Looking Over the Fence at Networks

A Neighbor's View of Networking Research

Committee on Research Horizons in Networking

Computer Science and Telecommunications Board

Division on Engineering and Physical Sciences

National Research Council

NATIONAL ACADEMY PRESS
Washington, D.C.

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

Support for this project was provided by core funds of the Computer Science and Telecommunications Board. Core support for the CSTB is provided by its public and private sponsors: the Air Force Office of Scientific Research, Defense Advanced Research Projects Agency, Department of Energy, National Aeronautics and Space Administration, National Institute of Standards and Technology, National Library of Medicine, National Science Foundation, Office of Naval Research, AT&T, Hewlett-Packard, Intel, Microsoft, and Time-Warner Cable.

International Standard Book Number 0-309-07613-7

Additional copies of this report are available from:

National Academy Press
2101 Constitution Ave., N.W.
Box 285
Washington, DC 20814
800-624-6242
202-334-3313 (in the Washington metropolitan area)
http://www.nap.edu

Printed in the United States of America

THE NATIONAL ACADEMIES

National Academy of Sciences
National Academy of Engineering
Institute of Medicine
National Research Council

The **National Academy of Sciences** is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Bruce M. Alberts is president of the National Academy of Sciences.

The **National Academy of Engineering** was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. Wm. A. Wulf is president of the National Academy of Engineering.

The **Institute of Medicine** was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Kenneth I. Shine is president of the Institute of Medicine.

The **National Research Council** was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Bruce M. Alberts and Dr. Wm. A. Wulf are chairman and vice chairman, respectively, of the National Research Council.

COMMITTEE ON RESEARCH HORIZONS IN NETWORKING

DAVID A. PATTERSON, University of California at Berkeley, *Chair*
DAVID D. CLARK, Massachusetts Institute of Technology
ANNA KARLIN, University of Washington
JIM KUROSE, University of Massachusetts at Amherst
EDWARD D. LAZOWSKA, University of Washington
DAVID LIDDLE, U.S. Venture Partners
DEREK McAULEY, Marconi
VERN PAXSON, AT&T Center for Internet Research at ICSI
STEFAN SAVAGE, University of California at San Diego
ELLEN W. ZEGURA, Georgia Institute of Technology

Staff

JON EISENBERG, Senior Program Officer
MARJORY S. BLUMENTHAL, Executive Director
JANET BRISCOE, Administrative Officer
MARGARET HUYNH, Senior Project Assistant
DAVID PADGHAM, Research Assistant

Preface

This report is the result of a new approach by the Computer Science and Telecommunications Board (CSTB) of the National Research Council to developing research agendas in key areas of information technology. Typically, only the members of a particular research community participate in defining an agenda for their future research activities. CSTB convened a small workshop in which more than half of the attendees were researchers in *other* fields. The premise behind this approach was that working together with a smaller number of network research insiders, these outsiders—people whose primary research interests were not in networking but who represented instead additional subdisciplines of computer science, as well as other disciplines such as Earth science, economics, and information studies—would function much like a visiting committee, providing a fresh perspective on research topics and directions and helping to stimulate development of a strategic view of future research directions. CSTB picked networking as the subject of this board-initiated project—the first in a planned series of workshops—since it is a field that has enjoyed both success and great attention due to its most visible creation, the Internet. As this report illustrates, it is also a compelling field in which to explore alternative research visions because that very success has constrained some avenues of research.

The presence of outsiders was critical to the dialogue at the January 2001 workshop. To kick off the discussion, both insiders and outsiders were asked to identify near-term and long-term research topics, as well as topics that should probably be deemphasized, at least for a while (Box P.1). In the discussions that followed, outsiders posed provocative questions that challenged conventional wisdom and suggested different research approaches drawn from their own research communities. They also brought the perspectives of experienced network users to the discussion and repeatedly expressed frustrations about the inadequacies of the Internet from their user perspective.

Workshop participants noted but did not elaborate on a number of topics that are of current interest in the network community. In some cases, topics were explicitly taken off the table. A shared sense among outsiders and insiders that topics other than network performance were at least as important and receiving less attention in the research community meant that workshop participants paid little attention to the issue of how to build higher-speed networks. Limitations of time and expertise precluded an in-depth examination of the implications of wireless and optical technologies, but participants did observe that such examination would be an important activity for the research community. In other cases, subjects arose in discussions but were not ultimately identified as areas meriting greater attention by networking researchers (e.g., last-mile access links). Other topics provoked mixed reactions; for example, some felt that multicast continues to be important while others felt that it should be abandoned as a research topic.

The workshop proved educational for everyone involved. The outsiders learned about some surprising characteristics of networking culture. For example, the research community is protective of the Internet; reviewers often reject papers that make proposals perceived as potentially deleterious to the Internet. Also, even when confidentiality is not at issue, network researchers are reluctant to identify by brand name specific products or services with features they find undesirable. The insiders, in turn, were surprised to hear that the outsiders were not very interested in seeing great efforts expended on more research to improve raw network performance (distinct from work on better characterizing network performance, which was of interest to some participants). There were other surprises: For example, outsiders were surprised by how mistakes made by a few people in the configuration of routing tables could bring a significant portion of the Internet to its knees.

BOX P.1 Questions Posed in Advance to Workshop Participants

1. What are three pressing problems in networking (that is, short-term problems that ideally would have been research problems 5 to 7 years ago)?

2. What are two fundamental research problems in networking (that is, things that would be important to put into practice in 5 to 7 years)?

3. What is one topic in networking that you would rather not read about again (that is, a topic that could be deferred to allow work on other problems)?

This report does not provide answers to these specific questions—the questions were posed as a way of stimulating discussions at the workshop.

The report that follows was written by the Committee on Research Horizons in Networking, composed of six networking researchers and four researchers from other areas in computer science, based on the 2 days of discussions among a larger group of workshop participants that was dominated by outsiders. The committee met immediately following the workshop and conducted a series of discussions by e-mail to formulate a fresh look at networking research, drawing on the workshop experience.

The report is organized around the three major themes, closely connected to the *process* of networking research, that emerged at the workshop—measuring, modeling, and creating and deploying disruptive prototypes. It is not a report that seeks to lay out a detailed research agenda per se. The issues raised in this report, which reflect in large part the concerns of the outsiders, would certainly require further consideration by the network research community to be translated into an actual research agenda that would help meet the needs of network users. For example, while outsiders bring a valuable fresh perspective, they can also miss obstacles that insiders see. The intent of this report is to stimulate such an examination.

