

International Computer Science Institute

ACTIVITY REPORT 2000

Contact Information:

Postal Address:

International Computer Science Institute
1947 Center Street
Suite 600
Berkeley, CA 94704-1198
USA

Phone: (510)-666-2900

Fax: (510)-666-2956

E-mail: info@icsi.berkeley.edu

Internet: <http://www.icsi.berkeley.edu>

Principal Sponsors:

AT&T
Intel
Nortel
Qualcomm
Förderverein (German industrial organization)
German Ministry of Research and Technology (via U. Munich)
Spanish Office of Science and Technology
Finnish National Technology Agency (TEKES)
National Science Foundation
DARPA

Affiliated Sponsors:

Cisco
European Media Lab
European Union (via Sheffield University)
IBM
Xerox

Corporate Officers:

N. Morgan (President and Institute Director)
S. Shenker (Vice President)
C. Papadimitriou (Secretary and Treasurer)

Board of Trustees:

Dr. H. Ahmadi, AT&T Labs-Research
Prof. E. Berlekamp, UC Berkeley
Vice Chancellor B. Burnside, UC Berkeley
Dr. J. Cohen, VoiceSignal Inc.
Dr. A. Goldberg, Neometron
Dr. T. Grandke, Siemens CRC
Prof. D. Hodges (Chair), UC Berkeley
Prof. R. Karp, ICSI and UC Berkeley
Mr. P. Lizcano, Spanish Office of Science and Technology
Prof. K. Melhorn, Max-Planck-Institute for Informatics
Prof. N. Morgan (Director)
Prof. C. Papadimitriou, CS Division Chair, UC Berkeley
Prof. D. Schütt, Siemens AG

Senior Trustees:

Prof. G. Barth, Dresdner Bank AG
Prof. D. Ferrari, U. Cattolica
Prof. J. Feldman, ICSI and UC Berkeley
Dr. R. Kay, former acting Director of ICSI (until 1988)
Prof. G. Goos, U. Karlsruhe
Prof. H. Schwärtzel, TU Munich
Prof. N. Szyperski, U. of Cologne; former head of the GMD; originator of ICSI.

Institute Overview – June 2001

The International Computer Science Institute (ICSI) is an independent, nonprofit basic research institute affiliated with the University of California campus in Berkeley, California. Its establishment was motivated by a recognition of the lack of an international facility for fundamental research in the field of computer science. ICSI was started in 1986 and inaugurated in 1988 as a joint project of the Computer Science Division of UC Berkeley and the GMD – Research Center for Information Technology GmbH in Germany. Since that time, Institute collaborations within the university has broadened (for instance, to the Electrical Engineering Division and the Linguistics Department). Additionally, Institute support has broadened to a range of international collaborations, US Federal grants, and to direct industrial sponsorship. Throughout these changes, the Institute has maintained its commitment to a pre-competitive, open research program. The goal of the Institute continues to be the creation of synergy between world leading academic and industrial research in an international environment through excellence in fundamental research in computer science and engineering.

The particular areas of concentration have varied over time, but are always chosen for their fundamental importance and their compatibility with the strengths of the Institute staff and the affiliated UC Berkeley faculty. ICSI currently has significant efforts in four major research areas: Internet research, including Internet architecture, related theoretical questions, and network services and applications; theoretical computer science, including applications to the modeling of both biological and internet-related phenomena; artificial intelligence, particularly for applications to natural language understanding; and natural speech processing.

The Institute occupies a 28,000 square foot research facility at 1947 center Street, just off the central UC campus in downtown Berkeley. Administrative staff provide support for researchers' housing, visas, computational requirements, grants administration, etc. There is an average of about eighty scientists in residence at ICSI including permanent staff, postdoctoral Fellows, other visitors, and affiliated faculty and students. The senior investigators are listed at the end of this overview, along with their current interests.

Institute Sponsorship – 2000 fiscal year (same as calendar year)

As noted earlier, ICSI is sponsored by a range of US Federal, international, an industrial sources. The figure below gives the relative distribution of funding between these different sponsoring mechanisms.

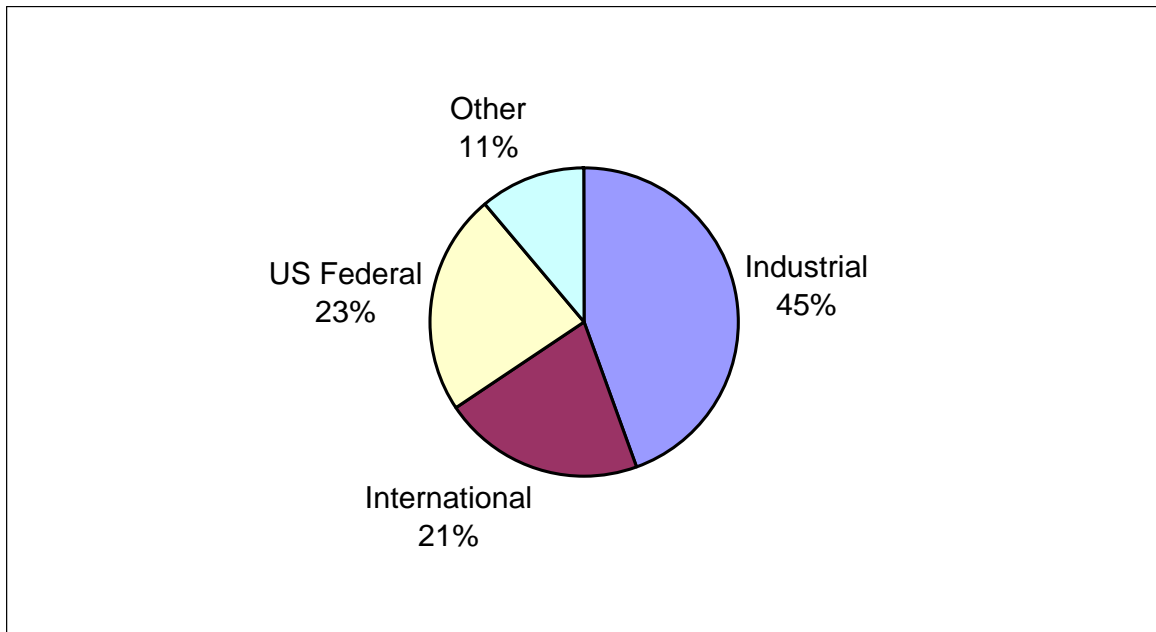


Figure 1: Distribution of sources of ICSI revenue for 2000.

The US Federal funding comes from a range of grants to support research across the Institute. Most of this funding comes from the National Science Foundation and DARPA. International support in 2000 was from government and industrial programs in Germany, the Ministry of Science and Technology in Spain, and the National Technology Association of Finland, with additional support from the European Media Lab and the European Union. Industrial support in 2000 was primarily from AT&T, but by the end of the year we had received commitments from Intel, Nortel, and Qualcomm with additional sponsorship from Cisco, IBM, and Xerox.

Revenues increased significantly in 2000 (about 32% over 1999), to roughly \$6M for the year.

Institutional Structure of ICSI

ICSI is a nonprofit California corporation with an organizational structure and bylaws consistent with that classification and with the institutional goals described in this document. In the following sections we describe the two major components of the Institute's structure: the Administrative and Research organizations.

Management and Administration

The corporate responsibility for ICSI is ultimately vested in the person of the Board of Trustees, listed in the first part of this document. Day-to-day operation of the Institute is handled by Corporation Officers, namely the President, Vice President, and Secretary-Treasurer. The President also serves as the Director of the Institute, and as such takes responsibility for ongoing Institute operations.

Internal support functions are provided by three departments: Computer Systems, Finance, and Administration. Computer Systems provides support for the ICSI computational infrastructure, and is led by the Systems Manager. Finance is responsible for payroll, grants administration, benefits, and generally all Institute financial matters; it is led by the Controller. All other support activities come under the general heading of Administration, and are supervised by the Administrative Manager; these activities include human resources, office assignments, proposal preparation and submission, and housing.

Research

Research at ICSI is overwhelmingly investigator-driven, and themes change over time as they would in an academic department. Consequently, the interests of the senior research staff are a more reliable guide to future research directions than any particular structural formalism. Nonetheless, through much of its history, ICSI research has been organized into Groups: the Networks Group (internet research), the Theory Group, the Applications/AI Group, and the Realization Group (parallel computer architecture and speech processing). In 1999 we added a fifth group, the AT&T Center for Internet Research at ICSI (ACIRI). Consistent with this history, the bulk of this report is organized along these lines, with one sub-report for each of the five groups. In 2001 we have begun the process of merging the two internet groups, so that future reports will be organized around the remaining four groups.

Across these groups, two strong themes can be seen, in Internet research and in Human Centered Computational Intelligence (HCCI), a term intended to encompass a range of topics related to the use of computing for improved human-machine interfaces and human-human cooperation. Other topics of study are pursued opportunistically. Some of these efforts will ultimately grow into a larger activity, given sufficient Investigator and Sponsor interest.

Senior Research Staff

The previous section briefly described the clustering of ICSI research into major research themes and working groups. This document describes a current research snapshot, and future work could be extended to new major areas based on strategic Institutional decisions and on the availability of funding to support the development of the necessary infrastructure. At any given time, though, ICSI research is best seen as a set of topics that are consistent with the interests of the Research Staff. In this section, we give the names of current (June 2001) senior research staff members at ICSI, along with a brief description of their current interests, the current group name, and the Research Group that the researcher is most closely associated with. This is probably the best current snapshot of research directions for potential visitors or collaborators. Not shown here are the range of postdoctoral Fellows, visitors, and graduate students who are also key contributors to the intellectual environment at ICSI.

Jerome Feldman (AI): neural plausible (connectionist) models of language, perception and learning and their applications.

Charles Fillmore (AI): building a lexical database for English (and the basis for multilingual expansion) which records facts about semantic and syntactic combinatorial possibilities for lexical items, capable of functioning in various applications: word sense disambiguation, computer-assisted translation, information extraction, etc.

Sally Floyd (ACIRI): congestion control, transport protocols, queue management, and network simulation.

Atanu Ghosh (ACIRI): extensible open source routing, active networks, protocols, multimedia, and operating systems.

Steven Greenberg (Realization [Speech]): Spoken language processing by humans and machines, analysis of spontaneous speech at the phonetic and prosodic levels, automatic labeling and segmentation of phonetic and prosodic material in spontaneous speech corpora, analysis of automatic speech recognition systems, computational models of auditory processing, hearing-aid design for improving speech intelligibility, auditory-visual speech processing.

Mark Handley (ACIRI): scalable multimedia conferencing systems, reliable multicast protocols, multicast routing and address allocation, and network simulation and visualisation.

Hynek Hermansky (Realization [Speech]): acoustic processing for automatic speech and speaker recognition, improvement of quality of corrupted speech, human speech communication. (Also with the Oregon Graduate Institute).

Richard Karp (Theory [Algorithms] and ACIRI) : mathematics of computer networking, computational molecular biology, computational complexity, combinatorial optimization.

Nelson Morgan (Realization [Speech]): signal processing and pattern recognition, particularly for speech and biomedical classification tasks

Vern Paxson (ACIRI): intrusion detection; Internet measurement; measurement infrastructure; packet dynamics; self-similarity.

Lokendra Shastri (AI): Artificial Intelligence, Cognitive Science, and Neural Computation: neurally motivated computational models of learning, knowledge representation and inference; rapid memory formation in the hippocampal system; inference with very large knowledge-bases; neural network models for speech recognition; inferential search and retrieval.

Scott Shenker (ACIRI): congestion control, internet topology, game theory and mechanism design, scalable content distribution architectures, and quality of service.

Elizabeth Shriberg (Realization [Speech]): Modeling spontaneous conversation, disfluencies and repair, prosody modeling, dialog modeling, automatic speech recognition, utterance and topic segmentation, psycholinguistics, computational psycholinguistics. (Also with SRI International).

Andreas Stolcke (Realization [Speech]): probabilistic methods for modeling and learning natural languages, in particular in connection with automatic speech recognition and understanding. (Also with SRI International).

Contents

	Research Group Highlights	1
1	AT&T Center for Internet Research at ICSI (ACIRI)	3
1.1	Overview	3
1.2	Research on Internet Architecture	4
1.3	Research on Internet Operations	14
1.4	Research on Internet Theory	19
1.5	Community Activities	21
2	Network Services and Applications	28
2.1	Overview	28
2.2	USAIA: Ubiquitous Service Access Internet Architecture	28
2.3	Mobile Network Architecture for Vehicles	29
2.4	Routing for Mobile Access Scenarios	30
2.5	Content Delivery and Source Specific Multicast Deployment	32
2.6	Further Activities	32
3	Speech Processing: Signals and Systems (the Realization Group)	36
3.1	Speech Signal Modeling	36
3.2	Speech Recognition	40
3.3	Speech Perception	44
3.4	Applications	44
3.5	CUSP	47
4	Artificial Intelligence and Its Applications	49
4.1	Language Learning and Use, Jerome Feldman, PI	49
4.2	FrameNet Project, Charles Fillmore, PI	52
4.3	Connectionist Modeling, Lokendra Shastri	55
5	Exploratory Areas in Computational Algorithms	59
5.1	Computational Biology	59
5.2	Theory of Computation	61

Research Group Reports

2000 was a good year for the Institute. Professors Richard Karp and Charles Fillmore ramped up their ICSI efforts, several new researchers began work here, and the Institute won several new competitive grants. New industrial partnerships were formed with Nortel and Intel, the latter based on the launching of an exciting new open source initiative. The annual AT&T review went well and the ACIRI/AT&T relationship was renewed for another 3 years. A new research partnership was begun with the European Media Lab. ICSI's participation in the German SmartKom project ramped up to a significant level, involving ICSI researchers, permanent staff, visitors, and students. Collaboration with OGI and Qualcomm on cellular speech processing also accelerated. New collaborations with campus faculty (both UC Berkeley and UCSF) were established, offering the potential for improved ties in the future.

While networking and human-centered computing continued to be strong areas of concentration, we are committed to other activities such as the work on fundamental computational algorithms.

In this report, we have described our research in terms of the 5 research groups in place during 2000: the AT&T Center for Internet Research at ICSI (ACIRI), Network Service and Applications (NSA), Realization, AI, and Theory. Future versions of this report will incorporate a modified group structure that we will be implementing this year, consisting of 4 major groups: Networking (merging the 2 network groups), Algorithms (theory plus applications of computational methods in biology), Speech (the former Realization group), and AI.

RESEARCH GROUP HIGHLIGHTS

The following are a selection of key achievements in our research groups for the year 2000. Although not a complete listing and, by necessity, quite varied given the differing approaches and topics of each group, it should nonetheless give the flavor of the efforts in the ICSI community for the last year.

Networking

- Explicit Congestion Notification (ECN) has become proposed standard
- Proposal for Aggregate Congestion Control as way to deal with flash crowd and denial-of-service attacks
- Work on Content Addressable Networks as a scalable way to build peer-to-peer applications.
- Complete rewrite of the PIM-SM specification
- TFRC, an equation-based congestion control protocol, has been proposed.
- Establishment of XORP project

- Sally Floyd elected Fellow of ACM, winning award for best paper in Transactions in Networking

Theory (Algorithms)

- Development of feature selection methods for supervised and unsupervised classification of biological samples using DNA microarray data
- Development of a powerful randomized algorithm for sequence prediction with incomplete information (Piccolboni and Schindelhauer)
- Successful proposal process leading to a large QB3 grant from the State of California to UCB, UCSF and UCSC for research and education in quantitative biology.
- Successful proposal process leading to NSF ITR Grant for research on optimization problems related to the Internet for collaboration between ICSI, UCB, and UCLA.

AI

- Major revision of FrameNet software to support expansion
- First major publication on Formal Cognitive Linguistics
- Publication of detailed model of Long Term Potentiation
- Initiating joint project with European Media Lab on “Even Deeper Understanding” (of language)
- Successful proposal process leading to NSF-funded FrameNet II project
- Successful proposal process leading to NSF-funded neural modeling research

Realization (Speech)

- Completion of efforts on discriminative approaches to confidence estimation on recognition systems
- New discriminative approaches and systems for robust multi-lingual speech recognition
- Collaboration with OGI and CMU, leading to best results on govt speech-in-noise ASR tests
- Key results in the study of prosodic and articulatory-acoustic patterns in conversational speech
- Preliminary English-language recognition system for German SmartKom project
- New funded projects in the study of multi-channel meeting recordings, with DARPA and IBM
- Integration into group of SRI visitors Shriberg and Stolcke

The remainder of this report consists of the 5 individual summaries for the ICSI 2000 groups.

1 AT&T Center for Internet Research at ICSI (ACIRI)

1.1 Overview

The year 2000 ushered in the second year of ACIRI's existence. We entered the new Millennium fully engaged in our mission to perform pioneering and relevant research in the area of Internet architecture. Before describing the content of our research, we first start with a brief overview of ACIRI's staffing, its three major research themes, and its community involvement.

1.1.1 Staffing

The ACIRI senior research staff consists of Sally Floyd, Mark Handley, Richard Karp, Vern Paxson, and Scott Shenker. We currently have two researchers with limited term appointments, Jitendra Padhye and Sylvia Ratnasamy. Eddie Kohler joined us in 2001 to begin a two-year appointment. Paul Francis joined our staff in 2000 and made significant contributions during his six month stay, but he has since left for FastForward. Luigi Rizzo, a professor from U. Pisa, is visiting ACIRI for several months, and ACIRI hosted numerous other visitors and student interns over the year. In addition, several non-ACIRI ICSI researchers (Brad Karp, a former ACIRI intern, among them) have augmented our research efforts.

1.1.2 Research Program

Research is ACIRI's primary mission. The research efforts at ACIRI can be roughly divided into three main areas:

- **Internet Architecture:** The basic Internet architecture, by which we mean the IP protocol family and closely related issues, is the fundamental technical underpinning of the Internet.
- **Internet Operations:** This area encompasses issues that are important in keeping the Internet operating efficiently and securely.
- **Internet Theory:** One of the goals in setting up ACIRI was to become a leader in developing the *science* of datagram networks. A necessary component of this is to bridge the gap between traditional theoretical computer science and the Internet community.

Sections 1.2-1.4 of this report describe ACIRI's major research projects in these areas. A list of publications is at the end of the report.

1.1.3 Community Participation

Active participation in the Internet community is another important aspect of ACIRI's mission. ACIRI members are active in several different organizations of the Internet community, most notably:

- **IETF**: The Internet Engineering Task Force is the standards body for the Internet community.
- **IRTF**: The Internet Research Task Force organizes research groups on topics of interest to the IETF.
- **NANOG**: The North American Network Operators Group (NANOG) provides a forum for the exchange of technical information, and promotes discussion of implementation issues that require community cooperation.
- **SIGCOMM**: SIGCOMM is the ACM special interest group on communications, and the conference it sponsors is the premier academic conference for the Internet community. The journal it co-sponsors, ACM/IEEE Transactions on Networking (ToN), is the premier academic journal for the Internet community.
- **Network Simulator**: The NS network simulator is based on the original simulator developed at UCB by Steve McCanne and Sally Floyd. It has now become the centerpiece of a larger community-wide effort to create and support a common network simulator environment for the Internet community. The goal is to encourage more synergy between different research groups by allowing them to share simulator modules.

Section 1.5 contains a more detailed list of ACIRI contributions to the research and standards activities in the Internet community.

1.2 Research on Internet Architecture

The main efforts in ACIRI's research program on Internet architecture fall into the following categories: congestion control (both host algorithms and router support), multicast, content-addressable networks for peer-to-peer content delivery systems, novel architectures, and the extensible open router platform.

