Report of the National Science Foundation Workshop on Fundamental Research in Networking

April 24–25, 2003
Airlie House, Virginia

REPORT TO THE NATIONAL SCIENCE FOUNDATION
DIRECTORATE FOR COMPUTER AND INFORMATION
SCIENCE AND ENGINEERING (CISE)

Report of the National Science Foundation Workshop on Fundamental Research in Networking

April 24–25, 2003
Airlie House, Virginia

An online version of this report is available at: http://www.cs.virginia.edu/~jorg/workshop1

The NSF Workshop on Fundamental Research in Networking was supported by the NSF under grant ANI-0328250.

Authors and Contributors

The following are the attendees of the NSF Workshop on Fundamental Research in Networking in Airlie, Virginia in April 24–25, 2003, who contributed to the writing of this report.

Name	Affiliation
Mostafa Ammar	Georgia Institute of Technology
Gustavo de Veciana	University of Texas, Austin
Magda El Zarki	University of California, Irvine
Anthony Ephremides	University of Maryland, College Park
Anthony Joseph	University of California, Berkeley
George Kesidis	Pennsylvania State University
Edward Knightly	Rice University
Jim Kurose	University of Massachusetts, Amherst
Monica Lam	Stanford University
Jörg Liebeherr	University of Virginia
Ravi Mazumdar	Purdue University
Biswanath Mukherjee	University of California, Davis
Klara Nahrstedt	University of Illinois, Urbana-Champaign
Brian Noble	University of Michigan
Larry Peterson	Princeton University
Martin Reisslein	Arizona State University
Henning Schulzrinne	Columbia University
Ness Shroff	Purdue University
Suresh Singh	Portland State University
Arun Somani	Iowa State University
R. Srikant	University of Illinois, Urbana-Champaign
Ion Stoica	University of California, Berkeley
Don Towsley	University of Massachusetts, Amherst
Nitin Vaidya	University of Illinois, Urbana-Champaign
Malathi Veeraraghavan	University of Virginia
Jean Walrand	University of California, Berkeley
Lixia Zhang	University of California, Los Angeles

Members of the organizing committee of the workshop were:

- Mostafa Ammar, Georgia Institute of Technology
- Jörg Liebeherr, University of Virginia *(Co-organizer)*
- Biswanath Mukherjee, University of California, Davis
- Ness Shroff, Purdue University *(Co-organizer)*
- Don Towsley, University of Massachusetts
- Nitin Vaidya, University of Illinois, Urbana-Champaign
- Hui Zhang, Carnegie Mellon University and Turin Networks

Acknowledgements

We wish to thank the NSF Directorate for Computer and Information Science and Engineering (CISE) for the support of the workshop. We thank the Division of Advanced Networking Infrastructure and Research (ANIR) and the program director of the Networking Research Program, Admela Jukan, for support and encouragement.

Table of Contents

Executive Summary

The workshop on Fundamental Research in Networking, supported by the NSF CISE ANIR program, was held in Arlie House, Virginia, on April 24-25, 2003. This was the fourth NSF sponsored workshop that addressed fundamental research in networking since the establishment of a networking research program at the NSF in 1987. Since the last of these workshops, held in 1994, the networking research community and its research agenda have significantly evolved, thus, creating a need to re-evaluate the needs of fundamental research in networking. The workshop consisted of 27 expert researchers from highly diverse backgrounds in networking. The goal of this workshop report is to predict, assess, and help shape the future research agenda in networking, and articulate the vision of the workshop participants of the critical role that fundamental research will play in the future of networking.

This Executive Summary contains the findings and recommendations from the workshop. The detailed context of the findings can be found in the main body of this report.

SUMMARY OF FINDINGS AND WORKSHOP RECOMMENDATIONS

The following summarizes the most important findings of this workshop report.[1] Based on these findings and extensive discussions, the workshop participants have made recommendations to the National Science Foundation.

Our first set of findings is a result of reflections on the past, present and future of networking research, and a discussion that attempts to specify success criteria of networking research.

FINDING: Growth of the networking community. The range of topics of networking research is continuously increasing. At this time, there is virtually no researcher in computer science or electrical engineering, who does not, directly or indirectly, include aspects of networking in his/her research agenda. The central role of networking research in computer science and engineering has led to a significant growth of the research community.

RECOMMENDATION: To sustain the enormous growth in networking research, continuous development and improvement of NSF's ability to foster visionary and far-reaching research in networking is needed.

FINDING: Impact of fundamental research. Fundamental research in networking has been and continues to be the driving force behind technological innovation in many emerging areas of communication networks (e.g., secure and robust communication, embedded networks, peer-to-peer networks, networks supporting high-quality services, etc.).

1 *The findings are highlighted in the text of the report, but may appear in a different order than listed here.*

FINDING: The state of networking research today.

1. The current body of knowledge of networking is significantly larger than the technology found in the Internet today.
2. The true promise and full potential of wide-area, ubiquitous networking requires significant research efforts well into the future.
3. In many respects, networking is still in its infancy with many exciting paradigms and possibilities for information collection, dissemination, and communication still ahead of us. Much of our understanding about networks will need to undergo substantial innovation and paradigm shifts.

RECOMMENDATION: Encourage research that seeks radical innovation and paradigm shifts. Allow multiple future research agendas, in which many technologies, paradigms and approaches are considered simultaneously and encouraged for their intellectual merit. Less emphasis should be placed on immediate deployment feasibility and short-term economic viability.

RECOMMENDATION: Encourage multi-disciplinary aspects of fundamental research in networking. New networking paradigms and applications bring with them many challenges that require fundamental research that brings together expertise from different disciplines, and this should be encouraged.

FINDING: Success criteria and models of success. Networking research has well-established criteria to measure success of outcomes of research agenda, including: (1) adoption in products/large user community, (2) solution to real-world problems, (3) development of new design principles, (4) creation and enabling of new modes of communication, (5) creation of new abstractions, models and tools with long-term value, (6) impact in the development and education of human resources. There are several models of success, each incorporating several success criteria.

RECOMMENDATION: Support a research agenda that is successful with respect to a broad range of success criteria. Evaluate research outcomes using a diverse set of success models.

FINDING: Success of the Internet carries the risk of the "network innovator's dilemma." The success of the Internet carries the risk of complacency about the need for true innovation and outside-the-box thinking. The networking research community must avoid the "network innovator's dilemma," where it gets too involved in improving the existing Internet technology, and fails to observe novel and disruptive technologies that take radical new shape.

RECOMMENDATION: Support a research agenda that looks beyond the success of the Internet and emphasizes radical innovation and outside-the-box thinking.

FINDING: Need for reproducibility of research outcomes. The maturing of networking research and its continued success require bolstering the scientific underpinnings in the area of benchmarking and reproducibility of results.

RECOMMENDATION: Foster efforts that increase reproducibility of networking research outcomes. Encourage the development of reference models or benchmarks to ease the burden of reproducible experiments on complex systems.

Our second set of findings underlines the essential role of basic research in networking, and points to the current situation of basic research funding.

FINDING: Networking has evolved as a scientific research field with its own identity. Basic networking research involves science from many areas and develops new scientific methods to fit the networking context, resulting in a field of networking with its own identity. While there is no particular networking science, the methodologies developed are a result of addressing specific networking problems and may not necessarily be of direct interest to the individual constituent science disciplines.

FINDING: Significant role of basic networking research. As the Internet increasingly becomes an integral part of ubiquitous and critical services, understanding the behavior of the Internet and learning how to improve its properties is a fundamental endeavor that will have an enormously beneficial impact on all facets of our economic and social life.

RECOMMENDATION: Promote the key role of basic networking research in the development of the Internet. The workshop participants observed that the role of basic networking research to the development of the Internet has not been articulated to the fullest extent, and should be promoted to a stronger degree.

FINDING: Basic networking research has played a key role in the development of the Internet. Developments that have been influenced by basic networking research include random medium access protocols (such as Wi-Fi and Ethernet), TCP congestion control, flow control algorithms, scheduling algorithms, distributed routing protocols, and many more.

FINDING: Need for increased role of government in support of basic networking research. The disheartening trends of shutting down or downsizing important industrial networking research labs means that the brunt of the basic research costs will have to be borne by government agencies such as the NSF.

RECOMMENDATION: Increase support for basic research in networking. Emerging applications and directions require fundamental understanding on how to design and control networks on an even grander scale than before. The changing landscape of the funding structure for basic research in networking has increased the role of the NSF in supporting such research.

The same trend requires universities to re-establish themselves as the main generators of know-how in networking research.

The workshop report emphasizes that applications are the raison d'etre of networks. The third set of findings and recommendations addresses the crucial role of applications in networking research.

FINDING: Important role of applications. Many of the important principles underlying today's networks have been strongly influenced by application-level considerations.

FINDING: Application requirements and challenges of applications will continue to influence networking far into the future. An application push for developing new networking mechanisms, protocols, and architectures will continue to influence networking far into the future.

FINDING: Relevance of killer application(s). Challenges in the future may not result from a single "killer application" but rather from the need to simultaneously support a large and diverse mix of applications.

RECOMMENDATION: Encourage closer ties between applications research and networking research. Often, the future requirements of networking are only apparent when the requirements of truly visionary applications are taken into consideration.

The last set of findings and recommendations identifies grand challenge problems in networking research that will play an important role in the next decade of networking research.

RECOMMENDATION: Focus research agenda on application-induced challenges. The workshop has identified a set of user-focused challenges relating to the user experience (robustness, invisibility, ease of configuration, exploiting storage and processing capabilities, power consumption), as well as network-centric challenges (heterogeneity and scale, manageability, evolvability).

FINDING: Meta-grand challenge of networking research. The current Internet architecture suffers from several shortcomings that will impede the development of the next generation of applications. The meta-grand challenge facing networking research today is to develop new network theories, architectures, and methodologies that will facilitate the development and deployment of the next generation of services and applications.