David A. Patterson, *Chair*
Committee on Research Horizons
in Networking

Acknowledgment of Reviewers

This report has been reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise, in accordance with procedures approved by the NRC's Report Review Committee. The purpose of this independent review is to provide candid and critical comments that will assist the institution in making its published report as sound as possible and to ensure that the report meets institutional standards for objectivity, evidence, and responsiveness to the study charge. The review comments and draft manuscript remain confidential to protect the integrity of the deliberative process. We wish to thank the following individuals for their review of this report:

Craig Partridge, BBN Technologies,
Larry Peterson, Princeton University,
Scott Shenker, AT&T Center for Internet Research at ICSI, and
James P.G. Sterbenz, BBN Technologies.

Although the reviewers listed above provided many constructive comments and suggestions, they were not asked to endorse the conclusions or recommendations nor did they see the final draft of the report before its release. The review of this report was overseen by Jerome H. Saltzer, Massachusetts Institute of Technology, appointed by the NRC's Division on Engineering and Physical Sciences, who was responsible for making certain that an independent examination of this report was carried out in accordance with institutional procedures and that all review comments were carefully considered. Responsibility for the final content of this report rests entirely with the authoring committee and the institution.

Contents

1 Introduction

The Internet has been highly successful in meeting the original vision of providing ubiquitous computer-to-computer interaction in the face of heterogeneous underlying technologies. No longer a research plaything, the Internet is widely used for production systems and has a very large installed base. Commercial interests play a major role in shaping its ongoing development. Success, however, has been a double-edged sword, for with it has come the danger of ossification, or inability to change, in multiple dimensions:

- *Intellectual ossification*—The pressure for compatibility with the current Internet risks stifling innovative intellectual thinking. For example, the frequently imposed requirement that new protocols not compete unfairly with TCP-based traffic constrains the development of alternatives for cooperative resource sharing. Would a paper on the NETBLT protocol that proposed an alternative approach to control called "rate-based" (in place of "window-based") be accepted for publication today?
- *Infrastructure ossification*—The ability of researchers to affect what is deployed in the core infrastructure (which is operated mainly by businesses) is extremely limited. For example, pervasive network-layer multicast remains unrealized, despite considerable research and efforts to transfer that research to products.[1]
- *System ossification*—Limitations in the current architecture have led to shoe-horn solutions that increase the fragility of the system. For example, network address translation violates architectural assumptions about the semantics of addresses. The problem is exacerbated because a research result is often judged by how hard it will be to deploy in the Internet, and the Internet service providers sometimes favor more easily deployed approaches that may not be desirable solutions for the long run.

At the same time, the demands of users and the realities of commercial interests present a new set of challenges that may very well require a fresh approach. The Internet vision of the last 20 years has been to have all computers communicate. The ability to hide the details of the heterogeneous underlying technologies is acknowledged to be a great strength of the design, but it also creates problems because the performance variability associated with underlying network capacity, time-varying loads, and the like means that applications work in some circumstances but not others. More generally, outsiders advocated a more user-centric view of networking research—a perspective that resonated with a number of the networking insiders as well. Drawing on their own experiences, insiders commented that users are likely to be less interested in advancing the frontiers of high communications bandwidth and more interested in consistency and quality of experience, broadly defined to include the "ilities"—reliability, manageability, configurability, predictability, and so forth—as well as non-performance-based concerns such as security and privacy. (Interest was also expressed in higher-performance, broadband last-mile access, but this is more of a deployment issue than a research problem.) Outsiders also observed that while as a group they may share some common requirements, users are very diverse—in experience, expertise, and what they wish the network could do. Also, commercial interests have given rise to more diverse roles and complex relationships that cannot be ignored when developing solutions to current and future networking problems. These considerations argue that a vision for the future Internet should be to provide users the quality of experience they seek and to accommodate a diversity of interests.

[1] Other instances of infrastructure ossification noted by networking researchers include challenges associated with deploying various flavors of quality of service and IPv6.

This report explores how networking research could overcome the evident obstacles to help achieve this vision for the future and otherwise better understand and improve the Internet. The report, which reflects interactions among networking researchers and outsiders (researchers from fields other than networking) at CSTB's January 2001 workshop, as well as subsequent discussion by the Committee on Research Horizons in Networking, stresses looking beyond the current Internet and evolutionary modifications thereof and aims to stimulate fresh thinking within the networking research community. Since it is not a formal research agenda (which would, among other things, entail a much more intensive effort than is afforded by an exploratory workshop such as this), the report does not, for example, review past literature and current research programs but instead briefly characterizes past progress, current efforts, and promising directions. It focuses on three key areas in which networking research might be invigorated: measuring the Internet, modeling the Internet, and making disruptive prototypes.

2 Measuring: Understanding the Internet Artifact

A remarkable creation, the Internet encompasses a diversity of networks, technologies, and organizations. The enormous volume and great variety of data carried over it give it a rich complexity and texture. It has proved difficult to characterize, understand, or model in terms of large-scale behaviors and a detailed understanding of traffic behavior. Moreover, because it is very difficult to prototype new networks—or even new networking ideas—on an interesting scale (see Chapter 4), data-driven analysis and simulation are vital tools for evaluating proposed additions and changes to its design.

Experimental science is an important approach in many areas of computer science and engineering, especially where the artifacts being studied are complex and have properties that are not well understood.[1] Central to the experimental method is the repeated measurement of observed behavior. Without acquiring such data it is impossible to analyze and understand the underlying processes, let alone predict the impact of a change to the environment being observed. Further, data often help suggest new theoretical approaches. Measurement is at least in part driven by a particular question at hand, and changing questions over time may well lead to different measurement needs.

However, there are also strong arguments for collecting data in anticipation of future use. Citing the heavy dependence of our knowledge and understanding of global climate change on a record of atmospheric carbon dioxide measurements that Charles David Keeling started on Mauna Loa in 1957, workshop participant Jeff Dozier observed that "good data outlives bad theory."[2] Hence a data set with typical days from the next 10 years of the Internet might be a treasure chest for networking researchers just as the carbon dioxide record has been to earth scientists. Also, outsiders at the workshop observed that in other areas of computer science, older versions of artifacts—old microprocessors, operating systems, and the like—are important as bases for trend analysis and before/after comparisons of the impacts of new approaches.[3] Archived Internet snapshots could provide an analogous baseline for evaluating the large-scale impact of both evolutionary and revolutionary changes in the Internet. Archived data could also be used by Internet researchers to determine if newly identified traffic phenomena (for example, a future equivalent of heavy-tailed behavior) existed in earlier instantiations of the Internet.

Unfortunately, the ability of network researchers or operators to measure the Internet is significantly limited by a number of interdependent barriers. The extreme scale of today's Internet poses a challenge to acquiring a representative set of data points. The Internet architecture itself also makes measurement difficult. Factors such as the end-to-end design,

---

[1]For more discussion of this issue, see Computer Science and Telecommunications Board, National Research Council. 1994. *Academic Careers for Experimental Computer Scientists.* National Academy Press, Washington, D.C.