1.2.1 Congestion Control: Host Algorithms

- **Equation-Based Congestion Control (ongoing)**

Because TCP is the most prevalent form of congestion control on the Internet, there is widespread consensus that any non-TCP congestion control algorithm must be *TCP-friendly* or *TCP-compatible*; that is, it must coexist with TCP without consuming an undue share of the bandwidth. ACIRI researchers, along with others, have been working on an approach to TCP-friendly congestion control called *equation-based* congestion control. Equation-based congestion control, rather than imitating TCP's window adjustment algorithm, seeks to match TCP's bandwidth usage equation in the longer run (the equation describing TCP's bandwidth usage resulting from a given packet-drop rate), while avoiding TCP's halving of the sending rate in response to a single packet drop. During this past year, a unicast version of this approach, called TFRC (TCP-Friendly Rate Control) has been developed and extensively analyzed,

and presented in both a Technical Report and a paper at SIGCOMM 2000 [13]. An internet-draft is also in progress to encourage experimentation with TFRC.

Coupled with the work on equation-based congestion control is a simulation-based comparison between equation-based congestion control and other TCP-friendly approaches to congestion control that also give a smoother sending rate than does TCP [68]. Other work explores the implications of TCP-friendly congestion control, particularly those variants with slower responses to congestion than that of TCP, in terms of fairness, responsiveness, and the potential for extended periods of congestion [66].

[Sally Floyd, Mark Handley, Jitendra Padhye, Scott Shenker, Deepak Bansal (intern)]

- **Equation-Based Multicast Congestion Control (new)**

Whilst we believe that our work on equation-based congestion control for unicast applications is now mature enough to consider standardization and real-world deployment, our research on equation-based congestion control for multicast applications is still ongoing. In addition to the problems associated with stable congestion management that emerge for unicast, with multicast there is a tension between providing rapid feedback, and preventing an implosion of feedback messages at the sender. Thus the control loop for multicast must necessarily function with larger delays. We are investigating a number of possible solutions for feedback suppression in the context of this control loop, and at the present time this work is looking rather promising. There are however still some potentially undesirable side-effects when the degree of statistical multiplexing is low and the number of receivers is large, and so this work is ongoing. Unlike with unicast congestion control, this multicast work is currently entirely simulation-based, as it is extremely hard to investigate such scalability issues in real networks.

[Mark Handley, Jörg Widmer (ICSI)]

- **Explicit Congestion Notification (ongoing)**

TCP senses congestion on the Internet by detecting packet drops. However, since the early DEC-bit era there has been interest in delivering Explicit Congestion Notification to end hosts. There is an ongoing effort in the IETF, led by Sally Floyd and K. K. Ramakrishnan, that advocates the use of a single bit in the packet header to indicate congestion. The basic idea is quite simple, but there are still some issues to arise from the initial deployment of ECN. In the past year, they have specified the treatment of ECN with respect to retransmitted TCP packets, in order to provide security against possible denial-of-service attacks using spoofed source addresses [40]. They have also specified the general use of ECN with IP tunnels, such as IP in IP [39, 38]. In addition to the technical work involved, considerable time has been devoted to dealing with obstacles to deployment presented by certain firewalls and other equipment in the Internet that block TCP connections attempting to negotiate ECN [70].

The Transport Area Directors of the IESG (Internet Engineering Steering Committee) have agreed that the next step for ECN is to combine all of the separate documents about ECN into a single internet-draft, and to submit this to the IETF to

be considered as a Proposed Standard. Thus, we expect ECN to become a Proposed Standard some time in the next year.

[Sally Floyd, K. K. Ramakrishnan (AT&T)]

- **Other Congestion Control Topics**

TCP Limited Retransmit (new): We have coauthored a proposed modification to TCP to more effectively recover lost packets when a connection’s congestion window is small [36]. The “Limited Transmit” algorithm calls for sending a new data segment in response to each of the first two duplicate acknowledgments that arrive at the sender. Transmitting these segments increases the probability that TCP can recover from a single lost segment using the fast retransmit algorithm, rather than using a costly retransmission timeout.

[Sally Floyd]

TCP Idling (new): Although TCP has been in use for many years, there are still a number of aspects of its behavior that are not well specified or that behave less than optimally. When our other research leads to insights into how to make TCP behave better or lead us to discover deficiencies in the specifications, we attempt to standardize remedies. In the past year, this has resulted in a IETF “Best Current Practice” RFC becoming standardized on how TCP should behave when the sending application goes idle or doesn’t fully utilize the available bandwidth. We are also looking at how to improve TCP’s retransmit timeout behavior by separating the concept of when the sender is allowed to send a packet after a timeout from when the missing packet is retransmitted.

[Mark Handley, Sally Floyd, Jitendra Padhye]

Stream Control Transfer Protocol (SCTP) (ongoing): SCTP is a reliable transport protocol being developed within the IETF, primarily for transport of telephony signalling, but with an eye towards serving as a general transport protocol for applications requiring multiple, concurrent communication streams. ACIRI’s primary contribution to the protocol is the design of its congestion control.

[Vern Paxson]

1.2.2 Congestion Control: Router Support

- **Aggregate-Based Congestion Control and Pushback (new)**

The research project on pushback and aggregate-based congestion control focuses on identifying when a rise in the packet drop rate is due to an increase in traffic from a traffic aggregate, as a subset of the traffic of the congested link. Such an aggregate might be traffic from a distributed denial-of-service attack, or a flash crowd of legitimate traffic to a web site related to a news-worthy event. In either case, if the router is able to identify a traffic aggregate largely responsible for the traffic increase, the router might want to preferentially drop packets from that aggregate, to protect

the rest of the traffic on that link from an overall increase in the packet drop rate. Coupled with this preferential dropping at the congested router, the router might invoke pushback to request the immediate upstream routers to also drop packets from this aggregate, and this dropping might propagate recursively. This use of pushback prevents an unnecessary waste of bandwidth by packets that will only be dropped downstream. In the case of a denial-of-service attack, pushback could help focus the preferential packet-dropping on the malicious traffic within the identified aggregate [67].

[Sally Floyd, Ratul Mahajan (intern), Vern Paxson, Scott Shenker, Steve Bellovin (AT&T)]

- **Self-Verifying CSFQ (new)**

Most Internet routers use FIFO packet scheduling. As a results, flows that use more aggressive congestion control algorithms obtain larger bandwidth shares. This requires the community to adopt a universal standard for congestion control (the TCP-friendly paradigm) and leaves the network vulnerable to misbehaving flows that don't adhere to this standard. One way to protect the network from such misbehaving flows, and to lift the requirement for a single standard congestion control algorithm, is to have the routers allocate bandwidth in a max-min fair manner. In the past, such max-min fair proposals have required complicated per-flow scheduling. The recent work on *Core-Stateless Fair Queueing* (CSFQ) avoided these complications but imposed the need for a core-edge distinction and left the network vulnerable to malfunctioning routers. In this project, we extend the original CSFQ approach to allow all routers to verify, through sampling, the accurate labelling of incoming packets. In this approach, hosts do the rate estimation themselves (whereas previously it was left to edge routers) and the network *trusts but verifies* these rate estimates.

[Scott Shenker, Ion Stoica (UCB), Hui Zhang (CMU)]

- **Approximate Fairness through Differential Dropping (new)**

Another approach to the problem posed above is to perform the rate estimation at each router. At first glance, this appears to require too much state and processing at the router. However, in this work we show how one can achieve high degrees of fairness (max-min fairness) without complicated scheduling or processing, and without keeping excessive amounts of state.

[Scott Shenker, Lee Breslau (AT&T), Rong Pan (Stanford), Balaji Prabhakar (Stanford)]

- **Controlling High-Bandwidth Flows at the Congested Router (new)**

One weakness of the FIFO scheduling typical of routers in the current Internet is that there is no protection against misbehaving flows that send more than their share, or fail to use conformant end-to-end congestion control. This work is an investigation of a mechanism to use the packet drop history at the router to detect high-bandwidth flows in times of congestion, and to preferentially drop packets from these high-bandwidth flows in order to control the bandwidth received by these flows

at the congested queue. Extensive simulations show the effectiveness of the proposed mechanism. We have completed a draft paper [69], and we expect to have a finished paper on this topic by the end of the year.

[Sally Floyd, Ratul Mahajan (intern), Scott Shenker]

- **Network-Assisted Congestion Control (new)**

Deployed congestion control mechanisms seem to exist at two extremes of a spectrum—TCP works entirely end-to-end with only a binary signal of congestion (packet loss or ECN) provided by the network, whereas ATM ABR service is entirely network driven with the switches indicating the desired transmission rate to the sender.

We are interested in congestion control mechanisms that lie in the middle of this spectrum, where the network elements can provide more explicit feedback to the end systems than occurs with TCP. This explicit feedback gives faster convergence to a fair state than TCP can manage, but where unlike with ABR, the network elements do not keep any per-flow state, but simply attempt to actively manage their queue and available bandwidth. The motivation for this work is to examine possible mechanisms for congestion control for future high-bandwidth network environments where the available per-flow delay-bandwidth product is large compared to the length of a typical transfer. In such an environment, the existing TCP mechanisms have difficulty utilizing the available capacity because most flows will terminate before finishing the slow-start phase.

[Mark Handley, Dina Katabi (intern)]

1.2.3 Multicast

- **PIM Sparse Mode (ongoing)** PIM Sparse-Mode is the most widely deployed multicast routing protocol on the Internet, but since its deployment, multicast service has become increasingly unstable. Part of the reason for this was that the PIM-SM specification was sufficiently poorly written that router vendors could not correctly implement from it, or verify that their implementations conformed to specification. To rectify these issues, in collaboration with Bill Fenner from AT&T, we have re-written the specification from scratch. This involved figuring out how the protocol was supposed to work, and during this process we discovered many aspects where the protocol could not previously have worked correctly. We now believe that many of these failures of the original design are responsible for the unstable behavior of the deployed multicast infrastructure.

We have completely re-designed the state model used by PIM-SM so as to rectify the deficiencies we discovered, whilst at the same time attempting to remain protocol-compatible with existing deployed infrastructure. This process is now almost complete, and the revised PIM-SM protocol is now specified in an Internet Draft[37] which should become a IETF Proposed Standard during the next year. Once PIM-SM implementations are modified to implement the new state model, we believe that multicast stability will greatly improve. Cisco have now implemented the new specification in their next-generation router operating system “IOS MX”.

The current trend for interdomain multicast routing is towards source-specific multicast (SSM), where the receivers join a multicast group and also specify the IP address of the source or sources from which they wish to receive traffic. The new PIM-SM specification is organized in such a manner that the strict subset of PIM-SM needed to support SSM is clearly separable from the additional mechanisms needed to support the full Internet Standard multicast model.

[Mark Handley, Bill Fenner (AT&T), Isidor Kouvelas and Hugh Holbrook (Cisco)]

- **Bidirectional PIM (ongoing)**

Whilst the trend for interdomain multicast routing is towards source-specific multicast (SSM), the full Internet Standard multicast model is still of great relevance intra-domain, and particularly within corporate intranets. In such contexts, multicast is often used for resource-discovery protocols, where a client attempts to discover the closest instance of a service. PIM-SM is not well-suited to such tasks because the initial traffic from a source is encapsulated to a router known as the Rendezvous Point (RP), before flowing to the receivers. Thus resource-discovery protocols tend to discover servers close to the RP instead of servers close to the client.

To address this deficiency, we have developed Bi-directional PIM, a variant of PIM-SM that constructs shared bi-directional trees (as opposed to PIM-SM's per-source unidirectional trees). Bi-directional PIM is also intended to copy well with groups with large numbers of multicast senders because it does not keep per-source state in the routers on the multicast distribution tree. This work is currently being implemented by Cisco Systems.

[Mark Handley, Isidor Kouvelas and Hugh Holbrook (Cisco)]

- **IGMP version 3 (ongoing)**

To deploy source-specific multicast (SSM), we require routers that support the SSM subset of PIM-SM, and we require that the hosts can indicate to the routers which sources they are interested in. The latter functionality is the task of IGMP version 3 (previous versions of IGMP didn't allow the source to be specified). To speed up the deployment of SSM and to verify the specification for IGMPv3, we have developed a complete implementation of IGMPv3 for the FreeBSD operating system. This code has been used to discover bugs in the alpha-release of Cisco's IGMPv3 router code, and to verify that the protocol works as designed. In the near future it will be rolled into the shipping FreeBSD codebase, and may also form part of Apple's MacOS X which is BSD-based.

[Wilbert de Graaf (ICSI), Mark Handley, Bill Fenner (AT&T)]

- **Yoid (new)**

If we take a broad view of the term "multicast" to mean any distribution of content to more than one machine, we find that multicast is proceeding along two distinct architectural tracks. On one track is IP multicast, which mainly targets realtime non-reliable applications, but for which hopes run high for reliable applications as well.

On the other are a plethora of open or proprietary host- or server-based approaches, each typically targeting a specific application or product line.

IP multicast suffers from a number of technical problems, lacks applications, and in general is having trouble reaching critical mass, especially regarding anything resembling a global infrastructure. Server-based approaches are valuable and widespread, but there is no synergy in terms of multiple distinct groups working within the same architecture. As a result, progress is not as fast as it could be, and consumers are strapped with multiple application-specific infrastructures to deal with.

This paper presents an architecture, called yoid, that aims to unify both tracks under a single umbrella architecture. Yoid attempts to take the best from both tracks—reliable and asynchronous distribution from the server-based track, and dynamic auto-configuration via a simple API from the IP multicast track.

A key component of yoid is that it allows a group of endhosts (the hosts where the content-consuming application resides) to auto-configure themselves into a tunneled topology for the purpose of content distribution. Yoid can run over IP multicast, but does not require it. This allows application developers to bundle yoid into their applications, giving their applications robust and scalable configuration-free out-of-the-box multicast. This is key to the initial acceptance of yoid and for allowing it to reach critical mass.

Yoid is not limited, however, to endhost-based distribution. It can also work in infrastructure servers (boxes that receive, replicate, and forward content but that are not consumers of the content). This allows improved performance for applications that require it. It can also provide other benefits such as better security. The endhost- and server-based modes of operation taken together, along with yoid's ability to utilize local islands of IP multicast, allow yoid to support the broadest possible range of applications.

[Paul Francis]

- **Reliable Multicast (ongoing)**

We are actively involved in the IRTF Reliable Multicast Research Group (Mark Handley is co-chair of the group) and the IETF Reliable Multicast Transport working group. After many years Reliable Multicast is now emerging from the research arena, and the standardization process has started in the IETF. We are involved both in shaping the way that this standardization is happening [59],[51], and in the design of the protocols being standardized[35].

[Mark Handley, Sally Floyd]

- **Hybrid Unicast/Multicast Applications (new)**

At least for the next few years we are unlikely to see ubiquitous IP multicast deployment. As a result, few application writers develop commercial multicast-based solutions despite the significant scalability advantages multicast can provide. This has led us to become interested in hybrid unicast/multicast applications, which use IP multicast where available, but which also build an application-level distribution

tree to reach unicast-only participants. Such applications not only include reliable-multicast mechanisms, but also require an application-level routing protocol (similar to a link-state protocol) so as to figure out the best hybrid path from each sender to all receivers. As a testbed for these ideas we are developing a shared-whiteboard application to allow real-world deployment and testing.

[Mark Handley]

1.2.4 CANs for Peer-to-Peer Content Distribution Systems (new)

The past few years have seen the dramatic (and surprising to many of us) rise of peer-to-peer content distribution systems. Napster and Gnutella are the most well-known and most popular examples of this genre, but new peer-to-peer content distribution systems seem to arise daily. However, most of these systems, including Napster and Gnutella, suffer from severe scalability problems.

ACIRI has been looking at how one might make such systems scalable. We have identified that a key component of scalable peer-to-peer distribution systems is a *Content-Addressable Network* (CAN). A CAN is much like a hash table; it maps "keys" onto "values" and so any host can store a file based on its (well-known) name, and then any other host can retrieve that file just by knowing its name. Note how this differs from the web, where to locate a web site you must now only know the name of the file, but the location (the web site).

While our initial motivation is scalable peer-to-peer distribution systems, we conjecture that many other large-scale, distributed systems could utilize a CAN as a core system building block. Large scale storage management systems like OceanStore and Publius, and distributed, location independent name resolution services are but two examples.

Our design of a Content-Addressable Network, based on the abstraction of routing through a logical grid of nodes, is scalable, highly fault-tolerant and completely self-organizing. We have studied the performance and robustness properties of our CAN design through simulation. In the near future, we shall look into different applications for this CAN technology and develop and test at least one such application.

[Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, Jitendra Padhye and Scott Shenker]

1.2.5 Novel Architectures

- **IPNL (new)**

We propose a solution to the IP address depletion problem. This solution does not require (unlike IPv6) any changes to the core routers; instead, it leverages the DNS (Domain Name System) to provide a routing and addressing architecture that retains the look and feel of IP to the extent possible. The key is to use FQDNs (Fully Qualified Domain Names) in an end-to-end routing mechanism that, in addition, provides benefits such as address-space separation between providers and subscribers (leading to better address-space aggregation at the provider, and better multihoming support for the subscriber), and safe mobility. Briefly, we use variable-length FQDNs as globally unique and persistent addresses to bootstrap communication between end points,

and two sets of aligned and fixed-length numeric addresses to facilitate fast routing from that point onwards. One set of addresses is locally unique and transient, while the other set of addresses comprises the usual IP address space. IPNL-routers (that exist only at the edges of the IP-network) interface between end-points which are assigned addresses drawn from the locally-unique, fast and transient address space (the first set) on one side, and the usual IP realm on the other side. Unlike other solutions, such as TRIAD, this scheme is architected to be lightweight (state at IPNL-routers only proportional to number of active local hosts, and not number of active connections) and scalable (transparent load-balancing between multiple routers exists), robust to router crashes and link failures (mechanisms for fail-over, fault-isolation, and recovery exist), and secure against a variety of attacks (such as address-spoofing, replay, and Denial of Service) against mobile hosts.

[Paul Francis, Ramakrishna Gummadi (intern)]

- **Host-Based Admission Control (new)**

The traditional approach to implementing admission control, as exemplified by the Integrated Services proposals in the IETF, uses a signaling protocol to establish reservations at all routers along the path. While providing excellent quality-of-service, this approach has limited scalability because it requires routers to keep per-flow state and to process per-flow reservation messages. In an attempt to implement admission control without these scalability problems, several recent papers have proposed various forms of endpoint admission control. In these designs, the hosts (the endpoints) probe the network to detect the level of congestion; the host admits the flow only if the level of congestion is sufficiently low. This paper is devoted to the study of such algorithms. We first consider several architectural issues that guide (and constrain) the design of such systems. We then use simulations to evaluate the performance of such designs in various settings.

[Scott Shenker, Lee Breslau (AT&T), Ed Knightly (Rice), Ion Stoica (UC Berkeley), Hui Zhang (CMU)]

- **Future Generation Internet Architecture (new)**

There is a belief amongst many Internet researchers that incremental development and the introduction of Network Address Translators, firewalls and layer-4 switching has significantly eroded the Internet architecture to the point where we are starting to see significant problems from feature interaction, and that some required future developments will be almost impossible to deploy. With researchers from MIT and USC/ISI we are involved in a project to redesign the Internet architecture from scratch. We do not believe that this is likely to result in completely new architecture that will actually replace the Internet, but rather that the exercise of working through the basic principles is likely to result in significant new ideas that may eventually prove useful in the existing or future Internet.

[Scott Shenker, Mark Handley, Bob Braden (ISI), Dave Clark (MIT), John Wroclawski (MIT)]

1.2.6 Scalable Routing for Wireless Datagram Networks (new)

Distributed shortest-path routing protocols for wired networks either describe the entire topology of a network or provide a digest of the topology to every router. They continually update the state describing the topology at all routers as the topology changes to find correct routes for all destinations. Hence, to find routes robustly, they generate routing protocol message traffic proportional to the product of the number of routers in the network and the rate of topological change in the network. Current ad-hoc routing protocols, designed specifically for mobile, wireless networks, exhibit similar scaling properties. It is the reliance of these routing protocols on state concerning all links in the network, or all links on a path between a source and destination, that is responsible for their poor scaling.

