Network Science and Engineering (NetSE) Research Agenda

A Report of the Network Science and Engineering Council
Release Version 1.1
September 2009

NetSE Research Agenda: Executive Summary and Recommendations

Over the past forty years, computer networks, and especially the Internet, have gone from research curiosity to fundamental infrastructure. In terms of societal impact, the Internet has changed the way we live, work and play, and altered our notions of democracy, education, healthcare, entertainment and commerce. In terms of its design, the Internet has shown a remarkable ability to adapt to, even inspire, changes in technologies and applications. In short, the Internet has been a powerful engine for technological innovation and societal evolution.

However, this is no time to rest on the successes of the past. To meet society's future requirements and expectations, networks in general, and the Internet in particular, will need to be better: more secure, more accessible, more predictable, and more reliable.

In 2008, the Computing Community Consortium (CCC) charged the Network Science and Engineering (NetSE) Council with developing a comprehensive research agenda that would support the development of better networks. The NetSE Council was to consider previous reports such as those produced by the Global Environment for Network Innovation (GENI) Science Council, as well as encourage new interdisciplinary participation. Over the summer and fall of 2008, the NetSE Council held a number of disciplinary and interdisciplinary workshops that, together with several GENI and pre-GENI workshops and documents, resulted in the network science and engineering research agenda detailed in this report. The NetSE-sponsored interdisciplinary workshops were structured to bring participants from closely related fields together with networking researchers to explore problems and opportunities in the intersection. The diversity of backgrounds of the workshop participants highlights the breadth of the intellectual space.

The work of the NetSE Council makes clear the following:

- Now is the time for the nation to significantly step up investment in the research that will lead to a better Internet, and improve network understanding and technology more generally. The urgency of this work derives from our society's heavy reliance on the Internet, clear and present threats to the current infrastructure, and competition from other nations making investments in the future Internet. Dramatic progress towards a better Internet is possible over the next decade, but significant effort will be required to achieve it.

  Recommendation 1: The funding agencies of the United States government must increase investment in research that will lead to a better Internet or risk a marginal future role.

- A key barrier to making near-term progress is the lack of adequate support for experimental research. The field of networking research began as an experimental discipline, and many fundamental networking insights have come from building and using prototype systems, as well as measuring existing systems. However, current

funding levels and funding consistency are inadequate for most research groups to build real prototypes, and there are no facilities for experimentation at realistic network speeds and scales. Measuring current systems requires access and technologies that are difficult for researchers to obtain on their own.

Recommendation 2: Funding agencies should rebuild the experimental capabilities of networking researchers, through funding individual systems-building efforts, providing adequate and persistent shared experimental infrastructure, and supporting research that leads to continued improvements in experimental methodology. Experimental work is expensive and long-term; typical NSF awards are insufficient, therefore either NSF will need to change its award portfolio or other agencies will have to play a significantly increased role.

- A key barrier to making long-term progress is the lack of a formal intellectual framework for understanding networking. Within computer science there is an effort to develop a "Theory of Networked Computing" and the nascent interdisciplinary field of "Network Science" provides a promising new theoretical perspective. These and other formal efforts will be critical in creating a fundamental science of networking that has the potential to underpin systematic engineering methodologies for the design of large-scale, robust, cost-effective and evolvable networked systems. Experimental and measurement-based research has a key role to play in the development of this discipline by providing opportunities to test and refine hypotheses, and to validate systematic design methodologies.

  Recommendation 3: Funding agencies should foster and support research activities relevant to network design within the theoretical computer science community, the new Network Science community, and other theoretical disciplines.

- The Internet is now as much a social phenomenon as a technological artifact, and thus its study requires a broadly interdisciplinary approach. Legal, social, or economic factors cannot be ignored when doing networking research; these other perspectives are crucial to the future of networking.

  Recommendation 4: Funding agencies should support a broad array of interdisciplinary research activities related to understanding the current Internet and designing future networks to include the Internet.

The remainder of this report consists of a (1) Synthesis Chapter that describes in more detail the imperative, challenge, opportunity and elements of a research agenda to improve the Internet and deepen our understanding of socio-technical networks and (2) the reports of six Workshops in Overcoming Barriers to Disruptive Innovation in Networking; Theory of Networked Computing; Science of Network Design; Behavior, Computation and Networks in Human Subject Experimentation; Network Design and Engineering; and Network Design and Societal Values.

The audience for the remainder of the report is the computing research community, the multi-disciplinary NetSE research community, and the funding agencies that support their work. We hope that researchers and funding agencies will use this report as a framework and foundation for pushing ahead with the various research agendas outlined herein.

Many people have been deeply involved in this effort over the last year. The NetSE Council would like to thank Jeannette Wing for recognizing the need for a NetSE effort and for convening the kick-off meeting at NSF that launched the work. We would like to thank the CCC, especially Ed Lazowska and Susan Graham, for their continuing leadership and strong support for this effort. We thank all the workshop co-chairs and participants for their important contributions to the intellectual conversation and for capturing the collective wisdom in the workshop reports that are appendices to this research agenda. Finally, we would like to thank David Clark, Helen Nissenbaum, Jennifer Rexford, Scott Shenker, John Wroclawski from the NetSE Council, Suzi Iacono from NSF, and several CCC members (notably David Tennenhouse, Anita Jones and Fred Schneider) for their collegial and constructive input during the course of the writing and especially for thinking about the format and tone of the synthesis chapter.

Ellen W. Zegura
Georgia Institute of Technology
20 July 2009

NetSE Council:

Charlie Catlett, Argonne National Labs
David Clark, MIT
Mike Dahlin, University of Texas - Austin
Chip Elliott, BBN Technologies and GPO (ex-officio)
Joan Feigenbaum, Yale University
Stephanie Forrest, University of New Mexico
Mike Kearns, University of Pennsylvania
Ed Lazowska, University of Washington
Helen Nissenbaum, New York University
Larry Peterson, Princeton University
Jennifer Rexford, Princeton University
Scott Shenker, UC Berkeley/ICSI
John Wroclawski, USC/ISI
Ellen Zegura, Georgia Institute of Technology

NetSE Research Agenda: Synthesis Chapter

The Imperative: A Better Internet

The Internet is the 21st century's fundamental societal infrastructure, comparable to the railways of the 1800s and roadways of the 1900s. The current Internet has been a remarkable success, providing a platform for innovation that far exceeds its original vision as a research instrument. The Internet and associated services have transformed the lives of billions of people in areas as diverse as democracy, education, healthcare, entertainment, commerce, finance and civil infrastructure. Yet these successes are seriously threatened by the increasing sophistication of security attacks and the rogue organizations that propagate them. A materially more secure Internet would be better. Indeed, it is vital to the advancement of society.

Societal reliance on the Internet for critical functions is increasingly disproportionate to the ability of the Internet to deliver high dependability. The Internet usually works pretty well, but every user has experienced inexplicable periods of degraded performance or outright non-function. The current Internet provides no visibility to end users and shockingly little visibility to network managers and operators to support understanding, adapting to and fixing reliability problems. Such limitations require lay people to spend their leisure time debugging problems with their home networks and companies to spend heavily on network operations. Further, the lack of performance reliability prevents the Internet from advancing to become a truly dependable, critical infrastructure. A more dependable, reliable and predictable Internet would be better.

The Internet implements a best-effort, point-to-point service model, well suited to applications between two easy-to-specify endpoints that can tolerate occasional performance degradation. The service model is strained for applications involving multiple endpoints, where identifying the endpoint(s) is difficult, and/or where performance degradation is unacceptable. Workarounds abound, but they are generally inefficient and fragile. A deep consideration of alternative service models and their associated economics would be better.

The Internet embeds societal values in ways that are often implicit and not well understood. For example, the Internet is "open", frequently taken to mean that anyone can join the network by implementing the public protocol IP. In principle, users can run any application on the Internet, without arbitrary limitation imposed by the network protocols. This openness enables organic growth but is not accompanied by mechanisms to vet participation, in particular those that threaten harm. Issues of trust and individual accountability are ongoing sources of concern, and the question of how to best address them remains not only technically difficult but also highly contested. An Internet that appropriately reflects individual and societal considerations in its design and engineering would be better.

Billions of people remain untouched by the advantages of the Internet; Internet World Statistics puts worldwide average Internet penetration at about 22% in mid-year 2008.

However, the reach of the Internet is accelerating as low-cost electronics and communication technologies, such as smart phones, are adopted everywhere, but especially in the developing world. There is, simply put, no scientific or engineering basis for assuming that today's Internet can support such growth. An Internet that effectively and affordably includes the other 78% of the world population would be better.

Interest in and demand for "a better Internet" extends from the most well connected regions of the world to the least; from scientific laboratories to government offices to corporate boardrooms to universities to homes and huts. Significant large-scale efforts to improve the Internet are underway across the world, in the European Union (FIRE program), Germany (G-LAB), and Japan (JGN2plus+Akari), to name a few. The United States must rise to the challenge with comparable investment or risk a marginal role in the future Internet. The US government was early to the table by funding efforts such as the original Internet and GENI, but others are now charging ahead.

The Challenge and the Opportunity

What makes it hard to design a better Internet? The Internet is unique. It is a federation of tens of thousands of networks connected by the IP protocol and distributed around the world. It is among the largest human-constructed technical artifacts. The Internet was, and continues to be, built in a distributed fashion, without centralized control. The Internet contains fundamental concepts of economic agency that distinguish it from traditional distributed computing: different parts of the Internet have different owners with different economic motivations, yet they cooperate to provide end-to-end service. The Internet spans a myriad of legal jurisdictions with vastly different laws, conventions and practices. The technologies available to build networks are constantly changing, as are the uses that humans dream up and invent. The list of desired properties is long, diverse, potentially conflicting, and personal.

The design and innovation process is further challenged by barriers to evaluating new ideas. In evaluation, the research and development community currently relies on analysis, simulation, small-scale experiments and larger over-the-Internet experiments. Each has its strengths, but also its significant limitations – analysis is often intractable; simulation can provide scale but generally insufficient realism; small-scale experiments provide realism on some dimensions but are unrealistic on the scale dimension. Larger experiments on the current Internet may be useful for evaluating new ideas that are compatible with the current Internet, but are less useful for evaluating ideas that depart significantly from current practice. All three of the current evaluation methods – analysis, simulation and experimentation – have room and requirement for significant intellectual advances, and at their best are used in complementary fashion to deepen understanding of the design problem as well as refine specific candidate solutions.

Lastly, even if the evaluation of an idea finds it promising along all dimensions, there is great difficulty in getting it deployed. In deployment, one must contend with the economic, social and political issues that accompany any new technology, but these are

more acute for a technology with as much heterogeneity and reach as the Internet. Some innovations are amenable to incremental deployment, but others only provide advantage with large-scale adoption. Researchers have had a complicated relationship with deployment issues – too much focus on ease-of-deployment may stifle innovative thinking while too little may promote work that is divorced from any current or future reality.

The challenge and opportunity, then, is critically important: to design, evaluate, and deploy new network technologies and services to operate in a complex and changing socio-technical-economic-political environment, while enabling the full range of societal and individual aspirations: in short, to realize better networks.

What Is Needed Now?

Now is the time for the nation to significantly step up investment in the research that will lead to better networks. The urgency of this work derives from threats to the current infrastructure as well as competition from other nations making investments in the future Internet.

What will it take to act on this challenge? Consider the three phases – design, evaluate and deploy.

Design. Designing a better Internet requires (1) identifying and articulating design goals and (2) engineering the appropriate components of a system that meets selected, achievable goals, based on reasoned and well articulated principles. Neither step is easy. High level goals such as "more secure" must be translated to specific, unambiguous constraints and capabilities that can be realized (or not) by specific system designs. The scientific community has considerable experience formalizing computational concepts that are easily quantified, but less experience formalizing human-centric goals. For example, we can formalize the computational difficulty associated with breaking an encryption scheme, but we do not yet know how to specify what it means for a system to have "sufficient privacy". We must increase the interaction between network technologists and social scientists to make progress on critical human-centered design goals.

Any discussion about goals quickly produces a long and formidable list. New concepts of feasibility, resources, and efficiency are needed to enhance and deepen our understanding of the power and limitations of networked environments. Some goals will be unachievable, by any system, even when considered in isolation. Some goals will operate in tension with one another, requiring an understanding of tradeoffs. The value placed on specific goals will differ among users, institutions and national governments. These last two observations point to the importance of understanding concepts such as universality and customization. We must address questions such as: What properties are and should be universal? What properties can and should be customized? What are the interfaces to customization and who has access to them? How is overall system stability assured when customization takes place?

The set of goals must include an explicit recognition of the unique environment in which the Internet operates, especially the fact that the requirements, uses and technologies change over time. The goal of evolvability – for an appropriate definition – must have first class status. Is "an Internet for the next 100 years" a grand challenge, too modest, or the wrong question? Thus, the first step in meeting the grand challenge is to understand and choose the set of goals.

Given a set of goals, the design process develops the appropriate components of a system to meet the goals. For any reasonable set of goals, this is a formidable challenge. Much of the research activity of the community is devoted to finding designs that meet various sets of goals, and these design activities will undoubtedly play the central role in acting on the challenge. As such, their continued support is essential.

One nascent but potentially promising direction in network design is the development of a new methodology centered on theoretically-derived architectures (components and their interfaces) with structuring principles derived from rigorous underlying theory and capable of adapting to unforeseen changes. That is, a science of network design to complement the traditional empirical design process.

Evaluate. We are on the cusp of an explosion in opportunities for evaluation, as measurement and computational capability provide the raw material for better models, greater fidelity, near real-time and real-scale simulations, and in-situ experimentation in testbeds of increasing size and realism (including the current Internet as testbed). Effort and balance are needed across all three legs of the evaluation stool – analysis, simulation and experimentation. Over time, the balance may change, but currently all three are valuable aspects of the networking design and evaluation process.

Taking advantage of these opportunities requires learning how to use evaluation tools to reach rigorous conclusions. To illustrate, the following questions are among those that must be addressed: What network models and abstractions should be used in simulations? How can one extrapolate the evaluation results obtained on smaller networks to predictions about larger networks? What is the meaning of worst-case analysis in a networking context? How should network users be modeled or otherwise included in testing?

Ultimately, the evaluation of any engineering artifact depends on its construction, deployment in prototype and eventually final form, and use. Realistic experimentation requires more than intellectual advances of the sort well served by traditional small-scale research grants: researchers must be able to build and study real prototypes and deployed systems. Unfortunately, most researchers do not have the resources to build full implementations of their designs, and have no facilities where they can experiment with those implementations on real networks, at significant scale, under reasonably realistic conditions. Without functioning instantiations of new designs, evaluation remains speculation, and firm conclusions are difficult to reach.

Deploy. The current Internet provides both a barrier to and an enabler of deployment of new technologies. It is hard to change the aspects of the network for which there is no economic or consumer stakeholder who directly benefits on a workable time scale. It is astonishingly easy to change aspects of the network that are amenable to incremental end-system deployment. One need only look at the deployment path for computer viruses or file sharing networks such as BitTorrent to be convinced of this. It is moderately easy to change aspects of the network that require a longer-term investment yet have a solid longer-term economic benefit to the investor, such as improvements in edge technologies.

We must develop a better understanding of deployment incentives, enablers, disincentives, disruptive transitions, and so on in large complex systems. This requires interaction among network technologists, economists, technology historians, and experts in business modeling. Of special importance is understanding and leveraging government investment that can facilitate the disruptive transitions that would not occur in a pure economy.

While some progress can be made independently in the design-evaluate-deploy phases, the most significant gains will come by acknowledging and leveraging the complex interactions. To wit, an initial idea is often refined during evaluation. An idea tested in the lab may be subjected to an initial deployment study to provide additional feedback in a real setting. A well-tested idea may be modified after deployment in controlled and uncontrolled ways. Occasionally a technology "escapes the lab" during testing and self-deploys. Of course there are also leaks in the process – many ideas are discarded during testing; a well-conceived and tested idea may never see deployment, for myriad reasons.

These are the challenges, and the community is poised to act. The time is right to advance and link across the design-evaluate-deploy challenges to realize better networks.

The Network Science and Engineering (NetSE) Effort

In January 2005 NSF held a workshop on "Overcoming Barriers to Disruptive Innovation in Networking" (OBDI) that recognized the need for a better Internet, and the many barriers to achieving that goal. The report particularly highlighted the need for a renewed emphasis on experimental networking research. Discussions subsequent to the workshop eventually led to the proposal for the GENI experimental infrastructure. The GENI Science Council was convened to define a Science Plan for GENI. The resulting science plan focused on networking problems that would require such an experimental infrastructure.

However, there is a much larger intellectual agenda related to networking, and the NSF created the Network Science and Engineering (NetSE) Council to articulate this broader research agenda by incorporating and extending the work of the GENI Science Council. The NetSE Council kicked off their effort in January 2008 with a meeting at NSF. In Summer and early Fall 2008, the Council ran five events focused on various parts of the NetSE design space:

June 11, 2008: Theory of Networked Computing (ToNC), Boston, MA
Organizer: Joan Feigenbaum, Yale

July 29, 30, 2008: Science of Network Design (WNSND), USC/ISI
Organizers: John Doyle, Cal Tech; John Wroclawski, USC/ISI

July 31, August 1, 2008: Behavior, Computation and Networks in Human Subject Experimentation (WBCN), Del Mar, CA
Organizers: Michael Kearns, Penn; Colin Camerer, Cal Tech

August 17, 18, 2008: Network Design and Engineering (WNDE), Seattle, WA
Organizers: Jennifer Rexford, Princeton; Ellen Zegura, Georgia Tech

September 24, 25, 2008: Network Design and Societal Values (WNDSV), Arlington, VA
Organizers: David Clark, MIT; Helen Nissenbaum, NYU

The Appendices contain the reports from the 2005 workshop and each of the five NetSE-sponsored meetings. The next section of this introductory chapter summarizes the 2005 workshop that started the community on this path, and the rest of this introductory chapter attempts to synthesize themes that emerged across the meetings and workshop by discussing four areas where research and innovation is necessary to act on the scientific challenge of designing better networks: challenges in network experimentation, new mathematical tools and frameworks, new disciplinary innovations, and new interdisciplinary conversations. We follow these with an exemplar – Security – that appeared in every meeting. The chapter concludes with a summary of what is needed to advance the agenda to realize better networks.

Each of the NetSE workshops was also attended by a representative from the GENI Project Office (GPO) with the aim of extracting requirements for the current GPO-led effort to develop an experimental research platform. A companion document – GENI System Requirements – describes the experimental requirements that the GPO group gleaned from these workshops and from the previous set of GENI-related documents. An important conclusion of our report is that requirements and objectives for experimental infrastructure be constantly reviewed as the NetSE research agenda evolves and experimental methodologies advance.

Challenges in Network Experimentation (See all Appendices, esp. OBDI)

In their attempts to move towards a "better" Internet, researchers should not be limited to designs that are small deviations from the current Internet, and which are consequently easy to build and experimentally deploy. To the contrary, it is imperative that researchers also pursue disruptive network architectures. However, paper designs, although thought provoking, are unconvincing, both to the companies that need to adopt them, and to the research community in evaluating ideas and in gaining insight into design tradeoffs. Thus, new architectures must be realized in functioning prototypes, and need to be evaluated experimentally, operating at scale, and under real-world conditions. This will entail building substantial systems, and then evaluating them in some experimental environment. Currently researchers are unable to do either of these.

- Traditional NSF support, while suitable for paper designs, is inadequate for building substantial systems.
- There is no shared infrastructure for testing designs under anything remotely resembling real-world conditions in either scale, speed, or traffic load.

Even if these barriers were overcome, challenges remain, as the design of good experiments can be as challenging as the design of good systems. Experimental research challenges include:

- Extension of real-world evaluation beyond the deployment-study model to include also a more structured and rigorous worst-case analysis model. As recent events in the financial sector have revealed, the worst case can happen, and models that fail to capture worst-case scenarios do not serve society well.
- Developing discipline and methodologies to validate models, measurements and data and to evaluate candidate technologies on testbeds, including those at universities and in industry, for example including WAN backbones, smart power grids and mobile sensor networks.
- Developing an experimental infrastructure for conducting large-scale behavioral experiments for situations in which underlying network structure, such as the number of immediate neighbors, strongly governs human interaction and strategy.
- Understanding the relationship of researchers and users in large-scale experimentation, including the role of consent in dynamic, unpredictable environments.
- Developing strategies for dealing with measurement data, including standards for collection, archiving and use. In addition, developing strategies for dealing with situations where there is a dearth of accurate and complete data as well as situations where the volume of data is overwhelming.

New Mathematical Tools and Frameworks (See Appendices WNSND and ToNC)

Realizing better networks requires a foundational mathematical theory of network architecture and design. Progress in this domain is essential to address a deeply troubling conundrum: Modern network technology promises to provide unprecedented levels of performance, efficiency, sustainability, and robustness across numerous domains of science, engineering and society. In many areas, it has already done so. Yet, as network scale, ubiquity and deployment grow, the problem of rare but catastrophic real-world failures and "unintended consequences" is, if anything, becoming worse. This "robust yet fragile" nature of complex systems is ubiquitous and fundamental. Foundational research, in the form of a theoretical framework to manage the complexity/fragility spirals of future network infrastructures, is essential to bring networked systems to the next level of reliability and robustness and fulfill the promise of networks as fundamental infrastructures of society.

Recently, study and comparison of networks across widely separated domains – ranging from the Internet to systems biology – has raised the possibility that a set of common sub-elements and abstractions might be found that together capture the essence of broad classes of networks. Chief among these are abstractions for network topology, for traffic or information generation, flow, and use, for the networked system's functional structuring and modularity, and for the control mechanisms that keep the network operating robustly and efficiently. While it is clear that the relative importance of individual elements within this set will differ for networks in different domains, a unifying theoretical framework that encompasses, builds on and integrates all simultaneously would have tremendously broad applicability. The development of this framework can be viewed as a "holy grail" of long-term network science research.

Key research questions in this area include:

- Developing a theory of *networked computing* by: formulating the definition(s) that a computational system must satisfy if it is to be called a "network"; identifying critical resources and bounding the consumption of those resources for networked computation to be considered "efficient"; formulating notions of "reduction" to prove that one networked-computational problem is as hard as another.
- Developing the foundations of a theory of *network architecture* that allows rigorous analysis and systematic design of complex networked systems, including organizational abstractions such as interfaces and layering.
- Developing new *network protocol design methods* to implement the protocol-related elements of networked systems, for applications such as cross-layer control in wireless networks, distributed data gathering, integrated network coding, and communication platforms for control of cyber-physical infrastructure.
- Developing and evangelizing a common *mathematical language* to broaden and deepen the contacts between and across networks in engineered systems and networks in the sciences, particularly biology.

New Disciplinary Innovations (See Appendices WNDE and WNSND)

Network design differs from classical engineering design. In classical engineering, the goal of the design is to produce an artifact given a set of requirements. In large-scale network design, there is no single entity that carefully plans and executes a design. Further, the design requirements change over the lifetime of the artifact. The design of a large-scale network instead emerges from conditions and decisions at micro-levels (e.g., within an ISP) to produce macro-level behavior (e.g., a robust and well-performing overall Internet). The conditions and decisions experienced at the micro-level are subject to change over time.

While satisfying any one design goal in isolation may be relatively straightforward, the underlying challenge in network design lies in reconciling trade-offs between seemingly conflicting design goals – for example, a secure network that enables innovation, or a reliable network at reasonable cost.

Key research questions in this area include:

- Creating a *new class of design methodologies and principles* that is concerned with steering collective, emergent behavior over time, rather than producing a final artifact.

- Understanding *tradeoffs and optimizations in the network design space* among multiple, possibly conflicting design goals and societal values, so that network designers can make tradeoffs explicit and network users can make appropriate choices within the tradeoff space. What regions of the design space are achievable? What tradeoffs are fundamental? What tradeoffs can be avoided with clever design? What tradeoffs can be exposed to users to support individual decision making?

- Revisiting *protocol layering* in the context of increasingly diverse applications, availability of programmability as a mechanism for innovation and customization, and recognition that network functions such as management are ill-served by strict layering.

- Appropriately reflecting *individual and societal considerations (including moral and political)* deeply in network design and engineering. For example, what are the appropriate approaches for defining identity? What values are at stake in alternative technical choices? Can and should application layer solutions such as reputation systems be embedded in a general network architecture? What sorts of collaborations are effective for exploring the technical-social space of identifiers?

New Interdisciplinary Conversations (See Appendices WNDSV, WBCN, ToNC)

As the security exemplar highlights, designing better networks will require network scientists and engineers to connect deeply to other disciplines including the social sciences. Interdisciplinary conversations are not yet commonplace for network domain researchers, though legal and economic considerations are, of course, dominant in the network service and technology industries. Network designers will need to increase their collaboration with economists, political scientists, behavioral scientists, philosophers and lawyers.

Network architectures and protocols embed assumptions about values, sometimes explicitly and often implicitly. The current Internet makes a strong assertion about the value of openness; in theory, anyone is free to implement on his or her own machine the open protocols that enable connection to the Internet. The current Internet makes an assumption that the appropriate policy for allocating bandwidth among competing connections is that implemented for technical reasons by TCP. Interestingly, it has taken many years and much research effort to even quantify the sharing policy that TCP implements; clearly it isn't explicit.

Key research questions in this area include:

- Developing an understanding of general moral and political obligations of and to various actors in a networked setting. Although traditional modes of analysis form an essential backdrop, they need to be adapted to the distinctive features of networked community and interaction.

- Developing workable definitions for security and privacy that allow progress on questions such as: what are the tradeoffs among security, privacy, and other values such as free speech? How does the value placed on security differ among users, institutions and national governments? How do we understand the type of privacy we want or need?

- Understanding how we can incorporate social values in the design of networks to promote and sustain sociality. Do social networks, based on voluntary associations among users, point toward a model for trust networks more generally? Is the trust that pervades networks durable over time and across platforms?

- Understanding how to develop, verify and evaluate incentive-compatible algorithms in the presence of rapidly changing or contradictory incentives and/or economically irrational actors.

- Unifying algorithmic and behavioral game theory so that behavioral models take into account computational considerations and algorithmic models are evaluated and improved based on experimental evidence.

Exemplar: A More Secure Internet

To illustrate the challenge of a better Internet and the readiness of the community to rise to it, we elaborate on the challenge of a more secure Internet. The security problems of the current Internet are the greatest impediment to its future. Security threats and attacks are increasingly rendering the Internet unusable for certain types of transactions, and if unchecked will threaten the development of future services. Further, security is not a single, identifiable "problem" to be solved once for all time. It is a collection of immediate and longer-term challenges that will continually change as the Internet and the security concerns it faces evolve.

Successfully addressing the security challenge is a formidable task that cuts across many fields of inquiry, including mathematical science, systems design, human-network interface design, economics, and political, social and ethical analysis. Potential solutions will surely leverage the formalisms and proofs that theory can provide, together with careful systems design, usable interfaces, conceptual clarification, and experiments in the wild that include human subjects. These approaches will all be needed to test the promise of proposed solutions.

Security is a science problem. The mathematical foundations of computer and communications security are ancient, dating back to the earliest human interest in sending secret messages. Cryptography provides the mathematical basis for protecting information from unauthorized use. Advances in cryptography have provided more provably secure algorithms, more flexible methods for encoding and decoding such as public-key cryptography, and impressive, powerful methods for using data without revealing it. Because security of a computational system can never be demonstrated experimentally (as opposed to insecurity, which is demonstrated every time an attack succeeds), formal specifications, analysis, and proofs are invaluable in the security world.