[2]Established initially with NSF support, the measurement program in the 1970s required cobbling together support from multiple sources and continuation of the measurement program was at risk. Today the carbon dioxide record is one of science's most famous data sets, and an illustration of the principle that "good data can outlast bad theory." The measurements that Keeling started, and maintained under some adversity, are the cornerstone of our analysis of the human effects on our climate. The data show that atmospheric CO_2 has increased about 1/2 percent per year since 1957. Measurements in other locations show a interhemispheric transport and seasonal variability. Comparison with data from ice cores allows us to extend the record backward more than 100,000 years. See Charles D. Keeling. 1998. "Rewards and Penalties of Monitoring the Earth," *Annual Review of Energy and the Environment*, vol. 23, pp. 25-82.

[3]The importance of archiving artifacts of complex software systems was discussed earlier in Computer Science and Telecommunications Board, National Research Council. 1989. *Scaling Up.* National Academy Press, Washington, D.C.

layering, and the statelessness of the basic datagram[4] make it hard to identify some types of flows. Factors such as routing asymmetry and multipathing make it hard to gather necessary information even about self-describing flows such as TCP. Also, business concerns and increased sensitivity to privacy limit the willingness of many stakeholders to participate in data collection and constrain the release of data to a wider research community. The resulting paucity of sound or representative data has severely limited the ability to predict the effects of even incremental changes to the Internet architecture, and it has undermined confidence in more forward-thinking research.

Progress in measuring the Internet artifact will thus require the effort, ingenuity, and unified support of the networking community. In other fields, grand challenges—such as mapping the entire human genome—have served to expose and crystallize research issues and to mobilize research efforts. Along those lines, a challenge that could stimulate the necessary concerted effort is the following: (1) to develop and deploy the technology to make it possible to record a day in the life of the Internet, a data set containing the complete traffic, topology, and state across the Internet infrastructure and (2) to take such a snapshot. Even if the goal were realized only in part, doing so would provide the infrastructure for establishing a measurement baseline.

A "day in the life" should be understood as a metaphor for a more precise formulation of the measurement challenge. For example, the appropriate measurement period might not literally be a single 24-hour period (one might want to take measurements across a number of days to explore differences between weekdays and weekends, the effects of events that increase network traffic, and the like) and, as discussed below, the snapshot might sample traffic rather than record every single packet. To achieve many of the goals, one would also measure on an ongoing basis rather than as a one-time event.

This ambitious goal faces many hurdles that together form the foundation for a valuable research agenda in their own right. Although the overarching goal is the ability to collect a full snapshot, progress on each of the underlying problems discussed below would be a valuable step forward toward improving our understanding of actual network behavior.

THE CHALLENGES OF SCALE

Accommodating the growth in link speeds and topology is a significant challenge for large-scale Internet traffic measurement. Early versions of equipment with OC-768 links (40 gigabits per second) are already in trials, and the future promises higher speeds still. Worse yet, each individual router may have many links, increasing the overall computational challenge as well as making per-link measurement platforms extremely expensive to deploy and difficult to manage. Addressing these problems presents both engineering and theoretical challenges. High-speed links demand new measurement apparatus to measure their behavior, and efficient measurement capabilities must be incorporated into the routers and switches themselves to accommodate high port densities. Even with such advances it may be infeasible to collect a complete record of all communication in a highly loaded router, and we may be forced to sample traffic instead. To do so effectively will require developing a deeper understanding of how to soundly sample network traffic, which is highly correlated and structured. An especially important statistics question is how to assess the validity of a particular sampling approach—its accuracy, representativeness, and limitations—for characterizing a range of network behaviors.

One particular challenge in measuring the network today is incomplete knowledge about the internal configuration of parts of the network, a reflection of network operators' reluctance to divulge information that may be of interest to their competitors. One way to cope with this impediment is the use of inference techniques that allow one to learn more about a network based

[4]For further discussion of these design issues, see Computer Science and Telecommunications Board, National Research Council. 2001. *The Internet's Coming of Age.* National Academy Press, Washington, D.C.

on incomplete, publicly accessible/observable information. For example, there has been research using border gateway protocol (BGP) routing table information to infer the nature of interconnection agreements between Internet service providers (ISPs). Inference techniques will not, in general, provide complete information, and more work is needed on how to make use of such incomplete information. Workshop participants noted that these statistical issues (along with the modeling issues discussed in the next chapter) would benefit from the involvement of statisticians.

A snapshot of an Internet day would contain an immense amount of data. Like other scientific communities faced with enormous data sets (for example, astronomy or the earth sciences), the Internet research community must grapple with analyzing data at very large scales. Among these challenges are effectively mining large, heterogeneous, and geographically distributed datasets; tracking the pedigree of derived data; visualizing intricate, high-dimensional structures; and validating the consistency of interdependent data. An additional challenge posed by measuring Internet traffic, which is also found in some other disciplines such as high-energy physics, is that data arrive quickly, so decisions about data sampling and reduction have to be made in real time.

MEASUREMENT INFRASTRUCTURE

In addition to the significant theoretical challenges, large-scale measurement of the Internet presents enormous deployment and operational challenges. To provide widespread vantage points for measuring network activity, even a minimal infrastructure will comprise hundreds of measurement devices. There is some hope that advances in remote management technologies will support this need, and lessons from several currently deployed pilot measurement projects could aid in the design of any such system. However, such an effort would also requires funding and personnel able to deploy, maintain, and manage the large-scale infrastructure envisioned here. In the long run, the value of this investment will be the creation of a foundation for watching network trends over time and establishment of an infrastructure available to researchers for new questions that are not adequately addressed by previous measurements.

Many of the challenges found in measuring today's Internet could have been alleviated by improved design, which underscores the importance of incorporating self-measurement, analysis, and diagnosis as basic design points of future system elements and protocols. This is particularly critical to providing insight into failures that are masked by higher layers of abstraction, as TCP does by intentionally hiding information about packet loss from applications.

NONTECHNICAL FACTORS

Although many of the challenges to effective Internet measurement are technical, there are important nontechnical factors—both within the networking community and in the broader societal context—that must be addressed as well. The committee recognizes that gathering this data will require overcoming very significant barriers. One set of constraints arises because the Internet is composed in large part of production commercial systems. Information on traffic patterns or details of an ISP's network topology may reveal information that a provider prefers not to reveal to its competitors or may expose design or operational shortcomings. A related set of challenges concerns expectations of privacy and confidentiality. Users have an expectation (and in some instances a legal right) that no one will eavesdrop on their communications. As a consequence of the decentralized nature of the Internet, much of the data can only be directly observed with the cooperation of the constituent networks and enterprises. However, before these organizations are willing to share their data, one must address their concerns about protecting their users' privacy. Users will be concerned even if the content of their communications is not being captured—recording just the source, destination, type, or volume of the communications can reveal information that a user would prefer to keep private.

