that level. What makes this work urgent is the observation that the present Internet architecture makes evolution more difficult than it should be. Too much detailed knowledge of how the network works is embedded in host software.

We are particularly pleased with the summer research program for graduate student interns. This three-month intensive program challenges the interns to complete the equivalent of a Master's thesis on a topic of interest to the intern and within the framework of the Institute's agenda. This program contributes substantially towards the work of the Institute while broadening the education of students with an interest in communications. It also strengthens our relationship with the academic community.

Fraser Research is an institute founded in the public interest with an understanding that new technology does not fly of its own accord from the pages of a research report into the market place. In the coming years we shall feel our way forward in conjunction with partners in industry and academia. Today we are fortunate to have a research relationship with six of the world's leading universities and we have started to develop industry relationships in certain areas of technology.

The past 18 months have been difficult for us as they have been for the entire communications industry. For the short term we have sought public funds, and are most grateful to the National Science Foundation for their support. Going forward we shall broaden our approach to include other sources of funds consistent with our work in the public interest. As progress is made, we will seek paths through which our work benefits society while contributing to an endowment that will sustain the research program for years to come.

Sandy Fraser
President


Thomas Sweet
ICE CREAM & CHOCOLATE
Now Hiring For The FALL SEASON!

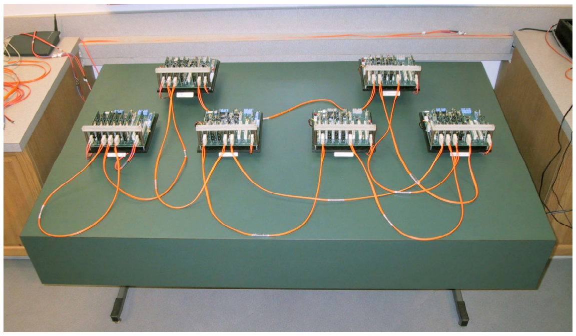

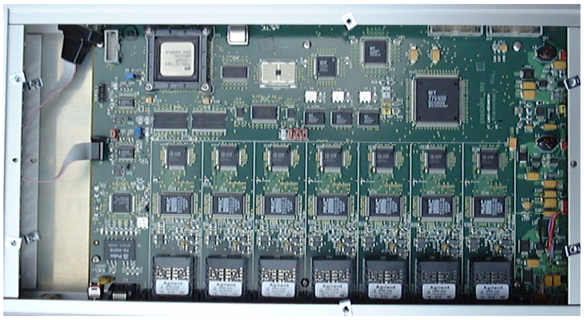

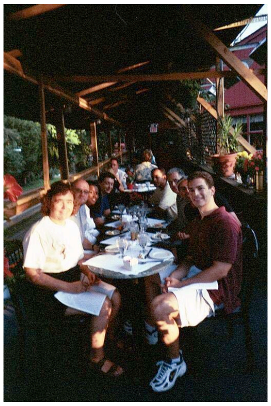

Overview

There is today a young generation for whom computers have always been a fact of life. They have insight, instinct and ambition that older generations lack. It is their energy which will be the agent for change. As they move through the phases of their lives the functions important to those phases will be transformed. One thing is already clear: the network must work reliably. Its performance must be predictable and without threat to their well-being.

Robustness and safety are best founded in the lowest layers of a communications system. While retransmission and encryption will at times be necessary they cannot be the foundation upon which all else is built. Responsiveness and robustness are best served when virtually every transmission succeeds on first try, and when the rigors and subtlety of keeping secrets are not visited on every network user every time the network is used.

The four topics introduced in this report explore network architecture at its most basic level: moving packets from one place to another over an infrastructure that is economically efficient, has high throughput for bulk traffic and low delay for interactive communications. Many factors have to be considered including ease of use, quality of service, privacy and security, and the ease with which the network and network applications can evolve. Infrastructural change will be continuous as it was for telephony. That change cannot depend upon the cooperation of billions of users - the user/network interface must be designed for independent evolution by both parties.

Our work over the next several years will contribute to successively higher layers of network design, operating system design and the design of applications. However, experience tells us that this cannot be a linear process. While building a solid foundation for the future we must develop and maintain a deep understanding for key applications requirements, user needs and service provider interests. Only by doing that can we avoid making a technically elegant mistake. Therefore we will also work with applications that are perceived to have important influence for the future. In recent years this motivation led to pioneering work on networked music, and an understanding for the impact that broadband web service can have when the network responds reliably in a human response time (about 100 msec.)

The summer program for student interns has been particularly helpful in probing new directions. Three months of a talented student's time can give much insight. The work described in this report builds substantially upon this method. Going forward the institute will assemble a permanent staff capable of following through and integrating the learning from these probes. We are also in the process of broadening the range of technical input through joint work with friends in business and academia. Over the next few years we expect to describe in detail a new target of network evolution, one which can have profound influence on the evolution of an information-based society. We foresee a high performance network that is robust, safe and easy to use for all people in the United States. It will be the foundation for a major advance in the quality of life, including medical practice, education, entertainment, employment, family life and personal affairs.

A Network for the USA

The Internet is at the threshold of change. The pioneering years are behind us; a mature service is now required. While there may be great flexibility in a network which has essentially one node type (a router), a minimalist payload type (the datagram), and complete freedom for anyone to send packets to anyone else, these characteristics do little to minimize cost, maximize performance, ensure robust communications, or provide a sense of safety. To illustrate this point we offer a design which has a clear structure that more readily supports a theory of its operation and encourages specialization in the design of its components.