FINDING: Grand challenges of networking research. The meta-grand challenge translates into a number of networking research challenges, which include the following areas:

1. **Need for an Information Theory for networks.** There is no unified, basic theory yet that extends information theory to networks. Such a theory is especially important in the case of wireless networks. It is imperative to understand the fundamental throughput and delay performance limits of wireless networks, with the goal of designing resource allocation algorithms that allow us to reach these performance limits.
2. **Overlay networks.** Overlays designed to support small-scale applications can utilize algorithms that would not scale to global size, but raises numerous challenging problems: How to share network resources between multiple overlay networks? How to build economic mechanisms that arbitrate among resource providers, overlay services and users? How to preserve the robustness of the shared infrastructure in the face of misbehavior or maliciousness?
3. **Resilient networking.** As the Internet grows in scale, it faces an increasingly adverse environment that the original design did not envision. Faults have become the norm rather than the exception, due to component failures, human operational errors, spreads of software viruses, and malicious attacks. It is important to develop a basic understanding of the rules and principles in network protocol design that facilitate the realization of resiliency in large-scale networks to these faults.
4. **Providing market incentives.** One of the impediments to the deployment of new services on the Internet is the lack of market incentives to improve network services and applications and to use them efficiently. Recent history has demonstrated that the bottleneck to improving the global network may be economics instead of technology. Thus it is important to develop a consistent economic theory for networks that provides mechanisms to enable the creation of efficient markets among users and providers of content, services, and applications. Such mechanisms require new algorithms and protocols to collect necessary measurements, calculate prices, charge users, and redistribute revenues.
5. **Sensorized universe.** In the future, remote sensing will become an integral part of our lives as we strive to use sensor technologies to monitor our surrounding space for a variety of applications such as security, health, education, comfort, environment, traffic, safety, and more. The design of a sensor network introduces daunting challenges in all aspects of the protocol stack of a sensor node.
6. **Virtual networks.** New users and applications introduce requirement for services with a higher level of abstraction that deal with user mobility, failures of servers and service providers, and automatic discovery of resources and routes, without a need for manual configuration.

RECOMMENDATION: Increase level of support and funding on grand challenge problems. Provide increased levels of funding on the meta-grand challenge of networking research and on the identified grand challenge problems.

1. Introduction

Since the establishment of a networking research program at the NSF in 1987, three NSF sponsored workshops have addressed fundamental research in networking. The networking research community and its research agenda have significantly evolved since the last of these workshops, held in 1994, and it is thus important that the community re-evaluate the needs of fundamental research in networking.

Two trends have strongly motivated the organization of a workshop on fundamental research in networking: the significant increase in the scope of networking research and the dramatic growth of the networking research community. A need for this workshop also arose from a growing sense that, support for fundamental networking research at NSF has not kept pace with the broadening of the scope of networking research, and the growth of the community. The problems faced by the network community are exacerbated, as industry and mission-oriented government agencies have broadly reduced, if not entirely eliminated, basic networking research from their agendas. Thus, there is a pressing need to better understand how to improve the ability of NSF to foster visionary and far-reaching research in networking.

The goal of this workshop report is to predict, assess, and help shape the future research agenda in networking. We hope to generate far reaching future initiatives that emphasize the importance of *fundamental research in networking*. The workshop report tries to identify major issues affecting fundamental research in networking for the future. There are a number of related workshops that have recently defined research problems and agendas in computer science and engineering [CRA 2002], computer networking research [CSTB 2001, Testbeds 2002] and networking education [Education 2002].We will describe the various salient attributes of fundamental research in networking and the enormous impact that such research has had, not only to the field of networking, but to commerce and society at large. This report will articulate the vision of the workshop participants, composed of experts in diverse areas of networking, of the critical role that fundamental research will play in the future of networking.

The term *fundamental research* or *core research* in networking refers to far-reaching, risk-taking research that is concerned with the development of new network theories, architectures, and methodologies that will facilitate the development and deployment of the next generation of services and applications. Fundamental research in networking includes both theoretical and systems-oriented networking research. It brings together methodologies and techniques from various disciplines of electrical engineering, computer and information sciences, and mathematics to address specific issues of relevance to networking problems. Fundamental research in networking spans a diverse array of topics from theory to practice.

This workshop was held in Arlie House Virginia on April 24-25, 2003. The participants were 27 networking researchers from highly diverse backgrounds. An organizing committee of six networking researchers set the agenda of the workshop. On the first day, the workshop consisted of two panel discussions, followed by four breakout sessions. Each breakout session was assigned a discussion leader and a scribe. The participants returned after the breakout sessions to report to the entire group the outcomes of their discussions and to receive feedback. The next morning the participants met again to plan the preparation of the report. All workshop participants were given writing assignments to be completed within a few weeks of the workshop. The details of the workshop program are available at http://www.cs.virginia.edu/~jorg/workshop1/agenda.html. Each group leader then assimilated the feedback from the various participants and sent it to the workshop organizers. The organization of the remainder of this report reflects the agenda of the workshop and the organization of the breakout sessions.

Section 2 is the result of an effort to describe the growth and evolution of the networking research community over the past several years, to describe the current state of networking research, and its outlook for the future. The section highlights past achievements in networking and provides models to benchmark future successes.

Section 3 discusses the significant role of basic networking research.

Section 4 highlights the importance of applications for networking research and presents application-induced challenges to networking research.

Section 5 identifies Grand Challenges that face networking researchers today—the solutions to which will have far reaching scientific, technological, and commercial impact.

While the workshop participants are aware of the dangers of making precise predictions about the future of networking, some consensus was reached, and a representative sample of future network applications and research challenges are included in the report.

2. Models of Success

REFLECTIONS ON PAST, PRESENT AND FUTURE OF NETWORKING RESEARCH

In addition to defining a technical research agenda, it is important for the networking research community to engage in a discussion of how the outcomes of such an agenda for future research should be judged, and to articulate criteria to evaluate the success of a networking research program. To that end the workshop participants have explored sets of criteria by which one can evaluate the success of various research threads in networking. The conclusions are: (1) that there are well-established criteria by which the networking research community judges its research and its product, and (2) that there are several models of success, each incorporating several success criteria—one size does not fit all.

2.1 SUCCESS CRITERIA

Networking research distinguishes itself from many other research disciplines in that it has well-established criteria by which the networking research community judges its research and its product. In the following, we briefly list some of these success criteria. It should be noted that there is no order implied in this list, and that research projects with the most impact succeed according to several of these criteria, albeit to differing degrees.

Adoption in products/large user community. The networking industry has been on a fast track in the last three decades. Research results have found their way into commercial products relatively fast. Hence, the adoption of research ideas into commercial products has become an important success criterion in the last 10 to 20 years. This is perhaps quite different from other disciplines where commercial viability of research output is often not easily discernible in the immediate term. We observe here that often deployment in products does not necessarily imply wide spread use of a product. For example, while multicast routing protocols can be found in most commercial router products today, the widespread use of multicast services has lagged behind.

Solution to real-world problems. Networking research is often motivated by real-world problems for which an immediate solution is needed. While most research efforts will typically target fundamental results with a broad scope of applicability, often results are directly applicable to existing and pressing real-world problems. This derives from the phenomenal growth of networking and its accelerated penetration of our daily lives. For example, the development of the naming architecture in the Internet (the Domain Name System) had a fundamental component in which it addressed general addressing, routing and naming architectures for large-scale networks. At the same time, this work provided a solution to a very real problem faced in the early days of the Internet.

Development of new design principles. Often networking research leads to the development of design principles that have broad applicability in a wide range of networking issues. For example, the research effort targeting the provisioning of quality of service (QoS) in wide-area networks has led to significant understanding of the various components of a QoS architecture, that is, a network architecture that distinguishes different types of traffic and gives a differential set of services guarantees to each type of traffic. The implications of such formulations of design principles go well-beyond applications within the Internet specifically and will continue to influence future research efforts.

Creating and enabling new modes of communication. The ubiquity and success of the telephone network popularized a communication paradigm that ultimately was taken for granted as a “natural” approach. With the advent of digital data communication and the development of wide-area data networks, there was increased interest in novel communication paradigms. Examples include the development of multipoint communication, anycast communication and location-aware services. Often an idea or a concept for a new communication paradigm is then followed by significant research activity that considers various aspects of the new paradigm ranging from fundamental to practical.

Creation of new abstractions, models, and tools with long-term value. As networking research matures, we are increasingly witnessing the development of results in what can be described as “networking science.” These results are characterized by their fundamental nature, broad applicability, and long-term value. For example, research into quality of service produced results in the area of scheduling, and in the development of completely new queueing paradigms including a network calculus, that provides methods to reason about delay and backlog in a network. As another example, research into mobile and ad-hoc networks is increasingly yielding a new fundamental understanding of network capacity.

Impact in the development and education of human resources. Networking research and its most visible output, the Internet, permeate all parts of society today. Because of this, there has been increasing demand for educational efforts in computer networking. This includes the education for networking scientists and engineers, and, since knowledge about networking technology has become an important part of most scientific and engineering education, the education of the broader community of future scientists and engineers. One of the important outcomes of networking research supported by NSF is its direct support for the training of the required human resources in this area and its direct and indirect support of curricular efforts in support of networking education.

2.2 CASE STUDIES OF NETWORKING RESEARCH SUCCESS

We now consider four long-term research threads that have occupied many researchers for at least the last 20 years. We provide a number of representative studies that illustrate the large variety of models of success for networking research outcomes.

2.2.1 Case Study 1: Multiple Access

Multiple access mechanisms are used to coordinate communication on shared channels, and are needed in most local area and metropolitan area networks, access networks, wireless networks, and satellite networks. Networking research on multiple access has a rich history, spanning over several decades, and has led to a large variety of MAC (media access control) protocol designs, which have come into widespread use. Several protocols have been standardized, and have led to widespread adoption, with Wi-Fi and Ethernet perhaps the biggest success stories. Research activity on multiple access mechanisms and MAC protocols has not only provided insights into fundamental issues affecting medium access, but has also led to practical protocols that have benefited the user community at large. New physical layer technologies (e.g., ultra-wide band communication) bring new challenges for multiple access. Practical solutions to these multiple access challenges will help launch the next wave of popular adoptions of new technologies.

2.2.2 Case Study 2: Quality of Service

Quality of service has represented a major research effort in the 1990s with contributions spanning theory, architecture, and protocol design. Broadly, QoS refers to traffic control mechanisms that seek to either differentiate performance based on application or network-operator requirements, or provide predictable or guaranteed performance to applications, sessions, or traffic aggregates.

QoS research has resulted in numerous innovations to packet-network technology. New packet scheduling algorithms were developed for differentiated and fair service among flows and aggregates, admission control algorithms were developed to ensure that offered loads are restricted to a level such that service performance objectives are satisfied. QoS mechanisms have permeated all components of a networked system, including routing protocols, web servers, and MAC protocols.