Security is a systems and engineering problem. Mathematical tools, while necessary, are insufficient for constructing secure systems. Security is an emergent property derived from the composition of system components, as well as the components themselves. Systems designers and engineers must decide the architecture of the system: what are the components, their functionality and their interfaces? Systems designers and engineers must navigate tradeoffs in the design space, aiming for the best possible balance when design goals conflict, or eliminating the conflict completely through technical advance. Stronger and deeper understanding of principles and paradigms to support system architecture decisions would be invaluable to addressing the security challenge.

Security is a network architecture problem. Network architecture provides opportunities to improve security by providing: (1) better control mechanisms for the transmission, modification and receipt of data, as might be used to thwart network-level denial of service attacks, (2) better mechanisms for attribution of past, present and future actions, and (3) more semantic transparency so that, for example, an intermediate network device could observe a packet stream and unambiguously understand protocol and end-host behavior. Changes in network architecture are neither necessary nor sufficient to solve

the security problem; however, appropriate changes can make a substantial difference in the problem and solution space.

Security is an economics problem. Security is ultimately about risk management, and decisions about "how much" security to employ must weigh the costs of providing security against the costs associated with breaches. Further, today's adversaries are sophisticated and well-funded as a result of a thriving marketplace built on illicit services ranging from bulk harvesting of identity information to remote control of millions of personal computers. Understanding the economic mechanisms and motivations of attackers is critical to understanding the frontiers of the security battlefield.

Security is a human problem. Humans are often called out as the weak link in secure systems, with their persistently poor security practices. This observation draws attention to the dire need for research into system design that reduces or eliminates tradeoffs between security and usability, as well as those between security and innovation. Human factors research also needs to consider approaches to making security issues and implications easy for lay people to understand. Advances in usable security are just beginning to show up on the HCI and security landscape; much more is needed. As Internet usage continues to spread, researchers must also address the needs and practices of diverse communities of users with diverse cultural practices and norms.

Security is a political problem. For political scientists and philosophers, security is an important and controversial concept. It calls into consideration the harms people and societies deserve to be protected against, and at what cost. Pursuit of security for some may leave others vulnerable; pursuit of security may also conflict with pursuit of other values such as freedom or privacy. In particular, it is relatively easy to design a closed, secure system. It is qualitatively much more difficult to design a secure system that is open, efficient, cost-effective, and preserves individual privacy. We are only beginning to understand the fundamental tradeoffs in this space as well as mechanisms to allow navigation of tradeoffs either system wide or by individuals.

Security requires experimentation. Because of the complicated component interactions and role of human behavior, experimentation is critical for developing better security solutions.

Summary

The Internet is remarkable, but it will need to be better. There are significant challenges that are emerging as a result of success in scale, capability, and the generative nature of the Internet. To meet society's requirements and expectations, to continue to inspire innovation, creativity and democratic engagement, the Internet will need to be materially more secure, reach all of humanity cost-effectively, realize performance, reliability, and predictability commensurate with societal reliance, and be adaptable to unforeseen changes in technologies, applications, and human behavior.

Our report calls out four important elements central to meeting this goal.

- First, we note the vibrant state of the research community and the research itself. NSF's networking, cybersecurity, and cyberphysical systems research programs, particularly the "clean-slate" FIND activity, have drawn strong interest and top-quality proposals from their respective communities. Similar activities in Europe, Japan, China, and Korea indicate that this interest is world-wide. The broad research agenda described across the NetSE Workshop reports and previous NSF activities outlines a compelling case for advance in the field.

- Building understanding is a key precursor to better design. To derive this benefit we must acquire a more fundamental understanding of the science underlying network design. Network Science and Engineering as a field will be greatly aided by the creation of an underlying discipline that can systematically formalize and explain large-scale, complex networked systems. Efforts in this direction have already begun, but there is much more to be done.

- Better designs and deeper understanding can only be achieved if a broader set of disciplinary perspectives are considered. Networking is no longer solely the province of system designers; legal, economic, and social considerations must play a crucial role.

- Evaluation is a key step in any engineering research process, but is particularly and crucially challenging in the case of networking research's complex, multidimensional problem space. Networking research depends on multiple models of evaluation – experiment, simulation, analysis, measurement – used appropriately and synergistically. To fully foster emerging research, researchers must have access to the full range of evaluation tools justified by the intellectual agenda.

This last requirement poses a significant challenge to progress, in light of the overwhelming disparity in cost between experimental evaluation and other forms of evaluation. Nonetheless, it is *crucial* that researchers be able to build realizations of their ideas and test them experimentally when required. The field is developing promising approaches to building flexible experimental environments that allow ideas to be tested under current and potential future scenarios, as well as allow ideas to move from lab settings into production settings. It is imperative that these activities be adequately supported, and that experimental environments be provided.

With grand challenges ahead, and new experimental, scientific, and interdisciplinary approaches in hand, networking as a field stands at the threshold of a great opportunity.

Appendix 1

Workshop Report on

Theory of Networked Computation

Theory of Networked Computing

The increasing prominence of the Internet, the Web, and large data networks in general has profoundly affected social and commercial activity. It has also wrought one of the most profound shifts in Computer Science since its inception. Traditionally, Computer-Science research focused primarily on understanding how best to design, build, analyze, and program computers. Research focus has now shifted to the question of how best to design, build, analyze, and operate networks. How can one ensure that a network created and used by many autonomous organizations and individuals functions properly, respects the rights of users, and exploits its vast shared resources fully and fairly?

The SIGACT community can help address the full spectrum of research questions implicit in this grand challenge by developing a Theory of Networked Computation (ToNC), encompassing both positive and negative results. Algorithms and complexity-theory research has already evolved with and influenced the growth of the Web, producing interesting results and techniques in diverse problem domains, including search and information retrieval, network protocols, error correction, Internet-based auctions, and security. A more general Theory of Networked Computation could influence the development of new networked systems, just as formal notions of "efficient solutions" and "hardness" have influenced system development for single machines. To develop a full-fledged Theory of Networked Computation, researchers will build on past achievements both by striking out in new research directions and by continuing along established directions.

The SIGACT community has identified three broad, overlapping categories of ToNC-research goals:

- **Realizing better networks:** Numerous theoretical-research questions will arise in the design, analysis, implementation, deployment, operation, and modification of future networks.
- **Computing on networks:** Formal computational models of future networks will enable us both to design services, algorithms, and protocols with provable properties and to demonstrate (by proving hardness results) that some networked-computational goals are unattainable.
- **Solving problems that are created or exacerbated by networks:** Not all of the ToNC-research agenda will involve new computational models. The importance of several established theoretical-research areas has risen dramatically as the use of networked computers has proliferated, and some established methods and techniques within these areas are not general or scalable enough to handle the problems that future networks will create. Examples of these areas include massive-data-set algorithmics, error-correcting codes, and random-graph models.

CISE's NetSE program (http://www.nsf.gov/pubs/2008/nsf08578/nsf08578.htm) welcomes proposals in all three categories. For more details about the ToNC-research agenda, see http://www.cs.yale.edu/homes/jf/ToNC.html.

Definitions and Models of Networked Computation

A broad range of theoretical research questions is likely to arise in the design, analysis, implementation, deployment, operation, and modification of future networks. Given our limited ability to model, measure, predict, and control today's Internet, we will need a more principled approach if we are to realize the ambitious goals now under discussion. What are the right primitives and abstractions with which to study networks? How should responsibility for essential network functions be assigned to various network components? How should state be allocated among components? What should the relationships be among naming, addressing, and routing; indeed, which objects in the network should have names that are meaningful network-wide?

In the systems-research community, these questions are representative of "network-architecture" research. From the SIGACT-community perspective, this type of question must be answered in the process of formally defining various types of networks and rigorously formulating models of networked computation.

From a ToNC perspective, one of the most basic unanswered questions is exactly what we mean by "a network" and by "networked computation." Clearly, networks have been in use for quite a while, and some of their computational capabilities and limitations have been formalized. However, existing definitions and models are not precise or comprehensive enough to enable us to prove the type of rigorous, general theorems about what can and cannot be computed on various sorts of networks that would constitute a rich and powerful "Theory of Networked Computation." Part of the difficulty is that the notion of a network has been a moving target, with new types of networks (such as sensor nets and wireless networks) gaining in prominence, making formal definitions a challenge. Our experience with networks is now sufficiently advanced that this difficulty can be overcome.

> **Research Goal:** Formulate the definition(s) that a computational system must satisfy if it is to be called a "network." Which critical resources are consumed in networked computation, and what upper bounds on the consumption of these resources must be satisfied for a networked computation to be considered "efficient"? Formulate notions of "reduction" that can be used to prove that one networked-computational problem is at least as hard as another or that two such problems are equivalent. Identify natural network-complexity classes and problems that are complete for those classes.

Multiple definitions and formal models may be needed, because "future networks" means more than just "next-generation Internet." The ToNC scope will also include theoretical aspects of the DoD's Global Information Grid [GIG], sensor networks, MANETS[1], closed "enterprise" networks, *etc*. Should each type of network be formulated independently, or is there one basic model with a few key parameters? What are the key properties that these parameters would have to capture? Open and evolving vs. closed and stable? Mobile vs. stationary? Designed vs. observed? Homogeneous vs. heterogeneous? Controllable vs. emergent? Is there a formal theory in which all of

[1] "MANET" stands for Mobile Ad-hoc NETwork.

these network types are actually distinct, and how does one prove that a given computational system falls into one particular category but not another?

These questions may seem overly ambitious, but similar theoretical frameworks have been developed and have proven highly useful in the related areas of parallel and distributed computing; examples include various PRAM models [Harr, Vish], Valiant's BSP model [Vali], the LogP model [CKP+], and Byzantine error models [LPS] .

> **Research Goal:** Develop a taxonomy of networks, with the goals of categorizing the important computational tasks that can and cannot be done efficiently on each network class and of classifying practical network designs.

References

[CKP+] D. Culler, R. Karp, D. Patterson, A. Sahay, K. Schauser, E. Santos, R. Subramonian, T. Eicken, "LogP: towards a realistic model of parallel computation," in *Proceedings of the 4th Symposium on Principles and Practice of Parallel Programming*, ACM Press, New York, 1993, pp. 1-12.

[GIG] Global Information Grid, http://www.globalsecurity.org/intell/systems/gig.htm

[Harr] T. Harris, "A Survey of PRAM Simulation Techniques," *ACM Computing Surveys* **26** (1994), pp. 187-206.

[LSP] L. Lamport, R. Shostak, and M. Pease, "The Byzantine Generals Problem," *ACM Transactions on Programming Languages and Systems* **4** (1982), pp. 382-401.

[Vali] L. G. Valiant, "A bridging model for parallel computation," *Communications of the ACM* **33** (1990), pp. 103-111.

[Vish] U. Vishkin, "A Case for the PRAM as a Standard Programmer's Model," in *Proceedings of the Workshop on Parallel Architectures and Their Efficient Use: State of the Art and Perspectives* (First Heinz-Nixdorf Symposium), LNCS Volume 678, Springer, Berlin, 1993, pp. 11-19.

Economic Agency in Networked Computation

Multi-agent systems have been extensively studied in both Economics and Computer Science, but the two communities have approached the topic very differently. The Economics literature traditionally stressed incentives and downplayed the design of algorithms and protocols, and the Computer-Science literature traditionally did the opposite. The emergence of the Internet has radically changed this state of affairs: Ownership, operation, and use by many self-interested, independent parties gives the Internet characteristics of an economy as well as those of a computer.

Economic agency appears on many levels in diverse types of networks. Internet domains (aka "autonomous systems" or ASes) are the subnetworks that directly serve users, *e.g.*, those run by companies for their employees, by universities for their students, or by commercial ISPs for their customers. They are organizationally and economically independent of each other (indeed some are direct competitors), and yet they must coordinate in order to enable interdomain communication. Nonetheless, re-examination of the autonomous-system concept is part of the clean-slate design agenda in network-architecture research:

> **Research Goal:** Are autonomous systems an essential part of Internet architecture? Are there more monolithic alternatives that could deliver significant advantages? If autonomous systems are essential, is the current hierarchical autonomous-system [GR] structure optimal?

On another level, individual users are self-interested, and they access networks through general-purpose computers that can be reconfigured in order to improve local performance; hence, network operators have to incentivize behavior that leads to good network-wide performance. In wireless mesh and ad-hoc networks, bandwidth is typically contributed and controlled by individual participating nodes; network performance could suffer dramatically if nodes fail to forward others' traffic in order to conserve local resources and are not penalized for this failure. To some extent, it is the centrality of economic agency that is now distinguishing the study of "networking" from that of parallel or distributed computing. For example, instead of studying agents who deviate from network protocols arbitrarily, as has commonly been done in distributed-systems research, it makes sense to consider agents who deviate from network protocols rationally in order to maximize their own utility.

The SIGACT community has focused intently on incentive issues in recent years, especially on the design of incentive-compatible algorithms. By building explicit payments to computational agents into the protocol, a system designer can incentivize the revelation of relevant private information and the choice of strategies that drive the overall system into a desirable equilibrium state. Substantial progress has been made in the design of incentive-compatible protocols for routing, multicast cost sharing, Internet-based auctions, peer-to-peer file distribution, and numerous other problems, but many questions remain open. General questions that form an important part of the ToNC agenda include:

Research Goal: Can one agent determine, through observation, modeling, and data analysis, whether another agent is responding to incentives or rather is behaving "irrationally" in the economic sense of this term?

Research Goal: Can incentive-compatible system designs handle agents with rapidly changing and apparently self-contradictory motivations and utility functions?

Research Goal: Are existing equilibrium concepts (such as strategyproofness, Nash, Bayes Nash, and ex-post Nash), together with randomized and approximate variations put forth recently, sufficient for the analysis of Internet-based computation, or are new, more fundamentally computational definitions needed?

Research Goal: Are standard algorithms concepts compatible with incentive analysis of networked computation? For example, because nodes and links fail, recover, join, and leave large networks frequently, the notion of a single problem instance on which a protocol either does or does not converge and, if it does, converges to a solution that either is or is not optimal may not be applicable. How should one evaluate incentive compatibility of a protocol that is carried out by a changing set of agents and that may never terminate?

Much of the recent work by the SIGACT community on incentive compatibility is covered in [NRTV].

References

[GR] L. Gao and J. Rexford, "Stable Internet Routing without Global Coordination," *IEEE/ACM Transactions on Networking* **9** (2001), pp. 681-692.

[NRTV] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, **Algorithmic Game Theory**, Cambridge University Press, 2007.

Networked Computation on Massive Data Sets

Robust technological trends (*e.g.*, the ever-decreasing cost of data storage, the ever-increasing ubiquity of computers and networks in daily life, and the accelerating deployment of sensor networks and surveillance systems) have led to an explosion of potentially interesting data. This situation has led people in many fields to observe that fresh thinking is needed about data privacy. The flip side of this observation is that these trends also strain our *algorithmic* ability to understand and use available data. Massive-data-set (MDS) computation will thus be a central theme of the ToNC agenda.

The SIGACT community has already taken up this challenge on multiple fronts. New computational models have been developed, including data streaming [FK+1, FK+2, Muth], external memory and cache obliviousness [ABW], and sampling, spot checking, and property testing [EKK+, GGR]. Applications have already been found in network measurement and monitoring (*e.g.*, [EV]). The emphasis has been on near-linear, linear, or even sub-linear time and/or space requirements, because the standard notions of polynomial time and space are inadequate when data sets are truly massive. Randomization and approximation are essential in many MDS tasks, and the fact that the SIGACT community has studied both in depth for many years will stand it in good stead.

Despite recent progress in MDS computation, much remains to be done. Indeed, no computational aspect of massive data is completely understood, and no concrete problem of interest has yet been completely satisfactorily solved. The Web-searching problem domain perfectly exemplifies both the great progress that has been made and the tough challenges that lie ahead. Who could have imagined a decade ago that the web would grow to its current size of tens of billions of publicly accessible pages and that, moreover, one would be able to search through this vast collection of pages in a split second? Despite these advances, most users have had the experience (all too often!) of searching for things that they have not found or of being unable even to express a query in the languages provided by today's search engines.

> **Research Goal:** Develop search techniques that work for images, video, audio, databases, and other non-text data on the web. Look for peer-produced structure in the web that can support search for non-text data in the same way that link structure [Klei] supports keyword search.

One research area that may greatly improve search but has only recently received attention by the SIGACT community is human-aided computing. Humans naturally provide feedback in many ways that could aid search; indeed, recent proposals (*e.g.*, [AD]) suggest creating games that, as a by-product, provide labels that could aid in the image-searching problem we've already highlighted.

Providing theoretical foundations for human-aided networked computation is a particularly novel ToNC challenge. Many observers have celebrated the "democratization" of the information environment that has been wrought by blogs, wikis, chatrooms, and, underlying it all, powerful search. More human input to the search process will make the information environment even more democratic, but it will also strain the algorithmic and mathematical foundations of correctness and information quality that have traditionally been present in the technological world. Trust, noise, and

scalability all play a part in human-aided networked computation, and these words mean different things when applied to humans from what they mean when applied to computers.

> **Research Goal:** Develop the theoretical foundations of human-aided networked computation; in particular, develop algorithms that allow networked computers to leverage and aggregate the results of a massive number of human actions. Explore the power and limitations of increasing human involvement in network-based search.

Generalizing from the search domain, numerous Web-based tasks have massive-graph computations at their core. Progress on MDS algorithmics will be an essential part of the solutions.

> **Research Goal:** Given a massive, evolving graph presented as a stream of edge-insertions and -deletions, are there one-pass, space-efficient algorithms to compute (or approximate) key graph properties, *e.g.*, conductance, eigenvalues, and bad cuts?

References

[ABW] J. Abello, A. Buchsbaum, and J. Westbrook, "A Functional Approach to External Graph Algorithms," *Algorithmica* **32** (2002), pp. 437-458.

[AD] L. von Ahn and L. Dabbish, "Labeling images with a computer game," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press, New York, 2004, pp. 319-326.

[EKK+] F. Ergün, S. Kannan, R. Kumar, R. Rubinfeld, and M. Viswanathan, "Spot checkers," *Journal of Computer and System Sciences* **60** (2000), pp. 717-751.

[EV] C. Estan and G. Varghese, "New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice," *ACM Transactions on Computer Systems* **21** (2003), pp. 270-313.

[FK+1] J. Feigenbaum, S. Kannan, M. Strauss, and M. Viswanathan, "An Approximate L^1-Difference Algorithm for Massive Data Streams," *SIAM Journal on Computing* **32** (2002), pp.131-151.

[FK+2] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang, "On Graph Problems in a Semi-Streaming Model," *Theoretical Computer Science* **348** (2005), pp. 207-216.

[GGR] O. Goldreich, S. Goldwasser, and D. Ron, "Property Testing and Its Connection to Learning and Approximation," *Journal of the ACM* **45** (1998), pp. 653-750.

[Klei] J. Kleinberg, "Authoritative sources in a hyperlinked environment," *Journal of the ACM* **46** (1999), pp. 604-632.

[Muth] S. Muthukrishnan, "Data Streams: Algorithms and Applications," http://athos.rutgers.edu/~muthu/stream-1-1.ps

Experimental Rigor in Networking Research

Until recently, most mainstream Computer-Science research has dealt with "man-made" or "designed" objects: Hardware and software systems were designed, built, programmed, and studied, using approaches and methods akin to those in engineering and mathematics. Today's large-scale networks (and even large, complex pieces of software) are in some ways closer to the "found" objects or natural phenomena studied by scientists: Detailed knowledge of the constituent components and processes of such a system is often insufficient for understanding and prediction of the system's aggregate behavior, because of the scale and complexity of the aggregate and the crucial role of exogenous forces, most notably the behavior of human users and operators. This presents abundant opportunity for mathematical modeling and analysis of network behavior.

One approach to modeling and analysis that has proved fruitful is to divide it into five stages [Mitz]: **Observe** (gather data about the behavior of the network), **Interpret** (explain the importance of these observations in context), **Model** (propose an underlying model for the observed behavior), **Validate** (find data to validate and, if necessary, specialize or modify the model, and **Control** (based on the model, design ways to control the network behavior).

Observation and interpretation have been proceeding apace for many years, and some consistent themes have emerged. For example, power-law and lognormal distributions[2] are observed almost everywhere that there is networked computation, both in Computer Science (file sizes, download times, Internet topology, the Web graph, *etc.*) and in other fields (income distributions, city sizes, word frequency, bibliometrics, species and genera, *etc.*). Despite their ubiquity in the study of network data, we do not yet fully understand how best to use these classes of distributions. In particular, it can be unclear whether observed data are more accurately modeled as a power-law distribution or a lognormal distribution. The distinction can be extremely important in some modeling contexts (*e.g.*, stock prices and insurance tables); when and why it is important in the modeling of network behavior is not always clear.

> **Research Goal:** Develop techniques for distinguishing empirically between power-law distributions and lognormal distributions. For situations in which they cannot be distinguished empirically, explore the implications of both modeling choices for validation of the model and subsequent control of network behavior.

Distinguishing empirically between power-law-distribution models and lognormal-distribution models is a specific case of the validation challenge. In general, there are many models of network behavior in the literature, but there are few effective techniques for validating that a model is the right one in order to predict and control future behavior. Some of the best work on model validation has actually resulted in model refutation [CCG+, LBCX]. Validation is inherently harder than refutation; in fact, it is not clear

---

[2] A *power-law distribution* is one that satisfies $Pr[X \geq x] \sim cx^{-\alpha}$. The random variable X is *lognormally distributed* if $\ln X$ is normally distributed.

exactly what constitutes convincing validation. Fleshing out this area is a basic ToNC challenge.

> **Research Goal:** Develop techniques for validating models of network behavior, *e.g.*, for proving that a probabilistic model is consistent with observed data or that one model is a "better fit" than another.

Ultimately, the goal of network modeling and analysis is the ability to predict and control network behavior. Accurate models should inform the co-design of networks and algorithms. They should also empower us to change various aspects of network design, use, or operation in ways that improve performance without unforeseen negative side-effects. Many other themes explored in this report, *e.g.*, incentive compatibility, network algorithmics, and networked-computational complexity, might be useful for control.

> **Research Goal:** Explore the feasibility of controlling networks for which models have been validated. In particular, explore the use of incentives (both with and without monetary transfers), limits on users' access to network resources (such as space and bandwidth), and limits on access to information about the network state.

Progress toward these goals will require significant advances in experimental networking research and facilities of a type and scale that are currently unavailable.

There are also purely theoretical problems that beckon in the area of analytical paradigms for networked computation. For example, the network analog of *smoothed* analysis [ST] would clearly be useful. Smoothed analysis, which has shed light on classic problems such as the running time of the simplex algorithm for solving linear programs, captures the fact that there can be uncertainty about in the input to an algorithm. This is quite relevant to network algorithms, where the uncertainty might come from, *e.g.*, unpredictable traffic congestion, unreliable network components, unpredictable user behavior, or intentionally supplied random bits.

> **Research Goal:** Expand the scope of network modeling and analysis. In particular, develop holistic models that capture many network features simultaneously and analytical methods that exploit uncertainty about the environment.

References

[CCG+] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "The Origin of Power Laws in Internet Topologies Revisited," in *Proceedings of INFOCOM 2002*, IEEE Computer Society Press, Los Alamitos, 2002, pp. 608–617.

[LBCX] A. Lakhina, J. Byers, M. Crovella, and P. Xie, "Sampling Biases in IP Topology Measurements," in *Proceedings of INFOCOM 2003*, IEEE Computer Society Press, Los Alamitos, 2003, pp. 332–341.

[Mitz] M. Mitzenmacher, "Editorial: The Future of Power Law Research," *Internet Mathematics* **2** (2006), pp. 525-534.

[ST] D. Spielman and S. Teng, "Smoothed Analysis: Why the Simplex Algorithm Usually Takes Polynomial Time," *Journal of the ACM* **51** (2004), pp. 385-463.

Theory of Networked Computation: Participants

This material was generated at two ToNC workshops during Spring semester 2006, and it was reviewed and updated at the June 2008 NetSE meeting in Boston. One of the 2006 workshops was held at the Nassau Inn in Princeton, NJ on February 16-17 and the other at the International Computer Science Institute (ICSI) in Berkeley, CA on March 16-17. Both workshops were attended by invited participants and by members of the Computer Science community who sent in successful applications. At both events, plenary talks were presented on important ToNC themes, and then participants formed "breakout groups" for in-depth discussion and problem formulation. Slides for all of the presentations can be found at http://www.cs.yale.edu/homes/jf/ToNC.html.

The Princeton ToNC workshop was chaired by Joan Feigenbaum and Jennifer Rexford. Breakout-group themes were Next-Generation Information Systems (Andrei Broder, chair), Next-Generation Network Architecture (Ashish Goel, chair), Next-Generation Network Protocols (Bruce Maggs, chair), Control of Personal Information in a Networked World (Rebecca Wright, chair), and Economic Approaches and Strategic Behavior in Networks (Michael Kearns, chair). The participants were Matthew Andrews (Bell Labs), Sanjeev Arora (Princeton), James Aspnes (Yale), Hari Balakrishnan (MIT), Boaz Barak (Princeton), Amotz Barnoy (Brooklyn College, CUNY), Andrei Broder (Yahoo! Research), Moses Charikar (Princeton), Nick Feamster (Georgia Institute of Technology), Joan Feigenbaum (Yale) , Michael Foster (NSF), Ashish Goel (Stanford) , David Goodman (NSF), David Johnson (AT&T Labs) , Howard Karloff (AT&T Labs) , Richard Karp (UC Berkeley and ICSI), Jonathan Katz (University of Maryland), Michael Kearns (University of Pennsylvania) , Vincenzo Liberatore (Case Western Reserve University), Bruce Maggs (CMU and Akamai), Stephen Mahaney (NSF), S. Muthukrishnan (Rutgers), Kathleen O'Hara (NSF), Jennifer Rexford (Princeton), Rahul Sami (University of Michigan), Alex Snoeren (UC San Diego), Daniel Spielman (Yale), William Steiger (NSF), Eva Tardos (Cornell), Robert Tarjan (Princeton), Sirin Tekinay (NSF) , Eli Upfal, (Brown), Avi Wigderson (IAS), Gordon Wilfong (Bell Labs), Tilman Wolf (University of Massachusetts), and Rebecca Wright(Stevens Institute of Technology).

The Berkeley ToNC workshop was chaired by Joan Feigenbaum and Scott Shenker. Breakout-group themes were Algorithmic Foundations of Networked Computing (John Byers, chair), Analytical Foundations of Networked Computing (Eva Tardos, chair), Complexity-Theoretic Foundations of Networked Computing (Russell Impagliazzo, chair), Economic Foundations of Networked Computing (Milena Mihail, chair), and Foundations of Secure Networked Computating (Salil Vadhan, chair). The participants were Moshe Babaioff (SIMS), Kirstie Bellman (Aerospace Corporation), John Byers (Boston University), Chen-Nee Chuah (UC Davis), John Chuang (SIMS), Luiz DaSilva (Virginia Poly), Neha Dave (UC Berkeley), Joan Feigenbaum (Yale), Michael Foster (NSF), Eric Friedman (UC Berkeley [on leave from Cornell]), Joseph Hellerstein (UC Berkeley), Russell Impagliazzo (UC San Diego), Matti Kaariainen (ICSI), Anna Karlin (University of Washington), Richard Karp (UC Berkeley and ICSI), Robert Kleinberg (UC Berkeley/Cornell), Richard Ladner (University of Washington), Karl Levitt (NSF),

Gregory Malewicz (Google), Milena Mihail (Georgia Institute of Technology), Christos Papadimitriou (UC Berkeley), Kathleen O'Hara (NSF), Satish Rao (UC Berkeley), Vijay Raghavan (UC Berkeley), Tim Roughgarden (Stanford), Amin Saberi (Stanford), Scott Shenker (UC Berkeley and ICSI), William Steiger (NSF), Ion Stoica (UC Berkeley), Eva Tardos (Cornell), Shanghua Teng (Boston University), Salil Vadhan (Harvard), and George Varghese (UC San Diego).