The architecture under consideration draws clear distinction between a backbone network, which serves a country or continent, and regional access networks each of which gathers and distributes traffic within a defined area. Distinction is also drawn between two layers of the network's architecture - a physical layer and a logical layer, each with clear and separately defined roles. The physical layer is focused on creating robust communication paths for the benefit of the logical layer, not individual users. Its design is tuned to support maintenance and fast service restoration. The nodes of the logical layer are regional switching centers that support the needs of individual users and user groups according to the style of service which they desire.

It is proposed that each country be divided into regions that are quite large, in the range 100,000 to a few million households. For example the USA might be divided into regions each 100 miles in radius. Regions of this size are expected to be economically efficient, and the number of regions would change only slowly so route tables would be small and static - much easier to maintain than they are today. Competitive service providers may operate overlapping regional networks and there can be competitive backbone service providers.

With many households per region comes the opportunity to aggregate traffic in large volume and to transmit it en masse over long distance backbone paths that use glass fiber and optical circuit switching to achieve high performance and exceptional reliability. This well understood architecture can be made very reliable, even at the highest transmission speeds. Most communications will pass through just one or two regional packet switches, contributing further to low delay and high performance. Furthermore, the small number of regional switching centers with fast service restoration provided by a circuit switched backbone make it reasonable to consider a deployment strategy with enderground installations that have a high level of security.

Three of the challenges posed by this design are highlighted in this report. For regional access networks the challenge has many components, including how to make the service very reliable while minimizing the costs of network construction and network operation. The regional nodes are challenged by the requirement to provide personal service to many households and businesses while carrying an unprecedented traffic volume. Finally, for the network as a whole, there is the challenge of ensuring predictable low delay service in what is expected to be a massive information flow.

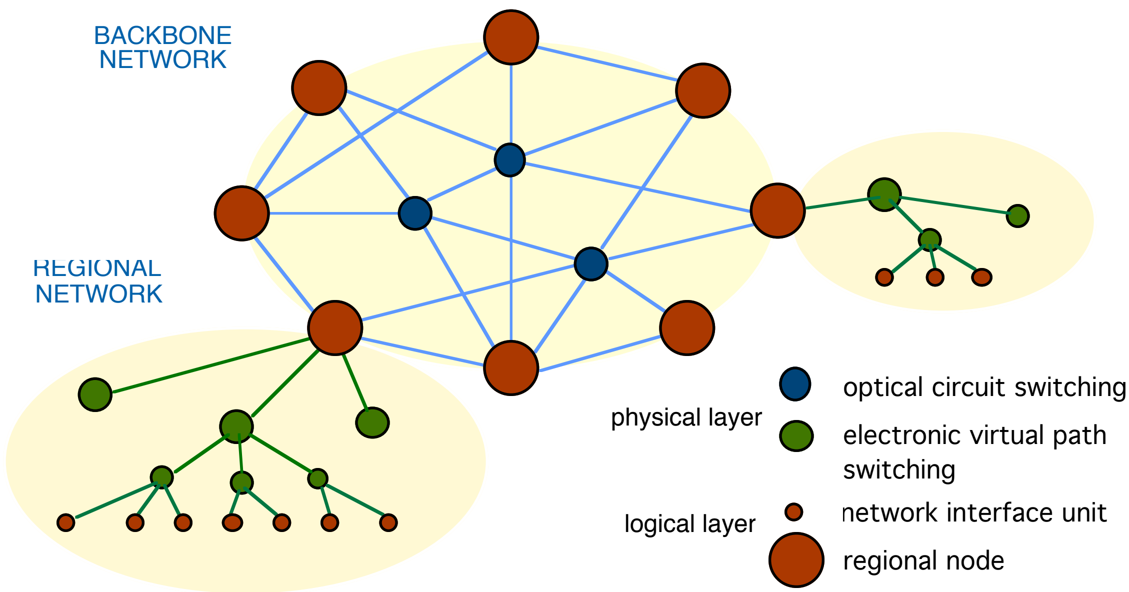

Tree structured access networks with automatic service restoration.
Full mesh logical backbone implemented with optical or virtual circuit switches.
Regional node provides packet switching and traffic aggregation.