We present Greedy Perimeter Stateless Routing (GPSR), a novel routing protocol for wireless datagram networks that uses the *positions* of routers and a packet's destination to make packet forwarding decisions. GPSR makes *greedy* forwarding decisions using only information about a router's immediate neighbors in the network topology. When a packet reaches a region where greedy forwarding is impossible, the algorithm recovers by routing around the *perimeter* of the region. By keeping state only about the local topology, GPSR scales better in per-router state than shortest-path and ad-hoc routing protocols as the number of network destinations increases. Under mobility's frequent topology changes, GPSR can use local topology information to find correct new routes quickly. We describe the GPSR protocol, and use extensive simulation of mobile wireless networks to compare its performance with that of Dynamic Source Routing. Our simulations demonstrate GPSR's scalability on densely deployed wireless networks.

[Brad Karp (intern, and ICSI)]

1.2.7 Extensible Open Router Platform (new)

IP router software currently differs from software written for desktop computers in that it tends to be written for a specific manufacturer's router family rather than being portable across systems from many different manufacturers. Whilst this makes sense in the case of code that interacts closely with high-speed forwarding engines, this observation also holds true for higher-level software such as routing protocols and management software. In addition, the APIs that would enable such higher-level software to be written by third parties are typically also not made public.

We plan to develop an open and extensible software platform for routers that might change this router software development model.

There are two main parts to such a platform:

1. Higher-level routing code comprising the routing protocols, routing information bases, and management software necessary to exist in today's Internet.
2. Low-level kernel code comprising an efficient forwarding path, together with the appropriate APIs to talk to routing code and to talk to additional higher-level functionality that is likely to be added later.

At the higher-level, the project needs to develop an architecture for interconnecting higher-level protocol modules in a manner that is efficient, but also modular and flexible. The

APIs between these modules need to be carefully designed and well specified to allow third-parties to contribute new modules that extend the functionality of the router. It is important that such extensions can be in binary-only form so that a wide range of third-party business models are possible. It is also important that a router can use well-verified trusted routing modules at the same time as experimental third-party modules without unduly compromising the stability of the router. In this way such a router platform would stimulate early experimentation and deployment of new network functionality. For good operational reasons, this is extremely difficult in the current Internet.

At the low level the architecture needs to be capable of spanning a large range of hardware forwarding platforms, ranging from commodity PC hardware at the low end, through mid-range PC-based platforms enhanced with intelligent network interfaces, to high-end hardware-based forwarding engines. Initially we plan to focus on a PC-based hardware platform as this allows for easy testing and early deployment, but great care will be taken to design a modular forwarding path architecture which can support some or all of the forwarding functions happening in hardware.

The primary motivation for ACIRI's involvement with this is that it is becoming increasingly hard for researchers to influence the development of protocols that require router support unless we can convince a router vendor that they have a business case for the protocol in question. However we also believe that if the project is completely successful, it may result in a change in the way that commercial router software is developed, allowing more competition and resulting in lower costs and better software for ISPs.

This project will be co-funded by Intel.

[Mark Handley]

1.3 Research on Internet Operations

The two main areas of research in Internet operations are (1) Measurement and Analysis, and (2) Network Intrusion Detection.

1.3.1 Measurement and Analysis

- **Measurement Infrastructure (ongoing)**

Historically, the Internet has been woefully under-measured and under-instrumented. The problem is only getting worse with the network's ever-increasing size. The National Internet Measurement Infrastructure project (NIMI), for which ACIRI provides technical leadership, aims to develop a scalable architecture for deploying and operating measurement infrastructures, i.e., a collection of measurement "platforms" that can cooperatively measure the properties of Internet paths and clouds by transmitting test traffic among themselves. The architecture emphasizes decentralized control of measurements; strong authentication and security; mechanisms for both maintaining tight administrative control over who can perform what measurements using which platforms, and delegation of some forms of measurement as a site's measurement policy permits; and simple configuration and maintenance of platforms. NIMI currently has 45 hosts.

[Vern Paxson, with Andrew Adams (Pittsburgh Supercomputing Center)]

- **Multicast Inference of Network Characteristics (ongoing)**

This project seeks to infer the properties of individual network links from the end-to-end drop and delay characteristics of multicast traffic. Such “tomography” approaches offer the intriguing possibility of pinpointing network performance problems without requiring ubiquitous measurement access to individual routers.

[Vern Paxson, with Nick Duffield (AT&T), Don Towsley (UMass), and numerous others]

- **A Benchmark for Content Distribution (new)**

In the last few years content distribution networks (CDNs) have emerged as a mechanism to deliver content to end users on behalf of origin sites. To date, there has been little work to soundly assess the performance of CDNs. This project strives to develop a sound methodology for benchmarking CDNs based on evaluating the different factors affecting the time for clients to load a “canonical page” from a CDN’s servers. The methodology focuses on the end user’s perceived latency and is suitable for CDNs using popular techniques, such as DNS redirection, that are visible at the Web browser client.

[Vern Paxson, Yin Zhang, with Balachander Krishnamurthy (AT&T), Craig Wills (WPI)]

- **Stationarity of Internet Path Properties (new)**

A basic question concerning how the Internet behaves is the extent to which network properties are well-modeled as stationary. Using data gathered from the NIMI measurement infrastructure, we have investigated the degree to which network routing, packet loss, packet delay, and TCP throughput are stable, using several different definitions of stationarity.

[Vern Paxson, Yin Zhang (intern), and Scott Shenker]

- **DNS analysis (new)**

This effort aims to update the groundbreaking DNS study by Danzig and colleagues, published in Proc. SIGCOMM '92. We have begun gathering traces from a wide variety of perspectives within the DNS, including dialup end-users, campus name servers, top-level domain servers, regional and backbone perspectives, and root name servers.

[Vern Paxson, Mark Handley, Brad Karp (intern, and ICSI)]

- **Impact of Routing Policy on Internet Paths (new)**

The impact of routing policy on Internet paths is poorly understood. Policy can *inflate* shortest-router-hop paths but, to our knowledge, the extent of this inflation has not been previously examined. Using a simplified model of routing policy in the Internet, we obtain approximate indications of the impact of policy routing on

Internet paths. Our findings suggest that routing policy does impact the length of Internet paths significantly. For instance, in our model of routing policy, some 20% of Internet paths are inflated by more than five router-level hops.

[Scott Shenker, Ramesh Govindan (USC), Deborah Estrin (UCLA), Hongsuda Tangmunarunkit (USC)]

- **TCP Behavior Inference Tool (new)**

There are a range of TCP congestion control behaviors in deployed TCP implementations, include Tahoe, Reno, NewReno, and Sack TCP, which date from 1988, 1990, 1996, and 1996, respectively. We recently asked the question “What fraction of the non-SACK TCP flows in the Internet use NewReno instead of Reno congestion control mechanisms? ”. One reason for asking the question “Who uses NewReno TCP?” is to better understand the migration of new congestion control mechanisms to the public Internet. A second reason to ask the question is to discourage network researchers from extensive investigations of the negative impacts of Reno TCP’s poor performance when multiple packets are dropped from a window of data, if in fact Reno TCP is already being replaced by NewReno TCP in the Internet. A third reason to investigate the TCP congestion control mechanisms actually deployed in the public Internet is that it is always useful to occasionally step back from our models, analysis, and simulations, and look at the Internet itself.

Inferring the congestion control behavior of a remote host can be done to some extent by actively initiating a transfer of data from a remote host over a TCP connection, and then passively monitoring the connection’s congestion control responses to packet drops on congested links within the network, if in fact packets are dropped. Passive monitoring was used by Paxson in previous work. However, it is difficult to determine TCP congestion control behaviors through passive monitoring, because one has to wait for the desired pattern of packet drops to occur.

We have developed a tool, called the TCP Behavior Inference Tool (TBIT) for answering this and related questions about TCP behavior. The tool uses some of the code developed by the Sting project. TBIT allows the user to control the receipt and sending of TCP packets at the local host, thus introducing specific packet drops at the host itself. Apart from testing for the congestion control “flavor”, TBIT is capable of testing for correctness of SACK implementation, use of timestamps, ECN capabilities and several other TCP behaviors. Both the TBIT source code and the experimental results are available from the TBIT Web Page [70]. TBIT has already helped detect and correct bugs in Microsoft, Cisco and IBM products.

We plan to continue the TBIT work in several ways:

- Better correlation of TBIT results with operating systems. We are currently running NMAP scans on over 27,000 hosts to correlate TBIT results with specific operating systems. This project has been made possible by kind assistance of Aaron Hughes of `bind.com`.
- Evolve TBIT into a complete TCP conformance testing tool. Such a tool should be of interest to OS vendors as well as ISPs that provide web hosting facilities.

Graduate students from Columbia University, working under the guidance of Prof. Daniel Rubenstein have expressed interest in helping us add more tests to TBIT to achieve this goal.

- Automatic generation of NS models. TBIT provides a characterization of TCP on a remote web server, and this characterization can be used to generate an NS model of that TCP variant.

[Sally Floyd, Jitendra Padhye]

1.3.2 Network Intrusion Detection

- **Bro (ongoing)**

The Bro project (LBNL and ACIRI) is a network intrusion detection system. It sniffs packets coming across a network link such as DMZ or a sensitive LAN and uses an event engine to analyze the traffic and extract from it events at different levels (e.g., connection attempted; user authenticated; FTP file retrieve request; new line of Telnet output). It then determines whether the traffic is consistent with the site's policy by running the series of events as input to a script that expresses the policy in a domain-specific language. The script can maintain and modify global state, record information to stable storage, synthesize new events, generate real-time alerts, and invoke shell scripts as a form of "reactive firewall".

Bro is currently operational at several sites (ICSI, LBNL, UCB, NERSC, ESNET, JGI). The code, significantly extended over this past year, is freely available.

[Vern Paxson]

- **Detecting stepping stones (new)**

One widely-used technique by which network attackers attain anonymity and complicate their apprehension is by employing "stepping stones": they launch attacks not from their own computer but from intermediary hosts that they previously compromised. We developed an efficient algorithm for detecting stepping stones by monitoring a site's Internet access link. The algorithm is based on the distinctive characteristics (packet size, timing) of interactive traffic, and not on connection contents, and hence can be used to find stepping stones even when the traffic is encrypted. We evaluated the algorithm on large Internet access traces and found that it performs quite well. However, the success of the algorithm is tempered by the discovery that large sites have many users who routinely traverse stepping stones for a variety of legitimate reasons. Hence, stepping-stone detection also requires a significant policy component for separating allowable stepping-stone pairs from surreptitious access.

[Vern Paxson, Yin Zhang (intern)]

- **Detecting backdoors (new)**

"Backdoors" are often installed by attackers who have compromised a system to ease their subsequent return to the system. We considered the problem of identifying a large class of backdoors, namely those providing interactive access on non-standard

ports, by passively monitoring a site's Internet access link. We developed a general algorithm for detecting interactive traffic based on packet size and timing characteristics, and a set of protocol-specific algorithms that look for signatures distinctive to particular protocols. We evaluated the algorithms on large Internet access traces and find that they perform quite well. In addition, some of the algorithms are amenable to prefiltering using a stateless packet filter, which yields a major performance increase at little or no loss of accuracy. However, the success of the algorithms is tempered by the discovery that large sites have many users who routinely access what are in fact benign backdoors, such as servers running on non-standard ports not to hide, but for mundane administrative reasons. Hence, backdoor detection also requires a significant policy component for separating allowable backdoor access from surreptitious access.

[Vern Paxson, Yin Zhang (intern)]

- **ITREX (new)**

ITREX (Internet Trap and Trace Experiments) is a collaborative effort to attempt to traceback attackers across the Internet, essentially a generalization of the stepping stone detection problem to the setting of multiple sites.

[Vern Paxson, Stuart Staniford (Silicon Defense), and Felix Wu (UC Davis)]

- **Normalizer (ongoing)**

There are many ways one can perturb traffic to make it difficult for a network intrusion detection system to unambiguously interpret the packet stream; for instance, a retransmission can contain different data than the original packet. This is a fundamental problem for network intrusion detection, in the sense that if it is not solved, network intrusion detection is doomed—because just as attackers today use a wide variety of toolkits to automate their attacks, so too will they use a wide variety of evasion toolkits to escape detection. Indeed, such toolkits already exists, and Phrack magazine, the effective journal of record for the attacker underground, has already published an article on the topic.

A traffic normalizer is a box that takes a packet stream and attempts to reshape the traffic flowing through it in order to eliminate ambiguities that can facilitate evasion, but while leaving the end-to-end semantics of the stream unperturbed. A draft paper has been written and further work is continuing.

[Vern Paxson, Mark Handley, Christian Kreibich (ICSI)]

- **Reflectors (new)**

A potentially devastating counter that distributed denial-of-service attackers can use to attempt to thwart proposed techniques for tracking back spoofed traffic flows to their source is to bounce their attack traffic off of "reflectors", i.e., any Internet host that when sent a packet will return one in reply. In this work we discuss the general problem and analyze the different types of reflectors available in the Internet to assess the degree to which we can defend against reflector attacks.

1.4 Research on Internet Theory

We have several Internet theory projects; they don't fall into obvious categories, so each project will be listed separately.

1.4.1 Randomized Rumor Spreading (new)

This project deals with the problem of spreading rumors in a distributed environment using randomized communication. In particular, we envisage the class of so-called epidemic algorithms which are commonly used for the lazy transmission of updates to distributed copies of a database. We introduce the *random phone call model* in order to investigate the possibilities and limits of this class of broadcasting algorithms. In this model, n players communicate in parallel communication rounds. In each round, each player calls a randomly selected communication partner. Whenever communication is established between two players, each one must decide which rumors to transmit. The major problem (arising due to the randomization) is that players do not know which rumors their communication partners have already received. In order to illustrate this problem, we will give a simple example of a commonly used algorithm in which each individual rumor is transmitted $\Theta(n \ln n)$ times.

In this paper, we investigate whether a large communication overhead is inherent to epidemic algorithms using randomized rumor spreading or can be reduced significantly. We show that there is an algorithm using only $O(\ln n)$ rounds and $O(n \ln \ln n)$ transmissions. We prove the robustness of this algorithm against adversarial node failures and inaccuracies in the randomized selection of the communication partners. Furthermore, we show that our algorithm is optimal among those algorithms in which the actions of the players do not depend on the addresses of their communication partners. Finally, we give a lower bound for general algorithms showing that time- and communication-optimality cannot be achieved simultaneously. In particular, we prove that any algorithm (based on the random phone call model) that distributes a rumor in $O(\ln n)$ rounds needs to send $\omega(n)$ messages on expectation.

[Richard Karp, Scott Shenker, C. Schindelhauer (ICSI), B. Vöcking (ICSI)]

1.4.2 Internet Topology (new)

In the past decade, traffic measurement and modelling has become an established and important tool in understanding the Internet. Much more recently, but in a similar manner, the measurement and modelling of Internet topology has started to receive significant attention. In this project, we are using the recently obtained empirical maps of the Internet to investigate how best to model the Internet topology. We find that a simple *power-law random graph* (PLRG) is the best model (among current contenders) for describing the Internet, and that such reasonably describe many other large-scale graphs (web, call-graph, airline routing, etc.).

[Scott Shenker, Walter Willinger (AT&T), Ramesh Govindan (USC), Sugih Jamin (U. Mich.), Hongsuda Tangmunarunkit (USC)]

1.4.3 Learning and Game Theory in Internet-like Settings (new)

Game theory has often been applied to the Internet, both for the design of low-level protocols and for the interaction of selfish agents over the Internet. Typically these treatments adopt the canonical notions of game-theoretic equilibria (Nash, etc.). However, the motivations for these equilibrium notions (full information, etc.) do not apply to the Internet. In this project we examine through experiments the implications of the Internet setting (low information, asynchronous play, etc.) to notions of equilibrium. For a dominance solvable version of a Cournot oligopoly with differentiated products, we find little evidence of convergence to the Nash equilibrium. In an asynchronous setting, play tends toward the Stackelberg outcome. Convergence is significantly more robust for a “Serial Cost Sharing” game (which resembles Fair Queueing), which satisfies a stronger solution concept of overwhelmed solvability. However, as the number of players grows, this improved convergence tends to diminish. This seems to be driven by high and correlated experimentation or noise and demonstrates that even when play converges, the convergence times may be too long to be of practical significance.

[Scott Shenker, Eric Friedman (Rutgers), Barry Sopher (Rutgers), Mike Shor]

1.4.4 Application Layer Multicast (new)

Multicast is a basic Internet function enabling any member of a group of Internet hosts to send a message to all the other members. Multicast is typically implemented in the IP layer by creating a tree of links and routers. A message can then be transmitted from any group member to all the others by transmitting the packets of the message across the tree, with each packet traversing each link in the tree once. IP multicast has been difficult to deploy because of issues of scalability, congestion control, error recovery and network management. Recently, it has been suggested (e.g., Francis’s work on Yoid or a recent paper by Chu, Rao, and Zhang) that multicast can be implemented in the application layer by configuring the group members into a ‘virtual tree’ and distributing messages by unicast transmissions among the nodes of this tree. Within a graph-theoretic abstraction Richard Karp and Claire Kenyon have studied the relative efficiency of this approach compared to IP multicast.

The vertices of the virtual tree T are the group members, and each edge represents a path in the Internet between the group members at the ends of the edge. A multicast message can be simulated by a single unicast along each edge of the virtual tree.

Several criteria are used to measure the performance of simulated multicast on the virtual tree T . $SIZE(T)$ denotes the total number of hops required to multicast to all group members. $DELAY(T)$ denotes the maximum number of hops between any two nodes in T , and $STRESS(T)$ is the maximum load that T imposes on any link in the Internet.

We would like to construct a virtual tree whose size and delay are not much greater than the size and delay of conventional IP multicast, and whose stress is small. We give

a construction which simultaneously achieves small size and small delay, and a second construction simultaneously achieving small size and small stress. These constructions require knowledge of the number of hops between any two group members, but do not require knowledge of the actual routes between the group members. We also show that there is an inherent antagonism between stress and delay; it is not possible in general to achieve both small stress and small delay, even if the routes between group members are known.

[Richard Karp, Claire Kenyon (University of Paris)]

1.4.5 Fairness in AIMD Congestion Control (new)

An Additive-Increase Multiplicative-Decrease (AIMD) transmission policy can be described by two parameters, the slope A and the drop fraction D ; A is the amount by which the flow rate increases per successful packet transmission, and D is the fraction by which the rate is reduced after a packet drop. A key question is how the average bandwidth available to a flow depends on these two parameters. Previous studies of this question have used simple stochastic or deterministic models of the occurrence of packet drops.

We have conducted the first investigation of AIMD performance under an adversarial model of packet drops. They consider a situation in which several concurrent AIMD flows with different parameters are subjected to packet drops at the same times. Their main conclusion is that, even under this worst-case adversarial model, it is not possible for one flow to deprive another flow of its fair share of bandwidth by aggressively increasing its slope or reducing its drop fraction. They show that, even if an adversary can choose the times of the packet drops, the average bandwidth available to a flow is proportional to its slope, and the ratio by which a flow can increase its average bandwidth by reducing its drop fraction from D_1 to D_2 cannot exceed an absolute constant times the ratio D_1/D_2 .

[Richard Karp, Michael Luby (Digital Fountain), and Avi Wigderson (IAS)]

1.5 Community Activities

ACIRI members engage in a wide range of community activities. While many of these activities are quite standard (program committees, editorial boards, journal reviewing, etc.) and won't be listed here, this report will focus on the activities that are more central to the Internet community.