If network providers could find ways of being more open while protecting legitimate proprietary or privacy concerns, considerably more data could be available for study. Current understanding of data anonymization techniques, the nature of private and sensitive information, and the interaction of these issues with accurate measurement is rudimentary. Too simplistic a procedure may be inadequate: If the identity of an ISP is deleted from a published report, particular details may permit the identity of the ISP in question to be inferred. On the other hand, too much anonymity may hide crucial information (for example, about the particular network topology or equipment used) from researchers. Attention must therefore be paid to developing techniques that limit disclosure of confidential information while still providing sufficient access to information about the network to enable research problems to be tackled. In some circumstances, these limitations may prevent the export of raw measurement data—provoking the need to develop configurable "reduction agents" that can remotely analyze data and return results that do not reveal sensitive details.

Finally, realizing the "day in the life" concept will require the development of a community process for coming to a consensus on what the essential measurements are, the scope and timing of the effort, and so forth. It will require the efforts of many researchers and the cooperation of at least several Internet service providers. The networking research community itself will need to develop better discipline in the production and documentation of results from underlying data. This includes the use of more careful statistical and analytic techniques and sufficient explanation to allow archiving, repeatability, and comparison. To this end, the community should foster the creation of current benchmark data sets, analysis techniques, and baseline assumptions. Several organizations have engaged in such efforts in the past (on a smaller scale than envisioned here), including the Cooperative Association for Internet Data Analysis (CAIDA)[5] and the Internet Engineering Task Force's IP Performance Metrics working group (ippm).[6] Future measurement efforts would benefit from the networking community at large adopting analogous intergroup data-sharing practices.

[5]See <http://www.caida.org>.

[6]See <http://www.ietf.org/html.charters/ippm-charter.html>.

3 Modeling: New Theory for Networking

The coming of age of the Internet has brought about a dual set of challenges and opportunities. The intellectual tools and techniques that brought us this far do not appear to be powerful enough to solve the most pressing problems that face us now. Additionally, concerns that were once relegated to the background when the Internet was small and noncommercial are now of crucial importance. In these challenges lies the opportunity for innovation:

- Understanding scaling and dynamics requires the development of new modeling methodologies and the undertaking of new modeling efforts (employing both well-known and newly developed techniques) to take our understanding beyond that afforded by today's models.
- Concerns of manageability, reliability, robustness, and evolvability—long neglected by researchers—are of critical importance and require the development of new basic understanding and theory.
- Even traditional problem areas, such as routing, must be addressed in a new context in light of how the global Internet has evolved.

PERFORMANCE

Even as the Internet has grown more complex, those who study and use it seek to answer increasingly difficult questions. What sorts of changes in the scale and patterns of traffic could lead to a performance meltdown? What are the failure modes for large-scale networks? How can one characterize predictability?

Researchers have worked for years to develop new theory and improved models. While this work has yielded many insights about network behavior, understanding other aspects of the network has proven a difficult challenge. Workshop participants encouraged the networking research community to develop new approaches and abstractions that would help model an increasingly wide range of network traffic phenomena. Simple models are more easily evaluated and interpreted, but complex models may be needed to explain some network phenomena. Queues and other resources cannot always be treated in isolation, nor can models always be based on simplified router-link pictures of the network. Small-scale, steady-state, packet-oriented models may not adequately explain all Internet phenomena. It is also well known that more sophisticated input models (such as heavy-tailed traffic distributions) are required to accurately model some behaviors. In other cases, the need is not for increased model complexity or mathematical sophistication but for just the opposite: new simple models that provide insights into widescale behavior. These may well require dealing with networking traffic at a coarser time scale or higher level of abstraction than traditional packet-level modeling. Here, theoretical foundations in such areas as flow-level modeling, aggregation/deaggregation, translation between micro and macro levels of analysis, and abstractly modeling the effects of closed-loop feedback and transients could be helpful. Simulation is another important tool for understanding networks. Advances in large-scale simulation efforts would aid model validation and permit higher-fidelity results to be obtained.

THEORY: BEYOND PERFORMANCE

Over the past three decades, several bodies of theory, such as performance analysis and resource allocation/optimization, have contributed to the design and understanding of network architectures, including the Internet. However, as the Internet has evolved into a critical infrastructure used daily by hundreds of millions of users, operational concerns such as manageability, reliability, robustness, and evolvability have supplanted performance of the data forwarding plane as the limiting factors. Yet theoretical understanding of these crucial areas is poor, particularly in comparison with their importance. The reasons for this disparity are many, including the lack of commonly accepted models for research in these areas, the difficulty of defining quantitative metrics, and doubts about the intellectual depth and viability of scholarly research in these areas.

As an example, consider the use of "soft state"[1] in the Internet, long hailed as a robust technique (when compared with hard-state approaches) for building distributed applications. Yet what, precisely, is the benefit of using soft state? The notions of robustness, relative "simplicity," and ease of implementation generally associated with soft state have not been defined, much less quantified. To take another example, the notion of plug-and-play is widely believed to make networked equipment more manageable. However, the implications of such factors as increased code complexity and the cost of reconfiguring default settings remain elusive.

At the heart of this endeavor is the seemingly simple but deceptively elusive challenge of defining the problems and an appropriate set of starting assumptions. The next steps include developing new concepts or abstractions that would improve present understanding of the infrastructure, defining metrics for success, and pursuing solutions. Because the basic understanding and paradigms for research here have yet to be defined, the challenges are indeed daunting.

APPLYING THEORETICAL TECHNIQUES TO NETWORKING

Outsiders observed that more progress on fundamental networking problems might come from greater use of theoretical techniques and understanding from algorithm design and analysis, complexity theory, distributed computing theory, general system theory, control systems theory, and economic theory. For example, routing has been well studied, both theoretically and practically, but remains a challenging and important problem for the networking community. Some of the open research questions relevant to routing noted by workshop participants include the following:

- Developing a greater understanding of the convergence properties of routing algorithms such as the border gateway protocol (BGP) or improvements to it. BGP has been found to suffer from much slower than expected convergence and can fail if misconfigured.
- Developing a better theoretical framework for robustness and manageability to inform the development of less vulnerable designs.
- Designing new routing algorithms that take into account real-world constraints such as the absence of complete information (and, often, the presence of erroneous information), peering agreements and complex interconnections among ISPs, and local policy decisions.
- Developing routing schemes that take into account the fact that the network is not simply composed of routers and links—network address translators, firewalls, proxies, underlying transport infrastructures, and protocols all come into play. Which of these elements are relevant and how should they be abstracted to better understand routing?
- Developing an understanding of the conditions under which load balancing and adaptive multipath routing work effectively and the conditions under which they can lead to instability and oscillation.

[1]"State" refers to the configuration of elements, such as switches and routers, within the network. Soft state, in contrast to hard state, means that operation of the network depends as little as possible on persistent parameter settings within the network.