Appendix 2

Workshop Report on

Behavior, Computation and Networks in Human Subject Experimentation

Behavior, Computation and Networks in Human Subject Experimentation

NetSE Workshop Report
Michael Kearns (Computer and Information Science, University of Pennsylvania)
Colin Camerer (Economics, Caltech)

Background: *On July 31 and August 1, 2008 we held a NetSE workshop on the topic of "Behavior, Computation and Networks in Human Subject Experimentation". This interdisciplinary workshop was designed to bring together a relatively small number of researchers with the following broad profiles:*

- *Researchers from economics, game theory and sociology whose interests include behavioral human-subject experiments.*
- *Researchers from economics, game theory and sociology interested in computational and algorithmic models for behavior, and algorithmic issues more broadly (such as equilibrium computation).*
- *Researchers from computer science with interests in game theory, economics and sociology, especially behavioral, experimental and simulation work.*
- *Researchers from all of the fields above with interests in social, organizational, technological and other networks, and how network structure and formation interact with individual and collective behavior.*

The remainder of this document describes the scientific agenda and research challenge areas emerging from discussions at the workshop and beyond.

1. Rationale for a New Research Agenda

Researchers in the computer science, economics, game theory and sociology communities have been engaged for some time now in healthy and vibrant interaction on a variety of *theoretical* topics. We assert that the natural and most important next frontier in this dialogue is the introduction of a *behavioral and experimental* component. Of particular interest are organizations and systems in which an underlying network structure strongly governs interaction and strategy.

Regarding the dialogue so far, between economics and computer science we have the well-established field of *algorithmic game theory*, whose work can now be found in multiple journals and conferences of both fields (ACM Conference on Electronic Commerce; Workshop on Internet Economics (WINE); STOC, FOCS and SODA of the theoretical computer science community; World Congress of Game Theory; Games and Economic Behavior; and many other examples); in an extensive recent edited volume from Cambridge University Press; in many notable publications co-authored by members of both communities; and so on. Similarly, by now there is a fair amount of contact between primarily mathematical topics within sociology (such as the diffusion of trends within a social network) and the theoretical computer science community.

If we are to take such interactions as more than theory for its own sake --- by which we mean that they might provide the foundation for an empirical science that is applicable to real problems and data, and able to make predictions (and potentially policy recommendations) --- then it is clear that we must begin to develop a heretofore missing behavioral and experimental component. In the same way that behavioral game theory and economics seek to adapt their theoretical counterparts towards actual human and organizational behavior (thus improving their applicability), we seek to build a behavioral and experimental discipline encompassing strategic settings important in computer science and technology, network science, and related fields.

As we shall discuss in the following section, building this discipline presents a number of significant conceptual, scientific and resource challenges to the constituent research communities. We believe the reward for meeting these challenges will be the creation of an important new field whose content will be widely applicable to the myriad modern problems in which strategic considerations, technology and behavior interact.

2. Emerging Research Challenges

Unifying Algorithmic and Behavioral Game Theory

In their own fashions, both algorithmic and behavioral game theory seek to "repair" classical game theory, arguably in the direction of "realism". Ideally game theory and related fields would provide accurate predictions of actual strategic behavior in individuals and organizations. Behavioral game theory seeks to reconcile theoretical models with empirically observed behavior in controlled experiments. Algorithmic game theory attempts to identify and rectify classical equilibrium notions by enforcing plausible demands on computational and other resources.

Ideally these two approaches should be unified and refined --- behavioral models taking more precise account of computational considerations, and algorithmic models evaluated and improved in light of experimental evidence. The mathematical, methodological and cultural chasms between the two communities are large, which is a significant part of the challenge. Algorithmic models will need to be refined in ways less related to traditional computational complexity (P vs. NP, and the various subclasses of P) and more related to cognition. But there is already promising work in this direction that crosses the disciplinary boundaries --- for example, recent work of Camerer and colleagues on "cognitive hierarchies" of varying levels of strategic behavior that directly account for computational limitations. We believe both algorithmic and behavioral approaches to game theory are sufficiently mature independently that the attempt to build a unified theory has arrived.

Network and Systems Infrastructure for Behavioral Experiments

At both the workshop and during a long series of dialogues between several participants, there has been discussion and excitement around the possibilities of building internationally shared networking and systems infrastructure for the conducting and

support of large-scale behavioral experiments in sociology, game theory, economics, and most recently, computer science. The fundamental observation is that the Internet, Web and other technologies have created the possibilities for (semi-)controlled experiments in these disciplines --- where small population sizes and the difficulties of human subject management have long been limiting factors --- on a large or massive scale. There are already numerous examples of such experiments, but each has employed highly specialized software and technology.

A significant portion of workshop time was devoted to discussing what such shared infrastructure might provide, how general it should be, what precedents there are, whether existing commercial platforms (such as Amazon Mechanical Turk) might suffice or at least serve as models, and many other issues. Perhaps the most important design issue is the trade-off between generality and ease of use, while the most important methodological concern is the maximal retention of experimental control and subject knowledge and management.

There is a strong sense among the stakeholders that (a) designed and implemented properly, such an experimental platform could have a transformative effect on behavioral research, and (b) its creation would be an extremely challenging and resource-intensive project requiring close collaboration between sociologists, economists, and computer scientists.

Theory and Design Principles for Peer Production

Related but distinct from the topic of shared experimental infrastructure is the phenomenon of recent "human peer production" --- systems in which massive numbers of distributed individuals voluntarily "solve" collective problems or build influential and useful artifacts. The diversity of such systems across several dimensions (problem solved, nature of individual contributions, incentives, etc.) is staggering and includes Wikipedia, the ESP Game, del.icio.us, Amazon Mechanical Turk, NASA click workers, Galaxy Zoo, prediction markets of many kinds, social networks, and many other examples. Yet there is essentially no theory about the design or behavior of such systems, including on basic questions such as contributor population size vs. quality of collective outputs, choice of incentive schemes, dealing with problems not easily decomposed into "modular" subtasks, and so on. The moment seems right to tackle such challenges with an interdisciplinary approach. (Duncan Watts and Kearns have held a series of DARPA workshops on this broad topic, with participation from many at the NetSE workshop as well as from industry and the military.)

3. Workshop Format

The workshop participants numbered approximately 20 active scientists representing multiple disciplines, each of them highly influential and at the forefront of their respective fields. We felt that assembling a group with this "gravitas" was important both

to obtain mature and accurate representations of the interests of the constituent areas, but also for follow-up evangelization of our nascent and emerging research agenda.

The workshop began with brief research presentations from each scientist. While the goal was to maximize the time allotted to semi-structured discussion and brainstorming, the diversity of backgrounds, interests and terminology was sufficient to warrant steeping the participants in work and viewpoints of the others. These talks were by themselves fascinating for both their individual content and collective variety and interconnections.

The bulk of the time, however, was devoted to informal and open discussion on a small number of central topics. The complete agenda, with participants, talk titles and discussion topics is provided in the following section.

4. Workshop Agenda and Participants

Behavior, Computation and Networks in Human Subject Experimentation
Thursday, July 31 and Friday, August 1, Del Mar Inn, Del Mar CA

Agenda

Thursday, July 31

9:00 Welcoming remarks: M. Kearns (Penn) and C. Camerer (Caltech)
9:15 Remarks from E. Zegura (Georgia Tech) on NetSE Council interest in our topics
9:30 Brief introductions
9:45 Brief research/project presentations (~15 minutes each):

V. Crawford (UCSD): *Studying Strategic Thinking by Monitoring Search for Hidden Payoff Information and Analyzing the Data in the Light of Algorithms that Link Cognition, Search, and Decisions*

M. Wellman (Michigan): *Software Agents and Empirical Game Analysis*

J. Ledyard (Caltech): *Agent-based models for repeated game experiments*

Break

C. Camerer (Caltech): *Evidence of algorithmic game theory from human experiments*

M. Kearns/S. Judd (Penn): *Behavioral Network Science and the Democratic Primary Game*

M. McCubbins/M. Paturi/N. Weller (UCSD): *Effects of Complexity, Incentives and Network Structure on Multi-Player Coordination Games*

J. Fowler (UCSD): *Eat, Drink, and Be Merry: The Spread of Obesity, Substance Use, and Happiness in a Large Social Network*

12:30 Lunch

1:30 Brief research presentations, continued:

S. Kariv (Berkeley): *A Normal Form Game Experiment of Trading Networks*

A. Pfeffer/K. Gal (Harvard): *Modeling the reasoning of people and computer agents in strategic settings*

B. Rogers (Northwestern): *Communication Networks: An Experimental Study of Influence*

2:30 Discussion of the morning's presentations: common themes and differences; marrying different approaches; what's missing; etc.
3:00 Brainstorming Topic 1: Algorithmic game theory and behavioral game theory/economics
4:00 Break
4:15 Brainstorming Topic 2: Relevance/incorporation of simulated agents in behavioral experiments
5:00 Adjourn
6:45 Informal dinner overlooking the Pacific, Martin Johnson House, Scripps Institute of Oceanography, La Jolla

Friday, August 1

9:00 Brief research presentation by D. Watts, Yahoo! Research/Columbia: *Virtual Labs: Using the Web to Conduct Human Subjects Experiments*

9:15 Recap of Thursday, discussion of new topics
9:30 Brainstorming Topic 3: "Scaling Up" behavioral experiments: use of the web, Amazon Mechanical Turk, peer production, etc. Do we need a "programmable infrastructure"?
10:30 Break
11:00 Brainstorming Topic 4: What are the applications of all this stuff?
12:30 Lunch
2:00 Brainstorming Topic 5: Where do we go from here?

Appendix 3

Workshop Report on

Network Science and Network Design

Summary Report

Workshop on Network Science and Network Design
July 29-30, 2008 at USC/ISI

Overview

In 30 years, computer and information networks have moved from research curiosity to fundamental and critical infrastructure of modern society. Numerous authors have described, from both technical and societal perspectives, the great promise and tremendous challenges created by this transition.

This report focuses on a research domain that workshop participants argue is central to both promise and challenge: a fundamental and mathematical theory of network architecture and design. Progress in this domain is essential to address a deeply troubling conundrum. Modern network technology promises to provide unprecedented levels of performance, efficiency, sustainability, and robustness across numerous domains of science and engineering. In many areas it has already done so. Yet, as network scale, ubiquity and deployment grows, the problem of rare but catastrophic real-world failures and "unintended consequences" is, if anything, becoming worse. This "robust yet fragile" nature of complex systems is ubiquitous, and foundational research, in the form of a theoretical framework to manage the complexity/fragility spirals of future network infrastructures, is essential to reverse the ominous trend and fulfill the promise of networks as fundamental infrastructures of society.

"Design" is a word widely used but rarely formalized. Despite the existence of multiple successful network architectures in technology and nature, there is not yet a systematic design methodology for building large-scale, robust, efficient and evolvable networked systems. Rather, current systems have arisen largely through a classic process of empirical design, deployment and evolution, with many, many solutions being tried, evaluated in service, and accepted or discarded.

It is absolutely critical to recognize the success of this methodology to date, and to note the depth and scope of intellectual insight already generated through its use. Further, we note that such deployment and evolution is the first step in the development of virtually all engineering disciplines throughout history. Nonetheless, we believe that the success of Network Science and Engineering as a field moving forward will increasingly be rooted in the creation of an underlying discipline that can systematically design, analyze, implement and maintain large-scale, complex networked systems. Further, we suggest that the timely convergence of several factors give confidence that research community can dramatically advance this agenda, at least in the context of complex computer and communication networks.

As a first starting point, we have today rich real-world examples of successful network architectures in both technical and natural systems, together with clear articulation by domain experts of precisely what these architectures entail and by what mechanisms they are implemented. While many of the details remain fragmented and largely unformalized, appropriate juxtaposition and translation reveals striking convergence in the underlying principles across different domains.

As a second starting point a nascent mathematical framework suggests that this convergence is not accidental but is the result both engineered and natural networks sharing common need for efficiency, robustness, and evolvability. These requirements drive the recurrence of certain ubiquitous architectural features, such as hourglasses and bowties, together with their role in protocols, layering, control, and dynamics. Initial success in formalizing these in the context of the Internet is most encouraging.

Within our overall objective a number of more specific goals stand out. Among these are:

- Developing the foundations of a theory of network *architecture* that allows rigorous analysis and systematic design of complex networked systems, including organizational abstractions such as interfaces and layering.

- Developing new network *protocol* design methods to implement the protocol-related elements of networked systems, for applications such as cross-layer control in wireless networks, distributed data gathering, integrated network coding, and communication platform for control of cyber-physical infrastructure.
- Developing and evangelizing a common *mathematical language* to broaden and deepen the contacts between and across networks in engineered systems and networks in the sciences, particularly biology.
- Developing discipline and methodologies to *validate measurements and data* and *to evaluate candidate technologies* on university and industry test beds, for example including WAN backbones, smart power grids and mobile sensor networks. Key here is the extension of real-world evaluation beyond the deployment-study model to include also a more structured and rigorous worst-case analysis model.

Connecting Network Science to Network Design

The title of our workshop, and the discussion above, reflect participants' belief in the importance of increasing connection between two fields that have to date not been tightly related: the evolving discipline of *network science* and the traditional field of *network design.* Nonetheless, this connection is fraught with definitional peril. We expand slightly on this point.

Network Science is a new and rapidly evolving discipline that has gained great visibility in the technical and popular literature over the past 10 years. Unsurprisingly, there is a divergence of views regarding the precise scope and definition of the discipline. Nonetheless, most views of the field to date share two basic properties:

1. The field is *analytic* in nature; much of the published literature is concerned with observing existing networks and building mathematical models that capture the properties and statistics of these networks.
2. The field hypothesizes that many properties of networks are *domain independent;* and aims to develop understanding and models that apply to networks in general, rather than to any particular problem domain.

Network Design is historically the exact opposite:

1. It is an activity *synthetic* in nature: the goal is to construct a new network to meet a need, rather than to observe and study an existing network.
2. It is typically *domain-specific* in nature: for example, the designers of the Internet were not particularly concerned about whether their work also applied to the design of railroad or energy networks.

Workshop participants argue that building and strengthening connections between these two fields will significantly benefit both. Much of this report is concerned with the potential of applying ideas from network science (defined broadly) to network design. But the converse is true as well. As examples, we note two points: that applying results from network science to network synthesis and evaluating the results can help to validate that those results capture fundamental properties, rather than observational artifacts; and that techniques for capturing, abstracting, and modeling a network's domain-specific properties and constraints within some general framework represent an important extension of network science.

Elaborations

This section summarizes discussion in the workshop that sought to expand on the fundamental problem statement described above. We briefly consider what is to be designed, the nature of the intellectual advance we seek, and the question of what, exactly, constitutes a network.

What is to be designed? The phrase "network design" can be interpreted at three levels. From the top down, these are:

1. Development of network *architecture*. The architecture of a network[1] captures the fundamental structuring properties of the system. Our challenge at this level, discussed further in Research Challenge 1, is to place complex architectural decisions within the scope of a theoretical and mathematical framework.
2. Development of network *protocols and algorithms.* This level of design focuses on the individual building blocks within the architecture. Our challenge at this level is to catalyze the development of theory and understanding that guides the development of protocols with provable robustness and performance properties.
3. Design of a specific network *instance.* This level is concerned with the design of a particular artifact – for example, the AT&T Internet backbone – within a framework guided by the results of levels 1 and 2 above – for example the IP/TCP architecture and protocols. While successful design at this level is the ultimate goal of the research we describe, our intellectual focus is centered at the architecture and protocol design levels. This is because architecture and protocols constitute reusable, high leverage *design patterns* or templates for the repeated realization of individual network instances.

Evolution of engineering understanding. In virtually all engineering domains, intellectual understanding develops over time – iteratively and sometimes nonlinearly – along a common progression. A key thesis of this report is that significant elements of our discipline, network design and engineering, are now advancing or poised to advance along this path, and that a fundamental research objective should be to intentionally and consciously target this advance. We briefly characterize the progression below.

1. **Verbal understanding** – at this level, intellectual understanding is expressed primarily in words. The understanding encapsulated may range from rudimentary to highly sophisticated – it is the *form* of the understanding, and thus the methodology by which it is applied to new problems, that we focus on. An example of well-developed reasoning primarily expressed at this level is the set of "End to End Arguments[2]" that are often viewed as underlying the Internet's architecture.
2. **Data and Statistics** – at this level verbal understanding is augmented by collection of data and statistics about a system. Such statistics are primarily *observational* – that is, they capture observed properties of a system under study, but may or may not capture fundamental

---

[1] Practitioners sometimes refer to the architecture of a specific network instance, speaking for example of "the architecture of the AT&T backbone" to describe basic design choices such as mesh versus ring topology or routed versus switched interconnects. We do not use the word for this purpose. We are referring here to the fundamental structuring and decomposition decisions within a protocol *family* – a design template for networks – rather than the structuring decisions for a particular network instance built using that family. However, the two uses are obviously related. It is to be hoped, though not always true in practice, that the structuring decisions of a specific network instance are conformant with the design architecture of the protocol family used to realize that network instance.

[2] The correct title from the original paper, but now often referred to as the "End to End Principle".

characteristics as opposed to mere artifacts. Data and statistics at this level are an often-necessary first step towards intellectual understanding, but should not be confused with understanding itself.

3. **Modeling and simulation** – represent a step beyond data collection and towards system and engineering *evaluation.* As with step 2, however, simulation models, and thus the results of using them, may but need not be rooted in fundamental and correct understanding of the system being modeled, and thus may or may not be limited or misleading in their application.
4. **Cause-based analysis** – represents the step beyond statistics and models that merely capture the behavior of a system "observationally" and towards statistics and models that explicitly capture behavior "for the correct reasons". Crucially, when the cause of a behavior is fully understood it becomes possible to reason successfully about the scope and validity of a behavioral model, and thus about an evaluation, simulation, or design activity based on that model.
5. **Verifiable synthesis** – represents the final step from causally based understanding, modeling, and validation of existing systems and designs to the development of intellectual frameworks that support synthesis of *new* designs with well understood, and ideally, formally provable behavioral characteristics and scope of applicability.

The fundamental elements of a network. Many early results from Network Science have focused on the topological properties of networks – both because data about these properties in existing real-world networks was thought to be easily obtainable and because conclusions about topology (and its evolution over time) were seen as potentially results about networks *per se*, independent of their domain of application.

It is apparent, however, that network design is about much more than topology, and that a corresponding view of Network Science must be equally as broad. Correspondingly, modern "networking research" is in fact a composite and integration of several sub-fields, with the various sub-fields at different stages of development and advancing somewhat independently along the evolutionary path described above.

Recently, study and comparison of networks across widely separated domains – ranging from the Internet to systems biology – has raised the possibility that a set of common sub-elements and abstractions of understanding might be found that together capture the essence of broad classes of networks. Chief among these are abstractions for network topology, for traffic or information generation, flow, and use, for the networked system's functional structuring and modularity, and for the control mechanisms that keep the network operating robustly and efficiently. While it is clear that the relative importance of individual elements within this set will differ for networks in different domains, a unifying theoretical framework that encompasses, builds on and integrates all simultaneously would have tremendously broad applicability. Several workshop participants see the development of this framework as a "holy grail" of long-term network science research.

To capture these ideas more concretely in the context of computer and communication networks, workshop participants utilized the "matrix" of Figure 1a below. Figure 1b shows a use of the matrix: capturing the progression over time in one specific sub-area of networking – understanding the statistics, causes, and network design implications of "heavy-tailed" Internet traffic. Significant discussion in the workshop focused around questions related to this matrix – the appropriateness and completeness of the column topics, the status of each column today along the understanding progression, and the identification of research directions that hold particular promise in advancing individual column topics in this progression or creating integrated theories that capture relationships *between* the columns.

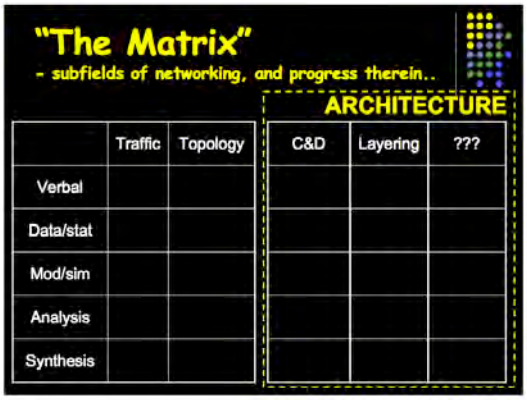

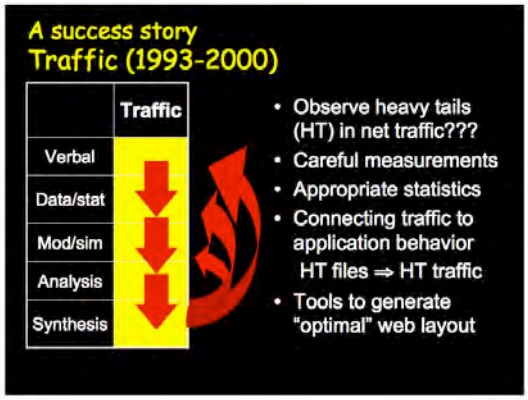

Research Challenges

The overview above suggests a number of research opportunities, both within the field of network science itself and in the interplay between network science and network design. This section of the report outlines two broad research challenges that workshop participants consider to be particularly relevant and particularly timely. It is interesting to note, however, that these challenges are not entirely independent or orthogonal. Early results suggest that each of these challenges may be approachable by drawing on methodologies and elements of a shared mathematical framework. This is excellent news, because it strengthens our sense that the developing theory may build understanding *across* columns in the networking research matrix, as well as deepening our understanding of the areas that comprise individual columns.

Challenge 1: Theoretically Derived Network Architectures

In modern usage, the term *architecture* focuses on the elements of a system's structure and organization that are most universal, high level, and persistent. The architecture of a complex, long-lived, decentralized system such as the Internet must address many goals, often in tension with one another. Such an architecture must facilitate system-level functionality as well as a number of properties that are currently less quantifiable, such as robustness and evolvability to uncertainty and changes in components, desired function and environment.

A well-formulated system architecture is typically conceived of at two distinct levels. First, there is a set of *structuring principles;* top-level design principles that are used to guide and bring coherence to decisions about system modularity and organization. Second, and separately, there is the *actual structure* of the system that emerges from the architect's application of these principles to specific modularity and organizational decisions. These two levels of system architecture act together to provide desired properties. Crisp, well articulated structuring principles provide a clear basis for reasoning about capabilities of an architecture, its ability to meet design objectives, requirements, and constraints, the tradeoffs being made, and the potential effects of architectural evolution (intended or otherwise) during the lifetime of the system.

To date the derivation and study of system architecture has been more art than science. In particular, the process of arriving at valid, useful, and understandable structuring principles has been one of intuitive study, based on experience, empirical observation, and evaluation of the properties of deployed systems. The research challenge we identify is the development of *theoretically derived architectures;* architectures with structuring principles that are derived from rigorous underlying theory, and thus that can provide stronger, more easily applicable, and more objective guidance for the overall architectural design of complex systems with predictable and well understood properties.

Current State of the Art: the Internet

The Internet provides an excellent example of how a well-framed architecture can address several criteria: facilitating evolution and robustness, as well as functionality and implementation. Indeed, much of the Internet's success has been a result of adhering more or less faithfully over time to a set of fundamental network design principles adopted by the early builders of the Internet. These principles, including layering, fate-sharing, and the end-to-end arguments, comprise the first, "structuring principles" layer of the classic Internet architecture.

The second, or "system modularity", layer of the Internet's architecture is manifest in the structure of the underlying TCP/IP protocol stack. At the overall level, the well-known hourglass concept creates a thin "waist" of universally shared data transmission and control mechanism (IP and TCP/UDP/SCTP/etc.) separating and decoupling the vast diversity of applications that sit above it from the equally vast diversity of hardware that may lie below it. Within the waist, additional modularity is evident. Roughly, protocols within the IP layer control routes for packet flows and thus, available aggregate bandwidth, while protocols at the TCP layer control individual flow rates and guarantee delivery.

From today's perspective, this architecture is remarkably effective in the choices that were made, but shallow in its connection to theoretical understanding of the full network design problem. Engineering design thus far has primarily been driven by reasoning and intuition, followed by considerable experimentation, either explicitly or by deployment in service. That is, the development of Internet technologies has followed from a largely empirical view, one in which validation of a design or protocol has been conducted via simulation or prototype.

The success of this approach has resulted in a scenario in which we are better at "trial and error via deployment" than at providing provable guarantees on performance, robustness, stability,[3] etc. However, as technological visions increasingly emphasize ubiquitous communications, computing, and cyberphysical system control, with systems requiring a high degree of autonomy and adaptation, but also robustness, evolvability, scalability, and verifiability, a more rigorous, coherent, and reasonably complete mathematical theory underpinning the technology is needed. Interestingly, recent progress in the development of theoretical underpinnings for network architecture has both confirmed the basic strength of the existing Internet architecture and suggested opportunities for dramatic improvement.

An Approach: Layering and Modularity as Optimization Decomposition

This section outlines one approach to the development of theoretically derived network architectures that appears quite promising.

The approach is rooted in a developing mathematical framework that views various protocol layers as carrying out asynchronous distributed computations to optimize a global objective function, subject to resource constraints in the network. Under this view, different layers iterate on different subsets of the decision variables using local information to achieve individual optimality. Taken together, these local algorithms attempt to achieve a global objective.

Such a framework exposes the interconnections between protocol layers as different ways to modularize and distribute a centralized computation. It formalizes the common practice of breaking down the design of a complex system into simpler modules, and provides a top-down approach to both the systematic design of layers and the tradeoffs between competing design objectives.

---

[3] The network architecture research community frequently refers to these properties collectively as "the ilities".

The broad outlines of the mathematical framework are as follows. The framework views the network as solving an appropriately defined general Network Utility Maximization (NUM) problem. The objective of the network becomes the optimization of a global cost function subject to all the physical and resource constraints in the network. Then network layering and protocol modularity can be understood as a *decomposition* of this large optimization problem into decentralized subproblems, with various protocol layers regarded as carrying out asynchronous, distributed computations to solve the subproblems. Different layers iterate on different subsets of the decision variables using local information to achieve individual optimality. Taken together, these local algorithms attempt to achieve a global objective.

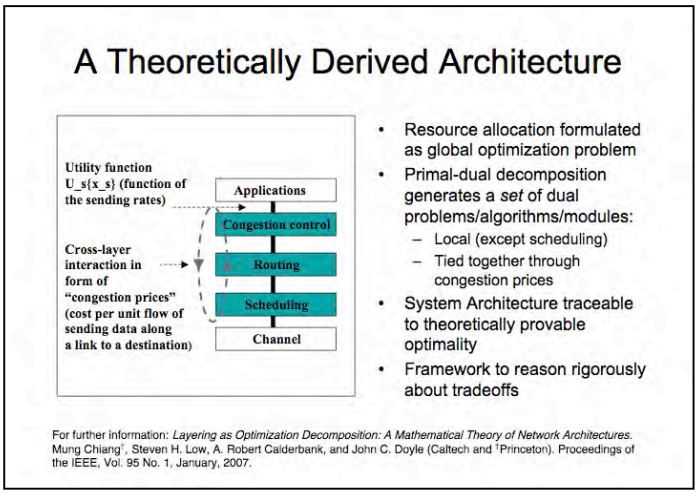

A key point is that there are many different ways to decompose a given problem, each of which corresponds to a different layering and modularity. These different decompositions have different trade-offs in efficiency, robustness, and asymmetry of information and control. Thus, some are "better" than others depending on the criteria set by the network users and operators, and the theory provides a vehicle to reason about how different decompositions address these criteria.