1.5.1 IETF

The Internet Engineering Task Force (IETF) is the most relevant standards body for the Internet. Sally Floyd, Mark Handley, and Vern Paxson serve on the Transport Area Directorate, an advisory body for the Transport Area, and are otherwise quite active in the IETF. Sally Floyd and Mark Handley attended the IAB Workshop on Wireless Internetworking in March. Vern Paxson chairs the working group on Endpoint Congestion Management, and finished his term on the IESG as area director for the Transport area in March. The TCP Implementation working group, which he co-chaired, concluded its work items and closed in October.

1.5.2 NANOG

The North American Network Operators Group (NANOG) provides a forum for the exchange of technical information among network operators. Sally Floyd and Jitendra Padhye attended the October NANOG meeting. Jitu gave a presentation entitled “TBIT: TCP Behavior Inference Tool” and Sally gave a presentation entitled “A Request for Information About the Deployment of RED.”

1.5.3 SIGCOMM and IRTF

The Internet Research Task Force (IRTF) is the research counterpart of the IETF. Sally Floyd and Scott Shenker are members of the IRTF End-to-End Research Group, and Vern Paxson and Mark Handley are frequent invitees to its meetings. Mark Handley is chair of the IRTF Reliable Multicast Research Group, and also a member of the Internet Research Steering Group (IRSG).

SIGCOMM is ACM’s special interest group on communications, and is the primary academic research organization for the Internet community. Sally Floyd is the current chair of the SIGCOMM Technical Advisory Committee, and Scott Shenker is the outgoing chair and continues to serve on the committee. Several members of ACIRI serve quite regularly on the SIGCOMM Program Committee.

1.5.4 NS simulator

The widely-used NS simulator [4] is based on the original simulator developed at UCB by Steve McCanne and Sally Floyd. Sally Floyd has been involved in the development of NS from the project’s inception, and continues to participate in the bi-weekly NS meetings and other general oversight activities. We are also actively developing new functionality and new validation test suites for NS.

References

Papers

- [1] A. Adams, T. Bu, R. Caceres, N. Duffield, T. Friedman, J. Horowitz, F. Lo Presti, S.B. Moon, V. Paxson, and D. Towsley. “The Use of End-to-End Multicast Measurements for Characterizing Internal Network Behavior”, *IEEE Communications*, special issue on network measurement, 38(5), May 2000.
- [2] M. Adler, J. Byers, and R. Karp, “Parallel Sorting with Limited Bandwidth”, *SIAM J. Comput.* 29(6): 1997-2015 (2000).
- [3] A. Ben-Dor, R. Karp, B. Schwikowski, and Z. Yakhini, “Universal DNA Tag Systems: A Combinatorial Design Scheme”, *Journal of Computational Biology*, forthcoming. (Preliminary version in RECOMB 2000.)
- [4] L. Breslau et al., “Advances in Network Simulation”, *IEEE Computer*, May 2000, pp. 59-67. (This is an outgrowth of the USC Technical Report 99-702 on “Improving Simulation for Network Research” reported last year.)
- [5] L. Breslau, E. Knightly, S. Shenker, I. Stoica, and H. Zhang, “Endpoint Admission Control: Architectural Issues and Performance”, *SIGCOMM*, August 2000.
- [6] L. Breslau, S. Jamin, and S. Shenker, “Comments on the Performance of Measurement-Based Admission Control Algorithms”, To appear in *Proceedings Infocom 2000*.
- [7] A. Condon and R. Karp, “Algorithms for Graph Partitioning on the Planted Partition Model”, *Random Structures and Algorithms*, to appear. (Preliminary version in RANDOM-APPROX 1999: 221-232.)
- [8] P. Dagum, R. Karp, M. Luby, and S. Ross, “An Optimal Algorithm for Monte-Carlo Estimation” *SIAM J. Comput.*, 29(5):1484-1496, 2000.
- [9] R. El-Yaniv, A. Fiat, R. Karp and G. Turpin, “Optimal Search and One-Way Trading Online Algorithms”, *Algorithmica*, forthcoming.
- [10] D. Estrin, M. Handley, J. Heidemann, S. McCanne, Y. Xu, and H. Yu, “Network Visualization with the Nam, VINT Network Animator”, *IEEE Computer*, 33 (11), pp. 63-68, November, 2000.
- [11] J. Feigenbaum, C. Papadimitriou, and S. Shenker, “Sharing the cost of multicast transmissions”, *STOC* 2000.

- [12] S. Floyd and V. Paxson, "Difficulties in Modeling and Simulating the Internet", *IEEE/ACM Transactions on Networking*, to appear. (This is an outgrowth of the paper "Why We Don't Know How To Simulate The Internet" reported last year.)
- [13] S. Floyd, M. Handley, J. Padhye, and J. Widmer, "Equation-Based Congestion Control for Unicast Applications", SIGCOMM, August 2000.
- [14] S. Floyd, "A Report on Some Recent Developments in TCP Congestion Control", *IEEE Communications Magazine*, to appear April 2001.
- [15] M. Handley, J. Crowcroft, "Internet Multicast Today", in *The Internet Protocol Journal*, Dec 1999.
- [16] T.E. Ideker, V. Thorsson and R. Karp, "Discovery of Regulatory Interactions Through Perturbation: Inference and Experimental Design", Pacific Symp. Biocomputing (2000)
- [17] B. Karp and H. T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless networks", MOBICOM 2000.
- [18] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking, "Randomized Rumor Spreading", FOCS 2000.
- [19] R. Karp, E. Koutsoupias, C. Papadimitriou, and S. Shenker, "Optimization Problems in Congestion Control", FOCS 2000.
- [20] R. Karp, C. Kenyon, and O. Waarts, "Error-Resilient DNA Computation", *Random Structures and Algorithms* 15(3-4): 450-466 (1999).
- [21] R. Karp and R. Shamir, "Algorithms for Optical Mapping", *Journal of Computational Biology* 7(1,2) (2000).
- [22] R. Karp, I. Pe'er, and R. Shamir, "An Algorithm Combining Discrete and Continuous Methods for Optical Mapping", *Journal of Computational Biology*, forthcoming. (Preliminary version in ISMB 159-68 (1999).)
- [23] B. Krishnamurthy, V. Paxson, C. Wills, Y. Zhang, "ABCD: A Benchmark for Content Distribution", in submission.
- [24] V. Paxson, A. Adams, and M. Mathis, "Experiences with NIMI", Proceedings of Passive & Active Measurement: PAM-2000, to appear.
- [25] V. Paxson, "On Calibrating Measurements of Packet Transit Times", in revision.
- [26] R. Rejaie, H. Yu, M. Handley, D. Estrin, "Multimedia Proxy Caching for Quality Adaptive Streaming Applications in the Internet" in Proceedings of Infocom 2000.
- [27] R. Rejaie, M. Handley, D. Estrin, "Layered Quality Adaptation for Internet Video Streaming", to appear in IEEE Journal on Selected Areas of Communications (JSAC).

- [28] M. Shor, E. Friedman, S. Shenker, B. Sopher, “Asynchronous Learning with Limited Information”, in submission
- [29] H. Tangmunarunkit, R. Govindan, S. Shenker, D. Estrin, “The Impact of Routing Policy on Internet Paths”, Infocom 2001, to appear.
- [30] D. Thaler, M. Handley “On the Aggregatability of Multicast Forwarding State”, in Proceedings of Infocom 2000.
- [31] W. Willinger, V. Paxson, R. Riedi and M. Taqqu, “Long-Range Dependence and Data Network Traffic”, in submission.
- [32] Y. Zhang, V. Paxson, and S. Shenker, “The Stationarity of Internet Path Properties: Routing, Loss, and Throughput”, ACIRI Technical Report, May 2000.
- [33] Y. Zhang and V. Paxson, “Detecting Stepping Stones”, Proc. 9th USENIX Security Symposium, August 2000.
- [34] Y. Zhang and V. Paxson, “Detecting Backdoors”, Proc. 9th USENIX Security Symposium, August 2000.

Internet Drafts

- [35] B. Adamson, C. Bormann, S. Floyd, M. Handley, J. Macker, “NACK-Oriented Reliable Multicast (NORM) Protocol Building Blocks”, Internet-draft draft-ietf-rmt-norm-bb-00.txt, July 2000.
- [36] M. Allman, H. Balakrishnan, and S. Floyd, “Enhancing TCP’s Loss Recovery Using Limited Transmit”, Internet draft draft-ietf-tsvwg-limited-xmit-00.txt, August 2000.
- [37] B. Fenner, M. Handley, H. Holbrook, I. Kouvelas, “Protocol Independent Multicast—Sparse Mode (PIM-SM): Protocol Specification (Revised)”, Internet-draft draft-ietf-pim-sm-v2-new-00.txt, 13 July 2000.
- [38] S. Floyd, David Black, and K.K. Ramakrishnan, “IPsec Interactions with ECN”, Internet-draft draft-ietf-ipsec-ecn-02.txt, December 1999. (This is a revised version of draft-ietf-ipsec-ecn-01.txt, October 1999, reported last year.)
- [39] Floyd, S. and K.K. Ramakrishnan, “ECN Interactions with IP Tunnels”, Internet draft draft-floyd-ecn-tunnels-01.txt, October 2000.
- [40] Floyd, S. and K.K. Ramakrishnan, “TCP with ECN: The Treatment of Retransmitted Data Packets”, Internet draft draft-ietf-tsvwg-tcp-ecn-00.txt, November 2000.
- [41] M. Handley, I. Kouvelas, L. Vicisano, “Bi-directional Protocol Independent Multicast”, Internet-draft draft-ietf-pim-bidir-00.txt, March 1, 2000.
- [42] M. Handley, H. Schulzrinne, E. Schooler, J. Rosenberg, “SIP: Session Initiation Protocol”, Internet-draft draft-ietf-sip-rfc2543bis-01.txt, August 6, 2000

- [43] M. Handley, S. Hanna, “Multicast Address Allocation Protocol (AAP)”, Internet-draft draft-ietf-malloc-aap-04.txt, June 2000.
- [44] M. Handley, J. Crowcroft, C. Bormann, J. Ott, “The Internet Multimedia Conferencing Architecture”, Internet-draft draft-ietf-mmusic-confarch-03.txt, July 2000.
- [45] M. Handley, V. Jacobson, and C. Perkins, “SDP: Session Description Protocol”, draft-ietf-mmusic-sdp-new-00.txt, Nov 2000.
- [46] M. Handley, J. Padhye, S. Floyd, and J. Widmer, “TCP Friendly Rate Control (TFRC):Protocol Specification”, draft-ietf-tsvwg-tfrc-00.txt,.ps, Nov 2000.
- [47] M.Luby, J.Gemmell, L.Vicisano, L.Rizzo, M.Handley, and J. Crowcroft, “Layered Coding Transport: A massively scalable multicast protocol”, draft-ietf-rmt-bb-lct-00.txt, Nov 2000.
- [48] M.Luby, L.Vicisano, J.Gemmell, L.Rizzo, M.Handley, and J. Crowcroft “The use of Forward Error Correction in Reliable Multicast”, draft-ietf-rmt-info-fec-00.txt, Nov 2000.
- [49] M.Luby, L.Vicisano, J.Gemmell, L.Rizzo, M.Handley, and J. Crowcroft, “RMT BB Forward Error Correction Codes”, draft-ietf-rmt-bb-fec-02.txt, Nov 2000.
- [50] K.K. Ramakrishnan, S. Floyd, and D. Black, “The Addition of Explicit Congestion Notification (ECN) to IP”, Internet-draft draft-ietf-tsvwg-ecn-00.txt, November 2000.
- [51] B. Whetten, L. Vicisano, R. Kermode, M. Handley, S. Floyd, M. Luby, “Reliable Multicast Transport Building Blocks for One-to-Many Bulk-Data Transfer”, Internet-draft draft-ietf-rmt-buildingblocks-03.txt, October 2000

RFCs (ordered by number)

- [52] M. Handley, C. Perkins, “Guidelines for Writers of RTP Payload Format Specifications”, RFC 2736, BCP0036, Best Current Practice, December 1999.
- [53] M. Handley, D. Thaler, R. Kermode, “Multicast-Scope Zone Announcement Protocol (MZAP)”, RFC 2776, Proposed Standard, February 2000.
- [54] S. Bradner and V. Paxson, “IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers”, RFC 2780, Best Current Practice, March 2000.
- [55] M. Handley, S. Floyd, J. Padhye, “ TCP Congestion Window Validation”, RFC 2861, Experimental, June 2000. (This was reported last year as an internet-draft.)
- [56] X. Xiao, A. Hannan, V. Paxson, and E. Crabbe, “TCP Processing of the IPv4 Precedence Field”, RFC 2873, Proposed Standard, June 2000.
- [57] S. Floyd, J. Mahdavi, M. Mathis, and M. Podolsky, “An Extension to the Selective Acknowledgement (SACK) Option for TCP”, RFC 2883, Proposed Standard, July 2000. (This was reported last year as an internet-draft.)

- [58] S. Floyd, “Congestion Control Principles”, RFC 2914, Best Current Practice, September 2000. (This was reported last year as an internet-draft.)
- [59] M. Handley, S. Floyd, B. Whetten, R. Kermode, and L. Vicisano, “The Reliable Multicast Design Space for Bulk Data Transfer”, RFC 2887, Informational, August 2000. (This was reported last year as an internet-draft.)
- [60] D. Thaler, M. Handley, D. Estrin, “The Internet Multicast Address Allocation Architecture”, RFC 2908, Informational, September 2000.
- [61] P. Radoslavov, D. Estrin, R. Govindan, M. Handley, S. Kumar, D. Thaler “The Multicast Address-Set Claim (MASC) Protocol”, RFC 2909, Experimental, September 2000.
- [62] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, and V. Paxson, “Stream Control Transmission Protocol”, RFC 2960, Proposed Standard, October 2000.
- [63] M. Handley, C. Perkins, E. Whelan, “Session Announcement Protocol”, RFC 2974, Experimental, October 2000.
- [64] V. Paxson and M. Allman, “Computing TCP’s Retransmission Timer”, RFC 2988, Proposed Standard, November 2000.
- [65] B. Whetten, L. Vicisano, R. Kermode, M. Handley, S. Floyd, and M. Luby, “Reliable Multicast Transport Building Blocks for One-to-Many Bulk-Data Transfer”, Approved by the IESG as an Informational RFC on October 27, 2000. (This was reported last year as an internet-draft.)

Work in Progress

- [66] H. Balakrishnan, D. Bansal, S. Floyd, and S. Shenker, “Exploring the Range between TCP-equivalent and TCP-compatible”, In progress, November 2000.
- [67] Steve M. Bellovin, Sally Floyd, Ratul Mahajan, and Vern Paxson, “Controlling High Bandwidth Aggregates Using Pushback”, November 2000.
- [68] S. Floyd, M. Handley, and J. Padhye, “A Comparison of Equation-Based and AIMD Congestion Control”, May 2000.
- [69] R. Mahajan and S. Floyd, “Controlling High-Bandwidth Flows at the Congested Router”, November 2000.
- [70] J. Padhye and S. Floyd, “Identifying the TCP Behavior of Web Servers”, URL “www.aciri.org/tbit/”, July, 2000.
- [71] V. Paxson, “An analysis of using reflectors to defeat DOS traceback”, in preparation.
- [72] V. Paxson, M. Handley, “Network Intrusion Detection: Evasion, Traffic Normalization, and End-to-end Protocol Semantics”, White Paper, in preparation.

- [73] Ratnasamy, S., Francis, P., Handley, M., Karp, R. and Shenker, S., “A Scalable Content-Addressable Network”, In preparation.

2 Network Services and Applications

2.1 Overview

After the broad research activities of the previous year, the Network Services and Applications Group has shifted its main activities to exploration of key issues in the area of mobility and related aspects of the Internet. The goal is to contribute significantly to the following network research areas:

- Internet protocols for mobile support
- mobility management
- applicability of existing networking Quality of Services approaches for mobility
- routing algorithms

But mobility aspects are not only related to these network topics or to technologies such as wireless or Mobile IP. Other topics like ubiquitous computing, Quality of Service (QoS), security, service provisioning, and even more application-relevant issues like e-commerce also have considerable impact on mobile Internet access. In addition to the network aspects, the NSA group also focuses on the following topics:

- ubiquitous computing, offering information access beyond provider/technology boundaries.
- security, because wireless access and mobile scenarios require a higher degree of security features (e.g. authentication, authorization, accounting, encryption, unobservability) than needed by wired access.

2.2 USAIA: Ubiquitous Service Access Internet Architecture

Recent initiatives to add mobility to the Internet and packet data services to third generation cellular systems are being considered by mobile service providers as candidate technologies for the delivery of IP data to mobile users. However, both of these candidate technologies have shortcomings. Mobile IP represents a simple and scalable global mobility solution but lacks support for fast handoff control, real-time location tracking, authentication and distributed policy management found in cellular networks today. In contrast, third generation cellular systems offer seamless mobility support but are built on complex and costly connection-oriented networking infrastructures that lack the inherent flexibility, robustness and scalability found in IP networks. Future wireless networks should be capable of combining the strengths of both approaches without inheriting their weaknesses.

This motivates us to focus on all-IP networks to provide an integrated solution for the ongoing convergence of IP and mobile telecommunications in wireless networks. The inclusion of IP-centric mobile telecommunications networks presents a number of challenges that we want to address with respect to QoS. In this project we have defined a framework that permits the provision of Quality of Service (QoS) independent of the movement of users. On each identified network level (cell, AD, global) we have defined QoS mechanisms in conjunction with appropriate mobility management schemes. The interrelationship of these network levels is depicted in the following figure.

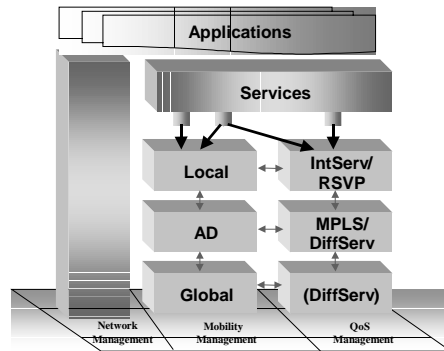


Figure 1: Interrelationship between the QoS and mobility planes of USAIA.

One of the principal challenges is to set up resources in advance (active and passive reservations, see next figure) on the cell level without overloading the air interface with control traffic. All necessary protocols between the mobile node and the base station have been defined with respect to simplicity without neglecting the end- to-end scope. For details and further information refer to [1].

Simulations, development of a Petri net for the cell level, and a prototypical implementation plan will be the final result of this project. The work is coordinated with similar activities at Siemens AG and the Georgia Institute of Technology in Atlanta.

2.3 Mobile Network Architecture for Vehicles

In this joint project with Daimler Chrysler Research Palo Alto, the current potential to provide seamless connectivity from vehicles to the global internet, assuming different air interfaces like GPRS, Richochet and WLAN, was investigated. The first steps defined an architecture and extent the network simulation tool ns with all necessary extensions [2]. The second step provided some simulations for all mentioned wireless technologies with respect to TCP as the transport protocol. We compared several TCP variants (e.g. Reno) and parameter settings (e.g packet size) to figure out the delay for handoffs between these different technologies, as seen in the next figure.

A detailed report on the achieved simulation results will be published soon.

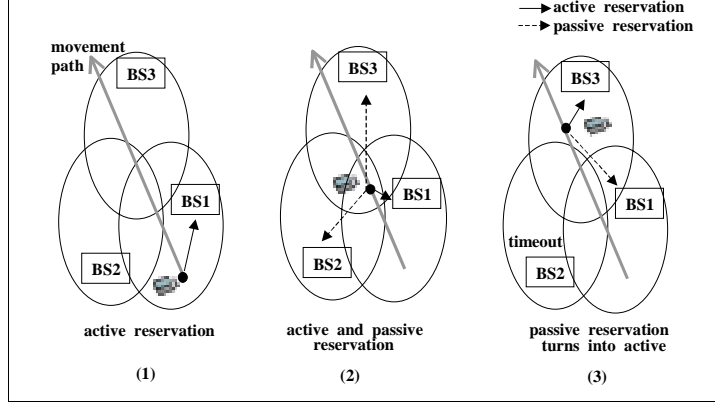


Figure 2: Active and passive reservation handling.