Perhaps the most widely studied scenario for QoS was that embodied by architectures such as IntServ and ATM's VBR service which sought to provide end-to-end per-flow guarantees on throughput and delay. If one were to consider adoption in products with a large user base as the measure of success, this effort has not succeeded in that hosts cannot today request bandwidth-on-demand from the routers and switches that compose its end-to-end paths. However, by other criteria, these efforts were highly successful. For example, guaranteed bandwidth over a packet-switched network is indeed a new mode of communication that was not previously possible. Thus, the question shifts from "can it be done?" to "should it be done?" given other factors such as economic viability. The value of QoS research often lies in that it explores the limits of what can be achieved. It is perhaps no surprise that today's commercial system are far from having achieved these limits, yet indeed feature a myriad of QoS features such as class-based scheduling and traffic shaping, voice-over-IP services, empirical load-based capacity planning vs. per-flow admission control, etc.

There is indeed a sense that end-to-end per-flow QoS has been well explored. However, it is increasingly clear that the lack of predictability in applications' performance remains a major hindrance in deployment of new network services and applications. New dimensions of QoS such as resilience and consideration of economic factors are clearly as critical as development of protocols and architectures. Thus, looking to the past, it is clear that significant progress has been made, yet looking forward, it is equally clear that significant challenges lie ahead. Perhaps, realization of a high-performance, predictable, resilient network will require a radical catalyst: a new network architecture, a new industry structure, or emergence of a new killer application.

2.2.3 Case Study 3: Optical Networks

The objective of optical networking is to exploit the capabilities of optical devices and technologies to create high-bandwidth transport facilities, which, in turn, may lead to the realization of new and unanticipated services. Today's state-of-the-art switching technology employed in optical networks is electronic circuit switching with sub-wavelength switching granularity, although other forms of switching deserve continued attention.

The development of the "lightpath" concept 15 years ago has resulted in a new branch of networking research in the 1990s, which is exclusively devoted to the design principles of optical networking. The optical networking research community has made tremendous contributions towards enhancing our knowledge and understanding on how to control resources in networks using optical technologies. Some of the important problems include the discovery of resources in the network and the rapid propagation of state information in the network, the method for setting up virtual bandwidth pipes between devices at the edge of the optical network, and the setup and maintenance of connection state tables. Starting in the mid-1990s, several startup companies were born to build commercial products around the developed concepts, eventually leading to (then) multi-billion-dollar publicly traded companies. Founders, CEOs, CTOs, and VPs, in several of these startups had before been active in the area of optical networking research.

Today, the field of optical networking is still in its infancy, and many significant open research problems need to be explored which can facilitate new services and applications, e.g., bandwidth on demand, pervasive high-speed access, new network architectures and algorithms to combat physical-layer optical impairments, optical packet/burst switching, etc. Optical networking technology will play undoubtedly a central role in future high-speed access networks to solve the "last mile problem," which is widely believed to be the major networking bottleneck in the realization of a low-cost next generation information infrastructure.

2.2.4 Case Study 4: Congestion Control

Congestion occurs in networks when offered traffic exceeds the capability of a link or router that is being traversed. Congestion tends to spread rapidly causing large segments of the network to shut down. The problem of congestion in the Internet first became evident in the mid 1980s when data networks transitioned from experimental, homogeneous systems to large heterogeneous systems with links and nodes of differing capabilities. Research on congestion control led to the simultaneous development of (1) new algorithms and techniques to deal with congestion and (2) a new network theory which attempted to understand this complex phenomena in order to better control it.

The investigation of the principles of AIMD (additive increase multiplicative decrease) fairness lead to the TCP congestion control mechanisms for the Internet. More recently, active queue management (AQM) algorithms, such as Random Early Detect (RED) at routers, as well as feedback schemes, such as Explicit Congestion Notification (ECN), enable end systems to respond to congestion before it occurs.

During the past decade, a simultaneous development in network theory has occurred which attempts to understand and model congestion in the Internet. Specifically, researchers have applied control theoretic formulations, fluid mechanics approaches, as well as more traditional queueing theoretic models to study the evolution of traffic behavior in the network as a first step in understanding and predicting congestion. Since congestion is so fundamental to the operation of networks, manufacturers and operators of the Internet have adopted many of the congestion control techniques developed by researchers. The new approaches being developed to model the Internet have fundamental scientific value today and will result in a more reliable Internet tomorrow as this new knowledge is incorporated into products and operating practices.

2.3 THE STATE OF NETWORKING TODAY

Much has been accomplished by networking research in the last 30 years. Networking research has covered architectural, technical, and engineering concerns in the design of such networks and has led to the instantiation of some of these ideas in the global infrastructure we know as the Internet. Networking research has witnessed tremendous success over the last several decades in large part due to the significant emphasis that NSF has placed on funding fundamental research. Fundamental research has been and continues to be the driving force behind technological inno-

FINDING: Impact of fundamental research
Fundamental research has been and continues to be the driving force behind technological innovation in the many emerging areas of communication networks (e.g., secure and robust communication, embedded networks, peer-to-peer networks, networks supporting high-quality services, etc.).

vation in the many emerging areas of communication networks.

Coinciding with, or due to, the success of networking research, the range of topics of networking research is continuously increasing, and now includes, among many other new topics, wireless sensor networks, embedded networks, peer-to-peer networks, optical networks, network security, Internet measurements, control theory applications in networking, and cross-layer optimizations. At this time, there is virtually no researcher in computer science or electrical engineering, who does not, directly or indirectly, include aspects of networking in his/her research agenda. The central role of networking research in computer science and engineering has led to a significant growth of the research community.

FINDING: Growth of the networking community
The range of topics of networking research is continuously increasing. At this time, there is virtually no researcher in computer science or electrical engineering, who does not, directly or indirectly, include aspects of networking in his/her research agenda. The central role of networking research in computer science and engineering has led to a significant growth of the research community.

Comparing the state-of-the-art of networking research with the deployed networking topologies, especially in the Internet, two aspects need to be emphasized:

1. The current body of knowledge of networking is significantly larger than the technology found in the Internet today.
2. Networking research has lead to an exploration of many of the basic dimensions of wide-area, ubiquitous networking but the true promise and full potential of such a technology requires significant research efforts well into the future. In many respects, networking is still in its infancy with many exciting paradigms and possibilities for information collection, dissemination, and communication still awaiting discovery and exploration.

Much of our understanding about networks will need to undergo substantial innovation and paradigm shifts, even though much of our knowledge today will continue to be relevant. Our children and grandchildren will recognize the Internet today as a "network" but a primitive one.

FINDING: The State of Networking Today

1. The current body of knowledge of networking is significantly larger than the technology found in the Internet today.
2. The true promise and full potential of wide-area, ubiquitous networking requires significant research efforts well into the future.
3. In many respects, networking is still in its infancy with many exciting paradigms and possibilities for information collection, dissemination, and communication still ahead of us. Much of our understanding about networks will need to undergo substantial innovation and paradigm shifts.

2.4 NETWORKING RESEARCH MOVES INTO THE FUTURE

While we feel that the long series of networking research successes according to the depicted models of success will continue to hold as networking research moves into the future, two additional criteria deserve highlighting, that will become especially important as the networking research area matures as a science and engineering discipline:

1. Appropriate understanding of networking research's own "Innovator's Dilemma" and its implication on our ability to incorporate disruptive technology and application demands into our research agenda.
2. Effective support for networking as a maturing research discipline through frameworks for reproducible research results and benchmarking.

2.4.1 The Network Innovator's Dilemma

In his well-known book, *The Innovator's Dilemma*, Clayton Christenson explains how enterprises can become too involved in continuously improving on a successful product or line of business that they often fail to observe novel technology that can disrupt the success of the enterprise. The reason this happens, it is argued, is that often the capabilities of such disruptive technology is not (yet) on par with the existing technology that has proven itself and continues to be improved. Ultimately, however, this new technology does catch up in capability and cause the established enterprise to fail.

We have seen this happen in the networking world when scientists and engineers working on the highly successful and growing telephone network were sometimes too busy to notice the less-capable packet switching as it emerged. Ultimately, ubiquitous wide-area data networking, developed and engineered by a different set of people, became a successful competing enterprise with the telephone network.

FINDING: Success of the Internet carries the risk of the "network innovator's dilemma"

The success of the Internet carries the risk of complacency about the need for true innovation and outside-the-box thinking. The networking research community must avoid the "network innovator's dilemma," where it gets too involved in improving the existing Internet technology, and fails to observe novel and disruptive technologies that take a radical new shape.

Are we about to witness the same development again? There is a good chance that the Internet as we know it today will become the network of the past tomorrow. What radical new shapes will wide-area, ubiquitous communication take in the future and how can we empower the networking research community to be the innovators behind this future? It is clear that an understanding of this innovator's dilemma as it relates to networking technology is a fundamental part of this.

Ultimately, the entire community should internalize this phenomenon and contribute to creating a culture in which true radical innovation is fostered. This culture will need to include:

1. A consideration for a multiple future research agenda in which many technologies, paradigms and approaches are considered simultaneously and encouraged for their intellectual merit.
2. Placement of less emphasis on immediate deployment feasibility and short-term economic viability. This is particularly difficult because of the critical infrastructure nature of much of the networking technology. However, it will give the right environment for disruptive technology to be fully investigated and developed.
3. A better understanding of the scalability concerns. Today's non-scalable technology may become scalable tomorrow. It has been argued that the networking community should not be interested in non-scalable solutions. This has often stymied consideration of truly innovative approaches that may not be scalable currently but may in the future.
4. Encouraging closer ties between applications research and networking research. Often, the future requirements of networking are only apparent when the requirements of truly visionary applications are taken into consideration.

2.4.2 Reproducibility and Benchmarking

One of the hallmarks of scientific and engineering pursuits, particularly in mature disciplines, is the generation of research data and results, which are reproducible. This is sometimes enabled by shared infrastructure facilitating reproducibility or the definition of (sometimes imperfect) benchmarks against which competing algorithms and designs can be tested. Benchmarks can serve the dual purpose of systematizing the goals for some of the key classes of problems at hand while facilitating efficient evaluation and comparison across various research efforts to quickly move towards best-in-class solutions.

FINDING: Need for reproducibility of research outcomes
The maturing of networking research and its continued success require bolstering our scientific underpinnings in the area of benchmarking and reproducibility of results.

The maturing of networking research and its continued success require bolstering our scientific underpinnings in the area of benchmarking and reproducibility of results.

The Current State of Affairs. The networking research community has been partly successful at fostering reproducibility and benchmarking. This is a difficult task given the complexity, heterogeneity, distributed nature, and important operational aspects of networked systems. A few examples and approaches are discussed below.