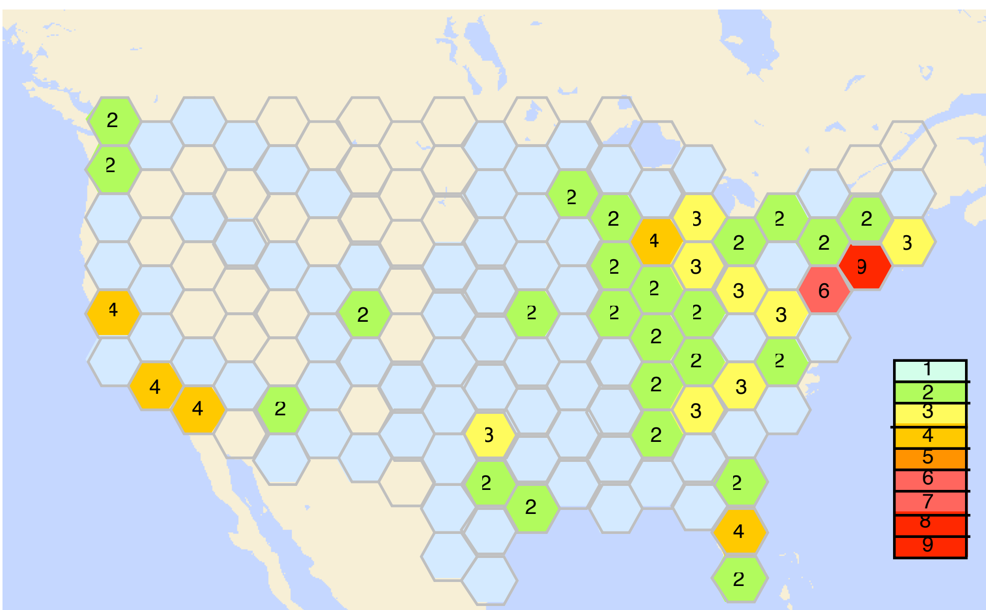

The United States divided into 109 regions shown in solid colors, each has 100 mile radius.
The number of households per region is rounded up to an integral number of millions.
1.5% of the nation's households are in the transparent hexagonal areas. Each of these areas has less than 100,000 households. In the above chart they are served by neighboring regions.
Population data per zip code was obtained from the US Census 2000.

Switched Fiber Access

Network access, referred to now as "the first mile", is a most critical part of the overall network structure because it represents a large part of the network's cost and it is where the customer meets the network. Access networks in the United States must connect to more than 100 million households and small businesses. Therefore the total cost is very sensitive to the unit cost per home. (Network access for large corporations and large public institutions is driven by different economics and historically has been treated as a separate business.)

Our research focuses on switched fiber access. Switched access means that a microelectronic switch is placed close to where the connection to a home joins with the fiber cable in the street. Fiber is the inevitable choice for new broadband networks because it has a low error rate, and low cost of maintenance. A switched fiber network is also an excellent foundation for wireless broadband service. It offers a symmetric communications service - one that provides an equally high data rate upstream and down. This means that any household or mobile device can operate a service and is not constrained by bandwidth to be merely a client on the network.

Local Exchange Carriers presently favor passive optical networks (PON) because that technology does not require powered electrical devices in the outside plant. This is seen as a saving in operating cost. However, we are investigating whether switches made of reliable low power electronics promise better value over all. The anticipated advantages include:

- Protection against fraud, a major source of lost revenue in today's CATV networks
- Automatic service restoration, which reduces service downtime and the cost of maintenance
- Active monitoring at all nodes in the outside plant, which further reduces maintenance cost
- Privacy, because one user's traffic does not enter another's premises
- Fast symmetric access with no scheduling delay for upstream flow
- Excellent support for multimedia communications to and from the home
- Lower first cost, because the more flexible topology requires fewer miles of glass fiber.

The switched fiber access is seen as a logical follow-on to PON and cable modem services.

At its inception, Fraser Research received in donation a partially complete research prototype network constructed with microelectronic switches manufactured by IDT. The switches implement virtual paths which connect individual homes to a regional switching center. In the past year the prototype hardware has been completed and an operating system, designed for the access network, is being tested on the prototype.

A strategy for automatic service restoration has been identified and simulated. Our goal is to restore service in 50 msec. Fiber layouts have been designed for Manhattan street networks of various sizes. It is observed that incorporating automatic restoration into an access network at the time of its design adds less than 10% to the cost. Software tools used for this study were also used to design PON networks. From this experience came the understanding that switched access networks may use fiber more efficiently and therefore have lower first cost than PON.

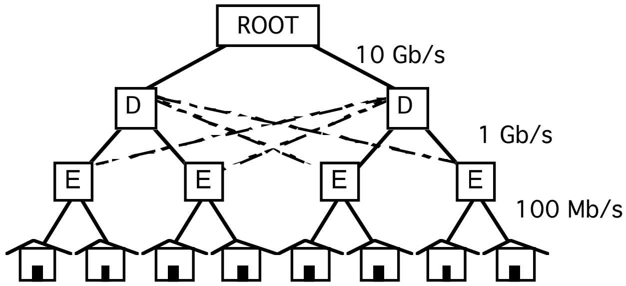

Each access network is a tree with extra links for redundancy
E is an edge node, D is a distribution node, ROOT is at a regional center