2.4 Routing for mobile access scenarios

2.4.1 Exploration of active routing

Conventional routing protocols work concurrently with data packet forwarding. These protocols exchange some information between routers in order to learn and/or to update the routes for data packets. Most of them have been designed to adapt to network traffic changes. Topological changes are also accounted for, but the assumption is that these changes will not happen very often. Active routing assumes there are not two different tasks such as routing and forwarding but only one. Routing functions are triggered by user data. Each active packet contains not only data but routing instructions. Therefore networks that use active routing do not have pre-established routing protocols but leave it to the users to decide what routing protocols to use. In fact, every user can opt for a different routing protocol. Active routers process the routing information contained in each active packet and interpret its routing commands. By exposing routing functions to network users, the most suitable algorithms for the desired quality of service can be chosen. At the same time, network users become aware of network load conditions.

The application of active routing is being considered in our project. However, because of security constraints we restrict the active routing zone only to the access network. The main advantage of our approach is that by using active routing, a mixture of several routing protocols can be obtained. Some routing protocols are more efficient for a certain class of traffic. A mapping of different quality of service levels to different routing protocols is under study.

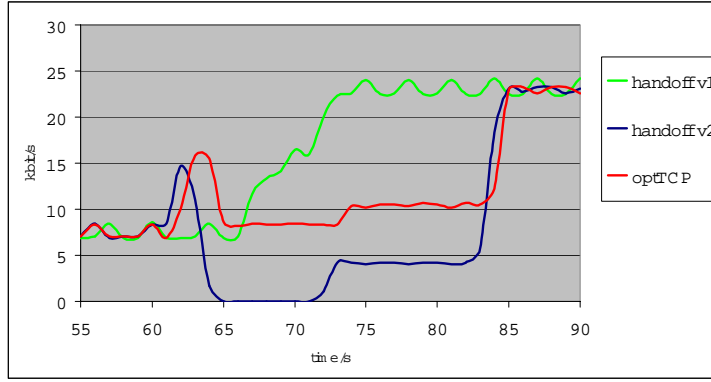


Figure 3: Throughput and handoff delay for different TCP variants.

2.4.2 Ad-hoc and infrastructure-based networks convergence

Wireless data networks are made of nodes, sometimes mobile nodes, and some forwarding mechanism that allows node to node communication even when source and destination nodes are far apart. Cellular networks are one successful example of what we call an infrastructure-based network. A set of wireless and linked stations, called base stations, provides end-node one-hop access to the network. Every transmitting node sends its data packets to the closest base station. The base station forwards the packet through the intra-base station network (either wireless or wire line) to the base station closest to the destination node. This base station then sends the packet to the destination node wirelessly.

In Wireless Local Area Networks (WLANs) the base stations are called Access Points (AP) and their function is quite similar. IEEE 802.11 working group developed a specification that provided 1 and 2 Mbps wireless short-range communication. The new 802.11b revision provides 11 Mbps. These networks provide wireless access to the company or campus wired network. Of course it is possible to have a WLAN not connected to any wire line network, just made of wireless nodes. It is also possible to have no Access Points (Aps) if nodes are set to ad-hoc networking.

But ad-hoc networks are only made of a set of nodes that can communicate wirelessly. No other forwarding elements are present. Therefore, each node must also act as a router. In this way any data packet can be forwarded to its intended destination through a route of several hops. In order to find routes, network nodes need a routing algorithm. Several proposals are being discussed at IETF's MANET group. Routing function will provide end to end connectivity whenever possible. But occasionally a node may go too far away from the rest and become isolated. This can also happen because of a node's motion; the network becomes partitioned so a set of nodes cannot communicate to nodes outside. The focus of this project is to evaluate the impact and feasibility of a hybrid operational mode for WLAN adapters. Complementary routing functions will be evaluated too.

2.5 Content Delivery and Source Specific Multicast Deployment

Within ICSI, KPN Research worked on a project called IP Spray. This project is concerned with content delivery and multicast on broadband ISP access networks. A result of this project is a blueprint for content delivery and multicast deployment, and the KPN ISP environment in particular as a use case.

A topic that has been addressed more specifically is Source Specific Multicast. Source Specific Multicast (SSM) is a special case of IP Multicast: the number of senders to a group is limited number and the receivers are able to inform the network about those senders. One of the most complicated issues with IP Multicast is the mechanism to get the traffic from all senders. This is not necessary with SSM because the receivers explicitly specify the senders. This relieves network operators from difficult configurations. Source Specific Multicast has been tested within this project. Within the SSM architecture, the IGMPv3 protocol is used by the receivers to inform the network about the senders. Within this project an implementation of the IGMPv3 protocol within the FreeBSD operating system has been developed and made available. This implementation has been used by a number of manufacturers such as Cisco, Lucent, Juniper and 3Com to test their products.

2.6 Further activities

2.6.1 Project Proposal: Joint activity of NSA group

The rapid growth and evolution of wireless networks and the uncontested success of the Internet will result in ubiquitous access to the Internet independent of time and location. An important issue in the Internet, and consequently in every network connected to it, is the support for multimedia applications. Wireless/mobile Internet access brings in a number of complications on the network level that are not well considered for the usage of those applications. These complications include:

- the available bandwidth is narrow (especially during movements) in comparison to wireline access
- communication conditions such as error rate change dynamically due to the effect of fading
- on small devices, the capabilities are restricted due to portability

In this context, the support of interactive multimedia applications has additional requirements:

- asynchronous traffic patterns for upstream/downstream
- real-time support for both, short-lived interactions and content delivery
- mostly one-to-many multicast communication pattern

Our project proposal focuses on network topics for mobile access to interactive multimedia applications (e.g. intelligent online classrooms, interactive games, interactive TV and radio, online customer care) with respect to multicast, mobility management, Quality of Service (QoS) support, routing, and their interrelationship on the network level, see the following Figure.

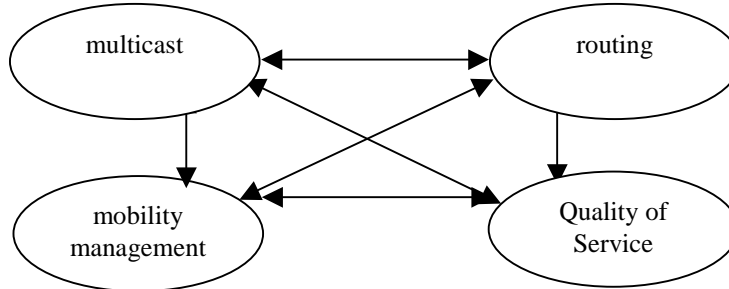


Figure 4: Networking topics and their mutual influence with respect to mobile multimedia applications.

The goals of this project are twofold. First, the networking requirements for seamless mobile access to interactive multimedia applications will be examined with respect to 3rd and 4th (and beyond) generation wireless networks. The applicability of IP centric functionality with respect to multicast, QoS support, mobility management, and routing is investigated to derive solutions, which could also be seen as building blocks for future generations of all IP-based wireless networks. Second, after prototypical implementations, a real scenario based on an intelligent classroom application is set up to verify our achieved results. For details refer to [3].

2.6.2 Multi area IP traffic engineering with MLPS

With the advent of MultiProtocol Label Switching (MPLS), technologies that seemed to be impractical in the Internet become possible. One of these is traffic engineering for IP networks. MPLS gives routing control to the traffic engineer, separating routing and forwarding on the network layer (OSI layer 3). In particular, it allows the use of explicitly routed paths which is fundamental for traffic engineering.

In this ongoing work we combine current intra-domain MPLS traffic engineering methods and intra-domain QoS routing concepts for novel routing algorithms between areas

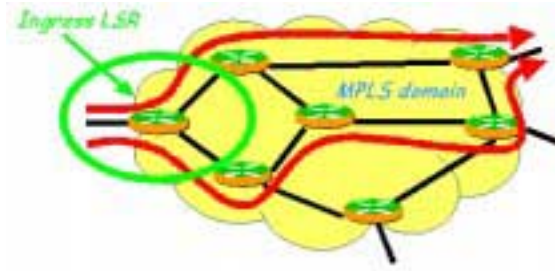


Figure 5: Different forwarding paths in a MPLS capable domain based on ingress LSR decisions.

in a single domain. Our investigations take into consideration the tradeoff between the necessary update intervals of network state information versus signaling overhead, and the degree of the provided QoS information. The main focus is the development of optimized method for QoS routing and Label Switched Path (LSP) setup over several areas within a single domain. We expect that these developed methods can also be used as a basis for inter-domain information exchange.

2.6.3 Modeling of Error Patterns in Transmission Channels

The use of Hidden Markov models for modeling transmission channels and the analysis and evaluation of different error coding techniques is a powerful method to ensure reliable transmission of data. Based on the models for mobile radio channel we extended their applicability to the transmission of video data. In this context, it seemed to be very auspicious to modify the previously developed channel model in a corresponding way; i.e. against the background of video transmissions the unpredictable disturbances or so-called channel noise has been adequately described on the different network layers. The error-correcting mechanisms being used on the different network layers could then be evaluated based on the modified channel model. In this context, it has been differentiated between so-called random and burst error-correcting codes depending on the error-correcting capabilities of the considered code. The residual error rate served as an appropriate criterion for the evaluation of the efficiency or the success of an error-correcting technique. While determining the residual error rate the different algebraic structures of different error-correcting codes have been considered.

In connection with video transmissions, the development of cell loss can be modeled by Markov models or even better by an appropriate autoregressive process. For an evaluation of error-correcting codes the developed channel model has been modified and adapted in a way appropriate to such an autoregressive error process. First steps in this direction have been taken, and an extension of the former underlying Markov model to an autoregressive error process has been further developed. Moreover, new implementations based on this modified approach confirm the results so far and demonstrate that if the correlation be-

tween transmission errors exceeds a certain level, then the use of a burst error-correcting code should be preferred.

For an evaluation of the efficiency of an error-correcting code, the residual error rate as a measure for the success of the considered code need not be exactly calculated because the computational expense of an exact calculation grows rapidly when the code length is increased. For this reason, an approximation formula has been developed. Former approaches for such a formula have been further developed. As a result, a closed expression has been proved in order to calculate the residual error rate of an error-correcting code. In this context, i.e. while developing this approximation formula, the different algebraic structures of error-correcting codes and therefore their different error-correcting capabilities have been taken into account. Under the conditions that

$$q_{GG}, q_{GB}, q_{BG} \quad \text{and} \quad q_{BB}$$

denote the transition probabilities of the underlying Markov model, the approximation formula for a burst error-correcting code results in:

$$\begin{aligned} \mathcal{R}_{BEC}(b) = & (q_{BB})^{n-1} \frac{q_{GB}}{q_{GB} + q_{BG}} \left(1 - \frac{n-b+2}{2^{n-b+1}} \right) \\ & + \frac{(q_{GG})^{n-3}}{4} \frac{q_{GB}q_{BG}}{q_{GB} + q_{BG}} \left(\frac{q_{BB}}{q_{GG}} \right)^b \\ & \left[\sum_{l=0}^{n-b-2} (2q_{GG} + q_{BG}(n-b-2) - q_{BG}l) \left(4 - \frac{3+l}{2^l} \right) \left(\frac{q_{BB}}{q_{GG}} \right)^l \right] \end{aligned}$$

In this context, the parameters n , b and l denote the code length, the length of a burst error which can be corrected by the considered code and the length of the transmission phase. By means of this approximation formula sufficient results can be achieved. The deviation from the exact calculation of the residual error rate is about 1% success of the considered code. Furthermore, the proof of this approximation formula is based on a complex relation between two hyper-geometric series. Here, a hyper-geometric series is defined as follows:

$$F \left(\begin{matrix} a_1, \dots, a_u \\ b_1, \dots, b_v \end{matrix} \middle| x \right) := \sum_{j \geq 0} \frac{a_1^{\bar{j}} \dots a_u^{\bar{j}}}{b_1^{\bar{j}} \dots b_v^{\bar{j}}} \frac{x^j}{j!}$$

The parameters $a_1, \dots, a_u \in \mathbb{R}$ as well as $b_1, \dots, b_v \in \mathbb{R}^{>0}$ and $a_k^{\bar{j}}$ and $b_l^{\bar{j}}$ denote for all $1 \leq k \leq u$ and all $1 \leq l \leq v$ the so-called upper Pochhammer symbols.

After an intensive study of hyper-geometric series it turned out that the above mentioned complex relation between a hyper-geometric series and its derivation was not known before in this form. After further detailed verifications the proof of this relation will be published in an appropriate mathematical journal. This relation will have an impact on further mathematical fields which make use of hyper-geometric series .

2.7 References

1. Sokol, Widmer, USAIA Ubiquitous Service Access Internet Architecture, 2000, subject to approval for CCR ACM Sigcomm
2. Widmer, Mobile Network Architecture for Vehicles 2000, ICSI Technical Report TR-00-009
3. Sokol et al Mobile access to interactive multimedia applications over the Internet Version 0.5, www.icsi.berkeley.edu/NSA

3 Speech Processing: Signals and Systems (the Realization Group)

The Realization Group at ICSI (renamed the Speech Group as of summer 2001) has been expanding its work in topics related to automatic speech recognition. For each area, a strong local program is augmented by strategic collaborations outside of ICSI, including a new DARPA-funded program conducted jointly with the University of Washington. The local efforts have been headed in 2000 by senior researchers Ellis, Greenberg, Hermansky, and Morgan, as well as by SRI visitors Shriberg and Stolcke. As has always been the case, major contributions have been made by students, postdoctoral Fellows, and visitors.

3.1 Speech Signal Modeling

The group has a number of projects focused on processing and modeling speech signals in order to provide fundamental information that can be incorporated in systems for recognition of spoken language and other related goals.

- **Discriminative acoustic processing for ASR based on frequency localized temporal cues.**

Conventional acoustic processing in current ASR is based on maximizing the likelihood of a Gaussian mixture model (GMM) that could have generated a very short (about 10 msec) segment of the speech signal with a given spectral envelope. We believe that this is not the way human listeners decode the linguistic message in the speech signal. While the best machine methods may not ultimately be the same as the approach that humans use, it seems illogical to ignore the characteristics of a system that works so much better than our current machine efforts. Therefore, we have pursued an alternative paradigm, where the goal is to independently optimize discriminability of rather long (about 1 sec) temporal patterns of frequency-localized spectral energies at different parts of the available speech spectrum.

Towards this goal we investigated mutual information between phoneme labels in hand-labeled speech and different regions in the time-frequency plane as a means for discovering the relative saliency of different regions of the time-spectral pattern of speech [25], [26], [27], linear discriminant analysis of short-term spectral vectors

[16], [19], of temporal patterns of critical-band spectral energies [23], [24], [13], and of two-dimensional time-frequency patterns [18]. The results of this research support the optimality of the speech code with respect to properties of human hearing [14], as well as practical RASTA-like FIR filters derived by the LDA analysis of temporal energy patterns [23], [20].

- **Multi-stream features**

In previous reports we have discussed our work in the design of speech recognition systems that incorporate multiple probability streams derived from different signal representations. In work concluded in 2000, we computed optimal signal representations for differing acoustic environments. In thesis work by Mike Shire, it was shown that combinations of front-end acoustic modeling components that have been trained in heterogeneous acoustic conditions can improve ASR robustness to acoustic conditions that were not seen during training. Discriminatively trained temporal filters, in addition to the discriminatively trained MLP probability estimator, were used to improve each stream's performance in a sample acoustic condition. We found that this method consistently improved phone classification at the frame level. For the task of word recognition, on tests with reverberation unseen during training, the multi-stream system having components trained in both clean and heavy reverberation produced results significantly superior to the system trained solely in one condition. Combination results appeared to be best when each of the streams was based upon different preprocessing strategies.

- **Models for Speech Confidence Estimation**

In Masters work by Berkeley student Andrew Hatch [12], discriminant approaches were applied to the estimation of confidence scores for recognition. These efforts were applied first to the ICSI hybrid HMM/ANN system, and then to a combined estimator incorporating both the ICSI system and the SRI Gaussian Mixture Model system. The work concluded that sizeable improvements in confidence estimation for the hybrid system alone could be achieved by discriminant methods, and in particular that posterior probability targets derived from forward-backward reestimation yielded improved scores over confidences derived from networks that were trained on phonetic labels corresponding to Viterbi alignments. However, when these estimators were combined with those from the SRI system, the extension from locally to globally derived ANN targets did not reflect the improvement seen in the single system. Nonetheless, on a task of significant scale (the Switchboard conversational speech transcription task), confidence scores from the hybrid system did yield improved confidence estimates in combination with the Gaussian-mixture-based system. Finally, we learned in the course of this work that the estimation of class priors (which is critical to the standard use of hybrid HMM/ANN systems) is best done using adaptive estimation from the test data itself, rather than using static estimates from the training data. Even when this estimation is done in a completely unsupervised fashion, both confidence estimation and speech recognition performance is substantially improved over the static case. Additionally, this adaptive estimation paradigm

appears to work particularly well on noisy tasks such as the Aurora cellular standards task, where the adaptation of priors not only appears to assist with particular speakers, but also with particular acoustic environments.

Warner Warren completed his Ph.D. thesis [28] during this last year, also focused on the task of confidence estimation. In this case, the focus was on the use of global posterior estimates, as defined by an earlier thesis by Konig; there the approach was designated Recursive Estimation of Maximum A posteriori Probabilities (REMAP). In the new thesis, REMAP's deviation from optimality was explored; in particular the issue of the global posterior estimates not summing to one over all the classes. Warren studies the deviation of the model over sequences from the reality, and also showed both practical and theoretical approaches to partially compensating for the errors. One simple approach that was important was the per-utterance normalization of the global posterior estimates. While this would be unnecessary for true posteriors, in the practical case it is in fact necessary due to the deviation between the model and the sequence itself, a deviation that increases with the length of the sequence. This approach led to much better performance, both for reducing false acceptances and rejections. However, Warren also designed a more elegant approach to optimize the model subject to a set of statistical constraints (such as the priors, as noted in Hatch's work), and directly reduce the modeling error. In preliminary experiments, this approach was shown to significantly reduce the word error rate on the Phonebook corpus. Overall, this thesis has explored in detail many of the underlying questions we have had about estimating class posterior distributions. It still raised more questions than it provides answers, but we now have a much better framework for future explorations.

- **Acoustic change detection**

We have been working on the problem of breaking speech into segments with relatively consistent acoustic characteristics, which is sometimes referred to as the acoustic change detection (ACD) problem. We have been investigating this in the context of broadcast news recognition. In particular, we have examined the improvements possible by clustering separated segments together and using cluster membership information iteratively to improve segmentation [8]. The interaction between the acoustic change detection and the clustering approach gave us substantial improvement over previously reported results on the 1997 Hub-4 Broadcast News test set; in particular, feedback of clustering information improved the Equal Error Rate (EER) of our ACD system from 26.5% to 18%.

- **Prosodic Stress**

A continuation of the automatic prosodic stress transcription project looked at various possibilities for pitch (f_0) variation being a more important factor in the perception of acoustic stress than we had previously concluded. Pitch variation within a syllable was specifically looked at in more detail, as well as using various scalings of pitch, all in an attempt to provide this parameter with every benefit of the doubt. The conclusions were the same as before - namely that pitch, while contributing to

the sensation of stress, is far less important than duration and amplitude. The only significant transform of the f_0 data that appeared to be highly correlated with stress was one in which segmental duration was conflated. When duration was neutralized (by taking the rate of f_0 change, rather than the magnitude of change) the effect all but disappeared, raising the possibility that pitch variation is subservient to duration with respect to stress marking.