Over the last decade the community has made substantial efforts towards developing tools and instrumenting networks so as to measure both the traffic, network state and dynamics of operational networked systems. The community has achieved this through the development of standardized measurement tools, university/industry consortia focused on collecting and distributing such data. In turn the data has found multiple uses, including for example, the study of traffic characteristics, growth trends, behavior of network protocols, e.g., addressing and routing problems. Further, the work driven by such measurements has led to fundamental research on appropriate models for traffic and network topology. These are a start towards developing reference models that capture the salient features of networked systems.

One approach towards facilitating reproducibility has been the adoption of a few openly available simulation tools, e.g., *ns-2*. The use of these tools has facilitated, to a great degree, the ability of researchers to evaluate and compare aspects of their work on a common simulation platform. Even with a common network simulation infrastructure, the task of reproducing simulated data can be daunting; yet, it is a necessary step towards a timely evaluation of research results.

Shared testbeds have provided another approach to enabling researchers to validate their results on a controlled yet common ground. As with simulation precise reproduction of experimental results on such testbeds, is often complex, and requiring an understanding of system configuration and software installations used for previous studies. See the Report of NSF Workshop on Network Research Testbeds [Testbeds 2002] for more discussion on this theme.

Desirable Future Capabilities. We believe that recognition that one of the methodological challenges in networking research is achieving better reproducibility of experiments, and validation of simulations in real world systems. Possible opportunities for continuing progress towards enabling these include:

1. Increasing the collection and availability of data and tying such activities to the development of specific reference models for networked systems, e.g., including types of traffic, network topology and/or configuration, and user or system responses, etc.
2. Encouraging research efforts that include the development and distribution of benchmarks targeted at specific classes of networking research, e.g., ad hoc routing, wireless MAC protocols, and congestion control.
3. Furthering the development of properly tested and continued improvement of middleware, prototype applications (e.g., there is no state-of-the-art video application), measurement, and simulation tools.

3. Basic Research in Networking

We will refer to basic research in networking as fundamental research in networking that can be widely seen as addressing theoretical aspects of networking. Basic research in networking can be best understood as discipline-oriented research. It brings together methodologies and techniques from applied mathematics, probability, optimization, automata theory, statistics, statistical mechanics and algorithms to address specific issues of relevance to networking and that are not addressed in any of individual constituent disciplines. The approach of basic research is to develop understanding by analyzing and studying traffic and network models that are generally a simplified version of the real system. The aim of basic research in networking is to develop insights, understand the performance limitations of existing or planned networks and develop methods to obtain estimates of their performance, suggest good policies, and guide the development of protocols and control mechanisms for networks to deliver performance and provide service guarantees required by the applications.

FINDING: Networking has evolved as a scientific research field with its own identity.
Basic networking research involves science from many areas and invents new scientific methods to fit the networking context, resulting in a field of networking with its own identity. While there is no particular networking science, the methodologies developed are a result of addressing specific networking problems and are often not of direct interest to the individual constituent science disciplines.

Basic networking research involves science from many areas and develops new scientific methods to fit the networking context, resulting in a field of networking with its own identity. While there is no, or, at least not yet a particular *networking science*, the methodologies developed are a result of addressing specific networking problems and are often not of direct interest to the individual constituent science disciplines. Networking is an engineering discipline but often requires significant extensions of basic underlying mathematical tools. Due to interaction with the physical environment, the nature of network traffic, and unpredictable human behavior, networking also has an important empirical component, which draws heavily upon statistics. Indeed even though networks are man made objects, created via engineering rules, the interaction of a vast number of entities as well as the human element make networks enormously complex. This has given rise to a healthy experimental and simulation component to basic research, which aims to understand the complex interplays and dynamics. In all these respects, networking research can be identified with a science or set of principles unique to networking.

Workshop participants identified the promotion of basic research in networking as an important activity under the purview of the NSF. Recent years have seen a drastic drop off in funding basic research from both industry and defense agencies. A lack of sustained funding in basic networking research will limit the further advances in developing versatile networks capable of offering reliable and high-performance transport capabilities especially because of increasingly demanding applications. The drop off in support of basic networking research has been fueled by the collapse in recent years of the telecommunications industry and the especially disheartening trends of shutting down or downsizing important industrial networking research labs. This means that the brunt of the basic research costs will have to be borne by government agencies such as the NSF, and that universities will have to re-establish themselves as the generators of know-how. It was also recognized that government funding should take into account both the long-term sustenance of on-going efforts as well as charting new avenues of basic networking research. Most methodological successes are due to careful and sustained long-term efforts.

FINDING: Need for increased role of government in support of basic networking research.
The disheartening trends of shutting down or downsizing important industrial networking research labs means that the brunt of the basic research costs will have to be borne by government agencies such as the NSF. The same trend requires universities to re-establish themselves as the main generators of know-how in networking research.

3.1 PAST IMPACT OF BASIC RESEARCH ON NETWORK TECHNOLOGY AND BASIC SCIENCE

Science, and scientific and mathematical methods have played a key role in the development of the Internet. Following, we provide specific examples of basic science leading to important innovations in networking are. The list, while not exhaustive, was identified in the workshop as areas where basic research provided the basis for the design of new generations of networks.

- Probabilistic models and analysis leading to increasingly efficient random medium access protocols from ALOHA to Ethernet to CSMA/CD to 802.x.
- Stochastic approximation theory and the concept of AIMD fairness lead the path to the *slow start* and *congestion avoidance* TCP congestion control techniques.
- Game theory, especially the cooperative and non-cooperative ideas of John Nash leading to an understanding of fairness of resource allocation and the precise functioning of congestion and flow control algorithms.
- Probabilistic and deterministic queueing theory and large deviations theory leading to leaky buckets, WFQ (weighted fair queuing) and other scheduling algorithms, network calculus, and effective bandwidths as mechanisms of traffic control and

FINDING: Basic networking research has played a key role in the development of the Internet
Developments that have been influenced by basic networking research are random medium access protocols, including Ethernet, TCP congestion control, flow control algorithms, scheduling algorithms, distributed routing protocols, and many more.

network provisioning.
- Information theory and optimization leading to capacity studies in wireless and wired contexts.
- Measurement, statistics and stochastic modeling leading to the observation of fractal and long-range dependent properties of network traffic.
- Control theory and queuing theory leading to RED/AQM and improved TCP performance.
- Centralized optimization algorithms leading to distributed routing protocols (Bellman-Ford and Dijkstra) and to BGP and OSPF routing protocols.

The flow of ideas from networking back to the underlying sciences has also been extremely influential. This has engendered new investigations in the corresponding disciplines that in turn have yielded tools to analyze even more complex systems. The following are just a few illustrative and high profile examples:
- The work on product-form solutions in queuing network models leading to a rich theory of quasi-reversibility of Markov chains.
- The primal-dual view of TCP-like algorithms leading to a framework for studying distributed algorithms for constrained optimization problems.
- The empirical observation of long-range dependence has given impetus to the study of queues with sub-exponential inputs and heavy-tailed characteristics.
- Teletraffic analysis has been the driving force in the development of Palm probabilities and their queuing applications leading to extensions of the underlying theory.
- Early work on connectivity of radio networks led to the development of continuum percolation theory that is now a thriving field.
- The stability analysis of fluid stochastic network models arising in networking has led to the development of Lyapunov techniques for studying stability issues for general classes of Markov chains.

Thus, it is clear that networking research has been in the best traditions of scientific inquiry and not just a simple application of ideas directly imported from the constituent fields. It has often been stated that the practical advances in networking have been de-coupled from basic research but a close evaluation of the situation reveals that it is far from reality. The increasing complexity of emerging applications, the huge end-user population and the enormous costs of deployment make cavalier, less formal approaches to networking research highly dangerous in terms of developing know-how, and will strengthen the role of basic networking research.

FINDING: Significant role of basic networking research. As the Internet becomes more and more an integral part of ubiquitous and critical services, understanding the behavior of the Internet and learning how to improve its properties is a fundamental endeavor that will have a enormously beneficial impact on all facets of our economic and social life.

As the Internet becomes more and more an integral part of ubiquitous and critical services such as banking and healthcare, understanding the behavior of the Internet and learning how to improve its properties is a fundamental endeavor that will have a enormously beneficial impact on all facets of our economic and social life.

3.2 FUTURE ROLE OF BASIC RESEARCH IN NETWORKING

Basic research in networking must involve core efforts in the basic methodologies and techniques such as control theory, queueing theory, information theory, distributed algorithms, and in simulation methodologies. However, it is felt that there will be a greater impact if research is pursued in conjunction with a focus area. Later in the report (in Section 5) we point out sources of fundamental problems, which can serve as the basis for developing appropriate methodologies and tools.

In the same way that basic research in networking has been successful in the past and continues to provide us the means to design better performing and more robust networks, recent research successes have indicated that there is much to be gained by studying the fundamental problems in networking and using the insights thus gained to design better networks. The emerging applications and directions require fundamental understanding on how to design and control networks on an even grander scale than before. Thus, the support of basic research in networking by agencies such as the NSF needs to be increased. The networking community continues to increase and with that we are seeing a much more multi-disciplinary flavor in the research. However, it is the purpose of basic research to identify the important tools and methodologies, which are necessary to provide solutions to problems. A failure to recognize the importance of basic research can lead to dramatic failures in terms of performance and stability of networks.

4. Vision of the Future

NEW APPLICATIONS AND PARADIGMS

Applications are the raison d'etre of networks. Given this importance of applications, it is not surprising that many of the important principles underlying today's networks have been strongly influenced by application-level considerations. From the basic principles of packet switching and statistical multiplexing (motivated by burstiness in application traffic), to today's congestion control algorithms (enabled by application elasticity) to QoS architectures (driven by QoS needs of interactive multimedia applications), much of the innovation in networking has been fueled by application-level needs. We believe that this application "push" for developing new network mechanisms, protocols, and architectures will continue to influence networking far into the future. Thus, when considering future directions in networking research, future applications—their requirements, and the new challenges that they pose—provide an important perspective with which to view this future.

FINDING: Important role of applications.
Many of the important principles underlying today's networks have been strongly influenced by application–level considerations.

FINDING: Application requirements and challenges of applications will continue to influence networking far into the future.
An application "push" for developing new network mechanisms, protocols, and architectures will continue to influence networking far into the future.