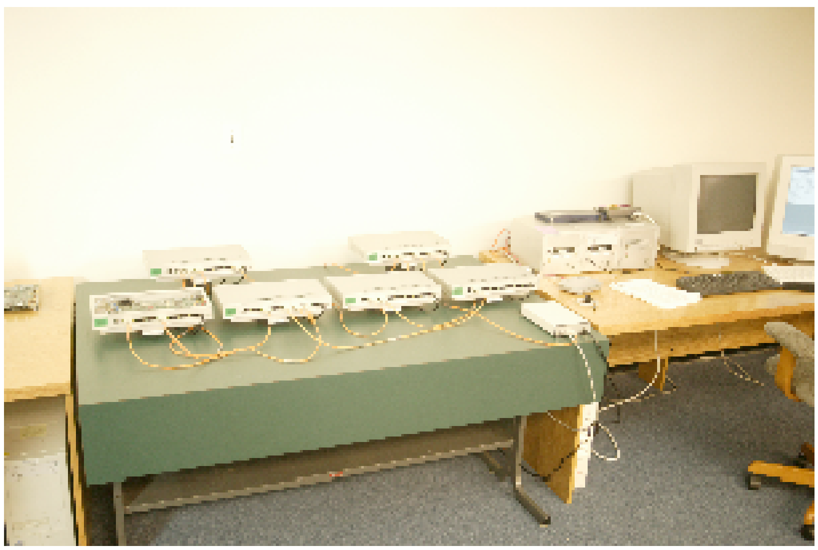

The “train set Mk2” - a prototype 155 Mb/sec switched fiber access network

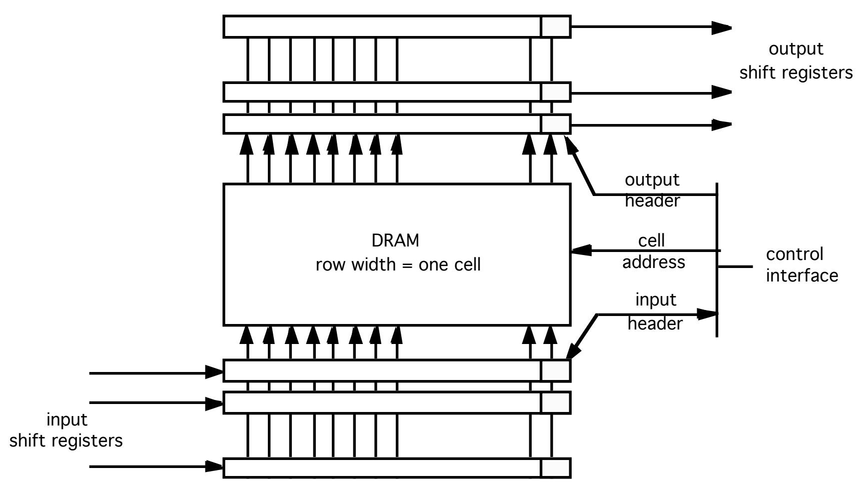

Data path of a microelectronic switch having a throughput of 1.2 Gb/sec.
Devices with 10 Gb/sec and 80 Gb/sec throughput are being investigated.

Switching Technology for a Regional Center

A region that serves in excess of one million homes and business at 100 Megabit/sec or more will require a switch with unprecedented capacity. Fraser Research is therefore intent on showing that it is possible to construct a switch that can handle traffic volume in excess 1 Petabit/sec.*

In 1956 Charles Clos described a solution to the Bell Telephone System's need for a high capacity switch. He showed how to construct a large switching system from a network of many small switches. He and others developed a theory for switching networks which has ever since been the foundation of large scale circuit switching. Unfortunately, large scale packet switches have proved much more difficult to design. There is insufficient theory upon which they can be based. Thus there is reason to be concerned about constructing the high capacity switches that will be required for the large regional networks which are anticipated for the future.

The main data path of a classic packet switch can be characterized as a multiplexer which concentrates traffic onto a central bus or memory, plus a demultiplexer which distributes the traffic from that bus or memory. The switch ingress ports feed traffic into the multiplexer. The switch egress ports are fed from the demultiplexer. Switches according to this design concept are reasonably well understood but cannot scale to large size because the central bus or memory becomes a bottleneck. Unfortunately, a Clos switching network, which can scale to have high capacity, has other troublesome properties. Depending upon the pattern of the offered traffic, traffic "jams" spontaneously occur within the switching network. Thus it was a surprise when we were able to show an equivalence between a Clos switching network suitably controlled and a multiplexer / demultiplexer switch having the same number of ingress and egress ports. During the past two years this concept has been explored with simulation, searching for a full understanding of how large switches of this design perform. We have growing confidence that the switch can be made to perform well and can scale to have capacity in excess of 1 Petabit/sec.* using technology that is presently within reach.

Contrary to present-day Internet practice, it is proposed that the switch must not lose packets in the course of normal operation, nor must the packets be allowed to arrive out of sequence. Packet loss wastes network capacity, causes delay or worse for the users whose packets are lost, and the practice of ignoring packet losses results in concealing the symptoms of other network problems. Therefore the switch must be expected to operate correctly and reliably under all conditions. It is particularly important that the network perform well in times of crisis when traffic patterns are anything but normal. Only then can there be the confidence which will allow network users to take full advantage of all that a worldwide broadband service can provide. Only then will we see the forecasted massive increase in gross national product and a revolution in medical care which will reach into every corner of the nation.

* 1 Petabit/sec = 1,000 Terabit/sec = 1 million Gigabit/sec = 1,000 million Megabit/sec.