Two papers were written describing the results [21][22], and a presentation was made to the NIST Speech Transcription Workshop.

- **Automatic phonetic transcription**

Shawn Chang, Lokendra Shastri and Steve Greenberg developed a system for automatically labeling and segmenting the OGI Numbers corpus at the phonetic-segment level. They first classified the signal in terms of articulatory-acoustic features (such as voicing, manner and place of articulation, using temporal flow model networks) and then mapped such feature clusters to phone labels (using MLPs). The results of these procedures were on the same level as a human transcriber (as assessed by interlabeler agreement). The ALPS (automatic labeling of phonetic segments) system achieves ca. 80% concordance with the human transcribers [3].

- **Articulatory-acoustic patterns in the Switchboard corpus**

The pronunciation patterns of the phonetically labeled component of this corpus were analyzed in terms of articulatory-acoustic features as a function of syllable position. The metric used was deviation from canonical pronunciation. It was found that there is a systematic relationship between syllable position and the articulatory-acoustic feature deviation patterns observed. At onset it is the manner of articulation feature that deviates from canonical more than others (by far). When other feature dimensions deviate it is always in tandem with manner deviations. A similar pattern obtains in coda position, but with the magnitude of deviations about twice as great as for onsets. Also, there are a significant number of segmental deletions in coda position (and relatively few at onset). The pronunciation patterns for vocalic nuclei is entirely different. Here, the manner dimension is highly stable, and the place of articulation feature most variable. In fact, there is a huge number of deviations from canonical place in the nucleus. From the pattern of deviations observed it was concluded that manner and place of articulation are primary feature dimensions (place being most stable in onset and coda, and manner in nucleus position). Voicing and rounding appear to be auxiliary dimensions in that they rarely if ever deviate from canonical states except when manner or place also vary (but not vice versa). This finding suggests that articulatory-acoustic features are not independent of each other (as has been proposed by others) but operate in concert (particularly voicing with manner, and rounding with place).

This material was presented at the Workshop on Patterns of Speech Sounds in Unscripted Communication Production, Perception, Phonology, Akademie Sankelmark, October 11, 2000 under the title: “Understanding Spoken Language Using Computational and Statistical Methods”, but has not yet been written up for publication.

3.2 Speech Recognition

The group has a number of projects focused on the recognition of speech per se. In this section we will focus on the more basic techniques as opposed to the efforts that are most closely tied to particular applications (see later section on Applications).

- **Tandem acoustic modeling**

Following on from our winning entry in the 1999 'Aurora' evaluation conducted by the European Telecommunications Standards Institute (ETSI) [20], we adapted our so-called 'Tandem' approach (in which the outputs of a neural network are used as input features for a conventional Gaussian mixture model (GMM)-based recognizer [15]) for the 2000 version of task. Following some criticism that last year's version of the 'noisy digits' task used the same 'noise' in both training and test, the 2000 version included three test conditions with increasing levels of mismatch from the training data. We found that the Tandem approach continued to perform very well, and was less adversely affected by the condition mismatch than the GMM baseline system using the HTK toolkit: for the lowest mismatch condition (similar to the 1999 test), a Tandem system based on two networks fed by PLP and MSG features made, on average, just 49% of the word errors observed for the HTK baseline; on the highest-mismatch condition, although the absolute errors were a few percent higher, the Tandem system made just 41.2% as many errors at the HTK system.

Since channel variation was one of the sources of mismatch, some kind of feature normalization was called for. To support real-time processing and recognition, it is not possible to rely on statistics of 'entire' utterances since recognition must commence before the utterance has completed. For this reason, we implemented simple online normalization of mean and variance within our standard speech recognition tools, and used these to repeat our Aurora tests. Tests to search for the optimal time constant to use with the running estimates of per-feature mean and variance revealed surprisingly little sensitivity, although the online normalization was clearly beneficial and confers a broad robustness to gain and channel variation on the system.

In addition to the work described above on noisy American English digit strings, we have also been investigating the applicability of the tandem approach to larger tasks. In collaboration with researchers at the Oregon Graduate Institute (OGI) and Carnegie Mellon University (CMU), we developed a Tandem system for the Speech In Noisy Environments (SPINE) task developed by the Naval Research Labs. This corpus involves a vocabulary of several thousand words, pronounced spontaneously by military personnel engaged in battlefield simulations, including significant and diverse background noise. The tandem system was a component of the best-performing recognizer in the evaluation, which combined the results of several different approaches at the hypothesis level.

A subsequent close comparison between the Tandem system and the conventional variants [4] showed that Tandem features reduced word error rates on this task by around 30% for the baseline, context-independent GMM-HMM system. The incorporation of context-dependent phone models and unsupervised adaptation (via

MLLR) into the GMM-HMM system improves the Tandem performance, but it has a greater benefit on the non-Tandem comparison, with the net result that for these best-performing systems the advantage of using Tandem features is reduced to only a few percent. However, these enhancements were optimized for the baseline (non-Tandem) system, and similar optimization was not performed for the Tandem system, but could improve the margin of advantage. Also, the SPINE task had a fairly limited training set size, and the network in the Tandem system would improve if more data were available. Overall, we were very encouraged that the Tandem approach continued to offer advantages on this very different and far more challenging domain. Work in this direction is continuing. Figure 6 shows the overall structure of our Tandem system for the SPINE task.

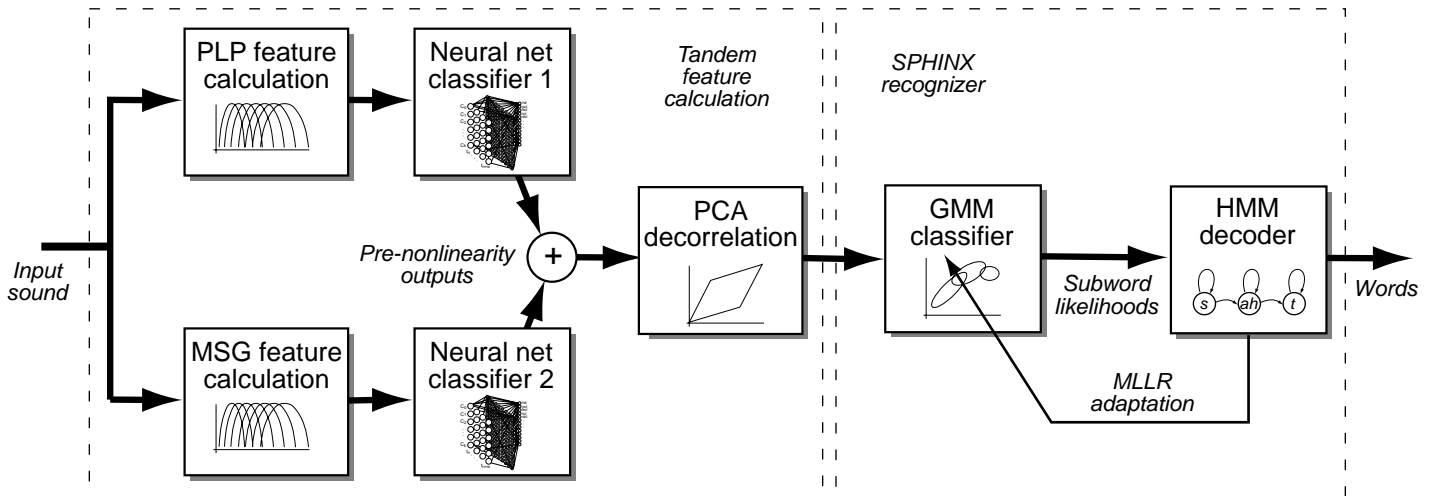


Figure 6: Block diagram of the tandem recognizer used for the SPINE task.

- **Combining information streams**

A major theme in our work to improve speech recognition accuracy has been to exploit combinations of diverse systems in ways to effect mutual corrections [6]; for instance, by combining two different feature extraction algorithms - Perceptual Linear Prediction (PLP) and Modulation-filtered Spectrogram (MSG), we routinely see a 10-15% relative reduction in word error rate. This concept is incorporated in the Tandem work described above, and we have made a detailed study of the differences in performance when the same basic signal features are combined in different ways, for instance by using a single statistical model, or dividing into several streams, training models on each, then combining the model results [7].

The complex pattern of performance exhibited by different combination strategies makes it difficult to find the best-performing recognizer configuration. In an effort to reduce this uncertainty, we investigated various indicators for their ability to guide combination system design. We reasoned that stream combination works best when the streams consist of complementary information, and thus we looked at estimates of the mutual information (MI) between the streams as a basis for deciding when

and how to combine them. We looked at the MI between feature streams before and after the classifier network. Our hypotheses were that (a) raw feature streams with high MI are showing correlated structure that may be removed within separate classifiers, so they should probably be modeled within a single large classifier, and (b) classifier outputs with relatively low MI are providing different information about the underlying speech stream, and are thus good candidates for combination at that level (assuming of course they both perform reasonably well). By correlating the MI measures with actual system performance, we found that (b) was well supported, although the evidence for (a) was more equivocal [5].

- **Multi-band recognition**

Frequency localized independent classifiers were introduced in our multi-band ASR technique. More recently, we started to study one extreme of the multi-band technique, classifiers of TempoRAI PatternS (TRAPS) [17]. Here we attempt to base phoneme classification on the temporal evolution of spectral energies in the vicinity of the target phoneme, rather than on the spectral shape of a short segment of the speech signal within the phoneme. The success of this paradigm has important consequences in our view of the role of spectral sensitivity of human hearing in decoding the linguistic message in speech. The technique was found useful in ASR of noisy speech in cellular telephony [20].

- **Handling multiple acoustic sources**

The standard approach to recognizing speech in the presence of background noise is to design features and classifiers whose output is least changed by the presence of the noise. Usually this involves assuming that the background noise is more or less stationary, which is often a poor assumption. An alternative perspective is to recognize the presence of many different energy concentrations in an acoustic signal, then to search among those pieces for a subset that correspond to a single target voice. This approach falls into the general category of ‘computational auditory scene analysis’ (CASA), the effort to build models of human sound organization.

We have been developing a speech recognizer based on this approach and also using the techniques of missing-data (MD) recognition. In MD recognition, the acoustic model is presented with only a subset of the acoustic features; the rest assumed missing due to corruption or masking e.g. by noise interference. The likelihood of a model fit can still be calculated based on the available data, for instance by integrating over all possible values of the missing dimensions.

The overall system, which we call the ‘multi-source decoder’ [1] to reflect its treatment of the acoustic signal as consisting of multiple sources, operates as follows: First, the signal is broken into several distinct energy concentrations in time-frequency. Then an incremental hypothesis search performs MD matches to the data represented by each possible subset of pieces. Because it is incremental in time, the search complexity is bounded by the number of simultaneously-present elements, rather than the total number of elements (which would quickly become intractable). Our current results show that this approach achieves an improvement of more than 25% relative in word

error rate compared to our best conventional MD approach. This is despite the fact that the current division of the acoustic energy into separate elements is extremely crude; we expect improvements in both performance and efficiency when a true CASA model is used to guide the segmentation.

- **Pronunciation modeling**

In work that was a continuation of Eric Fosler-Lussier's earlier efforts in our group, Mirjam Wester (a Dutch visitor) and Eric Fosler studied the impact on ASR of knowledge-based and data-derived modeling of pronunciation variations. This work was done on the Dutch database VIOS, which consisted of recordings of human-machine interactions in the domain of train timetable information. Using phonological rules led to a small improvement in word error rate (WER), while using a data-derived approach led to a significant improvement over the baseline [29].

Steve Greenberg and Eric Fosler-Lussier also wrote a paper making explicit the link between pronunciation variation and linguistic information [11].

- **Linguistic dissection of automatic speech recognition performance** Shawn Chang, Joy Hollenback and Steve Greenberg performed a diagnostic evaluation of eight separate speech recognition systems (from AT&T, BBN, Cambridge, Dragon, Johns Hopkins, Mississippi State, SRI, University of Washington) for a portion of the Switchboard corpus that had previously been phonetically transcribed. As part of the project the 54-minute portion of the corpus was transcribed as well in terms of prosodic stress. Based on the phonetic and stress labels (and phonetic segmentation performed using a quasi-automatic method) the speech material was analyzed on ca. 50 distinct acoustic and linguistic parameters (e.g., speaking rate, energy level, segmental duration, etc.). This analysis enabled us to correlate many properties of the speech signal with the word recognition results from the eight separate sites. In general it was found that the single best predictor of word recognition performance was the accuracy of phonetic classification (using the ICSI transcription material as the benchmark). Prosodic stress was highly correlated with word-deletion errors (most of these errors occurred in unstressed words). Syllable structure was also highly correlated with word recognition performance (except for AT&T), as was fast and slow speaking rates (as assessed using a linguistic measure - the acoustic measure of speaking rate did not correlate well with either recognition performance or with the linguistic metric of speaking rate). Other parameters of importance were the relation of the duration of the hypothesized word relative to the reference (ICSI) word duration and the magnitude of the articulatory-acoustic feature deviation from between the hypothesized and reference words. "Error" patterns were also analyzed with respect to correctly recognized words. There is far greater tolerance of "error" in the vocalic nuclei than for the onset and coda consonant components of a word. About 30% of the vocalic segments were "misclassified" in correctly recognized words, but only about 10% of the onset consonants were.

Two papers describing the results are available (www.icsi.berkeley.edu/~steveng) [9][10]. Four presentations were given on this material during the 2000 year - one

each at the workshops listed above, the third at AT&T in October and fourth at IBM in May. This project is being continued this current (2001) year, using the same material as in the competitive evaluation (as opposed to last year when we used a completely different set of materials that had been previously phonetically transcribed).

3.3 Speech Perception

Some of the work in this area continued in 2000.

- **The relation of intelligibility to the phase and magnitude components of the modulation spectrum**

Some of this work began in 1999, but the real analysis of the acoustic component of the sentence material used did not occur until January of 2000, so is included in this report. Takayuki Arai and Steve Greenberg locally time reversed segments (from the TIMIT corpus) of variable length (ranging between 20 and 180 ms) and assessed the intelligibility of the material. There is a sharp drop in intelligibility for signals whose reversed segments are longer than 40 ms. They developed a method of analyzing the modulation characteristics, called the complex modulation spectrum, in which they fold the phase and magnitude components into a single representation distributed across tonotopic frequency. They demonstrated that the intelligibility of speech locally time reversed can not be accounted for in terms of just the amplitude component of the modulation spectrum, but rather requires the phase component as well. The complex modulation spectrum (which essentially measures the phase dispersion of the modulation patterns across tonotopic frequency) is highly correlated with intelligibility.

The material was presented at the Midwinter Meeting of the Association for Research in Otolaryngology in February, 2000 (“What are the Essential Acoustic Cues for Understanding Spoken Language”) as well as at the Conference on Neural and Cognitive Systems in May (at Boston University).

- **Directionality of Consonant Confusions**

Both historical sound change and perceptual experiments show shifts in stop place (labial, alveolar, velar) that vary according to vowel context. Some of the stop shifts, such as [k] \leftrightarrow [p] in the context of [u], are symmetric. Some, such as [k] \rightarrow [t] in the context of [i], are unidirectional. In work by Madelaine Plauche, she attempted to account for the directionality of stop shifts based on their acoustic structure as well as listeners’ sensitivities to different parts of the speech signal. These methods include automatic feature extraction, machine learning algorithms, acoustic signal manipulation (filtering, splicing), and classic perceptual experimentation.

3.4 Applications

- **Meeting Recorder**

We have recently commenced a wide-ranging initiative to develop speech technology for use in conventional meeting situations. More obvious applications include information retrieval and summarization based on speech transcripts derived from acoustic recordings of meetings, and we are also interested in the deeper analyses of meeting structure that such recordings might support, as well as possible real-time interactive applications that would be possible if a computer were ‘listening in’ as part of the meeting.

The first stage in this project is the collection of a corpus of real meeting recordings so that we can work on the kinds of signal processing and recognition tasks that will arise in this scenario. We have ‘wired’ one of the meeting rooms at ICSI to be able to capture meetings, including very high quality sample-synchronous 16-channel hard disk recordings. We have already recorded tens of hours of meetings, in each case capturing on separate channels the individual speakers (via head-mounted microphones) as well as ambient signals collected from microphones on the conference table; ultimately, we are aiming at algorithms that require only these ambient signals to minimize the inconvenience to meeting participants.

Custom hardware and software was developed to make the data collection as convenient and robust as possible. In order to leverage our data collection activities, we are supporting the duplication of our setups at collaborating institutions the University of Washington (Seattle) and Columbia University (New York).

Our initial analyses of these data are taking multiple directions. Firstly, we are making hand transcriptions to provide training data for automatic classifiers. We have begun word-level transcriptions that will be used to train speech recognizers on the ambient acoustics, and we have also made some closer transcriptions of the many unusual and difficult acoustic entities that we are suddenly forced to address in this new domain. Speaker overlaps (where several people speak at the same time) are the most prominent of these, but other events such as extraneous noises, nonspeech (coughs, laughter etc.) and mic noise from physical contact are other categories that cannot be ignored.

Even before we have successful speech recognition for this material, we would like to be able to segment it into separate speakers; in fact, recognition is likely to be improved by knowledge of where speakers’ phrases begin and end. On the face of it, speaker turns should be relatively simple to extract, given that we have individual close-mic signals for each speaker, but in practice cross-talk, external noises and mic noise complicate the problem. We are looking at several ways to achieve automatic ground-truth labelling of speaker turns, including LMS coupling-function estimation and blind source separation.

Ultimately, we would like to be able to find speaker turns based only on signals recorded by a small tabletop device. If we have only a single channel available, this reduces to the conventional acoustic change detection (ACD) problem. We have been investigating this in the context of broadcast news recognition. In particular, we have examined the improvements possible by clustering separated segments together and using cluster membership information iteratively to improve segmentation [8].

If we can put several sensors on to the tabletop device, then we can also use inter-microphone timing to get direction-of-arrival information, and use this to help detect speaker changes. For closely-spaced microphones this timing difference can be rather small, but our initial investigations indicate that it can still be a useful correlate of the active speaker. We plan to extend this work by investigating more sophisticated direction-of-arrival estimation and tracking algorithms such as the ones that have been developed in underwater acoustics applications.

In 2001, we are beginning to get recognition and adaptation results with these data, but our 2000 efforts were focused on the infrastructural and exploratory aspects described above.

- **Speech Enabled Information Systems**

Our primary effort, in collaboration with the AI group, has been towards the development of an English language system for users to retrieve tourist information for the Public scenario of the German SmartKom project. We completed several milestones toward this goal during 2000:

1. We adapted our hybrid HMM/ANN recognizer to the Verbmobil software architecture. This integration was targeted as a stepping stone toward the forthcoming SmartKom testbed, which was expected to have similar interface properties. The functioning of ICSI' recognizer in the Verbmobil testbed was demonstrated successfully at the second Gesamt-Workshop on April 10-11, 2000, in Kaiserslautern. Several changes were made to the original ICSI recognizer to support the integration, e.g., the lattices were both changed in format and in pruning style. The software also was optimized for very fast startup and restarts (less than 1 second on a 300 MHz Sun Ultra-10 workstation). This capability is important to enable frequent restarts for multiple language models, etc. .
2. We further adapted the recognizer to the SmartKom architecture. The latter was determined by our German colleagues. Like the Verbmobil architecture, it made use of a Pool Communication Architecture (PCA) for communication between different subsystems. We developed an interface component that handles all PCA messaging and controls the startup, termination, and input and output from and to the recognition pipeline. The bulk of the development effort was in handling of the M3L/XML data formatting used in SmartKom.
3. In 2000 we also worked on porting the English recognition module to the tourist information domain. The fundamental challenge was that no in-domain speech data or transcripts were available. The acoustic model was based on the American English Broadcast News corpus, with further training from the Verbmobil domain. The vocabulary and language model of the recognizer were based on a combination of sources that are related to the tourist domain, including transcripts of dialogs collected at the European Media Lab (EML) in Heidelberg, totaling 5700 words. Other sources included phrase lists from guide books, Web pages with tourist information, hand-crafted word lists for numerical amounts etc., translations of the German SmartKom vocabulary, and a background of

frequent English words extracted from the Switchboard corpus. A class-based trigram model was developed using these sources and a manually determined set of classes. Missing pronunciations were generated by a phonetician, particularly for German place and streetnames. We are working this year on improving performance, though we are currently waiting for test data to provide objective evaluations.