What is encouraging about the layering-as-optimization-decomposition approach is that it has successfully evolved over time from origins in an abstracted and academic theory of the TCP protocol alone to a practically useful and significantly more general framework for layering and protocol design that is capable of addressing an increasing range of design issues. To date, the framework has been extended to address and integrate architectural elements well beyond its original focus on TCP and congestion control, such as routing, media access control (MAC) functions, power control and network coding.

Issues and Challenges

Despite this initial promise, a great number of issues and challenges remain if the line of inquiry described above is to lead to a true mathematical theory of network architecture that is capable of guiding design. Such a theory must encompass not just protocol layering and modularity, but also inform network control and dynamics, incentives, security, robustness and evolvability, and similar issues.

The challenges facing this emerging theory can be divided broadly into two categories. The first of these is the need to bring aspects or elements of network architecture that are not presently addressed under the theoretical umbrella. In many cases this may involve recasting or reformulation of an existing architectural concept into a format that allows it to be addressed within the theoretical framework. The second is advance to the theory itself, to broaden its applicability or improve its capabilities.

It should be apparent that these two types of challenges, and the research activities they engender, are not unique to the optimization theory we have outlined. Rather, these categories apply to *any* potential theoretical underpinning for network architecture. Further, the two types of challenge are closely related. There is a flexible and mutable boundary between the notion of extending mathematics to cover existing architectural concepts and that of extending architectural concepts

to fit within existing mathematics. It is interesting to note, however, that while the boundary is flexible the experience and domain understanding required to progress from each of the two sides is distinctly different. These factors suggest that the research domain is fundamentally interdisciplinary according to current taxonomy, and indeed the majority of recent progress has come from collaborative teams of experienced and intuitive architects and well-grounded theoreticians.

To provide further clarity to this subject we briefly discuss some examples from each category of challenge.

Incorporating new architectural elements and concepts

A key limitation of the optimization-based approach described above is that it is concerned with the optimization of a single global cost function subject to a set of constraints. Thus, any architectural requirement that is to be addressed within the theory must be expressed either as a) a cost, in the metric of the cost function, or b) a constraint.

For some architectural functions, such as capacity and resource management (which captures at least routing, scheduling and congestion control), this formulation is entirely natural. For others, such as security, it is at first glance less so. Methods must be developed to express additional architectural concepts within the theoretical framework. Approaches that appear useful include

- Subdividing or refactoring the problem so that one or more sub-problems can be expressed in terms of the optimization goals.
- Generalizing or otherwise redefining the optimization cost function so that it can be used to express a broader range of architectural objectives.
- Subdividing or refactoring the problem so that one or more sub-requirements can suitably be expressed as constraints on the optimization.

Extending theory to broaden applicability

Despite the success of the NUM framework, the line of research needs substantial extension along (at least) three lines: stronger understanding of global stability with delay in multi-layer protocol stacks; modeling of transient dynamics including flow and packet level dynamics; and capture of stochastics inherent in these dynamics. We describe the second of these in slightly more detail as an example.

The current theory focuses on convergence to a static optimal operating point. The primal-dual optimization model of TCP discussed in the sections above treats only the equilibrium rates of TCP flows, but says nothing about the transient trajectory. In practice, the dynamic nature of information within the networks and the evolution of the network itself necessitate analysis of transients and development of time-critical decision rules. Thus extending the static duality model to include transient dynamics using optimal control theory is essential, where part of the system dynamics become a constraint on the state trajectory over time. A promising approach with preliminary results interprets dual-based TCP congestion control as an optimal control law for the queuing dynamics, but much more remains to be done.

A second challenge for optimization based architectural decomposition theory lies in the multi-party, decentralized nature of the Internet and future “open” networks. A key requirement is to extend the current theory to better address issues of imperfect and local information. Extending the current theory to include concepts from game theory and mechanism design is critical to model distributed systems where agents have only local information and optimize their private objectives, and indispensable to understand the global behavior that will emerge from such interactions of local algorithms.

Challenge 2: Design by Constraint

Consider *design* as a verb. To design something is to carry out, with intent, an action or series of actions that lead to the creation of an artifact with certain properties.[4]

A *design methodology* is a model or approach to carrying out design. Much of classical engineering is concerned with the development of design methodologies and the analytic, modeling, and process tools that support them. A property shared across virtually all methodologies of classical engineering is that the outcome is the artifact itself. When an engineer designs a bridge, his job is to start with information about the requirements of the bridge-to-be, tables of material properties and so on, and the processes of his profession, and his aim is to produce a fully designed bridge.

We identify as a key challenge the creation of an entirely new class of design methodologies, different in scope than those of classical engineering. The unifying element of this class of methodologies is that they are concerned with *steering collective behavior over time* as opposed to *producing a final artifact.* Another way to say this is that what we are designing is an environment that causes the artifact we desire to come to be, rather than the artifact itself.

We refer to this class of design methodology as **design by constraint**. By "design by constraint" we mean the systematic creation of environmental elements and constraints such that the designed artifact that represents our final goal 1) emerges out of some ongoing collective process and 2) captures the designer's intent by exhibiting desired structural and behavioral features.

This concept of design goes well beyond the classical notion of engineering (e.g., bridge construction), yet it is far more grounded in domain-specific details than the types of emergent phenomena typically studied by the complex systems community. The research challenge is (1) to understand how the collective whole (in this case, a computer network such as the Internet) is shaped by forces and constraints acting at the local (i.e., microscopic) level and then (2) to develop the capability to *explicitly identify and specify* these local forces and constraints to achieve the desired outcome.

This objective creates considerable common ground with the broader network science literature, where a major emphasis in the study of "emergence" is exactly how local microscopic conditions give rise to the macroscopic collective whole. At the same time it goes significantly beyond the tenets of network science, which ignore the concept of intent. Within the network science community, one often encounters statements like "nobody designed the Internet, it simply emerged." From the perspective of classical engineering, where a single entity carefully plans and executes a design in support of a clearly articulated intent, this statement is correct. However, and contrary to the models of network science, the decisions of ISPs when building and operating their networks are not arbitrary or random, and suggesting that this process can be accurately represented as a sequence of (biased) coin tosses trivializes the very real engineering processes at work.

Example: Topology of the router-level Internet

To elaborate on what we envision, we consider as a concrete example the topology of the router-level Internet, which has been studied by both the Network Science and Network Engineering communities.

---

[4] Occasionally confusingly, *design* can also be a noun: the artifact that results from a design activity is itself a design.

The Internet is a federation of individual router-level[5] networks (i.e., where nodes are routers and links are connections between routers), each under their own administrative control. Each Internet Service Provider (ISP) is responsible for the design, provisioning, operation, and administration of one or more router-level networks, and ISPs interconnect these networks with one another in order to exchange traffic for a multitude of business reasons. Because there is no central authority on the Internet, the topology of the global router-level network is not directly controlled, engineered, or even known with certainty, making it a popular research topic for more than a decade.

Researchers in Network Science have primarily represented the router-level Internet as a *graph* (i.e., a mathematical object consisting of vertices and edges) and focused their efforts on identifying key structural features of this graph, as inferred from various sorts of measurement experiments. The stated intent of this class of research is to *abstract away* the domain-specific details of the system in order to isolate its most essential and presumably universal features.

For a complex system like the Internet, reducing all of the details to a simple graph means that there is not much to study other than connectivity properties, but this can be achieved using a combination of graph theory (to understand structure) and statistical mechanics (to understand evolutionary behavior). Thus, much of the work in the Network Science community has emphasized the identification of novel statistical features and signatures, along with the development of simple processes that can replicate them. Because simple generative processes often give rise to complex structures, this approach fits nicely with the belief that there are properties of the system as a whole that are not contained in the individual components and/or subsystems. Such properties are called "emergent", and major theme of the complex systems literature has been the study of emergent phenomena.

However, what is critical to recognize is that these simple models and generative processes have no notion of designer intent. The fact that simple intent-free models and generative processes can generate complex structures with the same statistical features and signatures as the real-world engineered structures they model can be (and has been) misinterpreted as implying that these simple statistical constraints and processes *can be used to design future versions of the real world structure.* But this is a significant logical fallacy. What is actually the case in this example is that the statistical features and signatures of the desired structure are an insufficient description, and thus that simple, intent-free models and generative processes appear to produce a good design result, without actually doing so.

An alternate approach, one that fits more closely with an engineering perspective, represents as constraints microscopic details sufficient to characterize the structure and behavior of the specific system of interest, as well as its underlying intent or purpose. For engineers, the articulated purpose often takes the form of a set of functional objectives and requirements. This approach allows one to ask questions of that particular system, typically in isolation. Models of larger systems require the inclusion of additional details of subsystems, and the overall system complexity quickly grows because of (1) a large number of components (complexity of size), (2) intricate relationships among components (complexity of interconnection), or (3) many degrees of freedom in the possible actions of components (complexity of interaction). A key issue in systems engineering is finding the "right" representation that manages the tradeoff between system fidelity and model utility. In the context of the router-level Internet, this means considering the functional objectives of the network, the technological capabilities of routers and the cables that connect them, along with the economic incentives of ISPs. Having gained

[5] The distinction between routers and switches is not significant to this discussion; we use "router" to encompass both.

sufficient understanding of a specific system and then validated it in detail, the challenge becomes to assess if and/or how the key features of that system generalize to other systems.

The actual artifacts that result from these two approaches can be dramatically different, even when they share the same statistical features. In our example, for graphs having highly variable (e.g., power law) statistics in their connectivity patterns, there is enough diversity such that two graphs exhibiting precisely *opposite* properties in terms of engineering attributes such as performance, reliability, and cost can nonetheless exhibit identical statistical properties. The figure to the left illustrates this point. Whereas one network is the most likely construction from a random process intended to match connectivity statistics, the other is the logical response to a need to provide high throughput while respecting the constraints on individual nodes. The fundamental goal of this research challenge can be stated as "understanding and controlling the full range of forces that would, in the real world, cause the desirable network on the left, rather than the undesirable network on the right, to emerge over time".

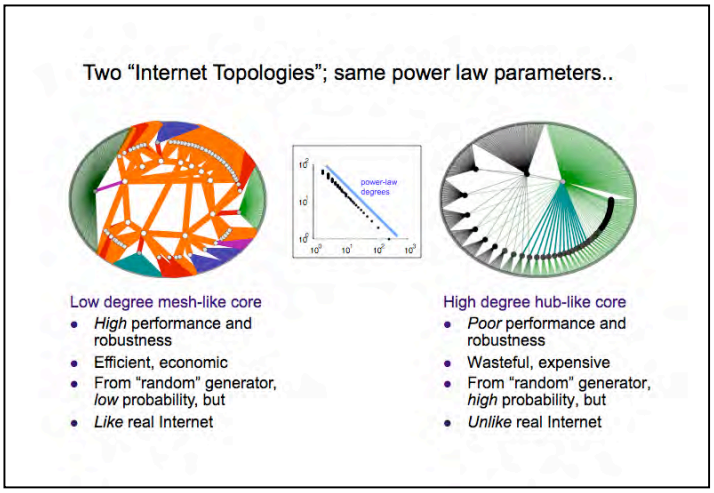

Validation

Critical to the research agenda described herein is the ability to accurately and realistically *validate results and conclusions.* This is particularly important because the agenda brings together different communities with very different traditions and perceptions of what useful "validation" entails. We call both for increased emphasis in establishing rigorous and accepted norms for validation, and for research specifically targeted at the development of *improved methodologies for validation* of conclusions that are intended to apply at significant scale and throughout the entire evolutionary lifecycle of an architecture. This problem is extremely challenging, because it requires acceptable validation of results across real systems that cannot yet be built and evolutionary events that have not yet happened.

It is easy to go wrong. It is ironic that "network science" has exploded in popularity in recent years while at the same time narrowing its meaning to something very different than what we envision. Unfortunately, the resulting errors and confusion have generated a backlash in certain communities that will have to be addressed if "network science and engineering" is to have legitimacy as a discipline. For example, four of the six most cited papers in ISI Web of Science searched with topics "internet and networ*" appear not in engineering journals but in Science, Nature, or Reviews of Modern Physics, and describe work primarily on scale free networks (SFN). The sensational claim that truly launched this genre (including the headline "Achilles heel of the Internet" on the cover of Nature) is that the router-level topology of the Internet is scale-free, and thus a few "hubs" (high connectivity routers) through which all packets must pass are crucial to the overall connectivity of the system and represent critical vulnerabilities (the unrecognized "Achilles heel") if attacked.

It is readily apparent that this line of research has little relation to the real Internet, where pragmatic engineers recognize that building "hubs" with millions of ports is not an appropriate engineering solution. Unfortunately, it did distract and confuse portions of the academic Internet research community for some time. Ultimately, research papers showing that SFN models are incompatible with both the real Internet and the underlying technology appeared in venues visible to the "networking research" community, and much of the confusion has receded. It is now well understood that alternative models are consistent with both data and engineering understanding, and that SFN's were a fiction of analysis error. Nevertheless, the "Achilles heel" remains for many physicists the canonical example of a discovery in network science. Put another way, the thesis and conclusions of the Nature article were judged as valuable and valid within one frame of reference, though they quickly failed this test once a more engineering-oriented perspective was applied. This example demonstrates crisply the need to build deeper links between researchers interested in intellectually strong "network science" and NSF/CISE's traditional "networking research" community.

Raising Standards for Measurement-Based Research

As a first step towards our goal of establishing rigorous and accepted norms for validation, workshop participants call for a renewed effort to raise the CISE networking community's standards for measurement-based research. We are certainly not the first to do so, but we note, as the example above illustrates, that the validity of measurement-based research becomes increasingly critical when interacting with scientific communities in which measurement and modeling are the primary intellectual activities, and even more so when its results are used to drive the progression of more fundamental understanding outlined in our matrix.

We argue that, within the larger context of "network science" and research on networks *per se*, the specific domain of computer and communications networking research has an exceptionally strong opportunity to contribute. Our argument is based on three points:

- The domain is accessible. The Internet and other computer networks provide a unique setting in which many claims can be unambiguously validated, albeit often with substantial effort.
- Because computer networks are both of great practical import and among the first highly complex networks to be subjects of serious academic study, the computer networking community has available a vast amount of domain knowledge accumulated over time.
- The computer networking community's engineering orientation has created a tradition of viewing domain specific knowledge as central, which provides an important balance to the existing "network science" perspective that domain-independent knowledge is more fundamental.

Building on these points, we argue for a two level objective. Our short-term call is to improve the quality of measurement-driven research in the computer networking community, by jumpstarting discussion involving the community as a whole, by developing canonical examples, and by establishing suggested sets of guidelines and standards for research of this type that will be easy to apply to specific measurement experiments. Our longer-term call is to use our field as a vehicle to improve the quality of network science overall, by providing an exemplar for the conduct of scientifically valid measurement based research in other domains of network science and by

demonstrating the value of bridge-building between domain-independent network science and domain-specific research areas.

Where to start?

This section of the report summarizes workshop discussion about the objective identified above.[6] We begin by identifying a basic and obvious, but too frequently ignored, question central to all research that is rooted in observation and measurement:

> **Q:** Do the available measurements and their interpretations, analysis, and modeling efforts support the claims and observations derived and presented?

To approach this question more structurally, we ask the related question "what are likely sources for errors that would make the above not true"? Measurement-driven research and its validation is largely a lesson in how errors of various forms occur and add up. We identify and discuss four sources for error: the measurement process itself, the analysis of the resulting data, the modeling work that is informed by this analysis, and the work undertaken to validate the resulting models.

We offer a set of questions that address these sources of error. These questions are intended to be trigger points for further discussion, methodology development, tool building, and similar activities. Summarizing, each question can be viewed as the root of an activity or challenge – the development of cultural standards, techniques, and tools that both allow and encourage individual researchers within the community to answer the question in the affirmative. We suggest that such activities be made explicit within the NetSE research program.

> **Q1:** Are the available measurements of good enough quality for the purpose for which they are used *in the present study*?

A particular focus of this question is to increase the community's focus on the creation and association of *meta-data* with any newly collected dataset. This suggestion is made frequently but seldom honored in practice. We believe it is paramount to revisit the meta-data concept, and develop it to the point where its utility becomes obvious and its implementation becomes as straightforward as possible. A first impediment is that there is today no "best practice" for collecting and maintaining meta-data with a dataset. Such guidelines should be developed, with the aim that meta-data description should include as much information as possible that is pertinent to the collection of the data and its future use by third parties. Ideally, this would include details about the measurement technique used, its shortcomings and limitations, and alternatives considered but not adopted. It should spell out in detail any issues concerning bias, completeness, accuracy, or ambiguity of the data that are known as a result of the data producer's in-depth understanding of the measurement and data collection effort. If possible, it should include any relevant information about the operating conditions of the network at the time the measurements were made that might impact validity of the data for subsequent studies by the producer or others (e.g., infrastructure or protocol-specific aspects, network usage, and application mix).

Such metadata is particularly relevant when data originally collected for one study is reused for another purpose. Providing convincing evidence that an existing dataset can be used for a very different purpose than was originally intended is a crucial responsibility of the user of such data.[7] At minimum, this requires a detailed account of the assumptions that are made about the dataset

---

[6] Discussion on this topic, and this summary, were driven by Walter Willinger's presentations at the workshop.

[7] Failure to do so has been the trigger for several spectacular mis-interpretations.

and a list of issues that a carefully crafted meta-data description of the measurements should or did address. While increased reliance on canonical datasets – common in other areas of science – would be clearly useful, we find that it is generally impractical today, both because our community lacks sufficient standards for the metadata that would make such datasets reliable, and because applicable situations may be less common in the Internet measurement field, where the underlying conditions tend to undergo constant change.

> **Q2:** Is the level of statistical rigor used in the analysis of the data commensurate with the quality of the available measurements?

After assessing the overall quality of available measurements, the next step towards improving measurement-based research as a whole concerns the quality of the analysis of the data. At issue is how to analyze datasets that are, in general, tarnished by various documented or undocumented types of errors and imperfections, yet contain some amount of useful information. Mining that information is at the heart of the problem and requires a data analytic approach that matches well with the quality of the measurements.

Clearly, it makes little sense to apply very sophisticated analysis techniques that are highly sensitive to inaccuracies in the data if the datasets have been identified to exhibit major deficiencies. In fact, the biggest take-away points from measurement studies are often in "broad rules of thumb" and not in details. For example, an observed Pareto-type principle or 80/20-type rule is often all that can be reliably and robustly inferred from high-variability data of questionable quality, and any attempt at fitting a specific parameterized model (e.g., a power-law type distribution) would be statistical "overkill". In this sense, the question concerning statistical rigor cuts both ways – reliance on statistically sophisticated methods in situations where the data don't justify their application should be as much frowned upon as avoidance of statistically rigorous approaches in cases where the data at hand justify a detailed and more elaborate analysis. Q2 is intended to raise the general awareness that there are important differences between analyzing high- and low-quality datasets, and that approaching the latter the same way as the former is not only bad statistics but bad science.

> **Q3:** Have alternative models that are also consistent with the available data been considered, and what criteria have been used to rule them out?
>
> **Q4**: Does the model validation go beyond showing that the proposed model is able to reproduce certain statistics of the data used to construct it?

For measurement-based research studies that include a substantial modeling component, an all too frequent description of the modeling element can be succinctly summarized as follows:

- Start with a given dataset, taking the available data at face value.
- Next, infer some distributional properties of the data, and determine the "best fitting" model (e.g., distribution, temporal process, graph) and corresponding parameter estimates. Here "best fitting" refers either to a subjective or "eyeballing" assessment of the quality of the fit, or to an evaluation involving some commonly used goodness-of-fit criterion.
- Lastly, argue for the validity of the chosen model by virtue of the fact that it reproduces the distributional properties of the data examined in the second step.

The commonly used recipe for network-related modeling described above has reduced this activity to an exercise in data fitting, a mostly uninspiring activity that creates little excitement and is generally detrimental to scientific advances. The reasons for this are all too clear. First, the approach gives little insight, as the recipe is guaranteed to produce *some* model. In fact, for the same set of distributional properties, there are often many different models that fit the data

equally well. Further, depending on the distributional properties of interest, the resulting models are likely to be significantly different, and rarely do there exist solid guidelines within this methodology for choosing among equally well-fitting models. Finally, given that more often than not the available measurements cannot be taken at face value, providing a precisely accurate encapsulated description of the data at hand is largely counterproductive.

The area neglected by this widely accepted approach is model validation. Models are generally declared to be valid by virtue of the fact that they reproduce the very same statistics of the data that played a key role in selecting the model in the first place. But being able to reproduce some statistics of the data that created it, while useful in some limited circumstances, is the most simple and uninteresting product of a good model. This capability alone gives little confidence that the model has captured any *fundamental* property of the system being modeled – and thus, that the model is of any value beyond describing the particular dataset from which it is derived.

To develop a more scientifically grounded and constructive model validation methodology, an initial suggestion is to make matching particular statistics of its generating data a non-issue in model evaluation. After all, a model that is "approximately right" (in the sense of Mandelbrot) can be expected to implicitly match most statistics of the data, at least qualitatively. A concrete procedure that would more strongly increase confidence in a proposed model is to examine it in terms of what new types of measurement properties and statistics it identifies that are either already available but have not been used in the model's generation, or that could be collected and used to check the validity of the model. Here, by "new" we do not mean "the same type of measurements, but more". Instead we look for completely new types of data, with different semantic content, that have played no role in the modeling process up to this point. A key benefit of such an approach is that the resulting measurements are only used for the purpose of model validation. While evaluating models according to this methodology does not reach the level of confirming "causal understanding" within our matrix, it is apparent that models that successfully predict the properties of data outside of their generating set are likely to be stronger than models for which this is not shown to be true.

Workshop Attendees

David Alderson
Naval Postgraduate School

David D. Clark
MIT CSAIL

Heidi Picher Dempsey
GENI Program Office

John Doyle
Caltech

Darleen Fisher
NSF

Fan Chung Graham
UC San Diego

C. Suzanne Iacono
NSF

Ali Jadbabaie
University of Pennsylvania

Will Leland
Telcordia Research

Dmitri Krioukov
CAIDA

R. Srikant
University of Illinois

Walter Willinger
AT&T Research

John Wroclawski
USC Information Sciences Institute

Ellen Zegura
Georgia Tech

Ty Znati
NSF

Workshop Agenda

Starting points

Framing Talk:

Leland – Army/NRC Network Science Study – topics, conclusions, open questions, issues raised.

Materials available

NRC Network Science Study (Executive Summary distributed, copies available)

M. Mitzenmacher's Internet Mathematics editorial (copies distributed)

Alignment and Analysis

(What's being aligned is the group's understanding of state of the art and key issues in several areas related to Network Science and Network Design. It would be ideal to have the Topology discussion first, but reordered so that Dmitri Krioukov could be here for it)

1) Control, Dynamics and Optimization

Framing Talk:

Srikant – Big picture CDO in the Internet Context

Jadbabaie – Flocks, swarms, and syncronization

Discussion

2) Traffic and Traffic modeling

Framing Talk:

Willinger – status and future direction

Discussion

3) Topology

Framing Talks:

Alderson – Review of progress, persistence of confusion

Graham – Challenges in the study of large graphs and algorithms

Discussion

4) Architecture

Framing Talks:

Clark – A methodical approach to modularity

Doyle – "architecture=constraints" framework for protocol-based architectures

Discussion

Synthesis

Framing Talk:

Doyle – Relating theory and design

Formulation of framing questions

Discussion of research challenges

Methodology and Infrastructure

Framing Talks:

Willinger – Raising the Standards for Measurement-based Networking Research

Wroclawski – Mind the gap – some troubling current assumptions about research infrastructure

Wroclawski – Evaluating Architecture

Discussion

Close

Appendix 4

Workshop Report on

Network Design and Engineering

NetSE Network Design and Engineering

Network design is unlike most other engineering practices. Rather than creating a single artifact that satisfies a set of requirements and obeys practical constraints, network architecture requires a *framework* that enables heterogeneous networks to address different requirements that evolve over time and in a distributed manner. Since the systems that operate and control these networks are inherently programmable, networks can continuously change to take advantage of new technologies and/or to address new societal needs.

The Internet today is the most complex network of networks ever constructed. The openness of its framework has allowed it to grow into a comprehensive global infrastructure that fulfills many of society's commercial, educational, health, communication and entertainment requirements. But many network goals remain elusive. Security, privacy, energy efficiency, and greater reliability are examples of network challenges that have not yet been solved.

As designers and engineers, how might we best move forward to achieve the numerous and often conflicting design goals that would make our future networks better? In this report, we identify four complementary research areas that we believe will lead to better network design and engineering and ultimately to better networks.

First, we must learn how to **satisfy competing design goals**: Designing a network that satisfies any one design goal---such as ensuring security or enabling innovation---is already quite challenging. Yet, the real task is to reconcile numerous seemingly conflicting goals into a coherent system. We must understand and precisely model the inherent trade-offs among competing goals to recognize when no solution can possibly satisfy them all. In addition, we must search for ways to eliminate false conflicts and expand the space of feasible solutions. And, where possible, we must create new solutions that strike a careful balance between different goals, or even enable a range of simultaneous solutions that prioritize different goals for different users or different applications.

Second, we must better **leverage the programmable nature of network infrastructure:** Many emerging technologies, such as Field Programmable Gate Arrays (FPGAs) and Photonic Integrated Circuits (PICs), enable much greater programmability, at reasonable performance, than ever before. We must find effective ways to harness these new capabilities, while still creating systems that are not too complex for us to model, build, deploy, manage, and evolve.

Third, we must **reconsider protocol layering:** From the early days of networking, protocol layering has been an invaluable tool for breaking complex problems into smaller, more tractable parts. Yet, protocol layers also stand in our way, leading to inefficiency particularly in the face of new technologies (such as wireless networks) and new applications (such as multi-player games) that violate old assumptions.

Fourth, we must better **address broad societal needs**: As the Internet becomes ever more ubiquitous, our design choices have increasingly serious consequences for the larger society. Future networks could become a major drain on global energy resources, or an important enabler for a greener planet. Networking could further widen the digital divide between rich and poor nations, or empower the developing world with access to critical information and a platform for untold innovation. Networking could put users' privacy at risk by concentrating sensitive information in remote locations, or enable a wide range of applications that simplify people's lives and improve their health. To play a positive role in society, the design of future networks must consider design goals that, while hard to articulate and measure, relate directly to basic human concerns.

Below, we discuss these four research areas in more detail. We describe the current dilemmas in each area of network design, emphasize the many goals that we must try to satisfy, and articulate some research directions that can lead to future networks that are increasingly trustworthy and beneficial to society.

Satisfying Competing Design Goals

While satisfying any one design goal in isolation may be relatively straight-forward, the fundamental challenge in network design lies in reconciling trade-offs among seemingly conflicting design goals. While there is a long list of competing design goals, we start the discussion here by selecting three pairs of conflicting goals that illustrate the particular tensions that need to be worked out. We also point to promising research that finds the "sweet spot" between the competing concerns.

Secure Networks that Enable Innovation

A common theme when securing a network is "that which is not expressly permitted is denied." This conservative approach reduces the odds of being surprised by the many threats that were not, or could not be, anticipated in advance. Yet, the Internet arguably owes its tremendous success to the ease of adding new applications, requiring only the cooperation of software running on two or more end-host computers. More importantly, the Internet lowered the bar for *who* gets to innovate, "democratizing innovation" by allowing anyone who can program a computer to create and deploy a new networked application. This model lies in stark contrast to special-purpose networks, such as the telephone network, where innovation requires central approval and global deployment.

The tension between security and innovation has been playing out on the Internet over the last two decades, as the network extends into mission-critical niches and as more and more sensitive data (financial, medical, and other data confidential to individuals or organizations) is transmitted across the network and stored on hosts attached to the network. Network designers and administrators require ways to ensure the availability, integrity, and confidentiality of their networks and data, pushed by high-profile incidents of customer information loss and privacy protection legislation such as HIPAA. And while such protection could, in theory, be

provided by a combination of end-to-end encryption and perfect host security, the complexity of modern applications and operating systems renders that approach impractical: a secure network today must greatly restrict the traffic that can cross it.