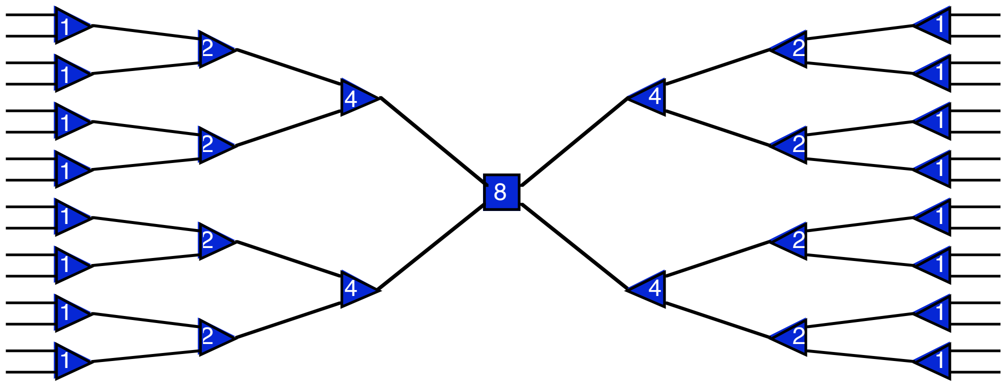

16x16 multiplexer / demultiplexer switch

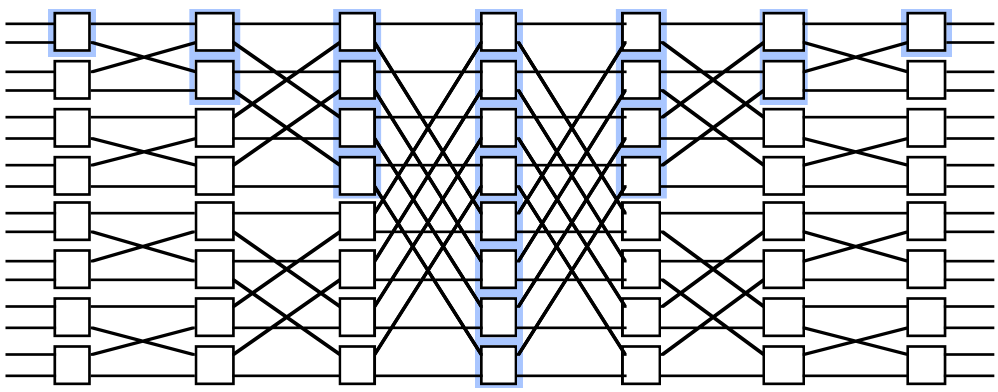

Multiplexer Stages Center Stage Demultiplexer Stages

A 16x16 Clos switching network with 7 stages and 2x2 nodes

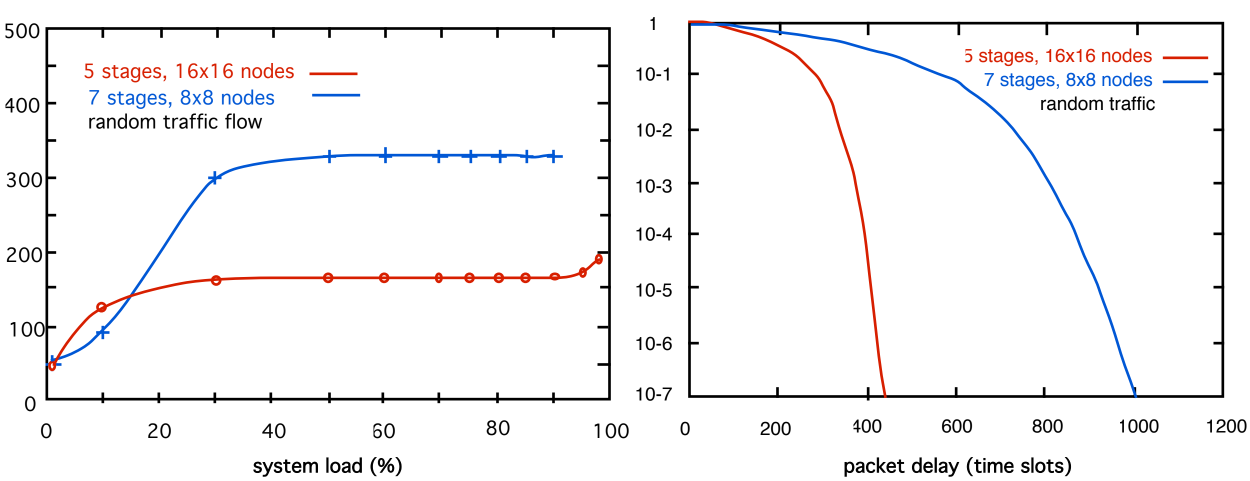

Delay (time slots) v. load Delay probability distribution for 75% load

Simulated performance of two 4,096 x 4,096 switching networks

Flow Control in a Large Scale Network

Flow control is the corner stone of network architecture. It defines network performance and underpins service quality. It greatly benefits from being considered during the earliest phases of network design. However it is not an easy subject. Only recently did Kelly and Willinger lay the foundation from which has emerged a theoretical basis upon which to build.

Flow control is about regulating traffic so that transmission lines are used efficiently while queues do not build to clog memory and cause significant delay and/or packet loss. Packet loss is bad, even for a flow with automatic retransmission, because retransmission adds to delay and wastes the network's capacity. More subtle but nevertheless harmful is the way that automatic retransmission conceals other networking problems.

Two additional objectives not met in today's Internet are being addressed as we do this work:

Host software must be decoupled from the details of network operation

> It is extremely difficult for the operator of a network deployed on a massive scale to introduce new technology and capabilities when those changes require cooperation from all customers. Such is the case for a change that requires new software in every host. The Internet in future must allow independent evolution of the network and the devices attached to it. This is a critical matter for network operators who intend to remain competitive.

Service quality in the network must not be conditional upon proper user behavior

> Recent experience teaches us that an internet vulnerable to erroneous or malicious user behavior cannot be regarded as robust, and most probably will become a national security risk. So, for example, it is not satisfactory to entrust congestion avoidance to software distributed in every host. While hosts must participate in end-to-end flow control, a network operator should be able to take full responsibility for ensuring the quality of service offered within its boundaries.