3.5 CUSP

The efforts described in the section on Speech Recognition have led to an expanded effort in small-medium vocabulary recognition for portable devices, particularly for cell phones. Towards this end, Nelson Morgan and Hynek Hermansky have begun an effort they refer to as the Center for Ubiquitous Speech Processing, or CUSP. The goal is to grow a new center of excellence in robust speech analysis for such an application area, funded by a combination of industrial and government sponsors. In May 2001 Qualcomm signed on as a sponsor, and discussions with other sponsors are currently under way. The year 2000 work described in this report is already being extended to multilingual recognition with a range of acoustic conditions. The Center also extends our previous collaborative arrangement with the Oregon Graduate Institute.

References

- [1] Barker, J., Cooke, M., and Ellis, D., "Decoding speech in the presence of other sound sources," Proc. ICSLP-2000, Beijing, October 2000.
- [2] Bourlard, H., and Morgan, N., *Connectionist speech recognition*, Kluwer Academic Publishers 1994.
- [3] Chang, S., Shastri, L. and Greenberg, S. "Automatic phonetic transcription of spontaneous speech (American English)." *Proc. Int. Conf. Spoken Lang. Proc.*, 2000, Beijing.
- [4] Ellis, D., Singh, R., and Sivadas, S., "Tandem acoustic modeling in large-vocabulary recognition," ICASSP-2001, Salt Lake City, May 2001.
- [5] Ellis, D., and Bilmes, J., "Using mutual information to design feature combinations," Proc. ICSLP-2000, Beijing, October 2000.
- [6] Ellis, D., "Improved recognition by combining different features and different systems," Proc. AVIOS-2000, San Jose, May 2000.
- [7] Ellis, D., "Stream combination before and/or after the acoustic model," ICSI Technical Report 00-007, April 2000.
- [8] Ferreiros, J., and Ellis, D., "Using acoustic condition clustering to improve acoustic change detection on Broadcast News," Proc. ICSLP-2000, Beijing, October 2000.

- [9] Greenberg, S., Chang, S. and Hollenback, J. "An introduction to the diagnostic evaluation of the Switchboard-corpus automatic speech recognition systems." Proceedings of the NIST Speech Transcription Workshop, 2000, College Park, MD.
- [10] Greenberg, S. and Chang, S. "Linguistic dissection of switchboard-corpus automatic recognition systems." Proceedings of the ISCA Workshop on Automatic Speech Recognition (2000): Challenges for the New Millennium, Paris.
- [11] Greenberg, S. and Fosler-Lussier, E. "The uninvited guest: Information's role in guiding the production of spontaneous speech," in the Proceedings of the Crest Workshop on Models of Speech Production: Motor Planning and Articulatory Modelling, 2000, Kloster Seeon, Germany.
- [12] Hatch, A., *Word-level Confidence Estimation for Automatic Speech Recognition*, Master's Thesis in Press, UC Berkeley, 2000.
- [13] Hermansky, H., "Should Recognizers Have Ears?," in Proceedings ESCA Tutorial and Research Workshop on Robust Speech Recognition for Unknown Communication Channels, pp. 1-10, France 1997.
- [14] Hermansky, H., Invited talk - conference on Neural Information Processing Systems, December 2000, Denver, Colorado.
- [15] Hermansky, H., Ellis, D., and Sharma, S., "Tandem connectionist feature stream extraction for conventional HMM systems," Proc. ICASSP-2000, Istanbul, III-1635-1638.
- [16] Hermansky, H., and Malayath, N., "Spectral Basis Functions from Discriminant Analysis," in Proceedings ICSLP'98, Sydney, Australia, November 1998.
- [17] Hermansky, H., Sharma, S., "Temporal Patterns (TRAPS) in ASR of Noisy Speech," in ICASSP'99, Phoenix, Arizona, USA, March 1999.
- [18] Kajarekar, S., Yegnanarayana, B., and Hermansky, H., "A study of two dimensional linear discriminants for ASR" ICASSP 2001, Salt Lake City, Utah.
- [19] Malayath, N., "Data-Driven Methods for Extracting Features from Speech" Ph.D. thesis, Oregon Graduate Institute, OR, USA, 2000.
- [20] Sharma, S., Ellis, D., Kajarekar, S., Jain, P., and Hermansky, H., "Feature extraction using non-linear transformation for robust speech recognition on the Aurora database," Proc. ICASSP-2000, Istanbul, II-1117-1120.
- [21] Silipo, R. and Greenberg, S. "Prosodic stress revisited: Reassessing the role of fundamental frequency." Proceedings of the NIST Speech Transcription Workshop, College Park, MD.

- [22] Silipo, R. and Greenberg, S., “Automatic Detection of Prosodic Stress in American English Discourse,” Technical Report TR- 00-001 (29 pages), International Computer Science Institute, Berkeley, 2000 (available from <http://www.icsi.berkeley.edu/techreports/2000>)
- [23] van Vuuren, S., and Hermansky, H., “Data-Driven Design of RASTA-Like Filters,” in Proceedings EUROSPEECH’97, Rhodes, Greece 1997.
- [24] van Vuuren, S., “Speaker Verification in a Feature-Time Space”, OGI Ph.D. Thesis, March 1999, Portland, USA.
- [25] Yang, H.H., van Vuuren, S., and Hermansky, H., “Relevancy of TimeFrequency Features for Phonetic Classification Measured by Mutual Information,” in ICASSP’99, Phoenix, Arizona, USA, March 1999.
- [26] Yang, H.H., and Hermansky, H., “Search for Information Bearing Components in Speech”. To appear in Advances in Neural Information Processing Systems, Vol.12, S.A. Solla, T.K. Leen and K.-R. Muller (eds.), MIT Press, 2000.
- [27] Yang, H.H., Sharma, S., van Vuuren, S., and Hermansky, H., “Relevance of TimeFrequency Features for Phonetic and SpeakerChannel Classification” Speech Communication, August 2000.
- [28] Warren, W., *Global Posterior Probability Estimates as Decision Confidence Measures in an Automatic Speech Recognition System*, Ph.D. Dissertation, UC Berkeley, December 2000.
- [29] Wester, M., and Fosler-Lussier, E., “A comparison of data-derived and knowledge-based modeling of pronunciation variation,” ICSLP-2000, Beijing, October 2000

4 Artificial Intelligence and its Applications

The Artificial Intelligence group continues its long term study of language, learning, and connectionist neural modeling. The scientific goal of this effort is to understand how people learn and use language. The applied goal is to develop systems that support human centered computing through natural language and other intelligent systems. Several shorter term goals and accomplishments are described in this report. There is continuing close cooperation with other groups at ICSI, at UC Berkeley, and with external sponsors and other partners. There are three articulating subgroups and this report summarizes their work.

4.1 Language Learning and Use, Jerome Feldman, PI

It has long been known that people would prefer to talk with computer systems in natural language if they could. The problem of communicating with machines is becoming increasingly important to society because computers will soon be embedded in nearly every

artifact in our environment. But how easy will it be for people of all ages and abilities to use them? In ten years or less, virtually every device in our environment will have a computer in it. This raises the specter of an embedded computing malaise—every device will have its own interface that the user has to learn. Most people today cannot program their VCRs or use more than half the functionality of their answering machines. In the world of embedded computing, there could be thousands of idiosyncratic interfaces to learn. Many people will not be in control of the devices in their own environments. In addition, ever widening aspects of society, from education to employment, depend upon everyone interacting with computational systems.

Natural interaction with computerized devices and systems requires a conceptual framework that can communicate about requests specified in ordinary language. Systems may well need to tell their human users what is going on, ask for their advice about what to do, suggest possible courses of action, and so on. The machines, and more especially the interactions among the machines, are getting to be so complicated and autonomous, and yet also so intimately involved in the lives of the human users, that they (the machines) have to be able to take part in a kind of social life. The central goal of this project is to provide a conceptual basis and a linguistic framework that is rich enough to support a natural mode of communication for this evolving human/machine society.

While the usefulness of natural language understanding (NLU) systems has never been questioned, there have been mixed opinions about their feasibility. Most current research is focused on goals that are valuable, but fall far short of what is needed for the natural interactions outlined above. We believe that recent advances in several areas of linguistics and computational theory and practice now allow for the construction of programs that will allow robust and flexible integrated language interaction (ILI) within restricted domains.

For many years, Jerome Feldman has studied various connectionist computational models of conceptual memory and of language learning and use. George Lakoff and Eve Sweetser have worked on the relation between linguistic form, conceptual meaning, and embodied experience. Over the past dozen years, the group has explored biologically plausible models of early language learning (Bailey et al. 97) and of embodied metaphorical reasoning (Narayan 97). About two years ago, we extended our efforts on modeling child language acquisition from individual words and phrases (Regier 96, Bailey 97) to complete utterances. This required us to develop a formal notion of what it means to learn the relationship between form and meaning for complete sentences. Many groups, including ours (Feldman 98) have worked on algorithms for learning abstract syntax, but we decided that it was time to look directly at learning form-meaning pairs, generally known as constructions. After an intensive effort by the whole research group, we now have an adequate formalization of constructions and are moving ahead with the project of modeling how children learn grammar from experience. This will form the core of a dissertation by Nancy Chang, a UCB doctoral student, currently on an IBM fellowship. But we also realized that our formalized notion of linguistic constructions that systematically links form to conceptual meaning is potentially a breakthrough in achieving robust and flexible NLU systems.

The most novel computational feature of the NTL effort is the representation of actions: executing schemas (x-schemas), so named to remind us that they are intended to execute when invoked. We represent x-schemas using an extension of a computational formalism known as Petri nets (Murata, 1989). As discussed below, x-schemas cleanly capture se-

quentiality, concurrency and event-based asynchronous control and with our extensions they also model the hierarchy and parameterization needed for action semantics.

Our goal is to demonstrate that unifying two powerful linguistic theories, embodied semantics and construction grammar, together with powerful computational techniques, can provide a qualitative improvement in HCI based on NLU. Over the last year we explored extending our existing pilot system to moderate sized applications in real HCI settings and developing the methodology needed for large scale realization of NLU interaction. This involves formalization and additional research in cognitive linguistics, development of probabilistic best fit algorithms and significant system integration. Much of the groups effort over the past year has gone into developing these formalisms and to producing a pilot version of the integrated language understanding system (Bergen, Chang, Paskin 2000).

For concreteness, we have chosen a specific task domain for the proof-of-concept demonstration of our research. We are constructing a system for understanding and responding to dialog with tourists, initially focused on Heidelberg, Germany. This applied project is being carried out in cooperation with a partner group at EML in Heidelberg, which has built an extensive data base describing their city (www.villa-bosch.de/english/research) and will implement the detailed actions for using it based on our natural language analysis. This cooperation will bring several benefits to the project and provides clear milestones for evaluating our effort. This project (called EDU for Even Deeper Understanding) has been in operation since July 2000, with multi-year funding from the Klaus Tschira Foundation. Robert Porzel, of EML, has joined our group for the calendar year 2001. This effort is also closely linked to the SmartKom project, which is discussed in the Speech section of this annual report. Another cooperation between the Speech and Language groups is the human interface section of the CITRIS proposal to the state of California, currently under legislative review.

The core computational question is finding best match of constructions to an utterance in linguistic and conceptual context. One of the attractions of traditional phrase structure grammars is the fact that the time to analyze (parse) a sentence is cubic in the size of the input. If one looks at the comparable problem for our more general construction grammars, context-free parsing becomes NP complete (exponential) in the size of the input sentence and thus impractical. But people do use larger constructions to analyze language and we believe that we have two insights that seem to render the problem of construction analysis tractable. The general computational point is that our task of finding a best-fit analysis and approximate answers that are not always correct presents a more tractable domain than exact symbolic matching. More importantly, our integrated constructions are decidedly not context-free or purely syntactic. We believe that constraints from both semantics and context will be sufficiently constraining that it will be possible in practice to build best-fit construction matchers of the required scale.

This sequence of operations: surface analysis, construction parse, SimSpec, simulation and inference is repeated for every clause. The current pilot system does not make use of extensive context or world knowledge, but these are central to our new design. There is currently a great deal of renewed effort to develop ontologies of words and concepts for a wide range of semantic domains (Fikes 1994). After analyzing these efforts, we have decided against committing to any one of the competing formulations and have instead defined an Application Programming Interface (API) that our system can use to access

information from any source. A preliminary version of this is used in the pilot system and we will evolve the API as experience requires. The current API has the usual commands for adding information and some special ones for retrieving ordered lists of concepts most likely to fulfill a request. This also facilitates our interaction with the EML project (EDU) and the German SmartKom effort.

Our HCI system also requires both situational and discourse context as well general knowledge. Despite a large literature on context (Hobbs 79, Kehler 93) there is currently no integrated theory or system that meets our needs. Again, we have built an API (which returns ordered lists of potential role fillers) for the pilot system and this remains an active research area Chang and Hunter 2000. We will also employ one of modern large scale parsers (Charniak 98, Collins 98). Although much needs to be done, we believe that we have identified all the main components of an ILI system and have at least preliminary versions of each subsystem operational and a sound overall design plan.

There is also a significant effort on related problems that elucidate or exploit our main results. Ben Bergen, A UCB linguistics doctoral student is using a statistical, corpus-based approach in combination with psycholinguistic experimentation, to explore probabilistic relations between phonology on the one hand and syntax, semantics, and social knowledge on the other. He and Nancy Chang developed a formal notation for an embodied version of Construction Grammar, which plays a crucial role in a larger, simulation-based language understanding system. They also devised an experimental means by which to test the psychological reality of construal, the variable, context-specific understanding of the semantic pole of linguistic constructions. Nancy Chang continued developing representations and algorithms useful for an embodied approach to language acquisition and use. She worked with colleagues to flesh out different aspects of a simulation-based approach to language understanding, including a formal representation for linguistic constructions (Embodied Construction Grammar, devised in collaboration with Ben Bergen and Mark Paskin). This formal representation was extended to handle idioms (with Ingrid Fischer) and integrated with a module (by Rob Hunter) that performs simple context-based reference resolution. A version of the formalism was incorporated into her thesis research, which focuses on the development of an algorithm that learns such constructions from a set of utterance-situation pairs.

Thus the NTL group has, over the last year, formalized and significantly extended its work on language learning and use based on deep conceptual semantics. Both the learning sub-task and the performance HCI system are moving ahead in collaboration with other efforts at ICSI and elsewhere.

4.2 FrameNet Project, Charles Fillmore PI

In the year 2000, the NSF-sponsored FrameNet project (NSF IRI-9618838 - "Tools for Lexicon Building") continued building the FrameNet Database, a lexicon which contain about 2000 words (made available on the World Wide Web), providing (1) a collection of semantically annotated examples for each sense of each word, (2) links to descriptions of the conceptual structures ("frames") which underlie each such sense, (3) details of the syntactic ways in which the semantic roles ("frame elements") in each such conceptual structure are grammatically realized in sentences containing the word (in a table of Frame

Element Realizations) and (4) records of the combinations of frame elements and their syntactic realizations expressed in the sample sentences (in a table of Valence Patterns).

The project makes use of the 100M word British National Corpus (provided by Oxford University Press), corpus-management tools developed in IMS-Stuttgart, MITRE-Corp annotation software which allows the introduction of XML tags into the sentences forming the examples database, and various in-house devices for viewing the database. Currently a MySQL-based web interface is used for making the database viewable and searchable by human users.

The manual analysis and annotation activities of the project are extremely labor-intensive, but it is hoped that the resulting data can be used as a training corpus for future automatic operations, these designed to serve to annotate sentences for that have not been manually annotated, for frames that have been analyzed, and to discover phrases that are most likely to have semantic relations to given target words, based on syntactic parsing of corpus sentences, in the case of frames that have not been analyzed.

The FrameNet project is in its second major phrase, having received \$2.1M in the year 2000 (NSF HCI: #0086132, FrameNet++: An Online Lexical Semantic Resource and its Application to Speech and Language Understanding) for expanding the lexical database itself and for pilot projects on a battery of NLP applications that make use of it. These applications will emphasize the areas of automatic word-sense disambiguation, automatic semantic role labeling, machine translation, information extraction, information tracking, and the design of meaning representations based on FrameNet annotations that can support automatic text understanding.

Charles Wooters joined the team in the early months of this phase to recast the entire resource into a MySQL database and Beau Cronin has taken on the job of maintaining the database and preparing the software needed for the applications. Subcontractors Dan Jurafsky of the University of Colorado, Srinivas Narayanan of SRI International, and Mark Gawron of the San Diego State University will participate in automatic labeling, semantic representation, and machine translation, respectively. Student members have been hired to participate in analysis of lexical information, preparation of sample materials for annotation, and doing the annotation. The actual production of lexicographic descriptions begins in February 2001.

The project's task, now as before, is to document from actual text data the varieties of uses of English lexical items, and this work has to begin by sorting the samples that provide this documentation according to the word's meanings. For each meaning of each word a semantic "frame" is postulated which presents the conceptual structure that underlies it. The descriptive terminology used in characterizing each frame provides labels for the annotation of sentences using words that belong to the frame, and both manual and automatic means are employed to provide a full account of the semantic and syntactic combinatorial possibilities of each word in each sense. The corpus on which these observations were based in the first phase was the British National Corpus (100M running words); we have now added an American Newspaper Corpus made available through the Linguistic Data Consortium (University of Pennsylvania), and we are actively participating in the new American National Corpus centered in New York University.

The newly designed database (Figure xx) makes it possible to map out accurate relations between inflected word forms and the lexical items that they belong to, between these

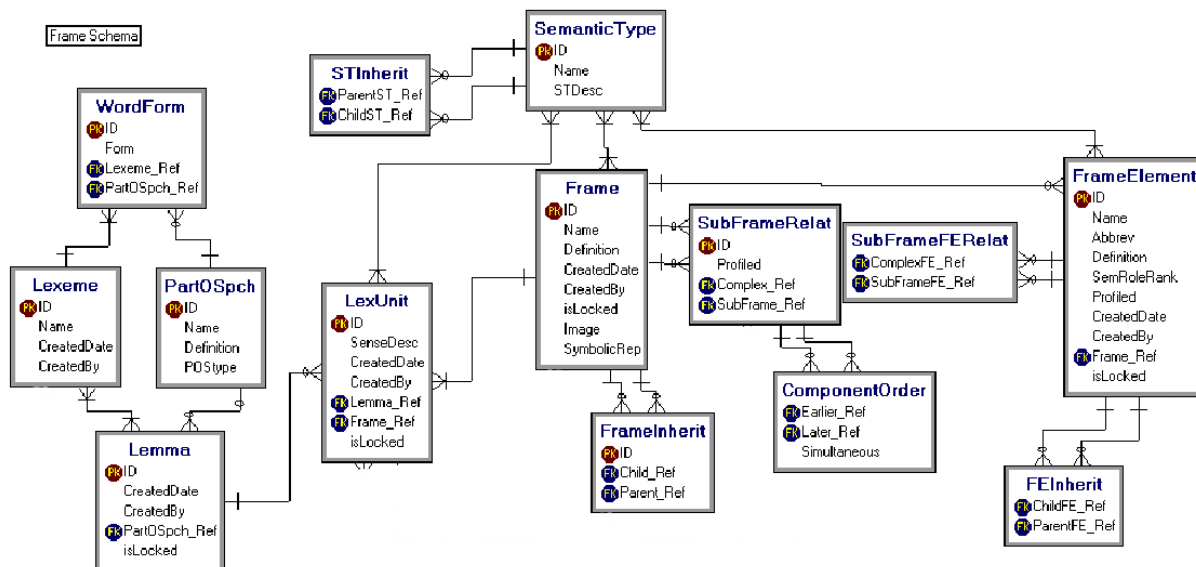


Figure 7: Database Tables for Frame Elements, Etc.

items and the (possibly multiple) semantic frames that they evoke, between frames and the other frames which are related to them through relationships of inheritance, co-existence, and composition, and between any and all of these properties (or combinations of them) and passages from the corpus which exemplify them. Computational linguists who have examined the facilities are generally agreed that the FrameNet resource, though it is currently quite limited in scope, carries deeper and more accurate information about lexical meanings than can be obtained from statistically acquired computer lexicons or from any of the existing machine-tractable electronic dictionaries.