4 Making Disruptive Prototypes: Another Approach to Stimulating Research

In addition to measurement and modeling, a third approach to stimulating continued innovation is to build prototypes. Very successful, widely adopted technologies are subject to ossification, which makes it hard to introduce new capabilities or, if the current technology has run its course, to replace it with something better. Existing industry players are not generally motivated to develop or deploy disruptive technologies (indeed, a good example of disruptive technology is a technology that a major network hardware vendor would not consider implementing in its router products). Researchers in essence walk a fine line between two slippery slopes: Either carry out long-term research that may be difficult to apply to the Internet or work on much shorter-term problems of the sort that would be of interest to a router manufacturer or venture capitalist today, leaving little middle ground in which to invent new systems and mechanisms. So it is no surprise that as the scale and utility of the Internet have increased, it has become immensely difficult to develop an alternative vision of the network, one that would provide important new benefits while still supporting the features of today's Internet, especially at the enormous scale of today's network.

The Internet itself is, of course, a classic example of a disruptive technology that went from prototype to mainstream communications infrastructure. This section considers how to enable a similar disruptive innovation that addresses the shortcomings of today's Internet and provides other new capabilities. Box 4.1 lists some research directions indentified by workshop participants as ways of stimulating such disruptive network designs. Research communities in computer architecture, operating systems, databases, compilers, and so on have made use of prototypes to create, characterize, and test disruptive technologies. Networking researchers also make use of prototyping, but the barriers discussed above make it challenging to apply the prototype methodology to networking in a way that will result in disruptive change.

CHALLENGES IN DEPLOYING DISRUPTIVE TECHNOLOGY

One important consideration in any technology area—a key theme of the book *The Innovator's Dilemma*[1]—is that a disruptive technology is likely to do a few things very well, but its overall performance and functionality may lag significantly behind present technology in at least some dimensions. The lesson here is that if innovators, research funders, or conference program committees expect a new technology to do all things almost as well as the present technology, then they are unlikely to invent, invest in, or otherwise encourage disruptive technologies. Thus (re)setting community expectations may be important to foster disruptive prototypes. Expectation setting may not be enough, however; a new technology must offer some sort of compelling advantage to compensate for performance or other shortcomings as well as the additional cost of adopting it. Those applications that do not need some capability of the disruptive technology will use the conventional Internet since it is larger and more stable.

Also central to the notion of developing a disruptive technology is suspending, at least temporarily, backward compatibility or requiring that technology developers also create a viable migration strategy. Outsiders observed, for example, that rigid adherence to backward compatibility would have made the development of reduced instruction set computers (RISCs) impossible.

Another key factor in the success of a disruptive technology is the link to applications. The popularity of many disruptive computer technologies has been tied to the applications that people can build on top of the technologies. One example is the personal computer. Early on, it

[1]Clayton M. Christensen. 2000. *The Innovator's Dilemma.* HarperCollins.

BOX 4.1 Some Potentially Disruptive Ideas About Network Architecture and Design

Workshop participants discussed a number of architectural/design issues that could stimulate disruptive network designs. The items that follow, though not necessarily points of consensus among the authoring committee, were identified as interesting questions worthy of further consideration and perhaps useful directions for future networking research.

Where Should the Intelligence in the Network Reside?

The traditional Internet model pushes the intelligence to the edge, and calls for a simple data forwarding function in the core of the network. Does this continue to be the correct model? A number of ad hoc functions are appearing in the network, such as NAT boxes, firewalls, and content caches. There are devices that transform packets, and places where the network seems to operate as an overlay on itself (e.g., virtual private networks). Do these trends signal the need to rethink how function is located within the network? What aspects of modularity need to be emphasized in the design of functions: protocol layering, topological regions, or administrative regions? Is there a need for a more complex model for how applications should be assembled from components located in different parts of the network? There was a sense in discussions at the workshop that the Active Networks research may have explored some of these issues, but that the architectural questions remain unanswered.

Is the End-to-End Model the Right Conceptual Framework?

The end-to-end model implies that the center of the network is a transparent forwarding medium, and that the two ends have fully compatible functions that interwork with each other. From the perspective of most application developers and, in some sense, from the perspective of users, this model is not accurate. There is often a lot of practical complexity in a communication across the network, with caches, mirrors, intermediate servers, firewalls, and so on. From a user perspective, a better model of network communication might be a “limited horizon” model, in which the application or user can see the detail of what is happening locally but beyond that can interact with the network only at a very abstract level. Could such a view help clarify how the network actually works and how application designers should think about structure?

How Can Faults Be Better Isolated and Diagnosed?

When something breaks in the Internet, the Internet’s very decentralized structure makes it hard to figure out what went wrong and even harder to assign responsibility. Users seem to be expected to participate in fault isolation (many of them know how to run ping and trace-route but find it odd that they should be expected to do so). This perspective suggests that the Internet design might be deficient in that it does not pay proper attention to the way faults can be detected, isolated, and fixed, and that it puts this burden on the user rather than the network operator. The fact that this situation might arise from an instance of the end-to-end argument further suggests that the argument may be flawed.

Are Data a First-class Object Inside the Network?

The traditional model of the Internet is that it moves bytes between points of attachment but does not keep track of the identity of these bytes. From the perspective of the user, however, the namespace of data, with URLs as an example, is a part of the network. The users view the network as having a rather data-centric nature in practice, and they are surprised that the network community does not pay more attention to the naming, search, location, and management of data items. Should content-based addressing be a network research problem?

Does the Internet Have a Control Plane?

The original design of the Internet stresses the data-transport function but minimizes attention to management protocols, signaling, and control. A number of ad hoc mechanisms supply these functions, but they do not receive the research attention and architectural definition that the data movement functions do. This seems out of balance and may limit what can be achieved in the Internet today.

Abstractions of Topology and Performance

The Internet hides all details of topology and link-by-link measures of performance (for example, bandwidth, delay, congestion, and loss rates) beneath the IP layer. The simple assumption is that the application need not know about this, and if it does need such information, it can obtain it empirically (by trying to do something and observing the results). As more complicated applications such as content caches are built, the placement of these devices within the topology of the Internet matters. Could a network provide an abstract view of its performance that simplifies the design of such systems? How could the various performance parameters be abstracted in a useful way, and would more than one abstraction be required for different purposes? What, for example, would it take for the network to provide information to help answer the question of which cache copy is most appropriate for a given user?

Beyond Cooperative Congestion Control

There seem to be a great number of papers that improve the current Internet scheme for congestion control. However, this scheme, which depends on the end nodes doing the right thing, seems less and less suitable in general as one can trust the end nodes less and less, suggesting that one needs to explore different trade-offs of responsibility between the users and the network. While some research is being done that explores alternatives to cooperative congestion control, this may be an area that deserves greater emphasis.

Incorporating Economic Factors into Design

It was noted that many of the constraints on the current network are economic in nature, not technological. Research, to be relevant in the immediate future, needs to take the broader economic, social, and governmental environment into account. One attendee noted that in many situations, the way to get people to behave how you want is to construct economic incentives, not technical constraints. This could be a useful way of thinking about network design issues.