A two-layer flow control architecture is being studied. Within each segment of the network there is control of the total traffic volume so that the bad effects of congestion are avoided. Flow is also controlled on an end-to-end basis so that individual flows are managed for the best tradeoff between low delay and good utilization of transmission capacity. Host applications can benefit by making the network aware of individual end-to-end flows. Given that every network has finite capacity, the root of each access network includes the means for denying or postponing recognition of a new flow when there is not the spare network capacity to handle it. In this way a network operator can offer and maintain a defined quality of service while achieving good transmission line utilization that is essentially free from packet loss.

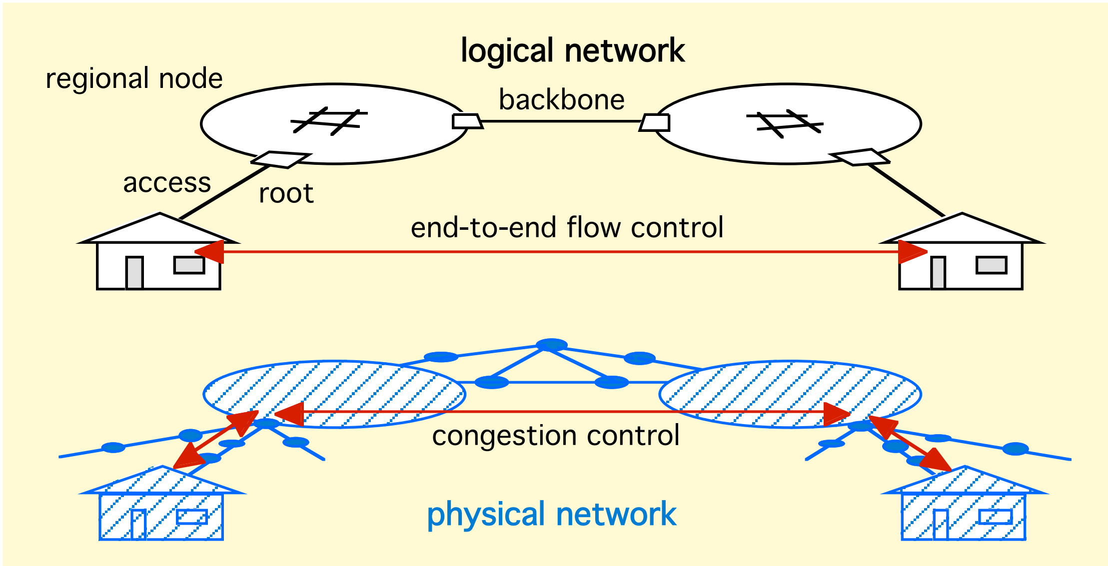

Two aspects of flow control for a wide area network

End-to-end flow control synchronizes application programs, and manages the rate of flow for each conversation so that queues are short and the network has the flexibility to provide a good quality of service.

Congestion control maintains a constant awareness of network capacity thereby providing a framework for conditionally admitting traffic and protecting service quality within the network.

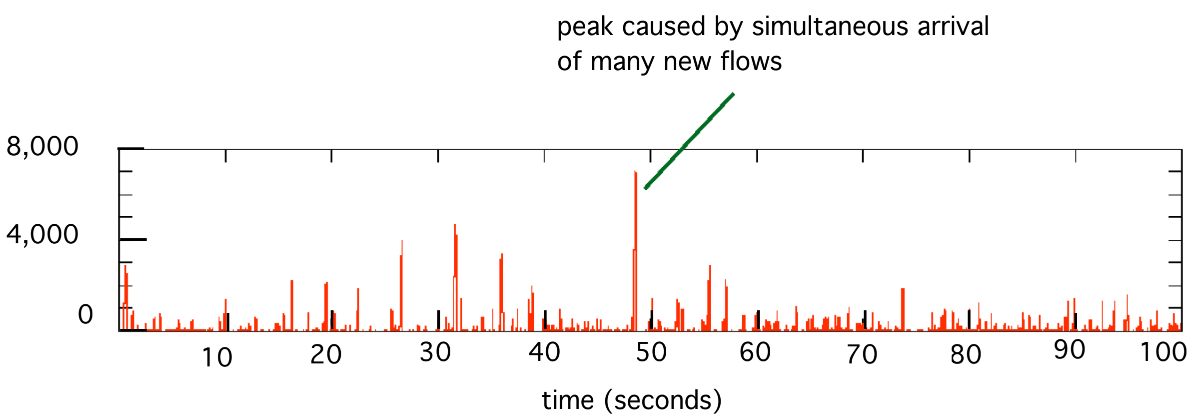

Queue length v. time
Rate based flow control in a simulated national backbone
One of several experiments with contrived "worst case" traffic.
Congestion control should be able address this situation.

Internship Program

Fraser Research hosts up to four graduate students for three summer months at its laboratory in Princeton, NJ. During this time each intern works closely with Sandy Fraser on a project of mutual interest that is compatible with the intern's research interests. The interns meet at least weekly for group discussions and informal talks, and at the end of the summer present their work to Board members and other visitors.

The interns for 2003 and 2004 were:

Alastair Beresford - A graduate of Cambridge University, England, Alastair was a fourth year graduate student in the Laboratory for Communication Engineering at Cambridge University. His research interests lay in developing techniques for improving privacy and security in the domain of Ubiquitous Computing. While at Fraser Research, Alastair simulated a new architecture to name network services and data flows in a nationwide network.