This conflict plays out in many forms. Oftentimes, users have to request changes to the configuration of a firewall to run emerging applications. In some cases, security software becomes a dependency that must be upgraded before new applications can function. Other applications, such as some peer-to-peer applications or the popular Skype voice-over-IP software, take the offense, masking their traffic or probing for ways to punch through firewalls. This state of affairs is *not* a solution: it frustrates users, slows the deployment of useful applications without effectively preventing the spread of malware, and creates an adversarial relationship between users and administrators---hardly a climate conducive to either productivity or security.

This tension raises a fundamental question for network research: How can networks provide both the assurance needed for them to be mission-critical, while providing the flexibility and openness to new applications that allow them to become critical in the first place? What is the allocation of responsibility between applications, host software, and the myriad network elements to achieve these goals, and how do these elements coordinate or trust each other---if at all---in doing so? Is there a level of abstraction at which to specify and enforce security policy that protects information or integrity instead of blocking applications in the hope of achieving that goal? Or, could we side-step the tension between security and innovation by supporting a range of solutions, each striking a different balance between the competing trade-offs, in parallel on a common network substrate?

Economically-Viable Reliability

The Internet is remarkably vulnerable to equipment failures, as evidenced by the serious disruptions in communications during the recent fiber cut in the Mediterranean Sea and other high-profile outages. The network also fails due to software bugs in the implementation of crucial protocols and, often, operator configuration errors. The routers in the network are slow to compute new paths that circumvent the failures, and sometimes suitable paths are never found despite the existence of spare resources. In addition, hidden dependencies (such as the multiple fiber-optic cables affected by the Baltimore tunnel fire a few years ago) often mean the network is less reliable than expected. Given the increasingly important role of communication networks, the future Internet should put a much greater premium on reliability and faster recovery from unavoidable failures.

Improving reliability is especially challenging because more reliable components and greater redundancy are both expensive, and different administrative domains often do not cooperate with each other. For example, from a systems perspective, the simplest way to improve reliability is to use more reliable components. Yet, highly reliable components are notoriously expensive, leading to the principle of building reliable systems out of unreliable components that (hopefully) fail

independently. However, economic motivations drive decisions that compromise reliability. To reduce equipment costs and operational expenditures, a provider may deploy the same model of equipment (with the same software bugs) at many locations. Also, different providers often co-locate their networking equipment in carrier "hotels" to reduce cost and simplify peering.

A deeper treatment of reliability needs to incorporate the dependencies across components and protocol layers, as well as the necessary economic incentives for improving reliability. For example, although different administrative domains are often business competitors, they could form more complex "peering agreements" for providing backup connectivity to each other during failures. Similarly, these networks can cooperate to identify shared risks and accurately model the influence of component failures on overall system reliability, without revealing proprietary information about their underlying network designs. In addition, network protocols could react more quickly to failures if appropriate information about the "root cause" of a problem were more readily available. Yet, any technical solutions must be coupled with appropriate economic incentives for the parties to participate honestly, or an accountability framework that can accurately identify the party responsible for a reliability problem.

Improving network reliability is an exciting and important research challenge for the future Internet, with ample scope for interesting interdisciplinary research that acknowledges the inherent tension between reliability and economic incentives.

Scalable Support for End-Host Mobility

With the proliferation of cell phones, PDAs, and laptop computers, end-host devices are increasingly mobile, and users increasingly demand seamless communication on the move. Sometimes large groups of hosts move together, as when a plane, train, ship, or tank changes locations in the network. In addition, "virtual" servers often migrate from one physical computer (or data center) to another to balance load and reduce energy consumption. Yet, most of today's networking technologies were designed to support communication between fixed end-point devices. Future networks should be designed with mobility as the norm, rather than the exception.

The desire to support seamless host mobility is seemingly at odds with the use of hierarchy to make large systems scalable. In particular, the Internet architecture ties end-point addressing to the host's location in the network topology, through the use of hierarchical IP addresses. Internet routing is based on 300,000 or so IP address blocks (or ``prefixes'') rather than hundreds of millions, or even billions, of individual host addresses. While local area networks often use so-called ``flat'' addressing (based on Medium Access Control, or MAC, addresses), these protocols do not scale to larger network configurations. Though researchers and practitioners alike have proposed new architectures that separate identity from location, we are still a long way from understanding the fundamental trade-offs between mobility and scalability, particularly in a wide-area setting.

Beyond the challenges of supporting end-host mobility, mobile hosts present new *opportunities* for delivering data. In wireless ad hoc networks, in particular, an end-host computer may simultaneously serve as a user device and as part of the underlying network infrastructure. While mobility makes routing protocols more complicated, the movement of the hosts provides an alternative way to transport data---on the computer itself! Recent theoretical work has shown that mobility can actually increase the capacity of a network, and such "opportunistic" networking is already seeing deployment in applications diverse as wildlife monitoring and communication in the developing world. Designing scalable protocols and practical systems that can exploit this extra "bandwidth," as well as incentives for hosts to contribute their resources to others, is an exciting opportunity for future research.

Leveraging a Programmable Infrastructure

Today's network stack came into being in a world where the lower layers were realized in hardware or subject to standards that were hard to change (or both). In this context, one role of the Internet Protocol (IP), and TCP and UDP above IP, was to provide a common, mediating interface between the dynamism of new applications (created in software, often overnight) and the comparative stasis of the lower protocol layers (created on the timescale of years). The lower layers had a fixed role to play, and IP's job was to make the different approaches taken at the lower layers have a common appearance to applications.

Times have changed. Just as the needs of applications have expanded, technological advances have substantially lowered the barrier to building programmable networks. The most visible aspect of this change is in the physical layer, where a combination of hardware innovations such as swiftly reprogrammable FPGAs, photonic integrated circuits (PICs), and improved DSP performance are enabling software-defined and software-determined implementations of the lower layers of the network stack (PHY, LINK, MAC layers). We are also seeing ever-increasing programming power inside middleboxes such as routers and firewalls, and massive parallelism in computing clusters is increasing the programming power at the edges of the network. Now that the underlying equipment is more programmable, IP (and TCP and UDP) has become a barrier that blocks new and innovative applications from capitalizing on new innovations in lower layer protocols.

Further complicating this picture, observe that the growth in computational power in middleboxes implies that we could conceivably run *multiple* protocol stacks side-by-side. In the extreme, each application could create its own custom network, using protocols that are tuned to the application's particular needs. In this environment, the whole notion of a "network protocol stack" disappears entirely---rather, a device simply runs the software package appropriate to the circumstances. Yet, programmability does not imply that interfaces are no longer necessary. The absence of interfaces is *not* an architecture; it is chaos. If every device has its own software environment, then deploying innovations would require dozens of vendors to update the software for all of their devices and then work among themselves to wrangle out the bugs and incompatibilities. Past history suggests this will inevitably

force innovations through a standards process that slows progress. To enable innovation at the pace of compilation, we need effective ways to run the same software across multiple platforms, as well as safe environments for running multiple customized protocol stacks in parallel on a shared substrate.

Reconsidering Protocol Layers

Designing protocols in layers, each offering a service to the layer above and each relying on services of the layer below, has been a fundamental principle in networking for decades. Protocol layering has been an immensely successful example of using modularity to manage complexity, by raising the level of abstraction for building networked services and enabling significant "code reuse." However, protocol layering has presented a number of challenges. It can lead to increasingly diverse applications that impose a wide range of *different requirements* on the network, begging the question of whether a "one size fits all" design is sufficient for the future. In addition, *network management*, a perennial problem facing the Internet, fundamentally requires the ability to look across layer boundaries, raising all sorts of important questions about how to design future networks that are inherently easier to manage. In this section, we discuss both of these issues.

Applications with Increasingly Different Requirements

Early networked applications, like bulk file transfer and electronic mail, were relatively natural fits with IP's model of best-effort packet delivery. Over time, the Internet has begun supporting a much wider range of applications with more diverse and challenging requirements. For example, interactive applications like Voice over IP (VoIP) and online gaming put a higher premium on low delay (as opposed to, say, high throughput) than earlier applications. In addition, these applications often involve communication between multiple parties, such as the many participants in a teleconference or the many characters in a virtual world, placing a higher demand on delivering the same data efficiently to many receivers. While these kinds of applications do run on today's Internet, they often experience serious disruptions that degrade the quality of the user experience; in addition, application developers face serious hurdles in "working around" today's architecture.

Similar to the traditional phone network, early applications supported communication between fixed end-points---telephones in the case of the phone network, and end-host computers in the case of the Internet. Increasingly, though, communication revolves around *content* that may reside on multiple computers spread throughout the Internet. For example, a Web page or MP3 file may be stored on many different computers, and the set of computers that can provide the content changes over time. These changes may occur because computers fail, or new peers arrive with their own copy of the content (or copies of portions of the content), or because load-balancing policies stop directing users to overloaded machines or network paths. The Internet architecture is surprisingly brittle in the face of these

kinds of dynamic changes to content-to-host mappings, something that arguably warrants serious revisiting in the future. Existing solutions map content names to hosts too early (as in DNS-based redirection) or too late (as in IP anycast).

In addition, many Internet protocols were designed under the assumption that the communicating hosts are both online at the same time, and have a path between them with a relatively low delay and loss. Existing Internet protocols perform quite poorly on paths with long round-trip times (e.g., as in satellite networks), or in wireless networks where the basic abstraction of a "link" does not really exist and interference can cause packet loss even in the absence of congestion. The Internet protocols are especially ill-suited to disconnected operation, and to entirely non-real-time tasks (such as nightly backups) that merely need to complete by the next morning. Most protocols explicitly create and maintain some kind of continuous session, or connection, between the communicating hosts, rather than trying to gradually move data closer to its ultimate recipient.

The traditional layering of protocols does not meet the needs of these applications. The link layer, by hiding details about the underlying physical medium, does not allow higher layers to adapt to network conditions. The network-layer routing mechanisms do not select different routes based on minimum delay (for some applications) and maximum throughput (for others). Assigning names and addresses to machines, rather than content, makes it difficult to handle churn in the relationship between content and its current location; in particular, the strict separation of naming and routing, while conceptually appealing, makes it difficult to adapt quickly when content moves to a new place. All of these concerns warrant a revisiting of the traditional layers in the network stack, with an eye toward customized stacks for different classes of applications.

Network Management and Protocol Layering

At some level, today's Internet manages itself. Routers automatically compute new paths when equipment fails, and end hosts adapt their sending rates automatically in the face of congestion. Yet, these mechanisms do not ensure that a network runs *well*. In practice, modern networks require an immense amount of day-to-day attention from chronically overwhelmed cadre of human administrators. In practice, the cost of the people and systems that run a network far exceed the cost of the underlying equipment; yet, despite the significant financial investment, more than half of network outages are caused by operator error. Clearly, part of the problem lies in the fact that innovations that reduce the costs of network management have not kept up with the Internet's growth. Another problem is that network management is, inherently, all about the things we don't (yet) know how to do, and requires looking across network elements and across protocol layers to diagnose problems and effect policies.

Network administrators monitor their networks, and tune the configuration of the network elements, to achieve a wide variety of goals. They need to balance load in the network, block unwanted traffic, make good on their promises to other

networks (such as their customers and peers), perform routine maintenance without disrupting existing applications, and make a profit along the way. All too often, these goals are a mismatch for the capabilities of the underlying protocols and mechanisms, forcing the network administrators to "work around" the limitations of the existing technology. Where layer boundaries are narrow to facilitate abstraction, network management requires visibility across the layers. Where each layer is meant to operate independently, network management needs to tune parameters (such as timers) to ensure an efficient interaction between the layers. Where layers hide the fact that multiple links share the same underlying physical risks, network management much expose this information to improve reliability.

As such, network management is often at odds with the notion of clean, simple layer boundaries. One solution is to keep the protocol layers simple, while having separate management systems collect and combine extensive measurement data at each layer, across different network elements. Another solution is to "widen" the boundary between layers to enable more adaptive network protocols that partially obviate the need for separate management functionality. The move towards greater programmability may be a help, or a hindrance. It remains unclear whether running several customized protocol stacks, each carefully designed for a particular class of applications, is inherently easier or harder than managing a single compromise solution that partially meets the needs of each application.

Addressing the Broader Needs of Society

Many of the most exciting challenges in network design arise when technical questions meet important human concerns. In this section, we touch upon three issues that are particularly relevant today. First, everyone is concerned about climate change and efficient energy consumption. How do we leverage networking technology to help reduce global energy consumption, and also reduce the energy the underlying network consumes? Second, universal access has not yet been achieved. How do we ensure that the developing world enjoys the many benefits of communication networks, rather than letting the "digital divide" be yet another example of the wide chasms between wealthy and poor nations? Third, the rise of cloud computing allows us to rethink long-term data management. How do we enjoy the convenience and lower costs inherent in cloud computing, without sacrificing user privacy? These and other similar questions require creating and evaluating new network designs under a much wider, and inherently subjective, way of judging whether one kind of is truly "better" than another.

Green Networking

Communication networks can play a major role in reducing global energy consumption by enabling remote collaboration to reduce the need for people to travel, or through distributed monitoring and control of energy use in buildings. In addition, information technology itself consumes as much energy as the automotive manufacturing industry, and will soon rival commercial aviation. Innovations in communication networks can enable selective "powering down" of data-center

servers and consumer devices---both notorious power hogs. In addition, the underlying network could consume less energy through new low-power equipment and techniques for selectively shutting down switches and routers. With the growing concern about global warming and energy costs, networking research can move beyond an early preoccupation "bigger and faster" to emphasize greater power efficiency.

While "going green" is clearly appealing, and early research suggests we may be able to become substantially more energy efficient without sacrificing performance, becoming energy efficient does introduce tensions with other important goals such as high reliability, low cost and predictable performance. For example, a centerpiece of energy efficiency is turning off underutilized components, but repeated powering down of equipment (whether end hosts or routers) often makes the components less reliable. In addition, powering down servers and routers, or consolidating functionality on fewer components, reduces the level of redundancy, lowering the reliability of the overall system. Similarly, to ensure that a network path or a server is available when needed, we often cannot turn components completely off (getting them restarted takes too long), rather we have to find ways to dynamically tune their power consumption to match demand (e.g. running a link at half speed when underutilized or turning on only some processors in a multiprocessor server). Gracefully shifting gears, from one power level to another, is a complex task and, if done wrong, can easily lead to painful (if often transient) performance problems . Finally, while "going green" can save money, particularly in terms of power and cooling, complex techniques for batching workloads, powering up and down components, and consolidating and migrating tasks can introduce additional equipment and management costs.

We still have a relatively limited understanding of the trade-offs between energy consumption and these other system goals. Future research can create accurate models of these trade-offs, and explore solutions that strike a careful balance between the competing goals. A crucial issue is to accurately predict future traffic demands, and to provision for expected variations in the traffic, to provision the appropriate amount of resources while allowing the rest to sleep. Otherwise, performance would suffer, or require frequent powering up and down of the equipment (which itself consumes power). Another important issue is to develop rapid techniques for powering up networking equipment, to allow networks to safely power down excess capacity without compromising responsiveness to unexpected failures. In addition, new network architectures could have better support for planned shutdown and migration of servers, even across data centers, to enable seamless service to end users while reducing energy consumption.

Networking in the Developing World

Much of the developing world remains disconnected from the Internet. In mid 2008, Internet World Statistics estimated a world-wide penetration rate of 21.9%; average penetration in Africa was just 5.3%. Internet penetration correlates with other measures of developmental progress. Indeed, developing regions are characterized

by scarcity on many dimensions: low education levels; little to no reliable, fixed infrastructure; geographic inertia and local livelihoods; low income; shared, low-capability endpoints; limited resources that are often hard to locate or obtain; frequent pressing needs, for example for food or medical supplies. These forms of scarcity imply that first-world solutions will often be unsuitable because they rely on resources (people, infrastructure) that are simply not present in the developing world.

For example, Internet routing is based on shortest- and single-path routing. Even in cases where topologies contain redundancy, that redundancy tends to be used for backups rather than as a normal mechanism to increase reliability or capacity. In developing regions, wireless is the most natural choice for rapid deployment with limited fixed infrastructure. Wireless links are well known for their varying reliability as well as their natural broadcast capability. Both characteristics suggest that multi-path routing is the appropriate fundamental approach, rather than single-path routing. This is a simple conclusion to reach, but the ramifications are considerable. The fundamental end-to-end reliability protocol TCP assumes in-order delivery with most losses due to congestion, not link quality. Hence a multi-path approach challenges the layering architecture that defines the current Internet.

The challenges of connecting the developing world are daunting, precisely because they span technology, policy, and society. A suitable technical solution cannot be developed in a sterile first-world lab; instead a close collaboration is required between technology experts and policy experts. Involvement of local people in local solutions is also a powerful, and perhaps even necessary, approach. We see opportunities for new algorithms, new architecture and new technology that are deeply coupled to an understanding of policy and societal norms.

Cloud Computing and Personal Information Management

Increasingly, we are seeing arguments for and opportunities to migrate storage, computation, and applications into a "cloud" of infrastructure maintained by third party providers. This infrastructure is attractive for users, who today generate large volumes of data they must manage, using cell phones, cameras, voicemail, digital video recorders, and desktop and laptop computers. A cloud-based infrastructure offers the hope of consistent and easy management of this data; today, each of these devices is an independent entity, with different formats, management interfaces, networking capability, and storage mechanism. Businesses, too, see potential benefits from cloud computing: managing a capital-intensive infrastructure as an expense, with outsourced management, and whose scale can be increased or shrunk at software timescales. With cloud computing, however, come enormous challenges of security and privacy; creating and maintaining the cloud infrastructure to meet demand; and even such factors as enabling a competitive market.

Computer users today are, often unwittingly, forced to act as amateur system administrators---performing backups, copying data from device to device,

upgrading software, and over the long term, determining ways to ensure that valuable personal data is accessible using new software and new computers. A seemingly simple question such as "how do I make sure I don't lose my digital photographs?" rapidly meets the cutting edge of research. Cloud computing offers a tantalizing answer to these questions: trust the cloud, and let an outsourced service manage it all for you.

Migrating to a cloud of computation and storage holds a number of potential benefits, including: consistent access to one's data from a range of devices, simplified sharing of documents and personal media, centralized administration of operating systems and application binaries, managed backup and archival of important data, and the ability to leverage tremendous computation on-demand without statically provisioning for peak requirements.

Unfortunately, it also offers significant challenges: personally-generated data has a wide range of sharing semantics, from fully public information (e.g., perhaps non-personal vacation photos), to that shared within myriad social networks and sub-networks, to highly restricted (e.g., financial or medical data). And while the online service may maintain backups, how does a user guard against the failure or shutdown of their chosen cloud services?

On the business side, consider a small company that wishes to be prepared for rapid growth in access from its user base or an established company that may see peak levels of demand that are an order of magnitude larger than its average case. In both cases, the companies must build out, manage, and power a substantial computing infrastructure that must necessarily sit idle for much of the time. The availability of third-party providers able to dynamically deliver computing and storage on-demand for a pre-determined fee holds the promise of significant efficiency, cost, and even energy gains.

This model is gaining traction, with product offerings from some of the largest companies, including Amazon, Google, and Microsoft. The primary factors for its increasing popularity include convenience and cost. From an end user perspective, it is simply more convenient to "outsource" the system management to a competent third party. In the face of application vulnerabilities, limited disk space, protecting against machine failure, and ever-changing application versions, simply running a personal computer can be a significant time investment. Certainly, corporations invest tremendously in their IT staffs. The advent of virtualization and increasingly ubiquitous wireless data access have been some of the last missing ingredients to enable third-party companies to run entire operating systems at a distance on behalf of remote users.

While there are significant opportunity and benefit from this emerging computing model, there are significant risks as well. First, we will be making all of our personal information and, in the case of companies, the personal information of others available to a third party, sacrificing privacy. Second, we may lose predictability in

resource availability. For instance, a company that wishes to leverage the cloud to absorb bursts in access demand may find that no additional resources are available because of the access patterns of other clients using the same provider. Similarly, verifying SLA's for long-term availability may be difficult to perform externally without understanding exactly how replication and network topology is managed internally within the cloud. Third, moving one's data and application into the cloud may result in lock-in to a single provider. Consider an exaggerated version of migrating one's cell phone number from one provider to another, but in the case of Terabytes of data, database information, application logic, and state. Finally, by outsourcing computation and storage to a small number of providers, we compromise on the ability of the research community and smaller companies to innovate. Much of the ability to provision, replicate, and manage the infrastructure is predicated on observing the access patterns of a large user population. This information will almost certainly be considered a trade secret, preventing others from developing novel systems and networking solutions.

The above challenges point to significant hurdles with the adoption of the cloud-computing model. However, since we are still in the relatively early stages of defining the architecture, use cases, and economics of cloud computing, there is also opportunity to carry out important research to ensure that the underlying cloud computing architecture delivers on the promised benefits while navigating some of the difficulties embedded within the current trajectory. As some initial examples, consider the ability to dynamically migrate computation and storage among a conglomerate of providers, with appropriate resource peering agreements in place to prevent lock in to a single provider and to enlarge the pool of potentially available resources in the face of demand bursts. Similarly consider an openly adopted API that enables trusted third parties (the equivalent of a Public Utilities Commission) to audit performance, security, and availability SLAs provided by individual cloud providers.

Workshop Attendees

David Anderson, CMU

David Clark, MIT

Aaron Falk, BBN/GPO

Nick Feamster, Georgia Tech

Darleen Fisher, NSF

Suzi Iacono, NSF

Dmitri Krioukov, CAIDA

Craig Partridge, BBN

Jennifer Rexford, Princeton

Amin Vahdat, UCSD

Ellen Zegura, Georgia Tech

Ty Znati, NSF

Appendix 5

Workshop Report on

Network Design and Societal Values

WORKSHOP ON NETWORK DESIGN AND SOCIETAL VALUES

September 24-25 2008
Arlington, Virginia

INTRODUCTION

Digital electronic networks have emerged as one of the most powerful and exciting technologies of the late 20th and early 21st centuries, embodying and promoting wide-ranging societal and individual aspirations to create, produce, communicate, buy, sell, organize, connect, associate, educate, learn, entertain, campaign, and collaborate on a local, community, national, and global scale.

One mark of a great technology is its capacity to transform and be transformed. This we have witnessed in the relatively short lifespan of digital electronic networks, as societies have reacted to them and, in turn, shaped and reshaped them in multiple iterative cycles of mutual transformation. For scientists and engineers, the challenges are legion. In this document, however, we report on some of the complex interactions between network science and technology and societal values, focusing on moral, political, and sometimes also cultural values.

BACKGROUND

The broad community of network scientists and engineers, in collaboration with the NSF Network Science and Engineering (NetSE) program, poses this challenge: to develop the fundamental principles and methodical knowledge that will help us understand large, complex networks, and help us better design such networks in the future. The scope of NetSE ranges from design and development of network technologies, to "network science" and to the relationships between and among both of these and with people and societies. As a step toward this ambition, scholars and researchers, both inside and beyond traditional science and engineering, have been invited by NSF and the NetSE Council (an external community organization helping to refine the NetSE objective) to participate in a series of workshops to think about key issues and approaches.

In this context, on September 24-25, 2008, the workshop on Network Design and Societal Values assembled a group of scholars and researchers in the humanities, social sciences, law and policy as well as scientists and engineers to identify promising research in the humanities, social sciences, law and policy, past and potential, that connects the study of moral and political values with computer and information system design, development, and deployment.[1] Although the focus of the workshop was specifically on networks, the workshop sought to bring to bear the wealth of expertise and past work on the complex, mutual interplay between design of technology and political and social life.

But identifying promising research in relevant non-technical fields was not the only goal. At least as important was to identify past and potential research that could speak beyond the communities of its authors' academic origins to network scientists and engineers as well – results, questions, approaches, literatures, cases, and issues that might be

[1] See Appendix I for a list of participants.

meaningful to scientists and engineers, that might even influence the design of computer and information networks (e.g., hardware and software). What might these be? How might decisions in network design usefully and systematically take them into consideration? And, by the same token, what hard problems in network science and engineering might successfully migrate onto the agenda of the humanistic, social, and political study of technology? How might these problems stir and energize these areas?

THE REPORT

The report is inspired by ideas emerging from brief presentations by Workshop participants and from the discussion following these presentations, where several salient themes gelled.[2] Going into the workshop, participants were asked to prepare remarks not only about their own work but reflecting a line of research or scholarship in which they conceived their work to fit. This placement did not need to track traditional disciplinary boundaries (e.g., the names of their home departments) but could be associated with a set of questions, a particular method, an object, or objects of study, a set of issues, an annual conference, etc. Participants were asked to reflect on how this line of work contributed to a landscape of study of networks and societal values, for scientists and engineers as well as the members of their communities. Participants from technical fields were asked to reflect on instances in which they had encountered problems that they understood to be socio-technical in nature, and prior collaborative experiences to address such issues.

Individual presentations and group discussion suggested that workshop findings would be more easily presented as a research landscape, characterized by several key dimensions, each defined by an open ended set of questions, rather than a research agenda, defined by a single list of questions. The dimensions that seemed best to capture relevant past research and exciting and valuable future research were: Methods and Approaches, Issues, Themes, and Integrated Case Studies. We also include a bibliography of additional readings.

It bears repeating that the workshop's horizon was not on all interesting and worthwhile research on network technology through the lens of humanities and social sciences but on research in both fields that offered exciting potential for mutual influence.

DIMENSIONS OF THE LANDSCAPE

RESEARCH METHODS AND APPROACHES.

Issues, questions, themes, methods, and cases raised at the workshop build upon a significant body of past and ongoing research across the disciplines. Workshop participants were asked to describe their own work, and to give some insights into their research methods.

Noshir Contractor, with a background in behavioral science, investigates factors that lead to the formation, maintenance, and dissolution of dynamically linked social and knowledge networks in communities. The goal of this research is to develop theories

[2] See Appendix II for the workshop agenda.

(network science) of network formation at this level, and to translate this theory into design principles that lead to more effective and useful networks among groups such as health care professionals, first responders, and other sorts of professional groups, as well as social contexts such as virtual worlds.

For this sort of work, the emergence of the Internet and the higher-level networks that form on top of it provide a source of observable data that can be used to test theories and designs. Networks that are embedded in technology are perhaps easier to study and analyze than networks that only manifest as social behavior and off-line records. But this fact hints at the possibility that with better instrumentation of today's networks, we may be able to extract data, apply theory, and help improve the operation of those networks.

As an example, peer-to-peer networks (P2P) are not designed and engineered by network operators. They just "happen", as individuals choose to have their machines join the network. The research in behavioral science and the factors that lead to the formation of human networks are similar to the factors that govern the formation of P2P networks. So perhaps a better understanding of network science at the human level can help us design P2P networks at the technical level that are more resilient, efficient and useful.

Yochai Benkler brings his background in law to the study of human collaboration, and in particular to the phenomenon he calls peer production. The efficiency and utility of networks such as the Internet make practical what was before perhaps too cumbersome to undertake: the creation of knowledge and content by the unmanaged cooperative contributions of many people. Wikipedia is perhaps the most recognized example of the peer creation of knowledge, but there are many other examples. Benkler believes that we, as a society, should place great value on this sort of collective endeavors, and studied the factors that make it practical and constructive. As an agenda for research, we must move from the rich, empirical observations we have of these systems to more abstract structures that can be studied and modeled. We need more knowledge of human behavior and the foundations of cooperation if we can make the design of peer systems a methodical process.