Finding Common Themes in User Requirements

Many user communities feel that they are expending energy trying to solve problems faced by many other groups of users in areas such as performance, reliability, and application design. These communities believe that their requirements are not unique but that the network research community does not seem to be trying to understand what these common requirements are and how to solve them. The tendency within the network community is to focus attention on issues at lower layers of the protocol stack even if significant, widespread problems would benefit from work at higher layers. One reason is that when networking researchers become heavily involved with application developers, the work becomes interdisciplinary in nature. Ongoing work in middleware development is an example of this research direction. Workshop participants noted that this sort of work is difficult and rarely rewarded in the traditional manner in the research community.

Using an Overlay Approach to Deploying Disruptive Technology

Along with specific disruptive ideas, workshop participants discussed the important implementation question of how one could deploy new technology using the existing network (to avoid having to build an entirely new network in order to try out ideas). The Internet is generally thought of as being composed of a core, which is operated by the small number of large ISPs known as the tier 1 providers; edges, which consist of smaller ISPs and networks operated by organizations; and endpoints, which consist of the millions of individual computers attached to the Internet.[1] The core is a difficult place to deploy disruptive technology, as the decision to deploy something new is up to the companies for which this infrastructure is the golden goose. Technical initiatives aimed at opening up the core might help, although ISP reluctance to do so would remain an issue. One of the successes of the Internet architecture is that the lack of intelligence within the core of the network makes it easy to introduce innovation at the edges. Following the end-to-end model, this has traditionally been done through the introduction of new software at the endpoints. However, the deployment of caching and other content distribution functionality suggest ways of introducing new functionality within the network near the edges. The existing core IP network could be used simply as a data transport service, and disruptive technology could be implemented as an overlay in machines that sit between the core and the edge-user computers.[2] This approach could allow new functionality to be deployed into a widespread user community without the cooperation of the major ISPs, with the likely sacrifice being primarily performance. Successful overlay functions might, if proven useful enough, be "pushed down" into the network infrastructure and made part of its core functionality.

---

[1] For a more detailed description of the Internet's design and structure, see Computer Science and Telecommunications Board, National Research Council. 2001. *The Internet's Coming of Age.* National Academy Press, Washington, D.C.

[2] The overlay approach has been used in several experimental efforts, including the Mbone (multicast), 6Bone (IPv6), and Abone (active networks).

was a low-cost computing platform for those who wanted to write programs. Like the Internet, the PC was immediately seen as valuable by a small user community that sustained its market. But it was not until the invention of the spreadsheet application that the popularity of PCs would rise rapidly. Similarly, in the networking world, the World Wide Web dramatically increased the popularity of the Internet, whose size went from roughly 200,000 computers in 1990 to 10 million in 1996, to a projected 100 million in 2001. Although the inventors of these applications were technically sophisticated, they were not part of the research community that invented the underlying disruptive technology. These examples illustrate an important caveat: It is hard to know up front what the "killer app" for new enabling technologies will be, and there are no straightforward mechanisms to identify and develop them. With any proposed technology innovation, one must gamble that it will be compelling enough to attract a community of early adopters; otherwise it will probably not succeed in the long run. This chicken-and-egg-type problem proved a significant challenge in the Active Networks program (as did failure to build a sufficiently large initial user community from which a killer application could arise).

There is a tension between experimentation on a smaller scale, where the environment is cleaner, research is more manageable, and the results more readily interpreted, and experimentation on a very large scale, where the complexity and messiness of the situation may make research difficult. A particular challenge in networking is that many of the toughest, most important problems that one would look to a disruptive networking technology to solve have to do with scaling, so it is often important to push things to as large a scale as possible. One-of-a-kind prototypes or even small testbed networks simply do not provide a realistic environment in which to explore whether a new networking idea really addresses scale challenges.

This suggests that if the research community is to attract enough people with new application ideas that need the disruptive technology, there will be a need for missionary work and/or compelling incentives for potential users. Natural candidates are those trying to do something important that is believed to be very hard to do on the Internet. One would be trustworthy voting for public elections; another, similar candidate would be developing a network that is robust and secure enough to permit organizations to use the public network for applications that they now feel comfortable running only on their own private intranets.

EXTERNAL DRIVERS

While focused on the disruptive ideas that could emerge from within the networking research community, workshop participants also noted the potential impact of external forces and suggested that networking (like any area of computer science) should watch neighboring fields and try to assess where disruptions might cause a sudden shift in current practice. The Internet is certainly subject to the possibility of disruptive events from a number of quarters, and many networking researchers track developments in related fields. Will network infrastructure technologies—such as high-speed fiber or wireless links—be such a disruption? Or will new applications, such as video distribution, prove a disruptive force? Workshop participants did not explore these forces in detail but suggested that an ongoing dialogue within the networking research community about their implications would be helpful.

5 Concluding Observations

A reviewer of a draft of this report observed that this proposed framework—measure, develop theory, prototype new ideas—looks a lot like Research 101. Why did this exploratory effort end up framing a research program along these lines? From the perspective of the outsiders, the insiders did not show that they had managed to execute the usual elements of a successful research program, so a back-to-basics message was fitting.

Both insiders and outsiders agreed that progress on each of these fronts would require effort, attention, and resources, and that each posed its own special challenges, and they also agreed that such investment could have significant payoffs. It is, to be sure, a daunting challenge, because the three dimensions of Internet ossification identified in Chapter 1 stifle the design of innovative alternatives. It is also possible that the workshop participants' enthusiasm about opportunities for change might be tempered by seeing new ideas realized. The outsiders seriously considered the words of Harry S. Truman: "I have found the best way to give advice to your children is to find out what they want and then advise them to do it."[1] If they had been Trumanesque, they would have applauded continuing research on higher Internet bandwidth, on quality of service protocols, and so forth. However, the outsiders expressed the view that the network research community should not devote all—or even the majority—of its time to fixing current Internet problems.

Instead, networking research should more aggressively seek to develop new ideas and approaches. A program that does this would be centered on the three M's—measurement of the Internet, modeling of the Internet, and making disruptive prototypes. These elements can be summarized as follows:

- *Measuring*—The Internet lacks the means to perform comprehensive measurement on activity in the network. Better information on the network would provide the basis for uncovering trends, as a baseline for understanding the implications of introducing new ideas into the network, and would help drive simulations that could be used for designing new architectures and protocols. This report challenges the research community to develop the means to capture a day in the life of the Internet to provide such information.
- *Modeling*—The community lacks an adequate theoretical basis for understanding many pressing problems such as network robustness and manageability. A more fundamental understanding of these important problems requires new theoretical foundations—ways of reasoning about these problems—that are rooted in realistic assumptions. Also, advances are needed if we are to successfully model the full range of behaviors displayed in real-life, large-scale networks.
- *Making disruptive prototypes*—To encourage thinking that is unconstrained by the current Internet, "Plan B" approaches should be pursued that begin with a clean slate and only later (if warranted) consider migration from current technology. A number of disruptive design ideas and an implementation strategy for testing them are described in Chapter 4.