Somdip Datta - A graduate of IIT, India, Somdip was a fifth-year graduate student in Engineering at Princeton University. His research interests lay in the area of routing and architectural issues in Optical Core networks. While with Fraser Research, Somdip simulated a new architecture for a high performance/large scale packet switch.

Nandita Dukkipati - A graduate of the Indian Institute of Science, Bangalore, Nandita was a third year graduate student in Electrical Engineering at Stanford University. Her research interests are broadly in the area of algorithms for high performance networks. Her current research is on congestion control. While with Fraser Research, Nandita investigated algorithms for backbone flow control in large-scale networks.

Anil Madhavapeddy - A graduate of Imperial College, London University, England, Anil was a second-year graduate student in Computer Science at Cambridge University, England. His research interests included operating system security and audio networking. While with Fraser Research, Anil explored techniques to make broadband home networks easy to use and configure.

Aaron Patzer - A graduate of Duke University, Aaron was a second-year graduate student in Engineering at Princeton University. His research interests lay in optical communications, switching, and systems architecture. While at Fraser Research, Aaron investigated algorithms for fast restoration of service in a broadband access network.

David Scott - A graduate of Cambridge University, England, David was a fourth-year graduate student in Computer Science at Cambridge University. His research interests lay in developing methods for abstracting security policies from application code at a high level in the domain of Ubiquitous Computing. While at Fraser Research, David worked on the dynamic creation of a virtual path through a local area network.

Alastair Tse - A graduate of the University of New South Wales, Sydney, Australia, Alastair was a second year graduate student in the Laboratory for Communication Engineering at Cambridge University. His research interests lay in indoor location systems for Ubiquitous Computing. While at Fraser Research, Alastair worked on a compact operating system for self-restoring network switches.

Collaborations
Funded by the National Science Foundation

100 Mb/sec to 100 Million Households

The 100x100 Project brings together economists, security and networking experts, network operators, and policy specialists to create blueprints for a network that goes beyond today's Internet. Drawing on technology trends and the experience of the past 30 years, these scientists are re-prioritizing the fundamental principles that underlie network design to craft networks that will be ubiquitous in scale, revolutionary in bandwidth, economically self-sustaining, resistant to attack, and tractable to manage.

Principal Investigators

AT&T Labs-Research	Albert Greenberg
Carnegie Mellon University	Raj Reddy, Mike Reiter, Hui Zhang
Fraser Research	Alexander Fraser
Internet2	Larry Landweber, Rick Summerhill
Pittsburgh Super Computing Center	Wendy Huntoon, Matt Mathias
Rice University	Ed Knightly
Stanford University	Nick McKeown
University of California at Berkeley	John Chuang, Pravin Varaiya
	Hal Varian, Jan Walrand

FAST Copper: Dynamic Optimization of Resources in Frequency, Amplitude, Space, and Time.

The goal of this 'FAST Copper' Project is to provide ubiquitous 100 Mbps broadband access to everyone in the U.S. who now has a phone line. This goal will be achieved through two threads of research: dynamic and joint optimization of resources in Frequency, Amplitude, Space, and Time to overcome the attenuation and crosstalk bottlenecks, and the integration of communication, networking, computation, modeling, and distributed information management and control for the multi-user twisted pair network.

Principal Investigators

Fraser Research	Alexander Fraser
Princeton University	Mung Chiang
Stanford University	John Cioffi

Alexander G. Fraser - A Biographical Sketch

Alexander G. Fraser, known as "Sandy", established the Fraser Research Institute after his retirement from AT&T in 2002. As Vice President for Research, he founded AT&T Labs Research in 1996 when AT&T split off its equipment manufacturing business into Lucent Technologies. He was appointed AT&T Chief Scientist in 1998.

Dr. Fraser arrived at AT&T Bell Laboratories in 1969 where he invented the Datakit Virtual Circuit Switch and the Spider ring network. Both were cell-based networks that anticipated the development of Asynchronous Transfer Mode (ATM). He created the UNIX Circuit Design Aids System which automatically produced wire-wrap circuit boards from schematic circuit diagrams. With S.C. Johnson, he developed a technique for computer instruction set optimization using a portable compiler which led to the design of a reduced instruction set machine. With W.T. Marshall and G.G. Riddle, he invented the Universal Receiver Protocol. With C.R. Kalmanek and R.C. Restrick, he created INCON, a cell-based network that operated at 2 Mb/s on home telephone wire. In 1982 he became the Director of the Computing Science Research Center, and in 1987 became Executive Director for Information Sciences Research.

Since the 1984 breakup of the Bell System, Dr. Fraser held responsibility for and participated in a wide range of research activities related to network infrastructure and network applications. He led collaborative research with universities and national labs using a 45 Mb/s nationwide network that included a 622 Mb/sec optically amplified trunk. In collaboration with an integrated circuit manufacturer he directed the creation of a 1.2 Gb/s packet switch which obtained low cost and low power by having its data path integrated into a DRAM chip. He also directed research on an experimental all-fiber access network in which 18 homes were directly connected to Bell Labs. From this work came recognition of the near-term importance of cable modems, and he launched a collaboration with cable service providers and manufacturers to create reliable cable modem systems. He proposed that packetized delivery of music and video might offset the projected decline of telephony, and he led a pioneering initiative to leverage Bell Labs' expertise in audio coding through MPEG 2 standardization to create the first music servers and portable players for music distributed via the Internet and around the home.