With a virtually unlimited number of potential future network applications, our discussion here is necessarily selective rather than exhaustive. A number of important network applications were discussed at the workshop, including gaming/virtual worlds; telephony and teleconferencing; telepresence; remote education; collaborative experimentation; e-science; environmental measurement/monitoring; P2P systems; home networking; wearable computing; sensor networks, and disaster recovery and emergency response networks. Of these, we chose three illustrative applications below as being particularly likely to play an important role in driving future networking research. In the following subsection we discuss these three applications, and the challenges they impose. In Section 4.2, we take the larger view, and discuss the user-oriented and network-centric needs and challenges induced by these and other applications.

4.1 THREE ILLUSTRATIVE NETWORK APPLICATIONS

4.1.1 Ubiquitous Tele-presence

Imagine your office or home where your walls and furniture consist of adjustable displays. The room listens to your senses and responds to your body gestures, touch, face expressions, spoken commands, eye movement, handheld controller commands via an avatar (personal digital assistant) by delivering desired information or connecting with requested person(s).

Imagine in this environment you are preparing a lecture on disappearance causes of rainforest decline and you need information from a digital library about the Amazon rainforest deforestation. Your personal assistant finds remote multimedia material about Amazon habitat in the form of HDTV video and HiFi sound together with three articles on the Amazon region and two related Earth Day conference talks, written and presented by your colleagues in Brazil, Canada and Thailand. On your behalf, the assistant contacts the authors to discuss the material. Your colleagues are available and join you in the virtual room for discussion of the multimedia material and articles that you want to use during your lecture. All see the same material retrieved by your personal assistant. As you discuss your lecture with your colleagues, and go through the retrieved material, you use a distributed environmental simulation that draws on model components and sources in input data drawn from around the world, to generate predictions of when the Amazon forest will disappear if no actions are taken. Your personal assistant records the conversation and simulations, annotates multimedia material with comments and composes the recording into a first draft of a lecture presentation. During your discussion, the personal assistant interrupts you and indicates to you that in 10 minutes you need to finish and go to another virtual room for your faculty meeting. As you finish the discussion with your colleagues, you preview your first lecture draft with your colleagues as well as you ask them if they would be available to join your class tomorrow at 3pm and answer questions from students. After your faculty meeting, you work jointly with your personal assistant to interactively refine and edit the content of your lecture.

The scenario outlined above illustrates a broad range of real-time, interactive multimedia services that will stress the underlying network infrastructure. (We note that in addition to the networking challenges, many of the implicit HCI capabilities contribute to the end-to-end complexity and may cause delays that are not negligible!). The needed distributed services in this scenario are: (a) real-time multimedia conferencing services integrated with virtual reality and other advanced display technologies (e.g., holography) among your colleagues, (b) real-time communication service with your personal digital assistant via different types of networks and through different types of information (e.g., gesture, voice, images), (c) real-time retrieval service of remote multimedia material from different sources (e.g., conference talks, articles, HDTV videos, auditory information on Amazon birds), (d) distributed real-time simulation, (e) real-time, reliable and secure interactive and distributed editing and information sharing, and (f) reliable and time-sensitive content distribution.

4.1.2 Sensing Everywhere

The environment—at a macroscopic scale—is of great interest both scientifically and practically. By understanding the proliferation of pollutants and how various species they affect, scientists gain important tools in providing a balance between our needs and our obligation to the environment around us. By monitoring and potentially controlling salient features of the environment, citizens gain advance warning of impending natural disasters, with the possibility of avoiding them or mitigating their effects.

The Mississippi River watershed is one example of such a macroscopic system. It is fed by an uncountable number of smaller rivers and streams, collecting water from a vast area of the country. Modern science has given us the coarse tools for understanding and predicting the first-order behavior of such large river systems, but lacks the fine-grained information needed to provide deeper predictive and prescriptive power. For example, if one were able to provide a distributed, decentralized sensor field across the entire watershed, one could plot the system's behavior with great accuracy, and well in advance of any danger. By instrumenting and adding distributed control to existing flood abatement mechanisms, such a sensor network can reduce or even eliminate the impact of floodwaters during and after exceptional weather systems. By coupling the observations of the watershed-sensing network with an atmospheric/meteorological sensing system, preventive measures may be taken in advance of exceptional weather systems.

Flood control is only one potential use of such a sensor field; one can also track the spread of pollutants through the watershed. Often, the release of pollutants goes undetected at the source, only becoming apparent over a long period of time at sites far from the original point of entry. A distributed sensor field can act as an early warning system, allowing cleanup and abatement work to happen closer to the pollution site, increasing effectiveness and reducing costs.

Sensors at scale have a role in the man-made world, as well. It is an oft-repeated maxim that engineers learn more from their failures than successes. For example, it is only when a structure fails that civil engineers can add to their ability to build structures that withstand ever-larger forces. However, one can consider building structures with many sensors embedded within them, measuring tension and compression in individual structural elements throughout the building. Then, as the building is subject to sub-critical stresses, the materials' response can be measured, cataloged, and understood. This in turn leads to the ability to build structures that are safer and more effective, and to do so at lower cost with less uncertainty in the design process. One could even envision embedding actuators in the building to adjust the tension applied to structural elements in reaction to external, settlement, or thermal forces, dampening their impact and improving the resiliency of the structure.

Constructing such sensor fields requires a number of fundamental technical advances in networking. For example, deployment of sensor networks at this scale requires self-configuration. The resulting networks must be fault-tolerant, as sensors are likely to fail over time, with only ad hoc replacement. If actuators are to be included, they must be controlled securely, to prevent malicious acts of destruction. Likewise, any system with real-world control capability requires strong guarantees about the timeliness of communication. We note that sensor networks exist at many scales, from small-scale motes that might perform species habitat monitoring to large-scale remote-sensing radars that monitor meso-scale weather events. An important architectural question is whether such apparently different sensing networks will demand individual “stovepipe” networks, or whether common architectural characteristics will result in one of more common layers of functionality, in much the same way that a common IP layer is able to sit above and abstract away the details of vastly different link layer technologies.

4.1.3 Emergency Response and Disaster Recovery Networks

Communication networks can play a major role in the event of a natural disaster such as an earthquake, tornado, or a fire, or a malicious attack such as arson, chemical/biological attack, or a bomb. They can be used to locate victims, assist, track and coordinate emergency rescue teams, provide communication capabilities to other assistance providers, and provide valuable forensic information on the disaster and design flaws that can be averted in the future.

Imagine a hypothetical disaster scenario such as a severe earthquake in downtown San Francisco, which destroys a significant part of the existing civil infrastructure. Emergency response teams that arrive at the scene can be immediately self-configured into a mobile ad-hoc network infrastructure. The ad-hoc network feeds video and data streams from the emergency sites through the wired infrastructure to multiple emergency management control rooms. Within only a few minutes after the earthquake strikes, several hundred emergency workers cover an area of almost a square mile, each feeding several hundred Megabits per second of audio, video, and diagnostic data to emergency sites, using wearable and handheld computers, digital cameras, data collection tools, and other devices. Millions of deployed sensors collect data and report on structural damage to buildings, fires, and survivors in the debris.

Hundreds of simultaneous high-resolution video feeds from rescue teams, satellites, and helicopters and aircrafts flying over the disaster sites are aggregated and transmitted in real-time across optical networks to several computing centers. There, the video data is processed and translated into a three-dimensional map of the catastrophe site. This map is accessed, again in real-time, at the emergency control centers to assess damage in various areas and zoom in the attention of the rescue efforts to the most needed areas.

The wired network infrastructure, which has been severely damaged during the earthquake, has dynamically reconfigured itself to bypass failed communication nodes. The basic communication infrastructure is robust to failures and still viable even after taking a lot of damage. High-bandwidth wireless network hubs are installed to connect areas where the wired infrastructure cannot be repaired through reconfiguration. The wireless hubs maximize the joint coverage area of the wireless network by trading data rate for increased transmission range. Unmanned aerial vehicles that are equipped with high-bandwidth switching equipment bridge communication across larger areas where the wired network infrastructure is paralyzed.

The global information flow through the disaster area is managed at the granularity of a micro flow. Prioritization of information distinguishes between regular communication and communication between emergency providers. Multicast communication is deployed to increase the efficiency of bandwidth usage for one-to-many transmissions. Access to the millions of different data streams is controlled by an access rights management system, which assures security and privacy, and provides a dynamic coalition management to share data between groups as needed to perform certain rescue tasks.

This emergency scenario entails a tight cooperation of heterogeneous sets of mobile and fixed networking technologies that can self-configure within minutes into a vast communications infrastructure that dynamically adapts to changes of the environment and the needs of application users. The complexity of such a network is beyond what can be planned or managed by human operators. Therefore, self-manageability and configurability of the emergency network infrastructure is essential. The applications in the depicted scenario range from data fusion from sensor data, to collaborative applications, and large-scale real-time visualization.

4.2 MEETING APPLICATION-INDUCED CHALLENGES

In our discussion of specific applications above, we have identified a number of particular challenges induced by each application. In this section we take the larger view and look broadly at the challenges posed in providing a network substrate that is sensitive to application, user, and network-provider needs. We will find it convenient to distinguish between those challenges associated with the user/application experience, and those that are more network-centric.

4.2.1 User-Focused Challenges

Robustness: Perhaps the most desirable attribute of a network is that of "robustness." Robustness implies having or exhibiting vigorous health, and being strongly formed and/or constructed. Robustness at the top level includes all attributes related to delivering network performance at the requisite level under both favorable and/or unfavorable (malicious or buggy) conditions. It includes various attributes such as reliability, availability, fault tolerance, security, privacy, anonymity, and performance. Since each of these attributes also is used independently, we will define them in brief here. To achieve robustness, all of these attributes need to be met.

A reliable and dependable network has the ability to provide its intended, expected, and agreed upon functions, behavior, and operations, in a correct and timely manner. The term "reliability" can be defined in many ways. A simple definition is the capability to operate as intended for as long as needed. A more formal definition states that it is the conditional probability to deliver uninterrupted satisfactory performance for a specified interval under specified operating conditions.

Availability, on the other hand is the probability of finding the network resources to deliver the required performance when needed. A secure system is the one, which is not prone to failures due to attacks and/or threats, and the vulnerabilities are difficult if not impossible to expose. Reliability and availability are achieved using various techniques, such as fault avoidance, fault removal, and fault containment, and fault tolerance.