A major effort is under way to create a multilingual version of FrameNet, with researchers on German, Spanish, Portuguese and Italian, centered in Stuttgart, Barcelona, Erfurt, and Pisa and including several corporate partners. The Berkeley team has agreed to collaborate in this "EuroFrameNet" project on a consulting basis. (Funding for this is being sought through the European Commission.)

The FrameNet project has both salaried and unsalaried participants. Volunteer programmers have included Qibo Zhu and Hiroaki Sato, visiting scholars from Guanzhou University in China and Senshu University in Japan, respectively, as well as Monica Oliver, UCB graduate student in Linguistics; Miriam Petruck, Linguistics Department visiting scholar has served as a volunteer annotator. The paid staff have included, over the three-year period, John B. Lowe, Collin Baker, Christopher Johnson, and Susanne Gahl (all graduate students who attained the doctorate while working for FrameNet); ICSI SysAd staff members Jane Edwards and Charles Wooters ; graduate students Josef Ruppenhofer, Tess Wood, Margaret Urban, Nancy Urban, Michael Ellsworth, Shweta Narayan and Paula Rogers; and undergraduate students Marianne Tolley, Christopher Struett, Peter Wong, Ursula Wagner, and Michael Locke, paid through an NSF REU grant. Students who have worked on the German-funded project while at ICSI include Susanne Gahl, Josef Ruppenhofer, Margaret Urban, and Hans Boas.

4.3 Connectionist Modeling

Lokendra Shastri's work on computational modeling has spanned three different representational and processing tiers of language processing. One tier focuses on high-level reasoning underlying language understanding. The second tier focuses on the formation of episodic memory whereby transient patterns of neural activity representing events and situations are rapidly transformed into persistent neural circuits (memory traces) capable of supporting recognition and recall. The third modeling effort concerns the extraction of syllabic segments from spontaneous and noisy speech. The results of the three efforts are summarized below.

4.3.1 A neurally motivated model of reasoning, decision making, and acting

SHRUTI, a structured connectionist model of reflexive inference, was extended to produce a unified connectionist architecture for representing beliefs and utilities. The extended Shruti system can propagate utilities and beliefs to make predictions, seek explanations, and identify actions that could make the world state more desirable. If the predictions and explanations suggest that undesirable things are going to happen then the system can seek actions that are likely to prevent them from happening (i.e., the system attempts to identify actions that would tend to maximize the expected future utility). This work is significant since it shows that a single causal structure (expressed as a neurally plausible network) can serve three purposes (i) understand the world, (ii) predict the future possibilities, and (iii) plan for a better future. Several key problems remain open. These include using the results of inference for making decisions, identifying complex plans, and learning the causal structure. Work on these problems is in progress.

An unsupervised learning mechanism, Causal Hebbian Learning, is being developed to learn a causal model with appropriate probabilistic link weights. This neurally plausible mechanism is a variant of hebbian learning where the precise timing of source and target activity is important to the learning process, and where different types of connections can respond to this timing in different ways.

The work on SHRUTI is being done with graduate student Carter Wendelken. The results of the above work are described in (Shastri & Wendelken, 2000; Wendelken & Shastri, 2000; Wendelken & Shastri, submitted). A new Java-based version of the SHRUTI simulator is also under development.

- **A biologically realistic model of episodic memory**

SMRITI is a computational model of episodic memory that demonstrates how a transient pattern of activity in cortical circuits representing an event can be transformed rapidly into a persist memory trace (i.e., a neural circuit) in the hippocampal system (HS) as a result of long-term potentiation.

Over the past year, work on SMRITI has further elaborated the role of the HS in the acquisition, maintenance and retrieval of episodic memory, and explicated how the HS might realize the episodic memory function in concert with cortical representations.

A quantitative analysis of the model was carried out to find out the type of episodic memory deficits that would result from cell loss in the HS and high-level cortical

circuits that project to the HS. The results of these analyses offer an explanation for the impact of semantic dementia and aging on the functioning of episodic memory.

A computational abstraction of long term potentiation (LTP) was developed that is simple enough to lend itself to rapid computation, but at the same time, complex enough to capture key temporal and cooperative properties of LTP (Shastri, To appear). This abstraction of LTP, in turn, is based on abstractions of cells and synapses that are significantly more detailed than those used in artificial neural network and connectionist models, but far simpler than those used in biophysically detailed simulations.

Results of the above work were presented as a featured presentation at CNS'00, the annual Computational Neuroscience Meeting, Brugge, Belgium.

- **Temporal flow models for speech processing**

Further progress was made in the development of a system for the automatic labeling of phonetic segments (ALPS). This system consists of several Temporal Flow Model (TFM) networks for Articulatory Feature (AF) classification followed by Multilayer Perceptrons (MLP) for performing AF integration and phone labeling. Over the past year, error analysis was carried out to understand the relation of phone classification errors and speech syllable structure, system performance was improved by using reduced and multiple temporal resolution, and the system was trained on mixed-noisy conditions to achieve robustness in the presence of noise. Some of the results of this work are published in Chang, Shastri & Greenberg, 2000.

We also worked on the Switchboard transcription project to obtain a full phonetic transcription, and at the same time, test the ALPS system on a more challenging corpus. The results on the 300 CV sentences were 88.23% hits and 7.57% insertions. When tested on a subset of 236 CV sentences that had good match between phonetic and syllabic transcriptions, the results were 92.5% hits and 5.69% insertions.

- **Training and human resources**

Two graduate students, Carter Wendelken and Shawn Chang, and one undergraduate student, Chetan Nandkumar, participated in further research efforts. The students continued developing representations and algorithms useful for an embodied approach to language acquisition and use. They worked with colleagues to flesh out different aspects of a simulation-based approach to language understanding, including a formal representation for linguistic constructions (Embodied Construction Grammar, devised in collaboration with Ben Bergen and Mark Paskin). This formal representation was extended to handle idioms (with Ingrid Fischer) and integrated with a module (by Rob Hunter) that performs simple context-based reference resolution. Most importantly, a version of the formalism was incorporated into their thesis research, which focused on the development of an algorithm that learns such constructions from a set of utterance-situation pairs.

4.3.2 Publications which appeared or were in press during 2000

1. Bergen, B. Probability in phonological generalizations: Modeling optional French final consonants. In Alan Yu et al. (Eds.), *Proceedings of the 26th Annual Meeting of the Berkeley Linguistics Society*. Berkeley: Berkeley Linguistics Society.
2. Bergen, B. Ramifications of phonology-syntax interactions for phonological models. In *Proceedings of the 24th Annual Penn Linguistics Colloquium*, 27-36.
3. Bergen, B. and Chang, N. (2000) Spatial schematicity of prepositions in Neural Grammar. Fifth Conference on Conceptual Structure, Discourse, and Language, Santa Barbara.
4. Bergen, B. and Chang, N. (2000) Prepositional semantics in Embodied Construction Grammar. Fifth Berkeley-Paris Cognitive-Enunciative Conference, Berkeley.
5. Bergen, B. and Chang, N. (2001) Semantic agreement and construal. 2001 meeting of the Linguistic Society of America. Washington, D.C.
6. Bergen, B., Chang, N. and Paskin, Mark A. Simulation-based language understanding in Embodied Construction Grammar. In Jan-Ola Ostman (Ed.), *Construction Grammar(s): Cognitive and Cross-language dimensions*.
7. Chang, N. and Fischer, I. (2000) Understanding idioms. KONVENS 2000/SPRACHKOMMUNIKATION. Ilmenau, Germany.
8. Chang, N. and Hunter, R. (2001) Reference resolution in child-directed speech: A deep semantic approach. International Workshop on Reference and Coherence in Discourse; Formal, Functional and Cognitive approaches. Utrecht, The Netherlands.
9. Chang, N. and Maia, T. (2001) Grounded learning of grammatical constructions. 2001 AAAI Spring Symposium on Learning Grounded Representations.
10. Chang, S., Shastri, L., and Greenberg, S. (2000) Automatic phonetic transcription of spontaneous speech (American English), *Proceedings of the 6th International Conference of Spoken Language Processing, ICSLP 2000*, Beijing, China, Oct., 2000.
11. Fillmore, C. J. and Baker, C. F. (2000) FrameNet: Frame semantics meets the corpus. Poster session at The Linguistic Society of America, January, 2000.
12. Fillmore, C. J., Wooters, C., and Baker, C. F. (2001) Building a large lexical databank which provides deep semantics. *Proceedings of the Pacific Asian Conference on Language, Information and Computation*. Hong Kong.
13. Gildea, D. and Jurafsky, D. (2000) Automatic labeling of semantic roles. *ACL 2000: Proceedings of ACL 2000*, Hong Kong
14. Jagota, A., Plate, T., Shastri, L., and Sun, R. (eds.), Connectionist Symbol Processing: Dead or Alive?, *Neural Computing Surveys* 2:1-40, 1999, 220 references. <http://www.icsi.berkeley.edu/~jagota/NCS>

15. Johnson, C. and Fillmore, C.F. (2000) The FrameNet tagset for frame-semantic and syntactic coding of predicate-argument structure. *Proceedings of the 1st Meeting of the North American Chapter of the Association for Computational Linguistics (ANLP-NAACL 2000)*(April 29-May 4, 2000):56-62.
16. Mani, D. and Shastri, L. (1999) Parallel inference with very large knowledge bases, In *High Performance Cluster Computing 2*, (Ed.) R. Buyya, Prentice Hall. 1999.
17. Shastri, L. (1999a) Advances in Shruti - A neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. *Applied Intelligence* 11:79–108.
18. Shastri, L. (1999b) Recruitment of binding and binding-error detector circuits via long-term potentiation. *Neurocomputing*, 26-27, 865–874.
19. Shastri, L. (1999c) Infants learning algebraic rules. *Science* 285(5434):1673. September, 1999.
20. Shastri, L. (1999d) A Biological Grounding of Recruitment Learning and Vicinal Algorithms. Technical Report TR-99-009, International Computer Science Institute, Berkeley, CA, 94704. April, 1999.
21. Shastri, L. (2000a) Types and quantifiers in Shruti — a connectionist model of rapid reasoning and relational processing. *Hybrid Neural Symbolic Integration*. S. Wermter and R. Sun (Eds.), Lecture Notes in Computer Science. Springer-Verlag, Berlin. (In Press).
22. L. Shastri (To appear) Biological grounding of recruitment learning and vicinal algorithms in long-term potentiation. L. Shastri. To appear in *Emergent Neural Computational Architectures Based on Neuroscience*, J. Austin, S. Wermter, and D. Wilshaw (Eds.), Lecture Notes in Computer Science, Lecture Notes in Artificial Intelligence Springer-Verlag, Berlin.
23. Shastri, L., Chang, S., and Greenberg, S. (1999) Syllable detection and segmentation using temporal flow neural networks. *Proceedings of the XIVth International Congress of Phonetic Sciences, ICPhS 99* 3: 1721-1724, San Francisco, CA, August, 1999.
24. Shastri, L., Grannes, D., Narayanan, S., and Feldman, J. (In Press) A Connectionist encoding of parameterized schemas and reactive plans. in *Hybrid Information Processing in Adaptive Autonomous vehicles*, G.K. Kraetzschmar and G. Palm (Eds.), Lecture Notes in Computer Science, Springer-Verlag, Berlin.
25. Shastri, L. and Wendelken, C. (1999a) Soft computing in Shruti: — A neurally plausible model of reflexive reasoning and relational information processing. *Proceedings of the Third International Symposium on Soft Computing, SOCO'99* Genova, Italy.(June, 1999):741–747.

26. Shastri, L. and Wendelken, C. (1999b) Knowledge fusion in the large – taking a cue from the brain. *Proceedings of the Second International Conference on Information Fusion, FUSION'99* Sunnyvale, CA. (July, 1999):1262–1269.
27. L. Shastri and C. Wendelken (2000) Seeking coherent explanations — a fusion of structured connectionism, temporal synchrony, and evidential reasoning. *Proceedings of the Twenty Second Annual Conference of the Cognitive Science Society* Philadelphia, PA. (Aug. 2000):453-458.
28. Wendelken, C. and Shastri, L. (2000) Probabilistic inference and learning in a connectionist causal network. *Proceedings of Neural Computation 2000* Berlin. May 2000.

5 Exploratory Areas in Computational Algorithms

In 2000, the Theory Group conducted research in computational biology and in the theory of computation. Additionally, much of the group effort has recently been focused on problems associated with internet research¹.

5.1 Computational Biology

The activities of living cells are performed and regulated by proteins. These activities include the maintenance of cellular structure, communication with other cells, cell division, programmed cell death, and the transport of chemicals within the cell and across the cell wall. The instructions for making proteins are encoded in the genes, but the steps required to manufacture and activate a protein are regulated by the environment of the cell and the proteins already active within the cell. These gene regulation processes involve an interplay among genes, RNA molecules, proteins and more complex molecular machines composed of these entities. High-throughput technologies for measuring cellular activity at the level of genes, proteins and molecular machines are developing rapidly, and there is a need for algorithms to interpret these measurements and construct mathematical models of cellular processes. Our ultimate goal is to contribute to the development of these algorithms and mathematical models.

The first step in making a protein is to transcribe its gene into messenger RNA (mRNA). DNA microarrays, also known as gene chips, enable the simultaneous detection and measurement of thousands of different species of mRNA, corresponding to thousands of different genes. Thus, a single microarray experiment can give a panoramic picture of the gene expression taking place in a tissue or collection of cells, at the level of mRNA production.

DNA microarrays can be used to measure gene expression in a set of biological samples which may correspond, for example, to different cancer specimens. These specimens typically fall into subclasses representing different phenotypes or clinical conditions. In the

¹R. Karp's research on the following topics is described in the ACIRI section under the subsections: CANs for Peer-to-Peer Content Distribution Systems; Randomized Rumor Spreading; Application Layer Multicast; and Fairness in AIMD Congestion Control. Several of R. Karp's publications are listed in the References for that Section.

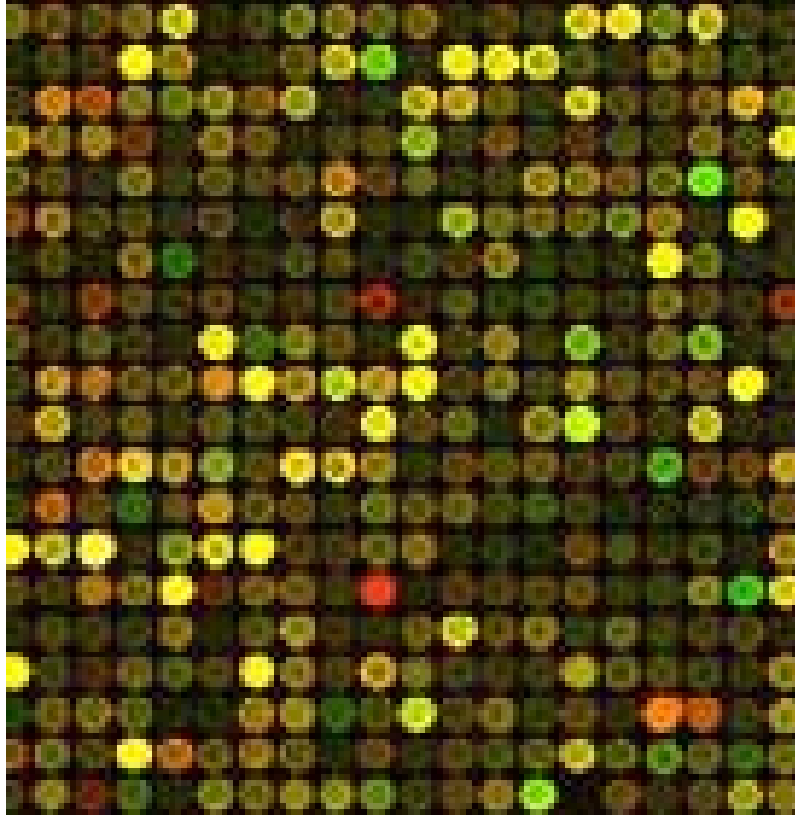


Figure 8: A DNA Microarray

case of *supervised learning* the subclass of each specimen is given (for example, by clinical observations), and the problem is to determine a genetic signature of each subclass, in order to classify future samples. In the case of *unsupervised learning* the classification is not given, and the task is to cluster the samples into homogeneous subclasses, and then to derive a genetic signature of each subclass.

Usually the number of samples is moderate, but the number of genes may be in the tens of thousands. Moreover, the great majority of genes may be irrelevant to the given classification problem, and even among the relevant genes there may be a great deal of redundancy, since many of the genes may have similar patterns of expression. These irrelevant or redundant genes merely introduce noise into the classification process. Thus we are faced with the *feature selection problem* of identifying the handful of informative genes that combine to determine the true classification and give a clear genetic signature of each subclass. In [1] and [2] we have introduced a battery of feature selection techniques and described how they can be used in combination to focus on a small explanatory set of genes. We have applied this method with excellent success to supervised and unsupervised learning problems in which the samples come from a microarray data set of 72 leukemia specimens.

References

- [1] Xing, E.P., Jordan, M.I. and Karp, R.M., “Feature Selection for High-Dimensional Genomic Microarray Data,” submitted, 2000.
- [2] Xing, E.P. and Karp, R.M., “CLIFF: Clustering of High-Dimensional Microarray data via Iterative Feature Filtering Using Normalized Cuts,” submitted, 2000.

5.2 Theory of Computation

Postdoctoral fellows Antonio Piccolboni and Christian Schindelhauer have obtained an important general result concerning prediction problems with limited feedback. They consider a game between a *predictor* P and a *sequence generator* S . At each step S generates a symbol s and, simultaneously, P makes a prediction p . P then incurs a *loss* $L(p, s)$ and receives feedback $F(p, s)$. The only information P receives, in addition to its own prediction, is the feedback $F(p, s)$; it is told neither s nor $L(p, s)$. Nevertheless, Piccolboni and Schindelhauer show that, for a wide class of loss and feedback functions, there is a randomized prediction algorithm whose asymptotic performance, in terms of cumulative loss, is at least as good as the loss of the best *constant* predictor (a constant predictor is one that makes the same prediction at every step). This result is an advance in a long line of research in statistics, information theory and economics, with applications ranging from predicting the outcomes of coin tosses to investing in the stock market. The problem setting was originally suggested by a bandwidth prediction problem arising in Internet congestion control.

Relevant references are given below:

References

- [1] Piccolboni, A. and Schindelhauer, C., “Discrete Prediction games with Arbitrary Feedback and Loss,” Technical Report SIIM-TR-A-00-18, Universität zu Lübeck, November, 2000.
- [2] Horn, G.B. and Karp, R.M., “A Maximum Likelihood Polynomial Time Syndrome Decoder to Correct Linearly Independent Errors,” to appear in Proc. International Symp. on Information Theory, 2001.
- [3] Ben-Dor, A., Karp, R.M., Schwikowski, B. and Yakhini, Z., “Optimally Multiplexed SNP Genotyping on Universal DNA Tag Systems,” submitted, 2001.