Before joining Bell Laboratories, Dr. Fraser was Assistant Director of Research at Cambridge University where he wrote the file system for the Atlas 2 computer, England's first time-sharing system, and developed file back-up and privacy mechanisms for that system. Earlier work included a language and compiler for commercial data processing written for the Ferranti Orion computer. Dr. Fraser has a B.Sc. degree in Aeronautical Engineering from Bristol University and a Ph.D. in Computing Science from Cambridge University.

Dr. Fraser was elected to the National Academy of Engineering in 2005, and in 2001 he received IEEE's Richard W. Hamming Medal. He received the 1989 Koji Kobayashi Computers and Communications Award in 1989, and the Sigcomm Award in 1992. He is an AT&T Fellow, a Fellow of the IEEE, and was a Fellow and council member of the British Computer Society.

Fraser Research Board of Trustees - Biographical Sketches

William F. Brinkman

Dr. Brinkman, currently a Senior Research Physicist at Princeton University, is internationally recognized for his fundamental research in the area of the physical sciences. He spent most of his career at Bell Laboratories working on theories of condensed matter. He retired as Vice President of Research. Among his many awards he is a member of the National Academy of Sciences, the American Academy of Arts and Sciences and in 1994, he received the George E. Pake Prize. He is Past President of the American Physical Society.

Edward E. David, Jr.

Dr. David, as a member of the Washington Advisory Group, currently advises governments, universities, and businesses on intellectual property, commercialization, and technology transfer. His early career was spent at Bell Laboratories where he worked in audio and visual communications and then in computing systems technology. In 1970 he became Science Advisor to the President of the United States. In later years he served as President of Exxon Research and Engineering until his retirement in 1986. Among his many awards he is a member of the National Academy of Sciences, the National Academy of Engineering and is a Fellow of the American Academy of Arts and Sciences.

James L. Flanagan

Dr. Flanagan recently retired as Vice President for Research at Rutgers University and Director of the Center for Advanced Information Processing (CAIP). Prior to his work at Rutgers he was at Bell Laboratories for many years where he made major contributions in the areas of signal coding algorithms and automatic speech synthesis and recognition. He invented auto-directive microphone arrays for teleconferencing and pioneered the use of digital computers for acoustic signal processing. His many scientific accomplishments have earned him Fellowships in the National Academy of Sciences, the National Academy of Engineering, IEEE, and the Acoustical Society of America.

Andy Hopper

Dr. Hopper, currently Professor of Computer Technology at the University of Cambridge and Head of the Computer Laboratory, is internationally known for his interest in networking, pervasive and sentient computing, mobile systems, and dependable infrastructure systems. He is equally well-known for his parallel pursuit of academic and industrial careers that has led him to co-found eleven companies. Until 2002 he was managing Director of AT&T Laboratories, Cambridge. He is currently Chairman of Cambridge Broadband and of Real VNC, and is a Director of Level 5 Networks and Ubisense. Among his many awards he is a Fellow of the Royal Academy of Engineering and a Trustee of the Institute of Electrical Engineers. He is a Fellow of Corpus Christi College, Cambridge.

Lawrence H. Landweber

Dr. Landweber, currently Research Professor and Professor Emeritus of Computer Science at University of Wisconsin – Madison is also a Senior Advisor to the Assistant Director of the National Science Foundation. Dr. Landweber played a large role in the evolution of the Internet from early research networks to the establishment of the first network gateways between the US and countries in Europe and Asia. Dr. Landweber served as President and Chairman of the Board of the Internet Society. As Vice President of Education for the Internet Society, he developed the Workshops for Developing Countries that were so important to the spread of the Internet to the third world. He is a Fellow of the Association for Computing Machinery**.**

David L. Roscoe, III

Mr. Roscoe, currently a senior executive at RiskMetrics, was employed from 1967-1999 at J.P.Morgan, from which he retired as a Managing Director. Mr. Roscoe has a broad understanding of the business strategies necessary for an organization to be successful. Mr. Roscoe received a B.A. with Honors in Economics from Yale University, and an M.B.A. in Finance from the University of Oregon.

Jeffrey R. Walsh

Mr. Walsh, currently Executive Vice President of WR Investment Partners, LLC and a partner with WR Capital Partners, LLC, has the ability to view the many aspects of a business and instinctively know what it will take to make it better and therefore, more valuable. Mr. Walsh joined Wesray Capital in 1987 as Vice President and was actively involved in the acquisition and management of numerous entities. Mr. Walsh is a board member of EAI Holdings Corp. and Vice Chairman of EAI Partners L.P and affiliated entities. He is Chairman of the Management Committee of Spinnaker Coating, LLC and a founding member of two independent alternative asset firms, Rumson Capital, LLC and ZAIS Group, LLC. Mr. Walsh is a Certified Public Accountant and a graduate of Villanova University.

Richard M. Wolf

Mr. Wolf worked in the telecommunication industry for 42 years and his knowledge is encyclopedic. His perspective on the future of the industry is perceptive and insightful. For twenty years he worked at New York Telephone and AT&T in various engineering positions and became involved in the Bell Systems antitrust case. He then joined Robert E. LaBlanc Associates, Inc. a telecommunication consulting firm where he held several positions and became the subject matter expert on Wireless Tower Sitings. Mr. Wolf graduated with a degree in Electrical Engineering from Cornell University and attended Stanford's Engineering Economics and Executive Programs.