Invisibility: Another user need is the desire to make the network invisible from the user perspective. This has serious implications from the point of view of meeting performance metrics. Simply put, a user should not feel any difference if they are communicating through a networked system or via a local system. It should be equally easy to connect and operate and see no noticeable difference in performance. This implies that the network control is such that it is easy to negotiate with, and obtain bandwidth on demand of any granularity, large or small. Networking relates to connectivity in a cost effective manner. It should be possible for many users to share the "wire" or achieve networking in a cost effective manner where an individual user is immune to delay, jitter, and performance constraints while the network service providers achieve higher utilization and can maximize their revenues for the services they provide. This implies that the protocols and control mechanisms effectively deliver data in a timely manner without significant performance overhead, delay, jitter, blocking, and loss of data. Thus all of these performance metrics are of importance.

Ease of Configuration: All kinds of users, from novice to expert, should be able to access the network. Connecting to the network should be as easy as plugging into a telephone jack or a power outlet. Even better, the network should detect a user's presence and automatically connect the user to the network (as with current wireless Wi-Fi cards). Self-configuration is a key component of ease-of-use. The same concept should be possible even when failures occur. Management of failures in a diligent autonomous way is a preferred attribute.

Scalability and Heterogeneity: The network and user base is expanding, using a variety of communication media, various types of devices, different protocols, and different bandwidth levels. Thus scalability and heterogeneity need to be managed. A protocol and/or control mechanism cannot break down with an increase in user base or a change in the network. Although, these issues can be considered as part of achieving robustness, it is important to consider them in their own right as well.

Exploiting Storage and Processing Capabilities: Users are interested in exploiting distributed storage and processing capabilities through networks in a seamless fashion. Often not all users will have sufficient computing and/or storage power on their end devices. Energy constraints may also require them to use other devices across the network. With the increasing (and largely unused) computing capabilities available at many nodes in the network, it should be possible for a user to access these underutilized computing facilities and/or storage as necessary. To facilitate such usage, appropriate protocols and capabilities need to be developed. Protocols are required for connection setup, processing capability reservation, intermediate processing, data acknowledgement, buffering and retransmission, flow and congestion control, ordered delivery, and security.

Power Consumption: Users desire anywhere-anytime connectivity to the network, often in locations without an external power supply. In such situations, reliance upon batteries makes power consumption of network interfaces an important issue. With the development of low-power devices, such as hand-held PDAs, network protocols and mechanisms should adapt to power availability.

4.2.2 Network-Centric Challenges

It was the sense of the workshop that many of the network-centric challenges in the future will result not from a single "killer application" but rather from the need to simultaneous support a large and diverse mix of applications. We identify several of the most important challenges below:

FINDING: Relevance of killer application(s).
Challenges in the future will result not from a single "killer application" but rather from the need to simultaneous support a large and diverse mix of applications.

Heterogeneity and scale. A single "converged" global network will be required to support many different types of traffic, ranging from delay- and jitter-sensitive high bandwidth telepresence applications to low-bit rate monitoring applications, to high-bandwidth point-to-point bulk data transfers in support of e-science. In addition to supporting a diverse mix of such applications, the network will also be required to support a large number of instances (i.e., many simultaneous users) of these applications. With such diversity, and such large numbers, control of traffic aggregates may be preferable to per-flow control mechanisms within the network. A large-number of high-bandwidth end-end applications will require core network speeds that dwarf today's gigabit core, requiring an architecture that will likely mix packet switching at the network edges with circuit switching in the network core.

Manageability. From home networks (with numerous "smart" network home appliances and devices) to enterprise networks (with firewalls, NAT boxes, caches and other layer-3 devices) to core networks (with a heterogeneous mix of devices, link technologies and speeds) to application-level networks (such as content distribution networks with thousands of caching servers), the number and types of network elements is rapidly increasing. Indeed, the rate of proliferation may soon—if not already—outpace our human capabilities to monitor, diagnose, and repair such systems. Thus, automated network management techniques, including autonomic, self-configuring devices and systems, and self-monitoring and self-healing architectures will become increasingly important. This will be particularly the case in high-stress rapid deployment systems (e.g., disaster recovery networks) and commodity systems designed to serve the technically non-savvy user.

Ability to Evolve. With the increasing reliance of the national economy on the Internet comes a concomitant difficulty in evolving this infrastructure to support future needs and applications. Concerns for backwards compatibility, investments in deployed infrastructure, and the need for ultra-reliable, unbreakable services result in an understandable resistance to change. In the face of this reluctance towards change, how can the needed changes be made to support future applications? How will the disruptive technologies/approaches of tomorrow be introduced, investigated, and incorporated? Two approaches are likely to be taken. Extensible network architectures (a.k.a. active networking) offer one possible path towards deploying new network services, with the ability to add new functionality and capabilities deeply within a device within the network. An alternative approach is to take the underlying infrastructure as being immutable, or only adopting change at very slow timescales. In this case, new functionality can be provided at the application layer via overlay networks. Here, the challenges include the interaction between control loops operating at the overlay and underlay levels. Another open question is to determine the right set of services to be provided by the underlay, in order to support services needed to build a diverse set of application-specific overlay networks.

Beyond the Internet. As noted above, sensor networks with devices ranging from small-scale power-limited motes, to large-scale atmospheric-sensing radars, to environmental monitoring devices, present a wide range of needs and constraints that are different from the dominant applications in today's Internet. And yet many of these networks will clearly want to be connected to the Internet. Is the end-to-end-IP model the right model for such networks? Similarly, for all optical core networks, how will the boundary between the circuit-switched core and packet-switched edge be defined and designed? In both cases the network of tomorrow (or pieces of this network) may be as radically different from today's Internet as the Internet was from the telephone network before it.

5. Grand Challenges

The current Internet has done an amazing job by providing a set of services that have led the development of the web, rudimentary real-time applications, and P2P file sharing applications. However, the current Internet architecture suffers several shortcomings that will impede the development of the next generation of applications such as ubiquitous telepresence, sensorized universes, and disaster recovery. The most important of these shortcomings are:

- The difficulty to make changes to the current architecture that will allow for the development of new services, such as differentiation of quality of service. Currently, specialized overlay networks have been developed to provide new services.
- The lack of appropriate market incentives to encourage the development of new services and applications.
- The lack of a theoretical foundation for the development of wireless- and mobility-based services. This is especially prominent because of the ever-increasing importance of wireless and mobile communications.

FINDING: Meta-grand challenge of networking research.
The current Internet architecture suffers several shortcomings that will impede the development of the next generation of applications.

The Meta-Grand Challenge facing networking research today is to develop new network theories, architectures, and methodologies that will facilitate the development and deployment of the next generation of services and applications.

These can be summarized in the following meta-grand challenge facing networking research today:

Meta-Grand Challenge: To develop new network theories, architectures, and methodologies that will facilitate the development and deployment of the next generation of services and applications.

It is sometimes said that many of the problems facing developers of new services can and will be resolved through the introduction of higher and higher bandwidths. This is not a panacea. First of all it is not always possible, e.g., in embedded systems. Second, throwing bandwidth at the problem may not be the most cost effective solution. Last, the presence of sufficient bandwidth does not necessarily translate into services without research into new mechanisms to map the bandwidth into services.

The above meta-grand challenge translates into a number of networking research challenges that include

- development of a network information theory,
- design of overlay networks,
- development of resilient networks,

- introduction of market incentives,
- development of architectures for embedded and pervasive networks, and
- development of virtual networks.

We now turn our attention to the networking challenges listed above.

5.1 NETWORK INFORMATION THEORY FOR WIRELESS NETWORKS

Classical information theory has been the fundamental basis for the major developments in communications, error control, and signal/image processing over the last half-century. In fact, today's advances in information technology would not have been possible without information theory. By contrast, while there have been many important developments in various aspects of communication networks (e.g., routing, congestion control, etc.), there is no unified, basic theory yet that extends information theory to networks. Such a theory is especially important in the case of wireless networks, where the random error filled nature of the channel makes the classical communication problem all the more challenging.

Wireless networks have experienced a phenomenal growth in the last decade, primarily through the exponential increase in the use of cellular phone and indoor wireless LANs. This growth has been fueled by technological advancements that grew out of fundamental research in physical layer communications on topics such as error-control coding, modulation, multiple antenna systems, etc., and in networking research on topics such as power control, traffic engineering, and medium access control. Further, in addition to the traditional cellular networks, a significant amount of research in wireless networking is being devoted to multi-hop radio networks, especially as the world becomes increasingly networked, with wireless connections being enabled to everyday appliances. Multihop radio networks would also be the foundation for networked sensor systems, with applications ranging from biomedical engineering and structural fault monitoring to security applications, and for community radio networks with small wireless base stations that communicate to each other in a multihop fashion. It is imperative to understand the fundamental throughput and delay performance limits of wireless networks, with the goal of designing resource allocation (power control, medium access, routing, etc.) algorithms that allow us to reach these performance limits.

FINDING: Grand Challenge: Need for a network information theory

There is no unified, basic theory yet that extends information theory to networks. Such a theory is especially important in the case of wireless networks.

It is imperative to understand the fundamental throughput and delay performance limits of wireless networks, with the goal of designing resource allocation algorithms that allow us to reach these performance limits.

Recently, a new line of research has been initiated, whereby the asymptotic sum throughput of a dense wireless network has been established as a function of the number of nodes in the network. Such results characterize the spatial diversity inherent in a wireless network. For example, if n sources share a wireline link, the throughput per node is obviously of the order of 1/n. On the other hand, if n spatially distributed wireless nodes share a common frequency band, the throughput per node is higher; it is of the order of $1/\sqrt{n}$. A constructive proof of this result also provides insight into resource allocation methods, showing that simple power control and MAC schemes, along with shortest path routing can achieve the throughput limit. These recent results on the achievable capacity in multi-hop wireless systems can be viewed as a first step in this direction. However, these results (just like classical information theory) do not involve any notion of delay or burstiness, both of which are basic and important elements of networking. So, a major challenge is to develop a network information theory that incorporates delay and burstiness. Some important related questions include whether we can develop *distributed* and *scalable* protocols and algorithms to achieve performance limits predicted by such a theory. What will be the effect of errors on performance and what will be the effect of feedback, which is inherently omnipresent in networks, are additional key questions.

Network information theory will also be instrumental in redesigning the current wireless network architecture. The current layered networking architecture supports transparency between layers at the cost of performance. While optimizing network designs and protocols within each layer of the network architecture help improve performance, it is increasingly clear that these improvements are insufficient to fuel major growth in wireless services. To achieve orders of magnitude gains in system performance, network protocols and designs will have to be engineered to exploit the capabilities of the underlying technologies (e.g., physical layer communication, VLSI design, etc.). For example, a network protocol could be improved by exploiting the time-varying channel conditions in wireless networks through intelligent coding and modulation, and network scheduling techniques (work on this problem has already motivated the development of commercial solutions such as Qualcomm's HDR); it could be further improved by tailoring the protocol to accommodate efficient channel estimation techniques; significantly more gains could be obtained by controlling and even creating channel variations in the system through smart antenna design techniques; further improvements in performance could be achieved by exploiting low power transmission capability to selectively reduce interference. The major challenge will be to quantify the tradeoff between complexity and performance, so that design of systems with target cost (i.e., complexity) and QoS (i.e., performance) constraints, will become possible.