He made the point that the idea of peer production does not only apply to the production of knowledge, but to the production of physical networks as well. P2P networks, mentioned above, are one example, and another is the creation of multi-hop (mesh) wireless networks out of devices contributed by the collective users. Such a network can only come into existence if the users choose to participate

Jinyang Li, an experimental computer scientist, echoed some of the above comments as she talked about the construction of distributed systems such as P2P systems. In decentralized systems with open membership, it is hard to create a stable system through technical constraints alone. Mechanisms that allow the maintenance of trust, such as identity systems and trust networks, seem to be important social building blocks of workable distributed systems. In balancing centralized and decentralized systems, technologists have observed that while centralized systems can be designed to be *technically* robust and resilient, they are prone to disruption at a higher level, for example legal. The balance should not just be seen in a space of technical tradeoffs, but in a larger space of social and legal considerations.

Judith Olson described her work in understanding successful and unsuccessful human collaboration from the perspective of psychology. Her work is empirical—extensive case studies of actual collaboration, and lab experiments involving humans tasked to solve real problems, has led to models of behavior that predict the outcome of collaboration. Her work provides a checklist of issues that will influence the success of collaboration: whether the intellectual structure of a field encourages competition or cooperation ("are you trying for a Nobel prize"), disjoint vocabularies, unnecessary heterogeneity in the technical tools of collaboration, physical distance, inherent modularity of the problem, and so on.

Her comment about the pragmatic barriers caused by the heterogeneity in the tools for collaboration should be a hint to the computer science community—we have not yet build network applications that can cover up the diversity in our systems and user interfaces. Tools for tele-collaboration are still awkward to use, prone to failure and disruption. We should step up and resolve these issues.

She, like several of our other speakers, used the example of collaborative activity, facilitated by the network, as an important goal. There are different words for the same essential idea: peer production, micro-contributions, distributed human computation, and so on.

Beki Grinter, with a background in computer science, described her research interest as interactive computing: the intersection of computing and humanity. Her research methods are anthropological and sociological. One thesis of her work is that human interaction is deeply local, the Internet is global, and the consequence is that society has now redefined "local". As an example of this phenomenon, she has studied the nature of online religion: the use of the Internet to allow participation at a distance in church services and the social fabric of the church. She has observed that large churches are now importing their services into the U.S. from abroad over the Internet, and these churches are an important part of the social linkage for new immigrants to this country. We both import and export religion over the Internet. Her work, like the work of others in this description, helps to shed light on those factors, both technical and cultural, that make the online social constructs effective and viable. She also observed that one must not be Internet-centric in this sort of analysis; mobile phones are an important part of the technology base that facilitates the creation and maintenance of these "local" groups.

Wendy Chun brings to the discussion the research method of critical theory, which has its roots in literary criticism but more generally invites us to think critically (and methodically) about both the process of interacting with the Internet and the content that is on it, but also to think critically about the framing of the Internet and its social implications. Critical theory reminds us that the act of "reading" is not just a one-way process where the reader is the recipient of the words of the writer. The reader, too, brings to the process a rich context, which participates in the construction of the meaning of that which is read. The reader is an active participant, as much as the writer. Critical theory reminds us, both as readers and writers, to think in a rigorous way about the context we bring to the process. Critical theory also reminds the reader to look "through" the presented media to the context, assumptions and motivations of the creator and the ways in which technology frames our language and actions.

In the context of a technology-rich environment like the Internet, critical theory would ask us to consider what aspects of technology shape our perception of media. Are the two

disconnected, or does the nature of Internet technology shape or limit our reaction to or perception of media? More generally, beyond the study of “media”, come questions about the modes of use of an artifact like the Internet, and the understanding of the local contexts of use, such as the example of online religion mentioned above. Critical theory raises questions regarding the relationship between technology, politics and society more generally.

We should also think critically when we consider language *about* the Internet. The Internet has been represented as a platform that fosters personal freedom and anonymous action, and also as the foundation of a global network of surveillance. We should consider the interplay between technological features and the context we each bring to the conversation in trying to understand such dichotomies.

The tools and discipline of critical theory will be particularly valuable as we try to *design* a different future, and must find ways to describe this future in terms that are both comprehensible from the varied contexts of the readers, and which invite a serious conversation about the values embedded in that future design. Words like “security”, “identity”, or “accountability” must be used with care, as they are rich in context and unshared assumptions.

Jon Kleinberg described the use of tools from graph theory and network algorithms to understand the structure of networks at all levels, from social to ``link and router''. The emergence of the Web as a vast mesh of linked objects triggered a change in his research discipline, since this large corpus allowed empirical study of the networks that have emerged. Certain aspects of these networks, such as their ``small-world'' structures, are well recognized at this point. But such observations point to deeper questions about why networks have this structure (especially networks that are not engineered but which just ``happen''). Does our understanding of these networks, and the nature of the forces that shape them, tell us anything about the formation of emergent networks at lower layers?

Another important topic is the study of online search. Search (and its results) are based on a ranking of sources, and there can be no general and neutral form of ranking. All ranking implies a value structure, which can be implicit or explicit, static or evolving. More generally, the nature of search offers a window into the collective minds of the searchers, and thus the mood of the time. Search, like reading, brings a great deal of user-specific context to the experience.

The emergence of large online social networks is another major shift in the landscape. Social networks, like other sorts of networks, can be analyzed and modeled to see what general properties they possess. One can also ask further questions about online social networks: for example do they provide a new platform for certain sorts of efficient search? Should we be trying to design networks that are optimized for search?

Helen Nissenbaum discussed research rooted in philosophy and ethics. She described an integrated approach formulated specifically for the task of analyzing design for values and approaches to guiding design practice taking values into consideration. The fundamental argument behind these efforts is that values are embodied, embedded, expressed, and reflected in design. Technology is not value-neutral: artifacts have politics, code is law, and technology has agency.

This approach includes “Values in Design” (VID), which generally refers to the study of fine-grain design characteristics for values embodied in them or promoted or afforded by them. Values-at-Play, Value-Sensitive-Design, and Reflective Design include heuristics for taking values into consideration during the design process, that is, for taking values into consideration in the design practice.

Larry Peterson is a computer scientist with interests in systems and networks. Recently he has spearheaded the development of a global platform for experimentation on distributed systems called PlanetLab. His discussion focused on a number of social and legal questions that have arisen as various researchers have built and deployed experiments on PlanetLab. Many experiments involve new applications that directly engage people, which raises the question of whether the deployment of a new application over the Internet constitutes performing an experiment involving human subjects. Some experiments do gather personally identifying information, which means that experimenters must be aware of and sensitive to issues of privacy and dignity. Many computer science experiments are now being reviewed by university internal review boards to confirm that they provide suitable safeguards for the people who might be involved in them. On the other hand, commercial players in the Internet deploy similar systems freely. This begs the question of what limits should be placed on academic researchers, relative to other actors who seek to understand and evolve the Internet.

Deirdre Mulligan, by training a lawyer, discussed both pragmatic and more fundamental issues. With respect to research (as discussed by Larry Peterson), she noted that in some fields such as health care, the research community protected itself legally by having language added to relevant legislation to add protections and exemptions to those doing research. The CS community, only recently coming to understand the deep ways in which their research intersects with social and legal issues, has not in the past sought out these protections. She raised the issue of more active and direct involvement with lawmaking in order to protect our ability to do research.

Paul Ohm provided some additional perspectives on the interplay of law and technology. First, he pointed out that the law (like some other non-technical fields) tends to look at technology as a static thing, but technology evolves rapidly. But noting that fact does not tell us how to model the future trajectory of something like the Internet. His practical experience at the Department of Justice, where he prosecuted criminal behavior, illustrates the range of stakeholders that bring pressure on technology to evolve. Their interests include surveillance and the gathering of forensic evidence. He noted with respect to surveillance, (e.g. observing what is sent over the network), law enforcement, academic researchers and network operators operate under three, very different sets of rules.

ISSUES

The workshop identified numerous issues. We acknowledge that those listed below are diverse in generality, size (of existing body of work), and scope, including some overlap. It is also important to note that disciplines vary in the ways they apply the identical label. Below are a sample of the issues that generated greatest interest and sometimes disagreement among workshop participants.

Security: We can study security from multiple perspectives. What does security mean to network researchers in computer science and engineering? For example, is it the perfection of private communication, or the inspection of communication by third parties to detect attacks? Is it possible to have both outcomes? What does security mean to political scientists, philosophers, or sociologists? What does it mean to the end-user, who must make sense of the rich network context and make decisions about safe and unsafe circumstances? How can we translate these definitions of security across fields? What happens, for instance, when different fields try to analyze an issue using the framework of security, or to justify a position by appealing to the goal of security? What are the tradeoffs between security and other values (e.g., free speech)? Among all the actors (or agents) on the network, including individuals, institutions, and governments, is everyone's security of equal value? And how does the value of security differ among users, institutions, and national governments themselves?

Identifiers and identities: The current Internet perhaps pays too little attention to how actors can and should be identified. For reasons of security (in its many guises, as discussed above), a future network may be designed to provide better tools for identifying users, services and other network components. But this objective in turn raises many important questions with rich social implications. Is there one or many approaches to defining identity? What's at stake in different (technical) choices? Why is identity often posed as a panacea? Can we test this proposition? How should identity online mesh with identity and identities in relation to other spheres of interaction, particularly in the relation of the individual to governments, financial institutions, and other corporate entities, such as merchants and service providers? Can we embed application layer solutions (e.g., eBay reputation systems, or social network identities) in a more general network architecture or design? Are there alternative approaches up and down the layers, and can these approaches successfully migrate? What sorts of collaboration would be effective in exploring this space and posing preferable approaches?

Openness: This is a term that has been used with various meanings in relation to networks (specifically, the Internet). One of the most important meanings, with technical and societal implications, is the capacity for everyone to *join* the network. In the case of the Internet, this means, at least in theory, that anyone is free to implement on their machine the protocols (e.g. TCP and IP) that will let that machine connect to the Internet. This contrasts with a scenario in which the protocols themselves are not open, in which case one would need the permission of a controlling or licensing authority to implement the protocols and/or connect to the network. Another important form of openness is that the users of the Internet are permitted to use any application that they choose: in principle neither the low-level protocols nor the Internet Service Providers limit what the user can choose to do. Important questions follow from this observation. What values are at stake? What does openness mean? What are the trade-offs in open and closed networks or protocols? Open networks may promote organic growth, but also suffer, in the case of malicious actors, from a lack of vetting or barring mechanisms. Open systems therefore often bring up issues of trust and individual accountability and thus identity online (as discussed earlier). Open systems can only invite users to participate. What are the mechanisms and conditions that encourage participation in open systems? Must the design of open networks take into account incentive structures? Systems that are open in the sense that anyone can use them, but closed in the sense that users cannot study the internal design, may force participation at a technical level in ways that the user did not expect. For example, a user of Skype may

be surprised to discover that they are relaying calls. Should such a system be called open?

Trust: A system that limits the opportunity for users to do harm to one another is not the same as one that achieves the same result based on trust. This is the difference between trustworthy technical artifacts (so-called trusted systems, as in secure banking) and technology that enables people to trust one another. How can we incorporate social values in the design of networks that actually promotes and sustains sociality? Do social networks, based on voluntary associations among users, point toward a model for trust networks more generally? Is the trust that pervades such networks durable over time and across platforms and between layers?

Mechanisms of regulation, control, or enforcement: Workshop participants agreed that behavioral constraints and affordances could be embedded in a network environment at various different junctures and layers. For example, they can be built into the technology, expressed in law and policy, through social norms, through incentives structures. The picture is even more complicated than this because even within these different junctures (or modes) there are various possibilities, and network design choices may produce unintended points of control. For example, technical constraints can be imposed at different layers (e.g., physical versus application) or following different strategies (e.g., through post-hoc auditing or front-end vetting). Reputation systems are an example of a socio-technical system that controls behavior, and yet disreputable people game these systems by putting in false scores or starting a new identity once the reputation is bad. Choosing mechanisms and points of control is a technical matter but ethical and political implications should be carefully considered. This issue covers a potentially huge terrain and offers great possibilities for collaboration among different approaches. It can be tackled thematically and also through detailed case studies.

Local and global: The Internet is touted as a global network but its value and meaning is often local (culturally, geographically). This requires study of networks embedded in a variety of contexts. Research might therefore address who appropriates a network for what purposes, and how different economic, social, political, and cultural contexts make such action possible. This research may draw from human-computer interaction, but may also adopt a more anthropological or sociological lens in examining the everyday and local uses of a network. Such insights may inform network design, particularly those attempts to develop a network that is sensitive to local variations in use and deployment. Can network design also build upon the general geographic distribution that tends to characterize social network membership? Can we develop networks that are optimized according to the spatial and social distribution of our likely network associates? Might we adjust our approach to network search, for example, given network information about geographic hotspots for certain query strings? How can or should scientists and engineers take local political contexts into consideration when designing the features of networks and network services?

Privacy: The issue of online privacy covers a universe of questions and issues. Do we agree on the meaning of the term; do we understand what other values it protects and what other values it clashes with? Can we identify technical mechanisms that might mitigate these conflicts? Do we understand how much do we want or need privacy, under any of its several definitions? What opportunities for monitoring and measurement do network design decisions create? Are privacy concerns inherent in the architecture of the network, or do secondary technologies (such a mass storage for data retention) play

a more important role? Whom do network architectures empower to monitor user behavior and information? The current architecture of the Internet, for instances, puts Internet Service Providers in a uniquely powerful position to monitor all the activity of its subscribers. What are the minimal features of a network that commercial service providers require (e.g., location-based IP)? How can we break apart information that we would find desirable for the network to reveal from that which it must necessarily produce?

Conditions of participation (related to several other issues above): This rather obscure title refers to a set of questions about expectations network users may reasonably have about the powers they have when they join a network and, conversely, what users may experience as part of normal participation in a network. To what kinds of activities are users legitimately expected to submit as a condition of participation? Specifically, researchers may wish to study whether traditional notions of real property have analogs in virtual worlds like Second Life. Do users have a right to object to unsolicited email as long as they have signed up for email, or to having their systems used as "zombies," or having search "bots" visit their websites? Could we imagine a network in which only consensual associates could exchange packets? To what degree do the conditions of participation of social networks already follow this model?

Motivations for participation: Why do people join and participate in a network? Noshir Contractor's work on the creation, maintenance, dissolution, and reconstitution of networks focused precisely on this role of motivations. Can we design networks that take into account the various motivations of their users? What defines a successful network from this perspective? Should networks adjust to users' motivations, and if so, how might networks determine or allow users to specify their respective motivations? What other criteria figure in the success of a network? Judith Olson's work on remote scientific collaborations, for instance, delineates the myriad factors that may obtain in pursuits supported by the network. On the other hand, what are the motivations for voluntary, collaborative online activities? What, for instance, are the social motivations of commons-based production on the Internet? If the degree to which certain network structures enable production of this sort has become clearer, there still remains much to explore about the micro-foundations of cooperation and collaborative production in general. For instance, can we develop networks that promote cooperation through solidarity rather than by reward or punishment? Can network design decisions (e.g. the nature and degree of revelation about identity) help cultivate voluntary participation and behavior that conforms to the norms of the community without recourse to punitive mechanisms or technical restrictions?

THEMES

Certain ideas seemed to crop up across discussion of several of the issues and case studies. They seemed more appropriately to be understood as themes, rather than as issues.

Visibility and transparency: The concepts of visibility and transparency are salient in two respects. The first we might describe as individual exposure and self-presentation on a network—that is the degree to which users can or must reveal information about themselves at different layers of a network (MAC address, IP address, application account, etc.). Trust, for instance, often requires some degree of exposure, as in reputation systems or social networking sites. Visibility in this sense may also refer to

the ability to communicate or reveal one's motivation for participation or collaboration (as discussed above). But these concepts have another meaning in a related context: the ability (or not) to examine the inner workings of a network design, protocol, or application. Take, for example, the design decision to allow Web users to view page sources, contrasted with the usually invisible algorithms that establish rank order for search results. Or technologies or software that are open, (in the terminology of the above discussion), and thus leave users free to tinker. Transparency of this sort has emerged as a political value among certain technologists and stakeholders, who argue that open access to software is an ethical virtue on the same level as the sharing of intellectual concepts.

Incentives: Understanding the structure of incentives can shed light on relationships between architecture or design, on the one hand, and behavior or outcomes, on the other. An integrated study of existing incentives through empirical, ethnographic, historical, etc. methods is an important way of understanding what is already in place. One may also wish to disrupt, shape, or take advantage of naturally occurring incentive structures in order to achieve certain ends, for example, security or privacy, in the context of networks or network transactions. How might we determine the generalizability of an incentive structure of a specific network or application? Are incentive structures from one network or application appropriate, legitimate, or effective in another?

Networks as Experimental Environments: The Internet and Web have emerged as hugely important environments for studying individual and social behavior. There is plenty of scope for thinking about the needs and requirements of research online. Network engineers are also engaged in experimentation in such activities as PlanetLab and potentially GENI inspired systems. What is the relationship between those who intentionally and inadvertently use these systems and the designers and developers of these systems? Must networks users consent to participation? Is there something importantly different in the responsibilities designers and engineers have to users when the systems they put out for use are "experimental?" To what ethical code should academic network researchers hold themselves, and how might such a code compare to the one, if any, that obtains in commercial research? How can network engineers best communicate the value of their research to those who are likely to be involved in the experiment or later affected? Or, alternatively, if large-scale experimentation is simply not possible with consensual parties, should we set a grand challenge for network engineers and designers which asks that they determine how to do research on networks that itself solves the problem of network monitoring?

INTEGRATED CASE STUDIES

There is an important place for integrated case studies. In general, these would be rich multidisciplinary studies of events, mechanisms, applications, architectures, etc. relating to networks.

Web search: One example discussed at the workshop was search, search in networks (social search, web search), including, for example, algorithm design and privacy. Why do we take for granted the current model? Are they the best we can manage? Must search algorithms tuned through machine learning be opaque to policy analysis? Is this a problem for values in design?

Technology adoption by government agencies: Under what circumstances do agencies view technology (e.g. RFID tags in passports) as a procurement or policy question? What determines the perspective different agencies take, and what are the effects of this decision on the primacy of values in the adoption process? Which procedures open up the most productive spaces for discussion of values?

Standards setting: Standard setting is an important site for determining socially relevant design features. There is often little reward for outsiders to participate in standard settings meetings. Why is this so and what about these meetings dissuades outside participation? How can outside stakeholder enter into or contribute to the debate? What are the social, bureaucratic, and epistemological conditions of participation?

*Self-organizing wireless network*s: We can imagine a study of the deployment of a wireless network in a municipality based on multi-hop or mesh technology that would call upon engineers, social scientists, and policymakers to shape the landscape of successful deployment. Researchers would consider the significance of local context and specific cultural, political, and motivational triggers.

Engineers' response to assertions about values: How do engineers articulate the values at play in their selection of and approach to a technical problem? How do they respond to the assertion that values figure in their work? Do they resist this idea? Under what conditions do engineers reflect on values in design, and how might these reflections lead to different design choices?

FOR FURTHER READING: METHODS OR THEORETICAL FRAMEWORK

This sample of books and papers written or recommended by workshop attendees will provide a deeper and broader window into the range of research methods and topics discussed here.

Books:

Behavioral Science

Monge, P, and N. Contractor, Theories of Communication Networks, Oxford University Press, 2003

Analytic, but also ethnographic:

Miller, D. and Slater, D., *The Internet: An Ethnographic Approach*. Berg, Oxford, England, 2000.

Legal:

Benkler, Y., The Wealth of Networks: How Social Production Transforms Markets and Freedom Yale University Press, 2007

Economics

Jackson, M. Social and Economic Networks. Princeton University Press, 2008.

Critical theory:

Chun, W.H.K., Control and Freedom: Power and Paranoia in the Age of Fiber Optics, MIT Press 2007

Historical approaches: not about the network per se, but about thinking about the relationship between infrastructure and society.)

Hughes, T., Human-built World: How to Think about Technology and Culture, University of Chicago Press, 2004

Nye, D., *Technology Matters: Questions to Live With*. MIT Press, Cambridge MA, 2007.

Nye, D.E., *Electrifying America: Social Meanings of a New Technology*. MIT Press, Cambridge, MA, 1990.

Rosenberg, N., *Inside the black box: technology and economics*. Cambridge University Press, Cambridge, UK, 1982.

Papers and other publications:

Analytic:

M. Flanagan, D. Howe, and H. Nissenbaum, "Values in Design: Theory and Practice" In Information Technology and Moral Philosophy Jeroen van den Hoven and John Weckert (eds.) Cambridge: Cambridge University Press, 2008

A. Barth, A. Datta, J. Mitchell, and H. Nissenbaum, "Privacy and Contextual Integrity: Framework and Applications," Proceedings of the IEEE Symposium on Security and Privacy, May 2006 (Showcased in "The Logic of Privacy," The Economist, January 4, 2007)

H. Nissenbaum, "Where Computer Security Meets National Security," Ethics and Information Technology, Vol. 7, No. 2, June 2005, 61-73 (Also, In Cybercrime, Eds Jack Balkin, James Grimmelmann, Eddan Katz, Nimrod Kozlovski, Shlomit Wagman and Tal Zarsky, New York, NYU Press, 2007

H. Nissenbaum, "Will Security Enhance Trust Online, or Supplant it?" In R. Kramer and K. Cook (eds.) Trust and Distrust Within Organizations: Emerging Perspectives, Enduring Questions, Russell Sage Publications (2004): 155-188

L. Introna and H. Nissenbaum, "Shaping the Web: Why the Politics of Search Engines Matters" The Information Society, 16(3):1-17, 2000

Law:

Schwartz, A.. Mulligan, D., Monda, I., Storing Our Lives Online: Expanded Email Storage Raises Complex Policy Issues, *I/S: A Journal of Law and Policy for the Information Society*

Mulligan, D., Reasonable Expectations in Electronic Communications: A Critical Perspective on the Electronic Communications Privacy Act," 72 *Geo. Wash. L. Rev.* 1557 (2004).

Economics:

Lian Jian and Jeffrey K. MacKie-Mason (2008), "Why Share in Peer-to-Peer Networks?", International Conference on Electronic Commerce (ICEC'08), Innsbruck, Austria, 19-22 August 20 Lian Jian and Jeffrey K. MacKie-Mason (2008), "Why Share in Peer-to-Peer Networks?", International Conference on Electronic Commerce (ICEC'08), Innsbruck, Austria, 19-22 August 20

Social Theory:

Olson, J. S., Hofer, E., Bos, N., Zimmerman, A., Olson, G. M., Cooney, D., and Faniel, I. (2008). A theory of remote scientific collaboration. in G. M. Olson, A. Zimmerman, and N. Bos (Eds.) *Scientific Collaboration on the Internet.* Cambridge, MA: MIT Press.

Empirical:

Qualitative:

Olson, J. S., Ellisman, M., James, M., Grethe, J. S., Puetz, M. (2008) Biomedical Informatics Research Network (BIRN) in G. M. Olson, A. Zimmerman, and N. Bos (Eds.) *Scientific Collaboration on the Internet.* Cambridge, MA: MIT Press.

Quantitative:

Bos, N., Shami, N. S., Olson, J. S., Cheshin, A., & Nan, N. (2004) In-group/out-group effects in distributed teams: An experimental simulation. *Proceedings of Conference on Computer Supported Cooperative Work.* 429-436.

Nan, N., Johnston, E. and Olson, J. S., Unintended consequences of collocation: using agent-based modeling to untangle effects of communication delay and in-group favor. *Computational & Mathematical Organization Theory*. Volume 14, Number 2 / June, 2008

Systems building:

Grinter, R.E., Edwards, W.K. Edwards, Newman, M.W. and Ducheneaut, N. The Work to Make a Home Network Work *European Conference on Computer-Supported Cooperative Work*, Springer, Paris, France, 2005, 469-488.

Chetty, M., Sung, J.-Y. and Grinter, R.E., How Smart Homes Learn: The Evolution of the Networked Home and Household *Proc. 9th International Conference on Ubiquitous Computing (Ubicomp 07)*, Springer-Verlag (2007), 127-144.

Shehan, E. and Edwards, W.K., Home Networking and HCI: What Hath God Wrought? *Proc. ACM Conference on Human Factors in Computing Systems (CHI 07)*, ACM Press (2007), 547-556.

APPENDIX I

Workshop Participant Bios
Co-chairs:

David Clark is a Senior Research Scientist at the MIT Computer Science and Artificial Intelligence Laboratory, where he has worked since receiving his Ph.D. there in 1973. Since the mid 70s, Dr. Clark has been leading the development of the Internet; from 1981-1989 he acted as Chief Protocol Architect in this development, and chaired the Internet Activities Board. His current research looks at re-definition of the architectural underpinnings of the Internet, and the relation of technology and architecture to economic, societal and policy considerations. He is helping the U.S. National Science foundation organize their Future Internet Design program. Dr. Clark is past chairman of the Computer Science and Telecommunications Board of the National Academies, and has contributed to a number of studies on the societal and policy impact of computer communications. He is co-director of the MIT Communications Futures Program, a project for industry collaboration and coordination along the communications value chain.

Helen Nissenbaum is Professor of Media, Culture and Communication and of Computer Science at New York University, where she is also Faculty Fellow of the Information Law Institute. Her areas of expertise include social, ethical, and political implications of computing and information technologies. Grants from the National Science Foundation, Air Force Office of Scientific Research, Ford Foundation, and U.S. Department of Homeland Security have supported research projects on privacy, trust online, security, intellectual property, and several projects investigating political values in information systems, including search engines, video games, and facial recognition systems. She has produced three books and over 40 research articles, which have been published in scholarly journals of philosophy, political philosophy, law, media studies, information studies, and computer science. Nissenbaum holds a Ph.D. in philosophy from Stanford University and, before joining NYU, served as Associate Director of Princeton University's Center for Human Values.

Participants

Yochai Benkler is the Berkman Professor of Entrepreneurial Legal Studies at Harvard, and faculty co-director of the Berkman Center for Internet and Society. Before joining the faculty at Harvard Law School, he was Joseph M. Field '55 Professor of Law at Yale. He writes about the Internet and the emergence of networked economy and society, as well as the organization of infrastructure, such as wireless communications. In the 1990s he played a role in characterizing the centrality of information commons to innovation, information production, and freedom in both its autonomy and democracy senses. In the 2000s, he worked more on the sources and economic and political significance of radically decentralized individual action and collaboration in the production of information, knowledge and culture. His work traverses a wide range of disciplines and sectors, and is taught in a variety of professional schools and academic departments. In real world applications, his work has been widely discussed in both the business sector and civil society. His books include The Wealth of Networks: How social production transforms markets and freedom (2006), which received the Don K. Price award from the American Political Science Association for best book on science, technology, and

politics, the American Sociological Association's CITASA Book Award an outstanding book related to the sociology of communications or information technology, the Donald McGannon award for best book on social and ethical relevance in communications policy research, was named best business book about the future by Stategy & Business, and otherwise enjoyed the gentle breath of Fortuna. In civil society, Benkler's work was recognized by the Electronic Frontier Foundation's Pioneer Award in 2007, and the Public Knowledge IP3 Award in 2006. His articles include Overcoming Agoraphobia (1997/98, initiating the debate over spectrum commons); Commons as Neglected Factor of Information Production (1998) and Free as the Air to Common Use (1998, characterizing the role of the commons in information production and its relation to freedom); From Consumers to Users (2000, characterizing the need to preserve commons as a core policy goal, across all layers of the information environment); Coase's Penguin, or Linux and the Nature of the Firm (characterizing peer production as a basic phenomenon of the networked economy) and Sharing Nicely (2002, characterizing shareable goods and explaining sharing of material resources online). His work can be freely accessed at benkler.org.