When contemplating launching a new agenda along these lines it is also worth noting, as workshop participants did repeatedly during the course of the workshop, that in the past, the networking research community made attempts at broad new research initiatives, some of which failed at various levels and others of which succeeded beyond expectations. There is little systematic process for learning from these attempts, however. Failures are rarely documented, despite the potential value of documentation to the community. Failures can be embarrassing to the individuals concerned, and writing up failures is unlikely to be considered as productive as

[1] Edward R. Morrow television interview, "Person to Person," CBS, May 27, 1955.

writing up successes. Accordingly, it would be useful to convene an "autopsy workshop" from time to time, perhaps even annually, devoted to learning from past history, at both the individual solution level and the larger research areas level. Documenting negative results will help avoid wasted effort (the reinvention of faulty wheels). A postmortem on larger research areas will help to guide future research by increasing understanding of the specific successes and failures, as well as the underlying reasons for them (economics, for example, or politics). Within research areas, the successes and failures of interest include attempts at disruption to conventional networking, which so far have met with mixed success. Two initial candidates suggested by discussions at this workshop would be quality of service and active networking.

Appendixes

Appendix A

Biographies of Committee Members

David Patterson, *Chair*, is the E.H. and M.E. Pardee Chair of Computer Science at the University of California at Berkeley. He has taught computer architecture since joining the faculty in 1977 and has been chair of the Computer Science Division of the Electrical Engineering and Computer Science Department at Berkeley. He is well known for leading the design and implementation of RISC I, the first VLSI Reduced Instruction Set Computer, which became the foundation for the architecture currently used by Fujitsu, Sun Microsystems, and Xerox. He was also a leader of the Redundant Arrays of Inexpensive Disks (RAID) project, which led to high-performance storage systems from many companies, and the Network of Workstation (NOW) project, which led to cluster technology used by Internet companies such as Inktomi. He is a member of the National Academy of Engineering and a Fellow of the Institute of Electrical and Electronics Engineers (IEEE) and the Association for Computing Machinery (ACM). He served as chair of the Computing Research Association (CRA). His current research interests are in building novel microprocessors using Intelligent DRAM (IRAM) for use in portable multimedia devices and in creating Intelligent Storage (ISTORE) to provide available, maintainable, and evolvable servers for Internet services. He has consulted for many companies, including Digital, Hewlett-Packard, Intel, and Sun Microsystems, and he is the co-author of five books. Dr. Patterson served on the CSTB committees that produced *Computing the Future* and *Making IT Better*.

David D. Clark graduated from Swarthmore College in 1966 and received his Ph.D. from the Massachusetts Institute of Technology (MIT) in 1973. He has worked since then at the MIT Laboratory for Computer Science, where he is currently a senior research scientist in charge of the Advanced Network Architecture group. Dr. Clark's research interests include networks, network protocols, operating systems, distributed systems, and computer and communications security. After receiving his Ph.D., he worked on the early stages of the ARPANET and on the development of token ring local area network technology. Since the mid-1970s, Dr. Clark has been involved in the development of the Internet. From 1981 to 1989, he acted as chief protocol architect for this development and chaired the Internet Activities Board. His current research area is protocols and architectures for very large and very high speed networks. Specific activities include extensions to the Internet to support real-time traffic, explicit allocation of service, pricing, and new network technologies. In the security area, Dr. Clark participated in the early development of the multilevel secure Multics operating system. He developed an information security model that stresses integrity of data rather than disclosure control. Dr. Clark is a fellow of the ACM and the IEEE and a member of the National Academy of Engineering. He received the ACM SIGCOMM award, the IEEE award in international communications, and the IEEE Hamming Award for his work on the Internet. He is a consultant to a number of companies and serves on a number of technical advisory boards. Dr. Clark is currently the chair of the Computer Science and Telecommunications Board. He chaired the committee that produced the CSTB report *Computers at Risk: Safe Computing in the Information Age*. He also served on the committees that produced the CSTB reports *Toward a National Research Network*, *Realizing the Information Future: The Internet and Beyond*, and *The Unpredictable Certainty: Information Infrastructure Through 2000.*

Anna Karlin is a professor in the Computer Science and Engineering Department at the University of Washington. After receiving her Ph.D. in computer science at Stanford University in 1987, she did postdoctoral work at Princeton University. She then joined Digital Equipment Corporation's Systems Research Center in 1988 as a research scientist and worked there until she came to the University of Washington in 1994. Her research interests include competitive analysis of online algorithms, design and analysis of probabilistic algorithms, and the design and analysis of algorithms for problems in operating systems, architecture, and distributed systems. She is currently a member of the Computer Science and Telecommunications Board.

Jim Kurose received a B.A. degree in physics from Wesleyan University in 1978 and his M.S. and Ph.D. degrees in computer science from Columbia University in 1980 and 1984, respectively. He is currently professor and chair of the Department of Computer Science at the University of Massachusetts, where he is also codirector of the Networking Research Laboratory and the Multimedia Systems Laboratory. Professor Kurose was a visiting scientist at IBM Research during the 1990-1991 academic year and at INRIA and EURECOM, both in Sophia Antipolis, France, during the 1997-1998 academic year. His research interests include real-time and multimedia communication, network and operating system support for servers, and modeling and performance evaluation. Dr. Kurose is the past editor in chief of the *IEEE Transactions on Communications* and of the *IEEE/ACM Transactions on Networking*. He has been active in the program committees for IEEE Infocom, ACM SIGCOMM, and ACM SIGMETRICS conferences for a number of years. He is the six-time recipient of the Outstanding Teacher Award from the National Technological University (NTU), the recipient of the Outstanding Teacher Award from the College of Natural Science and Mathematics at the University of Massachusetts, and the recipient of the 1996 Outstanding Teaching Award of the Northeast Association of Graduate Schools. He has been the recipient of a General Electric fellowship, an IBM faculty development award, and a Lilly teaching fellowship. He is a fellow of the IEEE and a member of ACM, Phi Beta Kappa, Eta Kappa Nu, and Sigma Xi. With Keith Ross, he is the coauthor of the textbook *Computer Networking, a Top Down Approach Featuring the Internet*, published by Addison-Wesley Longman in 2000.