5.2 OVERLAY NETWORKS

With our increasing dependence on the Internet, it is imperative that we are able to evolve it to adapt to changing application requirements, accommodate emerging technologies, and address new threats. Ironically, the very commercial success that has fueled our increased dependence on the Internet has also reduced our ability to evolve its underlying architecture to meet new demands. This has resulted in a new strategy by which services are introduced into the Internet based on the notion of *overlay networks*. This strategy adds a layer to the architecture, inserting special purpose overlay networks between the global infrastructure and the ultimate end nodes. Overlays constructed with service-level requirements in mind can make network-level decisions, such as routing, service model, and data manipulation, tuned to the specific service. Overlays designed to support small-scale applications can utilize algorithms that would not scale to a global size.

This potential has created great interest in the overlay model, with vigorous activity in the academic and commercial communities. Examples include file sharing and network-embedded storage, content distribution networks, routing and multicast overlays, QoS overlays, scalable object location, scalable event propagation, and anomaly detection mechanisms.

> **FINDING: Grand Challenge: Overlay networks**
> Overlays designed to support small-scale applications can utilize algorithms that would not scale to global size, but raises numerous challenging problems: How to share network resources between multiple overlay networks? How to build economic mechanisms that arbitrate among resource providers, overlay services and users? How to preserve the robustness of the shared infrastructure in the face of misbehavior or maliciousness?

As valuable as this work has been, however, it has focused on understanding and increasing the functionality of specific overlay networks and algorithms. This focus on overlay functionality has left unanswered an important question relevant to this approach, what should lie underneath? In fact, today's overlay solutions are developed independently under the implicit assumption that they will run in isolation. However, as more and more overlay networks are deployed, this assumption no longer holds. Since each overlay network will manage its resources and react to the congestion independently this can lead to instability. The classic example is when multiple overlay applications discover simultaneously that a link is lightly loaded and as a result decide to reroute their traffic on that link. This may lead to the link becoming congested, which in turn will prompt the overlay applications to reroute their traffic on another link. This will result in oscillations that will negatively impact the performance seen not only by the overlay applications but by the legacy applications as well. Hence, critical questions that need to be addressed are: What is the appropriate narrow, universally shared environment to underlay the overlays? What should the interface be between underlay and overlay networks, what information should be passed in both directions between the two networks? This environment, termed the *underlay layer*, must play three critical roles. First, it must *support* a variety of overlay structures and services; in the same way that today's Internet supports end-to-end applications. This implies, among other things, that the underlay must expose information about the underlying physical network that overlays need in order to do their job. Second, it must *protect* both overlays and underlying resources from damaging interactions, instabilities, and behaviors. Finally, it must support a level, open playing field, allowing technologies, services and participants to come and go while maintaining the basic integrity of the system.

Answering these questions involves addressing three basic research challenges:

1. How should resources (bandwidth, storage, computational cycles, shared information) be shared among competing overlays in an economical, computationally practical fashion? How might the underlay level best provide a flexible and powerful foundation on which to build economic mechanisms that arbitrate among resource providers, overlay services and users?
2. Virtualization, a basic technique for multiplexing shared resources, is needed to define the basic unit of network resources to be allocated. The underlay must provide and control access to the underlying network resources by overlay services, while also preserving the robustness of the shared infrastructure in the face of misbehavior or maliciousness. Reconciling the tension between making the Internet extensible and ensuring its integrity is a fundamental challenge for the next generation networking infrastructure.
3. The ability to gather, share, and use network information effectively for routing, overlay node placement, DDoS detection, fault recovery, and so on, is a key element of the overlay/ underlay architectural model. We envision a shared information plane (running above the transparent data plane) that creates stores, propagates, aggregates, and discovers observations about applications, overlays, and the underlying physical network. Developing such an architecture is a fundamental challenge.

5.3 RESILIENT NETWORKING

At a very fundamental level, all network services and applications rely on the Internet routing infrastructure for packet delivery. However, the Internet faces increasing challenges in providing *dependable* packet delivery. Today's Internet is a loose interconnection of tens of thousands of networks run by different Internet service providers, and accessed by hundreds of millions of users with diverse interests. Due to its sheer scale, faults have become the norm rather than the exception, and faults can include component failures, human operational errors, spreads of software viruses, and malicious attacks directed against the Internet infrastructure. Furthermore, due to its highly complex composition, faults in one part of the Internet often bring about unexpected consequences to the rest of the system.

> **FINDING: Grand Challenge: Resilient networking**
> As the Internet grows in scale, it faces an increasingly adverse environment that the original design did not envision. Faults have become the norm rather than the exception, due to component failures, human operational errors, spreads of software viruses, and malicious attacks. It is important to develop a basic understanding of the rules and principles in network protocol design that facilitate the realization of resiliency in large-scale networks to these faults.

The principles that guided the design of current network protocols focused on robust data delivery in the presence of physical failures. This design has led to a highly successful protocol architecture, and the existing Internet infrastructure and protocols effectively handle the frequent the occurrence of various scales of physical failures. This design however is based on the implicit assumption that all components abide by the protocol rules. An unexpected fault can lead to a system failure; past experience provides numerous examples dating back to the original design. A blackhole incidence of the old ARPANET distance-vector routing protocol was caused by memory corruption, and the "sequence-number fault" of the ARPANET link-state routing protocol is another *unexpected* system error resulting from the violation of design assumption. In the early days of the Internet, such unexpected faults were rare and damage was constrained due to the limited network user community and usage. But as the Internet has grown in scale, faults due to various causes have become the norm rather than the exception and the many faults go far beyond simple physical failures.

Although scalability is recognized as a research challenge, research in scaling has mainly focused on combating large address space, large routing tables, and higher traffic volumes. The research challenge raised by this unprecedented Internet growth brings far-reaching impacts that go beyond the sheer magnitude of changes. As the Internet grows in scale, it faces an increasingly adverse environment that the original design did not envision. As a result, the current system is incapable of effectively defending itself against new security vulnerabilities. Like any other human-engineered system, unexpected faults in the Internet lead to failures of various scales. While there have been many recent research efforts to combat various faults/attacks, each effort typically examines recent *specific* threats individually and proposes corresponding remedies. Overall, it is unclear how effective such separate, piecemeal efforts can be in preparing the Internet to defend itself in the long run—we do not know what kinds of new faults will occur in the future but are certain new faults will occur.

To provide the Internet with a resilient routing infrastructure requires an overall *systematic* approach. It is important to develop a basic understanding of the rules and principles in network protocol design that facilitate the realization of resiliency in large scale networks, and specific enhancements that can be readily applied to the existing Internet to effectively strengthen its defense ability.

5.4 PROVIDING MARKET INCENTIVES (THE ECONOMICS OF NETWORKING)

One of the impediments to the deployment of new services on the Internet is the lack of market incentives to improve network services and applications and to use them efficiently. The absence of service differentiation and congestion pricing results in network services that lag behind those that technology makes possible. This situation has resulted in an untapped potential, where network providers are unable to justify further investments in new services to support innovative and revenue-generating applications. In fact, recent history has demonstrated that the bottleneck to improving the global network may be economics instead of technology. The IETF has generated many protocols for better network services of the Internet, such as DiffServ, IntServ, and RSVP. However, most network providers do not deploy services that use these protocols, partly because of their complexity but, more fundamentally, because they lack economic incentives.

FINDING: Grand Challenge: Providing market incentives. One of the impediments to the deployment of new services on the Internet is the lack of market incentives to improve network services and applications and to use them efficiently.

Recent history has demonstrated that the bottleneck to improving the global network may be economics instead of technology.

Improved mechanisms should enable the creation of efficient markets among users and providers of content, services, and applications. Such mechanisms require new algorithms and protocols to collect necessary measurements, calculate prices, charge users, and redistribute revenues. The development of such algorithms and protocols requires research that must combine ideas from economics and networking. In peer-to-peer networks, proper incentives could stimulate users to share information and services. In Wi-Fi networks, pricing might compensate users for sharing their access points. In wired networks, the various participants should have the incentive to improve the utility of the network.

The last few years have seen a number of proposals for introducing such pricing mechanisms to better manage networks. However, in those proposals, pricing appears as a signal to control congestion and maximize the sum of the user utilities. Such a formulation, although useful by itself, does not adequately address the issues of market differentiation and incentives through revenue distribution among the providers of services and content.

It is widely accepted that the current Internet economic model is inefficient, in the sense that it fails to capture the potential utility of the network. Today's Internet does not support many potentially valuable applications. In fact, the progress of the Internet has been due to the development of applications that work reasonably well given the lack of dependability of the best effort service the network provides. However, the limitations of that best effort service severely constrain the applications it supports. The resulting loss of economic benefits, although hard to quantify even approximately, is probably substantial.

An analogy, although inadequate, of today's Internet is a post office without postage stamp where all the users would pay a flat fee every month, independent of usage. It is intuitively quite apparent that such a system would not provide the incentives for new services such as next day delivery or support the transport of large packages. The analogy makes it plausible that the Internet needs a postage stamp. However, to compound the difficulty, the Internet is not a single service provider. Instead, many independent commercial entities contribute to the delivery of information in the network. Some initial research has been conducted on pricing models in networks, such as Paris Metro schemes, differential charging for traffic classes, and charging for differential service guarantees. However, we believe that both the practical importance and the mathematical challenges of the problem justify that more work be encouraged in this area.

The objective is to design mechanisms and the corresponding protocols to enable the creation of efficient markets among the network users and the providers of services and contents. Examples include incentives for users to share information or offer services in P2P networks, micro-payments schemes in wireless hotspot networks, and charging mechanisms for new services in the Internet.
Such pricing mechanisms may be modeled as games among users and providers. Research problems concern the stability, convergence, scalability, incentive compatibility of the mechanisms, combinatorial auctions, evolutionary games, and behavioral models. An important issue is the extensibility of the mechanisms to different service models. Some users may want predictable quality whereas others may prefer predictable prices. Protocols for charging and revenue redistribution must be explored, with the relevant security issues.