Noshir Contractor is the Jane S. & William J. White Professor of Behavioral Sciences in the School of Engineering, School of Communication and the Kellogg School of Management at Northwestern University, USA. He is the Director of the Science of Networks in Communities (SONIC) Research Group at Northwestern University. He is investigating factors that lead to the formation, maintenance, and dissolution of dynamically linked social and knowledge networks in communities. Specifically, his research team is developing and testing theories and methods of network science to map, understand and enable more effective networks in a wide variety of contexts including communities of practice in business, science and engineering communities, disaster response teams, public health networks, digital media and learning networks, and in virtual worlds, such as Second Life. His research program has been funded continuously for over a decade by major grants from the U.S. National Science Foundation with additional funding from the U.S. National Institutes of Health (NIH), U.S. National Aeronautics and Space Administration (NASA), the Rockefeller Foundation, and the MacArthur Foundation. Professor Contractor has published or presented over 250 research papers dealing with communicating and organizing. His book titled Theories of Communication Networks (co-authored with Professor Peter Monge and published by Oxford University Press in English and scheduled to be published by China Renmin University Press in simplified Chinese in 2008) received the 2003 Book of the Year award from the Organizational Communication Division of the National Communication Association. He is the lead developer of CIKNOW (Cyberinfrastructure for Inquiring Knowledge Networks On the Web), a socio-technical system to enable networks among communities, as well as Blanche, a software environment to simulate the dynamics of social networks.

Wendy Hui Kyong Chun is Associate Professor of Modern Culture and Media at Brown University. She has studied both Systems Design Engineering and English Literature, which she combines and mutates in her current work on digital media. She is author of _Control and Freedom: Power and Paranoia in the Age of Fiber Optics_ (MIT, 2006), and co-editor of _New Media, Old Media: A History and Theory Reader_ (Routledge, 2006). She has been a fellow at the Radcliffe Institute for Advanced Study at Harvard, a Wriston Fellow at Brown, and Visiting Associate Professor in the History of Science Department at Harvard. She serves on numerous advisory boards of journals and is currently a co-PI on a Mellon Planning Grant to transform Visual Culture Studies. She

is also completing a monograph entitled _Programmed Visions: Software, DNA, Race_ (forthcoming MIT, 2010).

Rebecca E. Grinter (Beki) is an Associate Professor of Interactive Computing in the College of Computing at the Georgia Institute of Technology. Her primary research interests lie at the intersection of computing and humanity, exploring the human-centered problems of technology production and consumption. Her research has been published in Human Computer Interaction, Computer Supported Cooperative Work, Software Engineering, Security, and most recently Networking conferences. Before joining the faculty at Georgia Tech, she was a Member of Technical Staff at Bell Laboratories, Lucent Technologies (and briefly AT&T Bell Laboratories), and a Member of Research Staff in the Computer Science Laboratory of Xerox PARC. She holds a Ph.D. & M.S. in Information and Computer Science from the University of California, Irvine, and a B.Sc. (Hons) in Computer Science from the University of Leeds.

Jon Kleinberg is a Professor in the Department of Computer Science at Cornell University. His research focuses on issues at the interface of networks and information, with an emphasis on mathematical models for social and information networks, and algorithms for problems in search, data analysis, and network optimization. He is a member of the National Academy of Engineering and the American Academy of Arts and Sciences, and serves on the Computer Science and Telecommunications Board of the National Academies and the NSF CISE Advisory Committee. He has received a MacArthur Foundation Fellowship, Packard Foundation Fellowship, and Sloan Foundation Fellowship, NSF CAREER and ONR Young Investigator Awards, the Nevanlinna Prize from the International Mathematical Union, and the National Academy of Sciences Award for Initiatives in Research.

Jinyang Li has been an assistant professor in computer science at New York University since 2006. She is interested in distributed systems and networks, especially how to build reliable large scale systems. She received the NSF CAREER award in 2008. Her group is currently working on making peer-to-peer systems more trustworthy and applicable to a variety of applications such as censorship circumvention and cooperative storage systems. She received a Ph.D. from MIT in 2005 and was a postdoctoral researcher at UC Berkeley from 2005-2006. While at MIT, she worked on scalable lookup protocols for large distributed systems and multihop wireless routing.

Deirdre K. Mulligan comes to the UC Berkeley School of information from the Berkeley School of Law, where she was a clinical professor of law and the director of the Samuelson Law, Technology & Public Policy Clinic. She served previously as staff counsel at the Center for Democracy & Technology in Washington. Through the clinic, Mulligan worked to foster the public's interest in new computer and communication technology by engaging in client advocacy and interdisciplinary research, and by participating in developing technical standards and protocols. The clinic's work has advanced and protected the public's interest in free expression, individual privacy, balanced intellectual property rules, and secure, reliable, open communication networks. Mulligan writes about the risks and opportunities technology presents to privacy, free expression, and access and use of information goods. Professor Mulligan holds B.A. from Smith College (1988) and J.D. from Georgetown University Law Center (1994).

Paul Ohm joined the faculty of the University of Colorado Law School in 2006. He specializes in computer crime law, information privacy, criminal procedure, and

intellectual property. Prior to joining Colorado Law he worked for the U.S. Department of Justice's Computer Crime and Intellectual Property Section as an Honors Program trial attorney. Professor Ohm is a former law clerk to Judge Betty Fletcher of the U.S. Ninth Circuit Court of Appeals and Judge Mariana Pfaelzer of the U.S. District Court for the Central District of California. He attended the UCLA Law School where he served as Articles Editor of the UCLA Law Review and received the Benjamin Aaron and Judge Jerry Pacht prizes. Prior to law school, he worked for several years as a computer programmer and network systems administrator, and before that he earned undergraduate degrees in computer science and electrical engineering.

Judith Olson is the Donald Bren Professor of Information and Computer Sciences, with appointments also in the Paul Merage Business School and the School of Social Ecology at the University of California at Irvine. She was just recently the Richard W. Pew Professor of Human-Computer Interaction at the University of Michigan. She was a professor in the School of Information, the Business School, and the Psychology Department. She got her Ph.D. in Psychology at the University of Michigan then held a postdoctoral fellowship at Stanford University before returning to Michigan as a faculty member. Except for three years at Bell Labs and a year at Rank Xerox Cambridge, UK, and now at UC Irvine, she had been at Michigan her entire professorial life. Her research focuses on the technology and social practices necessary for successful distance work, encompassing both laboratory field study methods along with agent based modeling. She has served on a number of editorial boards and panels for both the National Research Council and the National Science Foundation. In 2001, she was one of the first seven inductees into the CHI Academy. In 2006 she and her husband Gary were awarded the 2006 CHI Lifetime Achievement Award.

Larry Peterson is the Robert E. Kahn Professor of Computer Science at Princeton University. He is also Department Chair and Director of the Princeton-hosted PlanetLab Consortium. Peterson is co-author of the best selling networking textbook Computer Networks: A Systems Approach (4e),* and chaired the initial planning efforts that led to NSF's GENI Initiative. His research focuses on the design and implementation of networked systems.
Professor Peterson recently served as Editor-in-Chief of the *ACM Transactions on Computer Systems*, he has been on the Editorial Board for the *IEEE/ACM Transactions on Networking* and the *IEEE Journal on Select Areas in Communication*, and he has served as program chair for SOSP, NSDI, and HotNets. Peterson is a Fellow of the ACM. He received his Ph.D. degree from Purdue University in 1985.

Ellen W. Zegura received the B.S. degree in Computer Science (1987), the B.S. degree in Electrical Engineering (1987), the M.S. degree in Computer Science (1990) and the D.Sc. in Computer Science (1993) all from Washington University, St. Louis, Missouri. Since 1993, she has been on the faculty in the College of Computing at Georgia Tech. She served as Interim Dean of the College for six months in 2002. From February 2003 to 2005, she was an Associate Dean, with responsibilities ranging from Research and Graduate Programs to Space and Facilities Planning. Starting in August 2005, she has chaired the School of Computer Science of the College of Computing. She is the proud mom of two girls, Carmen (born in August 1998) and Bethany (born in May 2001), whose pictures had never made it onto the web until the advent of photo sharing web sites. Prof. Zegura's research work concerns the development of wide-area (Internet) networking services and, more recently, mobile wireless networking. Wide-area services are utilized by applications that are distributed across multiple administrative domains

(e.g., web, file sharing, multi-media distribution). Her focus is on services implemented both at the network layer, as part of network infrastructure, and at the application layer. In the context of mobile wireless networking, she is interested in challenged environments where traditional ad-hoc and infrastructure-based networking approaches fail. These environments have been termed Disruption Tolerant Networks.

Scribes

Solon Barocas is a doctoral student in the Department of Media, Culture, and Communication and Student Fellow at the Information Law Institute at New York University. His research focuses on the implications of predictive technologies, such as profiling and personalization, in news media, politics, national security, and social welfare provision. Solon has worked with the Stanhope Center for Communication Policy and Research, the Center for Global Communication Studies, the Berkman Center for Internet and Society, and the Russell Sage Foundation. He obtained his MSc in International Relations from the London School of Economics and graduated from Brown University with a BA in Art-Semiotics and International Relations, where he worked on the Information, Technology, War, and Peace Project at the University's Watson Institute for International Studies.

Erika Shehan Poole is a PhD student in the Human-Centered Computing program at Georgia Tech. Her research interests broadly focus on how end-users make sense of networked computing in domestic settings and in more advanced ubiquitous computing environments. Her dissertation work focuses on understanding the causes of digital complexity in the home, as well as how householders seek help from third parties in overcoming these difficulties. Erika holds a BS degree in computer science from Purdue University and an MS in computer science from Georgia Tech. She is a member of ACM and IEEE, and is actively interested in research ethics and public policy issues related to computing.

GENI Project Office (GPO) participants

Brig "Chip" Elliott is the Principal Investigator and Project Director and Chief Engineer for the GENI Project Office. As Project Director, he will assume overall responsibility for timely completion of GENI's planning, including development and management of the GPO itself and its dependent working groups and sub-contracts. Chip has nearly thirty years of experience in leading large, technically-challenging projects, both in industry and in academia, with particular expertise in routers, wireless Internet technology, mobile "ad hoc" networks, quality of service issues, advanced optical techniques, and novel routing architectures. As Chief Engineer at BBN Technologies, Chip has led the design and successful implementation of secure, mission-critical networks based on novel technology for the United States and its allies, with aggregate value above $3 billion. From 2001 to 2006, Chip served as Principal Investigator for the DARPA Quantum Network, in which he led the design and build-out of the world's first quantum cryptography network. It became fully operational in October 2003 in BBN's laboratory, and since May 2004 has operated non-stop between Harvard, Boston University, and BBN.

Aaron Falk is GPO's Engineering Architect and Lead System Engineer. Aaron works closely with the community to ensure that GENI's end-to-end architecture is fully defined, that it satisfies the community's research requirements.

Aaron is a degreed system engineer with a strong background in building and managing networking projects. An IETF leader for over ten years, Aaron managed the DCCP, PILC, and TCPSAT working groups as they developed standards-track Internet protocols and advisory documents. He received his BS, Electrical Engineering in 1992 and MS, System Engineering in 1994 from University of Maryland College Park, MD.

APPENDIX II

Agenda
Wednesday September 23

5:00–9:00 PM

Presentation on the NetSE program
Ellen Zegura, Chair, NetSE Council

Co-chair introductions, review of meeting objectives, scope, and candidate report outline
Helen Nissenbaum and David Clark

Thursday September 24

8:30-9:00 Breakfast

9:00-noon Extended introductions and identification of issues

9:00-10:15 Contractor, Olson, Grinter, Chun, Kleinberg

10:15-10:45 Break

10:45-12:00 Peterson, Li, Benkler, Ohm, Burk, Mulligan

12:00-1:30 Lunch (provided)

1:30-5:00 Discussion of issues and case studies

1:30-3:00 Identify key issues for report

3:00-3:30 Break

3:30-5:00 Charge to report writers and Conclusion

Appendix 6

Workshop Report on

Overcoming Barriers to Disruptive Innovation in Networking

Report of NSF Workshop on

Overcoming Barriers to Disruptive Innovation in Networking

January 2005

The NSF Workshop on Overcoming Barriers to Disruptive Innovation in Network was supported by NSF under grant CNS-0439842.

Table of Contents

1. Executive Summary

There is little argument that the Internet faces many challenges, including both correcting vulnerabilities that arise from society's increasing dependence on it [PRE05], and capitalizing on opportunities that arise as new applications. It is critical that the network research community be engaged in addressing these challenges.

The research community typically pursues one of two paths when trying to affect the Internet. The first is to incrementally evolve the network to address new vulnerabilities and opportunities as they occur. The research community, in conjunction with the commercial players that define today's Internet, have successfully followed this path for nearly 30 years, resulting in point-solutions of narrow scope, many of which step outside the original Internet architecture. The second path is to create a new Internet architecture that better addresses the many challenges on the horizon. This approach potentially involves a clean-slate design, and so is likely disruptive.

While there is no way to be certain that the incremental path will ultimately fail to address the challenges facing the Internet, there are two reasons to be concerned. The first is that the point-solutions incrementally applied to the Internet result in increased complexity. The Internet's once clean architecture has become muddied by patches, which makes it hard to reason about the network as a whole. This increased complexity makes the Internet harder to manage, more brittle in the face of new requirements, and more vulnerable to emerging threats. The second is that are architectural limits that may eventually result in a dead-end for the current incremental path. This report identifies five such limits, which we express in actionable terms:

1. **Minimize trust assumptions:** the Internet originally viewed network traffic as fundamentally friendly, but should view it as adversarial;
2. **Enable user choice:** the Internet was originally developed independent of any commercial considerations, but today the network architecture must take competition and economic incentives into account;
3. **Allow for edge diversity:** the Internet originally assumed host computers were connected to the edges of the network, but host-centric assumptions are not appropriate in a world with an increasing number of sensors and mobile devices;
4. **Design for network transparency:** the Internet originally did not expose information about its internal configuration, but there is value to both users and network administrators in making the network more transparent; and
5. **Meet application requirements:** the Internet originally provided only a best-effort packet delivery service, but there is value in enhancing (adding functionality to) the network to meet application requirements.

Considering the risks of solely pursuing the incremental path, the workshop participants believe it is important that the research community also pursue the design, evaluation, and deployment of disruptive network architectures. This path is not without its own risks, however. First, researchers need more realistic evaluations of architectural proposals. New architectures need to be evaluated experimentally, operating at scale, and under real-world

conditions. Second, there must be a plausible deployment plan for any new architecture. Expecting global agreement about (and uptake of) a new network architecture is not realistic in an environment dominated by commercial considerations.

Despite these risks, there is a new approach to experimental network testbeds that both permit realistic experimental evaluations and have the potential to lead to wide-spread deployment. The key features of the new approach include (1) an overlay infrastructure with global reach that can be shared among multiple candidate network architectures; (2) interposition mechanisms that allow users to opt-into new architectures on a per-user/per-application basis, thereby providing real user traffic and facilitating incremental deployment; and (3) a high-performance substrate that provides sufficient capacity to make successful architectures viable on a larger scale.

In light of the current situation and opportunities, the workshop participants make the following recommendations to the National Science Foundation:

Recommendation 1: Immediately initiate a research program on experimental architectural research in networking. If successful, the potential benefits are enormous, easily justifying the modest initial outlay. A dedicated research program will pull together the research community to tackle the broad scope of thorny architectural questions that we have outlined in this report, and that must be addressed for the program to be successful.

Recommendation 2: Foster experimental validation of new architectural research in networking. Paper designs, although thought provoking, are unconvincing, both to the companies that need to adopt them, and to the research community in evaluating ideas and in gaining insight into design tradeoffs. Thus, to maximize our chance of success, NSF must foster an expectation within the experimental architectural research program that research ideas should normally be validated under real use

Recommendation 3: Fund the development and deployment of suitable testbeds. Since experimental validation is an important component of this research program, it is essential that researchers have access to suitable testbeds NSF should therefore endeavor to build a meta-testbed that reduces the barrier to entry for new architectural ideas. To meet short-term needs, NSF should support an initial meta-testbed that can be deployed immediately. At the same time, NSF should initiate a deliberative process through which the community can identify long-term solutions to its meta-testbed requirements.

Recommendation 4: Start a process that will lead to substantial increases in funding for a broad multi-disciplinary effort in this area over the next few years. To design, construct and widely deploy a new architecture for the Internet is an enormously difficult and, at the same time, an enormously important undertaking. To be successful, we will need to enlist the efforts of distributed systems researchers, e-scientists, application developers, computer architects, and network hardware technologists.

Recommendation 5: Find ways to promote synergy and convergence among architectural visions. Academic research focuses on novelty and, in so doing, often accentuates

differences rather than identifying commonality. The past success of the Internet strongly suggests that we will be the most successful if we can coalesce around common architecture features. Architecture, by its very nature, "defines that on which we must agree." Thus, to be effective, architectural researchers should seek convergence rather than divergence.

Recommendation 6: Help the community learn from industry. Disruptive architectural research should not be fettered by today's problems and practices, but it must be informed by them if we are not to simply repeat the mistakes of the past. The large gap between the research and commercial communities often prevents effective communication between the two, to the detriment of both. To bridge this gap, NSF should facilitate interactions between researchers and practitioners.

The workshop participants recognize that different outcomes are possible. One possibility is that multiple promising architectures bloom, but that over time, there is convergence on a single new architecture for the Internet. Ideally, the incremental deployment story proves successful, bringing the new architecture to the verge of commercialization. Another possibility is that many valid architectures emerge, but there is no consensus as to a single correct architecture. Instead, the experimental testbed that supports multiple architectures emerges as the substrate for a future global communications infrastructure. A third possible outcome is that the ideas developed as a part of this program provide new insights and architectural clarity, but these ideas can be incrementally retrofitted into today's Internet architecture. This possibility suggests that pursuing the second path (a disruptive architecture) actually improves the odds that the first path (incremental evolution) succeeds.

2. Problems, Opportunities, and the Impact Imperative

The Internet has, in a remarkably short period of time, radically transformed the world's information infrastructure. This success is in no small part due to its innovative architecture that, in several dimensions, broke with the conventional (and largely telephonic) wisdom. The architecture now accommodates a wide variety of network technologies, spans an enormous gamut of speeds, supports a broad range of applications, withstands a substantial number of failures, and scales to hundreds of millions of nodes. Moreover, the same architecture that facilitated organic and decentralized growth during the Internet's formative years has endured, without modification, the painful transition to a commercial enterprise with many competing providers. In both technical and commercial terms, the Internet architecture has succeeded beyond anyone's wildest dreams.

However, in the thirty-odd years since its invention, new uses and abuses, along with the realities that come with being a fully commercial enterprise, are pushing the Internet into realms that its original design neither anticipated nor easily accommodates. These problematic issues include: the awkwardness with which host mobility, host multi-homing, data migration and data replication are handled; the lack of protection from unwanted or harmful traffic; the increasing complexity and fragility of inter-domain routing; and the impact of radically diverse edge devices, including sensor networks. Such problems are numerous, and the Internet's emerging centrality has made these flaws all the more evident and urgent. As a result, it is now widely believed that the Internet architecture is in need of substantial change.

Unfortunately, there is increasing pessimism about the possibility of change. Adopting a new architecture not only requires modifications to routers and host software, but given the multi-provider nature of the Internet, also requires that ISPs jointly agree on that architecture. The need for consensus is doubly damning; not only is agreement among the many providers hard to reach, it also removes any competitive advantage from architectural innovation. This discouraging combination of difficulty reaching consensus, lack of incentives for deployment, and substantial costs of upgrading the infrastructure leaves little hope for fundamental architectural change. Thus, many believe that the Internet architecture, which began as a radical experiment, has now ossified into an unalterable status quo [PET04].

Freezing forevermore the current architecture would be bad enough, but in fact the situation is deteriorating. The inability to adapt to new pressures and requirements has led to an increasing number of ad hoc work-arounds, many of which violate the canonical architecture (e.g., middleboxes). While derided by architectural purists, these modifications have (usually) arisen to meet legitimate needs that the architecture itself could not. These architectural barnacles—unsightly outcroppings that have affixed themselves to an unmoving architecture—may serve a valuable short-term purpose, but significantly impair the long-term flexibility, reliability, security, and manageability of the Internet. Thus, the collision between the improbability and the necessity of change has resulted in expedient but eventually harmful architectural liberties.

While the commercial world applies point-solutions and work-arounds to the existing Internet, the research community is facing its own dilemma. A network architecture is a subtle thing that defies rigorous analysis or satisfying simulation, and is best understood through extensive live experimentation. However, current testbed paradigms are inadequate to this task. Traditional testbeds can be roughly categorized as production-oriented or research-oriented. Production testbeds, such as Internet2 [I2], support real traffic from real users, often in large volume and across many sites. As such, they provide valuable information about the operational behavior of an architecture. However, the users of such a production testbed have no choice about whether or not to participate in the testbed and usually do not even realize that their traffic is part of an experiment. They thus expect the performance and reliability to be no worse than the standard Internet. Production testbeds must therefore be extremely conservative in their experimentation, using well-honed implementations of incremental changes.

Research testbeds (such as DETER [DET]) do not carry traffic from a wide variety of real users but instead are typically driven by synthetically generated traffic and/or a small collection of intrepid users. This allows them to be much more adventurous, capable of running first-cut implementations of radically new designs. Unfortunately, this lack of real traffic also renders the results much less indicative of real operational viability. As a result, neither kind of testbed—production or research—produces the data needed to adequately evaluate new architectures. It is therefore difficult to make a compelling case for new architectural designs based on a testbed evaluation. In addition, because they utilize dedicated transmission links, both categories of testbeds involve substantial cost, and so are prohibitively expensive to operate at very large scale. Thus, they are typically of small geographic extent and arise only with substantial funding support. Given the limitations mentioned above, traditional testbeds offer far too little bang for their buck, and clearly cannot lead us into the future.

The preceding paints a depressing picture of the status quo, with an architecture incapable of change and a research community unable to validate its designs. However, within this bleakness there are seeds of hope. After roughly a decade where incremental research held

sway, there has been a resurgence of interest in more fundamental architectural questions. This architectural research is still at an early stage and needs significantly more support in order to reach fruition, but these initial architectural sprouts are very encouraging. While not providing any definitive solutions, they suggest that many of the challenges facing the Internet can be adequately addressed through architectural innovations.

In addition, there is now a promising alternative to the traditional testbed approach. Two recent trends, virtualization and overlay networks, can be combined to create effective and inexpensive testbeds. Overlay networks have often been used to augment the current Internet and deploy experimental designs. Overlay networks, in contrast to the traditional physical testbeds, are not limited geographically: in fact, overlay networks can be accessed by any user through packet-redirection implemented by host proxies. The decision about whether or not to use an overlay network can be made on a per-user, and even a per-application, basis. If the overlay network fails, the user's traffic can default back to normal Internet service. The lack of geographic limitations, the ability of fine-grained opt-in, and the presence of automatic fail-over suggest that experimental architectures could likely attract a sizable pool of volunteers willing to supply live traffic. This breaks the old dichotomy of experimental versus production testbeds; these overlays can now be both.

Moreover, overlay networks don't require significant investment in bandwidth. However, such networks require a great deal of effort to deploy and manage, and this overhead of deploying a single-purpose overlay is well beyond the means of most researchers. Fortunately, the advent of highly virtualized infrastructures, like PlanetLab [PET02, BA04], provides a solution to this problem. Virtualization allows each overlay node to emulate the actions of many logical "routers," and thereby enables such infrastructures to support many concurrent architecture experiments, each running on its own set of logical routers. The burden of running an overlay is thus shared among a large set of experiments, bringing the overhead imposed on any individual researcher to a much more manageable level. Thus, these virtualized testbeds offer new hope that large-scale live experimentation with new architectures is within reach of most researchers.

All such experimentation would be meaningless without a plausible deployment path. As argued earlier, the need for consensus and the consequent lack of competitive advantage, along with the sizable investment needed to upgrade the deployed infrastructure, makes it doubtful that the current ISPs will deploy a next-generation architecture. Thus, deployment of new architectures may rely on new entrants to the service provision market. Given the high capital costs and low operating margins of this industry, a market-entering foray by a traditional infrastructure-based ISP seems unlikely. Overlays, however, are a more cost-effective way to enter this market. A new-generation service provider could deploy an overlay supporting a new architecture and distribute proxy software that allows anyone, anywhere, to access that overlay. This deployment path would be further enhanced by a highly virtualized overlay infrastructure. Just as commercial web hosting facilities allow individual companies to easily establish production-grade web sites, a commercial overlay hosting facility could greatly lower the barrier facing entering service providers. In fact, this virtualized infrastructure need not be an overlay and could instead be based on a set of dedicated links and (virtualized) routers. As we discuss later, this would be especially relevant if a sizable market for the development and deployment of new architectures (and infrastructure-based services) develops.

While the status quo is good reason for pessimism, the new developments described above provide much hope for the future. There is growing interest in new architectural approaches,

and some of the early results are promising. Virtualized overlay infrastructures can allow extensive yet inexpensive live experimentation with these new designs, and eventual deployment may proceed through the same virtualization approach. Thus, the seeds for success are already present.

This opportunity will not be realized easily. The research community must rally around the grand challenge of designing new network architectures and following them through to deployment. This is no small task. Not only will it require abundant time and effort, it will also require a change in the community's culture. Researchers must move beyond merely academic models of success and rededicate themselves to making an impact.

Many in the community already feel this impact imperative. But they will require substantial support in order to succeed. A greater focus on architectural research would broaden the pool of interested designers and interesting designs. This endeavor will also require a greater focus on impact and a recognition of the nature of support such efforts require.

3. Challenges and Options for Meeting Them

It is clear that the Internet faces serious challenges, from improving the security and robustness of its core packet delivery service, to accommodating an explosion in the number and diversity of devices that connect to it, to enabling a new generation of applications. While a perfectly valid response to this situation is to identify the attributes an ideal Internet of 2015 might aspire to [CL05], the research community believes it is also important to re-evaluate the architectural decisions that underlie today's Internet. This research agenda involves identifying the key limitations and assumption of the current architecture and pursuing the opportunities made possible by removing these barriers, with the goal of converging on a new set of architectural features that provide the foundation for the global communications infrastructure. While there is also value in doing research that leads to incremental improvement of today's Internet, these architectural barriers must be taken head-on to fully address the challenges we face.

This section identifies seven specific architectural limitations or assumptions that the research community believes warrant investigation. The following subsections do not correspond to seven different network architectures, but rather, they identify "vectors" for architectural research that the community is already pursuing. Note that these vectors are not orthogonal; they revolve around five themes:

1. **Minimizing trust assumptions:** the Internet originally viewed network traffic as fundamentally friendly, but should view it as adversarial;
2. **Enabling user choice:** the Internet was originally developed independent of any commercial considerations, but today the network architecture must take competition and economic incentives into account;
3. **Allowing for edge diversity:** the Internet originally assumed host computers were connected to the edges of the network, but host-centric assumptions are not appropriate in a world with an increasing number of sensors and mobile devices;

4. **Designing for network transparency:** the Internet originally did not expose information about its internal configuration, but there is value to both users and network administrators in making the network more transparent; and
5. **Meeting application requirements:** the Internet originally provided only a best-effort packet delivery service, but there is value in enhancing (adding functionality to) the network to meet application requirements.

3.1. Security

Unlike the original Internet, in which the user community was a close-knit group of experts running relatively simple applications, today's user population and applications increasingly means that network traffic must be viewed as adversarial rather than cooperative. This fundamental shift makes security a major concern. In particular, the scale and heterogeneity of the network has increased dramatically to span scores of nations, thousands of network providers, and millions of users. Unfortunately, few of today's protocols are designed to minimize trust or even to recognize trust boundaries. To take one example, a single mistyped command at a router at one ISP recently caused widespread, cascading disruption of Internet connectivity across many of its neighbors. At the same time, a broad range of applications-including critical infrastructure, commerce, education, personal productivity-now depend on the Internet infrastructure. This raises both the incentives for malicious users and the consequences of successful attacks. Because of the Internet's ossification, any new security flaw in a protocol can take decades to address, handing malicious attackers a significant advantage.