Edward D. Lazowska is professor and chair of the Department of Computer Science and Engineering at the University of Washington. Lazowska received his B.A. from Brown University in 1972 and his Ph.D. from the University of Toronto in 1977. He has been at the University of Washington since that time. His research concerns the design and analysis of distributed and parallel computer systems. Dr. Lazowska is a member of the NSF Directorate for Computer and Information Science and Engineering Advisory Committee, chair of the Computing Research Association, a member of DARPA ISAT, and a member of the Technical Advisory Board for Microsoft Research. Dr. Lazowska is currently a member of the Computer Science and Telecommunications Board. He served on the CSTB committee that produced the report *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure*. He is a fellow of the ACM and of the IEEE.

David Liddle is a general partner in the firm U.S. Venture Partners (USVP). It is a leading Silicon Valley venture capital firm that specializes in building companies from an early stage in digital communications/networking, e-commerce, semiconductors, technical software, and e-health. He retired in December 1999 after 8 years as CEO of Interval Research Corporation. During and after his education (B.S., E.E., University of Michigan; Ph.D., computer science, University of Toledo, Ohio), Dr. Liddle spent his professional career developing technologies for interaction and communication between people and computers, in activities spanning research, development, management, and entrepreneurship. First, he spent 10 years at the Xerox Palo Alto Research Center and the Xerox Information Products Group where he was responsible for the first commercial implementation of the Graphical User Interface and local area networking. He then founded Metaphor Computer Systems, whose technology was adopted by IBM and which was ultimately acquired by IBM in 1991. In 1992, Dr. Liddle cofounded Interval Research with

Paul Allen. Since 1996, the company formed six new companies and several joint ventures based on the research conducted at Interval. Dr. Liddle is a consulting professor of computer science at Stanford University. He has served as a director at Sybase, Broderbund Software, Metricom, Starwave, and Ticketmaster. He was honored as a distinguished alumnus from the University of Michigan and is a member of the national advisory committee at the College of Engineering of that university. He is also a member of the advisory committee of the school of Engineering at Stanford University. He has been elected a Senior Fellow of the Royal College of Art for his contributions to human-computer interaction. Dr. Liddle currently serves on the Computer Science and Telecommunications Board.

Derek McAuley is director of Marconi Labs, Cambridge, England. He obtained his B.A. in mathematics from the University of Cambridge in 1982 and his Ph.D. addressing issues in interconnecting heterogeneous ATM networks in 1989. After 5 years at the University of Cambridge Computer Laboratory as a lecturer, he moved in 1995 to the University of Glasgow as chair of the Department of Computing Science. He returned to Cambridge in July 1997 to help found the Microsoft Research facility in Cambridge, a once-in-a-lifetime opportunity. In January 2001 he was presented with a second once-in-a-lifetime opportunity as founder of Marconi Labs. His research interests include networking, distributed systems, and operating systems.

Vern Paxson is a senior scientist with the AT&T Center for Internet Research at the International Computer Science Institute in Berkeley and a staff scientist at the Lawrence Berkeley National Laboratory. His research focuses on Internet measurement and network intrusion detection. He serves on the editorial board of *IEEE/ACM Transactions on Networking* and has been active in the Internet Engineering Task Force (IETF), chairing working groups on performance metrics, TCP implementation, and end-point congestion management, as well as serving on the Internet Engineering Steering Group (IESG) as an area director for transport. He has participated in numerous program committees, including SIGCOMM, USENIX, USENIX Security, and RAID; co-chairs the 2002 SIGCOMM networking conference; and is a member of the steering committee for the SIGCOMM Internet Measurement Workshop, 2001. He received his M.S. and Ph.D. degrees from the University of California, Berkeley.

Stefan Savage is an assistant professor in the Department of Computer Science and Engineering at the University of California, San Diego. Prior to joining the faculty at UCSD in 2001, his doctoral work was at the University of Washington. Dr. Savage's current research interests focus on wide-area networking, reliability, and security. Previously he has worked broadly in the field of experimental computer systems, including research on real-time scheduling, operating system construction, disk array design, concurrency control, and performance analysis.

Ellen W. Zegura received the B.S. degrees in computer science and electrical engineering (1987), the M.S. degree in computer science (1990), and the D.Sc. degree in computer science (1993), all from Washington University, St. Louis, Missouri. She has been on the faculty at the College of Computing, Georgia Institute of Technology, since 1993. She is currently an associate professor and assistant dean of facilities planning. Her research interests include active networking, server selection, anycast and multicast routing, and modeling large-scale internetworks. Her work in topology modeling is widely recognized as providing the best current models to use in simulation-based studies of Internet problems. A software package implementing these models is in frequent use by other research groups and has been incorporated into one of the leading public domain software tools for Internet simulations. Her work in active (or programmable) networking is among the earliest in this relatively new field. Her focus on applications of active networking—and rigorous comparison of active solutions to traditional solutions—distinguishes her work from that of the many other groups who have focused on enabling technologies. Dr. Zegura is currently leading a DARPA working group on composable services and applications and editing a document intended to serve as a foundation for research and development efforts in the DARPA community. Her work in the server selection area focuses on techniques for the effective use of replicas of servers distributed over the wide area. She has developed an architecture to support server selection on a wide range of performance and policy criteria, thus supporting diversity in the service or in client preferences. She has also explored the implications of server selection in the presence of emerging technologies, including multicast-capable servers and QoS-capable networks. Dr. Zegura has served on a variety of NSF award selection panels and numerous conference program committees, as well as three conference executive committees (as Publicity Chair for ICNP'97 Student Travel Grant, Chair for Sigcomm'97, and Tutorials Chair for Sigcomm'99). She was co-chair of the 2nd Workshop on Internet Server Performance, held in conjunction with ACM Sigmetrics. She also served on the selection committee for the CRA Distributed Mentor program. Dr. Zegura's work has been funded by DARPA, the NSF, the CRA, and NASA. She has also received industrial grants from Hitachi Telecom, USA, and internal university grants from the Packaging Research Center and the Broadband Telecommunications Center.

Appendix B

List of Workshop Participants

CHRISTINE BORGMAN, University of California, Los Angeles
DAVID D. CLARK,* Massachusetts Institute of Technology
DAVID CULLER, University of California, Berkeley
JEFF DOZIER, University of California, Santa Barbara
ANNA KARLIN, University of Washington
JIM KUROSE,* University of Massachusetts, Amherst
DEREK MCAULEY,* Marconi Research
JOHN OUSTERHOUT, Interwoven
DAVID PATTERSON, University of California, Berkeley
VERN PAXSON,* AT&T Center for Internet Research, International Computer Science Institute
SATISH RAO, University of California, Berkeley
STEFAN SAVAGE,* University of California, San Diego
HAL VARIAN, University of California, Berkeley
ELLEN ZEGURA,* Georgia Institute of Technology

Additional input, in the form of answers to the three questions posed in Box P.1, was provided by Andy Bechtolsheim* (Cisco), Eric Brewer (University of California at Berkeley), Stephanie Forrest (University of New Mexico), Ed Lazowska (University of Washington), and Tom Leighton (MIT).

---

*Indicates a networking insider.