Deployment issues include the possibility of introducing such schemes incrementally and the compatibility with existing systems. Should such pricing schemes be implemented as overlay networks or is it possible and preferable to modify current protocols?

5.5 SENSORIZED UNIVERSE

Networking between computers, in the early days, began by connecting a number of terminals to a costly mainframe and sharing knowledge within the office. In the 1990s, with the emergence of the Internet came the advent of a network age in which individuals could exchange information by personal computer. Today we have entered the age of personal web computing in which people use mobile phones and other devices besides personal computers to access networks. In the future, information terminals will become integrated into every aspect of daily life powered by sensors with different modalities and a sophisticated information network environment will evolve allowing information to be used anytime, anywhere, and by anyone. In such an environment, computational intelligence can be regarded as being embedded into the users' surrounding space rather than into specific devices, such as their notebook PCs, leading to a notion of smart or sensored-room, office, home, car, train, building, factory, campus, park, school, city, nation, globe, and universe. Remote sensing will become an integral part of our lives as we strive to use sensor technologies to monitor our surrounding space for a variety of applications such as security, health, education, comfort, environment, traffic, safety, etc.

FINDING: Grand Challenge: Sensorized universe.
In the future, remote sensing will become an integral part of our lives as we strive to use sensor technologies to monitor our surrounding space for a variety of applications such as security, health, education, comfort, environment, traffic, safety, and more. The design of a sensor network introduces daunting challenges in all aspects of the protocol stack of a sensor node.

The design of a sensor network is influenced by many factors, which include fault tolerance, scalability, production costs, operating environment, sensor network topology, hardware constraints, transmission media and most importantly power consumption. Each of these enforces specific requirements on all layers of a sensor node protocol stack. Modulation schemes, strategies to overcome signal propagation effects and low power hardware design are expected from the physical layer design. Determination of lower bounds on the energy required for sensor network self-organization, error control coding schemes, power saving modes of operation and taking care of mobility is a challenge for data link and MAC protocols. Addressing greater topology changes, scalability and interface with other networks is something that is expected from network layer protocols. The transport layer protocols should allow for diversity in end-to-end communication to capitalize on the variation in the communication channel characteristics. Also, one should not ignore the daunting task of collecting the data, often termed as data fusion, getting some meaningful information out of it, sensor query, and data dissemination.

5.6 VIRTUAL NETWORKS

The Internet infrastructure is much as it was in 1993, at the advent of the world-wide web. Ten years later, an influx of new users and applications is placing new requirements on the network. Services should work seamlessly as nodes travel across networks; nodes on different networks should be able to collaborate securely; protocols must quickly adjust to failures of not only links but servers and service providers; there needs to be less manual configuration and more automatic discovery of resources and routes.

FINDING: Grand Challenge: Virtual networks.
New users and applications introduce requirement for services with a higher level of abstraction, which deal with user mobility failures of servers and service providers, and automatic discovery of resources and routes, without a need for manual configuration.

All these different requirements can benefit from a semantically higher level of abstraction. More specifically, clients today no longer want to attach to a specific computer or IP address; they instead wish to access a resource, such as a service or a piece of content. The network should route users to nearby resources that are available and uncongested. Also, the mapping of resources to computers may change, even during a session; applications need to be able to survive such changes. The client is yet another resource which can move between computers or networks. Finally, the naming and transport of resources need to be authenticated and made secure against malicious parties.

In response to these new requirements, many ad hoc and often proprietary mechanisms have been proposed recently, many of which are not scalable and do not take advantage of all the underlying capabilities of the network. It is necessary that we identify those mechanisms that are useful across different applications and incorporate them into the network so as to speed up the development of new applications and improve the efficiency of the network.

References and Further Readings

[Christensen 1997] Clayton M. Christensen. The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail, Harvard Business School Press. June 1997.

[CRA 2002] Grand Research Challenges in Computer Science and Engineering. Computing Research Association, 2003.

[CSTB 2001] Looking Over the Fence at Networks: A Neighbor's View of Networking Research, Computer Science and Telecommunications Board (CSTB), National Academies Press, 2001.

[Education 2002] ACM SIGCOMM Workshop on Computer Networking: Curriculum Designs and Educational Challenges, *http://www.net.cs.umass.edu/sigcomm/education/workshop1.html*, August 2002.

[Networking 1994] Research Priorities in Networking and Communications, Report to the NSF Division of Networking and Communications Research and Infrastructure. May 1994.

[Testbeds 2002] Report of NSF Workshop on Network Research Testbeds. November 2002.

NSF Workshop on Fundamental Research in Networking

AGENDA

Dates: April 24-25, 2003
Airlie House, Virginia

Wednesday, April 23

6:30–7:30pm	Registration and Cocktail reception	*(Garden Room)*
7:30–8:30pm	Dinner	*(Airlie Room)*

Thursday, April 24

7:00–8:30am	Breakfast and Registration	*(Airlie Room)*
8:30–8:45am	**Welcome and introductions** N. Shroff, J. Liebeherr	*(East Room)*
8:45–9:15am	**Past and Present of Network Research at the NSF** A. Jukan	*(East Room)*
9:15–10:45am	**Panel 1: Grand Challenges in Networking** Chair: D. Towsley	*(East Room)*
10:45–11:00am	Break	*(2nd Floor, Airlie House)*
11:00am–12:30pm	**Panel 2: Network Applications of the Future** Chair: A. Ephremides	*(East Room)*
12:30–1:30pm	Lunch	*(Airlie Room)*
1:30–4:00pm	**Breakout sessions**	*(East, South, West, North Rooms)*
4:00–4:30pm	Break	*(2nd Floor, Airlie House)*
4:30–6:00pm	**Report back, summarization, report planning**	*(East Room)*
7:00–9:00pm	Dinner	*(Airlie Room)*

Friday, April 25

8:00–8:30am	Continental Breakfast.	*(Airlie Room)*
8:30–10:00am	**Discussion and report planning**	*(East Room)*
10:00–10:15am	Break	*(2nd Floor, Airlie House)*
10:15am–12:15pm	**Writing breakout groups**	*(East, South, West, North Rooms)*
12:15–12:45pm	**Wrap up, homework, and report completion schedule**	*(East Room)*
12:45–1:45pm	Lunch	*(Airlie Room)*

BREAKOUT GROUPS (AND WRITING GROUPS):

• Breakout Group 1:	**Grand Challenges**	*(East Room)*
• Breakout Group 2:	**Network Science**	*(South Room)*
• Breakout Group 3:	**Vision of the Future: New Applications and Paradigms**	*(West Room)*
• Breakout Group 4:	**Whence we came from and where we go**	*(North Room)*

Breakout Group 1: Grand Challenges

Using the panel session as a point of departure, this group formulates a fundamental research agenda for grand challenge problems. The breakout group can add to (or remove from) the grand challenge problems discussed in the panel. Ideally, the grand challenge problems cover a wide range of activities in networking research, as opposed to a laundry list of people's favorite research topics.

Suggested topics:

- What qualifies a particular problem as a grand challenge problem (high impact, currently intractable, solution can be advanced by fundamental research)?
- What have been some grand challenge problems that have been solved in the past and what impact have they had?
- What are the open grand challenge problems that we believe will have far reaching implications for the future of communications networking?
- How can grand challenge problems help to evolve and transform the networking research program?
- What is the time horizon for each of the grand challenges? How do we measure success?
- What are the research opportunities? What are the opportunity costs for not investigating these problems?
- How can the grand challenge programs impact other research areas, industry, and society at large?
- What must be done to enable or facilitate a grand challenge problem that is being investigated?

Breakout Group 2: Network Science

This group lays out an agenda for developing a Network Science, that has is as important and necessary as physics, astronomy, etc. A Network Science must necessarily borrow from many disciplines, control theory, statistics, algorithms, communications theory, signal processing. The development of a Network Science would result in a better understanding of how current complex systems work. Further, and more importantly, its goal would be to provide clear guidelines on how to better design and control networks. Multidisciplinary projects should be encouraged as long as the networking science objective is clearly identified and networking expertise is at the core.

Suggested topics:

- What can constitute and justify the development an Network Science? Are there laws that describe the various interactions that take place within the network? Are they backed by experiment?
- Is there a theory to guide the construction of stable control loops in a network where networks are layered on top of other networks, e.g., content delivery networks over traditional networks.
- There is some understanding of how to design a system built out of reasonably homogeneous components. How about the construction of an Internet consisting of heterogeneous components?
- How to model large, heterogeneous networks? How to develop such models? How to interface them with measurements?
- Is there a network information theory (perhaps one that includes delay)? In mobile networks, particularly, in sensor networks, the development of a network information theory has already begun. Is there a theory that one can develop and apply to other network management functions?
- Is there a theory to describe the evolution of networks? Such a theory ought to relate the introduction of disruptive applications such as the web and P2P to the growth in bandwidth and introduction of new technologies.
- Is there a systematic science that can be developed to model the interactions that take place over several layers and time-scales? Will such a theory help us significantly develop better networks.

Breakout Group 3: Vision of the Future: New Applications and Paradigms

This group continues the discussions of the panel and formulates a vision for network applications and paradigms in the future. The time horizon can be from 5 to 50 years.

Suggested topics:

- What role does fundamental networking research play in these applications?
- How will networking research change in the light of these applications? Which methodologies will be needed?
- How will these applications transform society?
- To what degree are new applications and paradigms enabled by networking research?
- Which new fields in networking research will emerge as a result of these applications?

Breakout Group 4: Whence we came from and where we go

The goal of this group is to describe the growth and evolution of the networking research community over the past several years, describe the current state of the area, and its outlook for the future. The group should document the contributions and accomplishments of basic networking research and discuss the development of the community. The group may want to use the report of the 1994 workshop (on the web) as a starting point.

The workshop announcement has stated that, at present, there is hardly a researcher in computer science or electrical engineering, who does not, directly or indirectly, include networking aspects in his/her research agenda. This has raised the relevance and importance of the field.

Possible topics:

- What are the major accomplishments and contributions of networking research?
- What has been the impact of this research on other research areas?
- How has networking research, and especially, fundamental research, evolved?
- Which new fields in networking have evolved in the last 10 years (security, wireless, active networks)?
- Which new areas are likely to emerge in the next 5, 10, 20 years?
- How has networking research influenced research in other areas?
- How have the methodologies of networking research changed or evolved?
- What strategies should we use to convince congress that growth in research funding for networking research is imperative. For example, networking continues to be vital to the U.S. (world) economy and that advances in research will spur significant economic growth?