Fundamentally changing the Internet architecture to assume adversarial rather than friendly use has the potential to yield dramatic benefits. Imagine, for example, a world where the Internet is a trustworthy network absent of attacks, where sensitive information is communicated safely, where corporations can rely on the Internet for their businesses without fear of disruption, and where governments can rely on it for their critical infrastructures.

Given the paramount importance of the Internet, a security-aware architecture that minimizes trust assumptions is necessary. For example, architectural support for security could (1) improve network robustness through protocols that work despite misbehaving participants, (2) enable security problems to be addressed quickly once identified, (3) isolate ISPs, organizations, and users from inadvertent errors or attacks; (4) prevent epidemic-style attacks such as worms, viruses, and distributed denial of service; (5) enable or simplify deployment of new high-value applications and critical services that rely on Internet communication such as power grid control, on-line trading networks, or an Internet emergency communication channel; and (6) reduce lost productivity currently aimed at coping with security problems via patching holes, recovering from attacks, or identifying attackers.

Several architectural approaches show promise for addressing security issues. One important thread are architectures that prevent denial of service by allowing a receiver to control who can send packets to it. Another is making firewalls a fully recognized component of the architecture instead of an add-on that is either turned off or gets in the way of deploying new applications. A clean specification for security that makes clear the balance of responsibility for routers, for operating systems and for applications can move us from the hodge-podge of security building blocks we have today to a real security architecture. A careful design of mechanisms for identity can balance, in an intentional way rather than by accident, the goals of privacy and accountability. Ideally, the design will permit us to apply real world

consequences (e.g. legal or financial) for misbehavior. This may require that the architecture be aware of such real-world attributes as boundaries of jurisdiction.

The main research challenges in defining a more secure network architecture include balancing accountability versus privacy, balancing processing overheads versus security guarantees, and determining what network information and processing to expose to the network infrastructure and to network users.

3.2. Economic Incentives

The original design of the Internet did not take into account the economic structure of the industry that would emerge to support it. The very early view of the Internet was an undifferentiated cloud of routers, with no recognition of the points where Internet service providers connect. In contrast to the telephone system, which has two kinds of phone calls: sender pays and receiver pays ("800" calls), the Internet has no equivalent of a call, and nothing to signal the direction of value flow. This lack of attention to value flow, and architectural mechanisms to underlie the flow of payments across the Internet, represents a barrier to future investment in the Internet, and a barrier to the overall economic health of the infrastructure sector. While many mechanisms have emerged in response to industry needs and in particular to the problem of bilateral connection among ISPs, it can be argued that lack of an overall architectural framework for flow of payments has hindered the deployment of inter-provider Quality of Service, of multicast, and of consumer broadband. A failure to attend to larger economic issues around the competitive nature of the industry structure can also be seen as one of the causes of poor security in the Internet, and the failure of the Internet to address larger social needs (public sector needs) such as emergency preparedness.

A future design for an Internet should take into account that a network architecture induces an industry structure, and the economic structure of that industry. The architecture can use user choice (to impose the discipline of competition on the players), indications of value flow (to make explicit the right direction of payment flow), and careful attention to what information is revealed and what is kept hidden (to shape the nature of transactions across a competitive boundary). The "architecture of economics" surrounding a new Internet must also reflect the necessity of governments to inject into the design functional objectives that do not necessarily align with the features that emerge through private sector, profit-seeking investment.

3.3. Address Binding

The way in which endpoints are identified for the purpose of directing traffic toward them is one of the most fundamental aspects of any network architecture. In today's Internet, endpoints are addressed with topologically-dependent IP addresses, and it is precisely the structured nature of the address space that enables scalable packet forwarding.

More precisely, endpoints in the Internet architecture are simply network attachment points—locations where a network device plugs in. IP addresses were not intended to say anything about the machine connected at that point. Unfortunately, due to the tight coupling between IP addresses and end hosts during the initial 20 years of Internet deployment—i.e., one of the unstated assumptions has been that machines rarely moved between attachment points, and attachment points were rarely shared between machines—IP addresses were reused as host

identifiers—that is, an IP address came to be far more than an ephemeral routing locator: *it was a machine's identity*.

A critical issue is that the use of IP addresses as end host identifiers creates a number of problems when end hosts move. Mobile IP [PER02] was developed to address this exact problem: it allowed machines to take their IP address with them, but resulted in an inefficient forwarding mechanism. A number of systems have proposed efficient mechanisms for intercepting packets near their source and forwarding them to a mobile host's current location. However, the Internet architecture effectively limits us to intercepting packets at a mobile host's home. Similarly, the design of a network that assumes that most hosts are mobile may be grossly different from either the current Internet architecture or any existing proposal. The deployment of more efficient mobile host support might greatly change the way that typical portable devices (e.g. laptops, PDA, cell phones) connect to the Internet.

In the interest of expediency, today's end hosts change IP addresses each time they move. However, operators have implemented policies that make implicit assumptions about the machines using particular IP addresses. For example, IP addresses are often used to specify security and access policies as in the case of ingress filtering to alleviate denial-of-service attacks [FE98]. To the extent that IP addresses change, we loose any identity or accountability that might have been keyed to addresses.

The current reality, then, is a mess: Internet addressing is neither secure, efficient, nor architecturally clean. Further, neither end hosts nor network attachment points are sufficient to describe the end points in today's Internet. Depending on the situation, end points may be applications (that may move between machines), sessions (that may move between applications), users (that may move between applications and machines), or data (that can exist almost anywhere).

A new architecture needs to remove the coupling between topological location and endpoint identity present in IP addresses. One proposal is that endpoints need to be given a topology independent identifier, and routing and addressing cannot depend on the identity of the endpoints. A number of proposals have explored possible avenues, each with their own strengths and weaknesses. The Host Identity Protocol (HIP) [MO05] provides each end host with a cryptographically secure identifier, which can be used to anchor transport end points, as well as input to security policies. Routing and addressing continues to be performed by traditional IP, but IP addresses are treated only as ephemeral locators. Another architectural possibility is that end-points (as equated with physical machines or operating systems) need not have any globally known identity at all. Instead, application level entities have shared identities that they use to confirm each end to the other, and higher level name spaces such as a re-designed DNS are used to give global names to services, so that they can be found.

While it's not yet clear exactly what the properties of an endpoint identifier should be, nor precisely what constitutes an endpoint, it is clear that IP addresses and network interfaces are not the right abstractions. A complete redesign of the architecture for location, global naming and shared identity will enhance the security, efficiency, ease-of-use, and flexibility of the basic forwarding infrastructure.

3.4. End Host Assumptions

The current Internet architecture makes several assumptions about the hosts that connect to its edge—that they are usually connected, that they do not move very often, that they are best

identified by relatively static names/addresses rather than more dynamic properties, and so on. This has made it difficult to incorporate devices such as sensor motes or functions such as delay tolerant network routing directly into the Internet infrastructure. Most commonly, intermediate nodes—e.g., sensor base stations, proxies, and home agents—are used to allow these devices to participate on the Internet. Unfortunately, this translation typically incurs some loss in functionality and performance.

Internet routing is based on destination address. But sensor nets often route data based on its value. Algorithms such as diffusion routing are used to build data-driven routing patterns that allow for an application-specific integrated pattern of routing and processing. Te extend sensor nets across the Internet, what is needed is support for a *sensor overlay* that allows a set of agents to recreate schemes such as diffusion routing on top of Internet connectivity. By simplifying the direct attachment of sensor networks to the Internet, we could enable a global-scale mesh of sensor networks, thus supporting a wide-variety of natural science research. In addition to the associated naming and routing architectural changes, a global sensor mesh may also require new security infrastructure. For example, while access to current sensor networks is limited by physical proximity, a global mesh of sensor networks would require enforcement of policies for access to and use of the collected sensor data.

Another limitation of the current architecture is the assumption that nodes are connected in a way that permits near instantaneous communications. Staged or delayed delivery is a part of some applications, such as email, but is not recognized as a problem at the Internet level. The correct solution to this requirement may involve a *delay tolerant* overlay, of the sort being developed by the DTN project. But it is also possible that the Internet architecture itself should better take account of nodes or regions that are poorly and intermittently connected. The deployment of transport and routing protocols that support such long-delay links would enable Internet access in a variety of impoverished and poorly connected regions. Unfortunately, the best design for supporting high delay links remains an open issue. In addition, it is unclear whether a common suite of the algorithms, protocols, and applications could support both interactive and delay-tolerant operation.

3.5. User-Level Route Choice

The current Internet architecture performs routing in an opaque manner that does not allow users control over the paths taken by traffic to and from them. In this context, a *user* could be an actual human, an application program, their Internet provider, or even an overlay service running on their behalf. This limitation restricts several desirable goals. For example, a user cannot express the choice of their ISP beyond their selection of an access provider, or direct traffic along links that have higher availability than the default path.

Relaxing this restriction provides several potential benefits, both technical and economic. Because users know whether or not a particular path is actually working for them, choosing between multiple paths in the network based upon whether they are functional can lead to improved availability and performance. Such path selection can create an enhanced competitive landscape by allowing users to easily switch between packet carriers based on their performance, cost, or availability, a choice that does not exist today. Permitting users to express their routing preferences in a more fine-grained manner may permit ISPs to offer increased service differentiation: Instead of applying a "one-size-fits-all" policy to their traffic, ISPs could perform routing and traffic engineering based upon the user traffic preferences in addition to

their own metrics and policies, or even offer unique policies such as keeping all traffic within the continental United States for security reasons.

Permitting users to control their routes opens a variety of issues. Foremost among these is resolving the conflicts between the preferences of multiple users and of the ISPs who carry their traffic. It is important that the architecture ensure the stability of the network despite the route changes induced by user choice. This selection creates a more complex economic environment; it offers potential rewards in user choice and competition, but requires solutions to issues of accounting, pricing, billing, and inter-ISP contracts. Because this architectural change involves the user or some proxy thereof, implementing a more flexible routing architecture involves changes to the entire network—including hosts. Finally, it is necessary to seriously consider the security implications of any proposed architecture (such as source routing) to ensure that they do not create additional vulnerabilities.

It is our hope that an architecture that enabled some level of user control over routes would lead to an increase in the reliability of the Internet, and an improvement in the sets of features offered to users. First, today's Internet often lacks the necessary reliability to serve as a basis for emergency services, real-time control, or for particularly time-sensitive transactions. Systems based upon user control of routing may be able to increase this availability sufficiently for the Internet to encompass a wider variety of critical services. Second, in today's Internet environment, it can be difficult to determine what party is responsible for poor (or superior) performance, and to reward them by choosing to use them for your Internet service. This inability keeps from the Internet many of the benefits of increased competition—lower costs, more efficient practices, and a rich set of services that differentiate providers from one another. User control of routes could help move the Internet in this positive direction.

3.6. Control and Management

The original Internet architecture focuses on best-effort reachability among cooperative users, which results in a primitive control/management infrastructure that bundles the reachability logic and data fowarding functions in each individual router. Today's networks, owned by competing entities and operated in different environments (data center, access/metro, enterprise, ISP) are called upon to meet far more sophisticated network wide-objectives: dependability, policy, traffic engineering, security, ease of management, cost-effectiveness, and so on. As new network-wide objectives need to be accomplished, the original box-centric control architecture (tightly coupled decision making logic and data plane in one box) forces point solutions to be invented, and then retro-fitted onto the network. This has resulted in significant complexity, with diverse and local decision logic distributed across multiple network elements. This is a fundamental reason for the fragility of the Internet, where a single local event can cause a network-wide meltdown. In short, operational complexity plagues today's Internet.

The trajectory of existing incremental efforts is to incorporate more point solutions into the control plane, which only exacerbates the problem of management complexity. If network management is re-architected to explicitly consider multiple network-wide objectives, there is the potential to reduce the fragility of today's networks and lower the complexity of the network. In general, such a change would enable rapid innovations management functions by explicitly separating the implementation of control logic from the routers that implement data

plane functions. The research community is actively pursuing this agenda [HJ00, TU01, KO00, HA02, DO02, YA04, FA05, HS03].

One of the key barriers to progress is the relative opaqueness of the network, meaning that components do not support communication of operationally relevant information to each other. Such information could be aggregated and analyzed [CL03], thereby facilitating load balancing, fault diagnosis, anomaly detection, application optimization, and other traffic engineering and network management functions [SA99, CA00, AR92, RE99, CR03, HO93, PO97, FE00a, FE00b, LA04, SH99, FE02a, FE02b]. This lack of transparency, even for components within the same administrative domain, is framed by not only technical reasons (e.g., a router can only export information about its best known routes, not all known routes, rendering it impossible to realistically simulate what-if scenarios) but also by competitive business realities (operators have a disincentive to reveal operational details about their infrastructure).

The opaqueness of the routing system merits particular consideration. It is impossible to realistically model routing behavior more than 1-hop away from a given node since the policy-rich features added to BGP (e.g., MEDs) have further removed what little transparency originally existed, and thus fatally hindered the ability to logically reason about the routing system. More fundamentally, the current market structure of the Internet promotes information hiding, and when those building and maintaining infrastructure consider opaqueness a feature rather than a limitation, an architectural position that favors transparency also needs to consider how to enforce that transparency on a market that will hide whatever it chooses, even at the expense of operational efficiency.

3.7. Meeting Application Requirements

One architectural decision of the original Internet stands out as playing a critical role in its success: the adoption of a *narrow-waisted hourglass model*. A minimal and carefully chosen set of global capabilities at the mid-level of the architecture allows both higher-level applications and lower-level communication technologies to coexist, share capabilities, and evolve rapidly. The narrow-waisted model is critical to the Internet's ability to adapt rapidly to new user demands and changing technologies.

However, new application classes place demands on core IP capabilities that many argue cannot be met within the current model. Moreover, the growing scale and increasing heterogeneity of the Internet increases the perceived value of placing functionality *within the network*, to better take advantage of localized knowledge and optimization opportunities.

There are several technical responses to these developments. One is to *widen the waist of the hourglass*, to augment the Internet's core forwarding service to include additional functionality. Proposals to satisfy new application requirements by adding new capabilities—e.g., QoS control, multicast, anycast, policy-based routing, data caching, and so on [PO81, SH97, BR97, BE00, QU01, HI03, ST93, AL99, BY98, AKA, DIG]—to the core IP protocols have dominated the last ten years of Internet networking research. Some of these have been deployed in specific circumstances; some have failed to be deployed at all. Whether or not we can deploy any specific enhancement, there is a risk to the stability and coherence of the Internet architecture if we keep adding function to the basic forwarding layer. It is a widely held belief that flexibility, deployability, and evolvability are achieved only when the truly *universal* portions of the architecture are also truly *minimal*.

A second response is to *add a layer* to the architecture, inserting purpose-tuned overlay networks between the global infrastructure and the ultimate end nodes. This strategy offers many potential advantages. Overlays constructed with application-level requirements in mind can make network-level decisions, such as routing, service model, and data manipulation, tuned to the specific application. Overlays designed to support small-scale applications can utilize algorithms that would not scale to global size. Moving functionality from shared infrastructure to multiple overlays increases decentralization by more cleanly modularizing responsibilities and administrative operations. This potential has created great interest in the overlay model, with vigorous activity in the academic and commercial communities [AN01, SU02, BA02, RA02, RO01a, RO01b, ST01, CH02, SP03, DA01, KU00, PA04, ZH04].

As important and valuable as this work is, however, virtually all of this activity has focused on understanding and increasing the functionality of specific overlay networks and algorithms, leaving unanswered the single most critical question relevant to this approach: what is it that lies underneath? What is the appropriate narrow, universally shared environment to support the overlays?

This environment, which we term the *overlay substrate,* must play three critical roles. First, it must *support* the different overlay structures and services that it underpins, in the same way that today's Internet supports end-to-end applications. This implies that the underlay must expose information about the underlying physical network that overlays need 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.

A third approach is to move the narrow waist to a lower level of the protocol stack, that is, define a *network substrate* consisting of a collection of physical resources (nodes and links) on top of which multiple, alternative network architectures could co-exist. Fundamentally, this implies that virtualization would become a first-class feature of the network architecture, allowing for on-going diversity and renewal at the network layer. In this world, IP would become just one of potentially many network architectures. Others might provide alternative security or robustness properties, or simply be tailored for certain classes of applications.

A diversified network could create a range of new opportunities for current stakeholders. It would allow providers of the physical infrastructure to focus on virtual networks as their primary "customers", allowing them to distinguish themselves through the quality of their infrastructure and the support services they provide to virtual networks. Equipment vendors would have the opportunity to create new classes of equipment and provide design services to virtual network providers. Shifting the provision of end-to-end services to virtual networks would create a whole new class of business opportunities, potentially sparking a wave of entrepreneurial innovation.

Before the diversified Internet concept can be put into practice, there is a range of open issues that will need to be explored through on-going research efforts. These include: (1) defining the nature of the resource provisioning interface between substrate providers and virtual networks; (2) developing mechanisms that enable virtual networks to easily use resources provided by the substrate to implement innovative new services; (3) design systems that allow virtual networks to co-exist on a common substrate without interfering with one another, while still allowing them to interact where such interaction is desired; (4) extending

host operating systems to allow users to conveniently use the services of multiple virtual networks; and (5) develop strategies for implementing access links that would allow the access link resources to be flexibly re-allocated among multiple virtual networks, while still allowing for predictable performance.

Although these last two approaches address the problem at different layers, they both focus on designing a suitable substrate on top of which multiple network architectures and services can run. In both cases, the goal is to identify the critical balance of functionality, minimality, stability, evolvability, and deployability that will allow a shared virtualized infrastructure to spread globally, while supporting a rich and changing architectural ecosystem. Moreover, understanding which layer is the most appropriate "new waist of the hourglass" is one of the most interesting questions facing the research community.

4. Experimental Deployment of Architectural Innovations

To be effective, a research program that seeks to promote architectural innovation must enable researchers to create, deploy, and evaluate novel architectures. These architectures must both run at scale, and carry traffic from real users. This calls for the creation of a testbed of global reach and diverse capabilities. This section outlines the goals and design principles that shape this testbed.

4.1. Goals and Scope

The testbed must provide an environment in which multiple new network architectures and services can be deployed. This means there should be as few restrictions as possible on the architectures that operate on the testbed and on the capabilities provided by the testbed. Toward this end, the testbed should include a diversity of links and nodes (both physical and virtual), and permit connection of arbitrary edge devices.

The testbed should be capable of bridging the gap between so-called production testbeds, which constrain research, and research testbeds, which constrain users [KU02]. It must be capable of attracting and supporting users of its services beyond the research community. This is essential for allowing new architectural innovations to be evaluated at scale, and for creating a population of users whose demonstrated interest in a new capability can stimulate technology transfer to the commercial Internet.

To meet these goals, the testbed need not have a single architecture in the traditional sense. Instead, its role is to provide an environment in which a diverse set of experimental networks—each with its own distinct architecture—can operate. In this sense, the testbed is really a *meta-testbed* that hosts a heterogeneous collection of testbeds within it. Each of these individual testbeds is allocated a portion of the meta-testbed's resources. The meta-testbed should constrain the hosted activities hosted to the minimum extent possible, and provide for varying degrees of isolation and interconnection among these activities. The common part of the meta-testbed, which we refer to as the *substrate*, provides the mechanisms for allocating and configuring resources and ensuring the necessary isolation.

The meta-testbed should be viewed as a dynamic artifact: the physical resources, management capabilities, governance processes, implementation, and even the substrate design will evolve with time. The physical resources in the testbed may include a mix of dedicated physical links and nodes, virtual components contributed on a permanent or temporary basis

by testbed users, and resources leased from third-party providers or consortia such as NLR [NLR]. The substrate and management infrastructure should incorporate standard service policies and interfaces to enable organic growth, provide incentives to contribute, and manage dynamic resources available to the meta-testbed on a temporary basis under various terms.

4.2. Key Concepts

There are several key architectural concepts that we expect to play a central role in the design of the meta-testbed. The meta-testbed will consist of links, nodes and edge devices. Links may be implemented in a variety of ways, including direct physical links, MPLS paths, and IP tunnels. The meta-testbed links can be shared by different experimental networks running within the meta-testbed, using well-known virtual link multiplexing techniques.

The meta-testbed nodes provide a collection of memory, processing, and storage resources. They might correspond to virtual machines running on commodity processors; dedicated general-purpose processors; both dedicated and virtualized network processors and programmable hardware; and virtualized routers. The meta-testbed provides mechanisms to configure these resources for use by different experimental networks and to provide isolation between experimental networks.

Edge devices (including traditional hosts) may participate in multiple networks running within the meta-testbed. In some cases, this will require that edge devices implement separate protocol stacks, although a key to the success of the meta-testbed will be the development of mechanisms that make it easy for users to "opt-in" to experimental networks that offer some value-added capability.

Each experimental network will run on some subset of the meta-testbed resources. We call the substrate resources bound to a particular experimental network a *slice*, borrowing the term from PlanetLab [PET02, BA04]. Each slice will include some number of nodes (including both physical processors and virtual machines multiplexed shared hardware) connected by links (including both physical links and virtual links). The main responsibility of the meta-testbed management software will be to provide mechanisms that can be used to allocate resources to slices, and ensure that slices do not interfere with each other.

We note that different users of the meta-testbed will require varying degrees of isolation, connectivity, dynamism, and control in their slices. Slices that require full isolation from other slices (including traffic and performance isolation) should have a means to acquire it, subject to the availability of the required resources. At the same time, it should be possible to connect different slices to one another, where that is appropriate and mutually agreed upon. While it is likely that the meta-testbed will initially incorporate a narrow range of resources and simple assignment policies, this range should advance over time.

4.3. Design Principles

The design and development of the meta-testbed will require decisions on a wide range of issues. The workshop did not provide sufficient time to fully explore these issues, but there was substantial agreement on some core design principles, which participants felt should guide the design process and the subsequent operation of the meta-testbed. These are summarized below.

- ***Service/architecture neutrality***. What is most important for research in network architecture is that the level of abstraction be low enough to permit full experimentation at layer 3 and above. Different slices of the common testbed may reflect different layer 3 architectures at the same time. In particular, networks running in different slices may use different packet formats and service models.
- ***End-system diversity***. The meta-testbed should enable heterogeneity in the end systems that connect to it and participate in the experimental networks running within it. In particular, it should enable the connection of limited functionality end-systems (such as wireless PDAs and sensor motes).
- ***Ease of user access***. Mechanisms are needed to make it easy for users to join one or more experimental networks running in the meta-testbed, and to transparently fall back to the standard Internet whenever the experimental network cannot provide the requested service. In some cases, this can be accomplished using transparent re-direction mechanisms [KA04]. In other cases, it may require the installation of new protocol stacks in hosts.
- ***Sustainability and incentives***. To ensure the sustainability of the meta-testbed, it should be possible for participating institutions to join by contributing resources in return for access to the resources of the meta-testbed as a whole.
- ***Inter-slice composition***. The testbed infrastructure must enable interconnection among slices by mutual consent, and between slices and the external Internet. This permits slices to host network services with external users, and/or to act as transit networks. Nothing should prevent a researcher from inter-connecting a network running within a slice with another network. This other network could be running within another slice of the meta-testbed, or it could be the commodity Internet or another custom network (or testbed) that runs over standard IP protocols.
- ***Policy and governance***. Since the meta-testbed will comprise shared infrastructure, there must be a governance process to guide allocation of resources to slices, and a software architecture that implements and enforces the policies. Some slices will likely require strong performance isolation, which will make some form of admission control necessary.

There were additional issues raised for which there was not a broad consensus. While these issues should be explored further, the workshop participants felt that design and development of the meta-testbed should not be delayed until these issues can be fully resolved.

First, there was a discussion about how much effort should be focused on performance, at least when considered relative to the need to design and evaluate new functionality. For example, there is an opportunity to incorporate high performance backbone links into the meta-testbed, using fiber optic facilities available through the National Lambda Rail [NLR]. NLR links can operate at 10 Gbps, allowing the meta-testbed to carry large traffic volumes and making it possible to evaluate experimental networks operating at high speeds. For many network research purposes however, this capability is not strictly necessary, and fully exploiting this capability will require the development of high performance testbed nodes. It is difficult to know if the benefits provided by such high performance nodes would justify the cost of their development; it may be possible to accomplish most of the objectives using clusters of general-purpose processors connected by COTS switches. The question, then, is one of priorities: should sufficient funding be available, work on functionality and performance should

proceed in parallel; if not, designing new functionality that address the many challenges facing the Internet is the highest priority.

Second, while the general consensus among workshop participants was that the meta-testbed must use packet transport—thereby allowing existing network access mechanisms (primarily Ethernet) to be used to connect end users to the meta-testbed—there was an acknowledgment that the broader network community might not agree, favoring instead a circuit-based approach. However, the workshop participants believe that experimental networks running within the meta-testbed could offer circuit-like services within their own slice, for example, by using virtual links with reserved bandwidth and implementing per virtual link smoothing buffers to convert packet links with a small amount of jitter into constant delay links. Such an approach does not allow for high performance circuit switched elements (such as optical cross-connects) to be incorporated into the meta-testbed, but there was no clear suggestion for how such elements could be included, given the likely resource constraints on the meta-testbed as a whole.

4.4. Departure Point

The meta-testbed we envision is ambitious, but the networking community has a strong track record of creating testbeds and testbed technologies, and using them to evaluate and demonstrate new research ideas. Some examples of current efforts that provide subsets of the capabilities needed for the proposed meta-testbed required include:

- PlanetLab [PET02, BA04], which focuses on node virtualization and global resource management, and is widely used for research in network services.
- X-bone [TO01, TO03] and 6-bone [SXB], which define core (L3) capabilities for network virtualization and are supported by multiple operating systems;
- Emulab [EMU,WH02] and Netbed, which allocate and configure heterogeneous end-system resources and network resources (using L2 virtualization) together;

The scale and presence of these testbeds have proven to be significant enablers of new research, with strong momentum and community involvement. They serve as incubators and proofs of concept for many of the architectural ideas outlined earlier in this report, as well as for the proposed meta-testbed itself. What is needed now, however, is a more comprehensive effort that incorporates a broad range of resources and capabilities.

PlanetLab, in particular, offers a starting point that can be leveraged immediately. It provides a shared overlay infrastructure that spans over 525 nodes distributed across over 250 sites and 28 countries. It currently hosts over 350 slices—each running a different network architecture, service, or application—on the shared infrastructure. PlanetLab also includes software that allows end users to seamlessly connect their desktop machines to services they want to employ, resulting in network traffic to over 500k unique IP addresses every day.

However, this infrastructure is not sufficient by itself. It needs to be enhanced with a richer set of link technologies (e.g., by adding MPLS paths and dedicated circuits to the currently supported IP tunnels) and a more diverse set of node configurations (e.g., by adding dedicated processors and customizable hardware to the currently supported commodity processors). Over time, the resulting infrastructure will become the meta-testbed that meets our objectives. Much work remains to be done to realize the full scope of the meta-testbed, but the end-goal is clear: a meta-testbed that combines the *global reach* of overlays with the *performance realism* of physical

links and programmable routers. In making progress toward this vision, there must be an open and inclusive community process for designing the mechanisms of the meta-testbed, and this process must carefully balance competing demands and the needs of the community.

5. Recommendations

This report has explained the need for dramatic improvements in the security, robustness, manageability, flexibility, and cost-performance of the Internet as a critical piece of our societal infrastructure, and the inability of evolutionary research to effect the needed architectural improvements. Thus, we argue that the National Science Foundation should take a new approach to fostering disruptive innovation in networking. With the right support from and partnership with NSF, the network research community is poised to define the next generation of Internet technology, fundamentally and permanently addressing the problems the Internet has today. Further, a revised architecture has the potential to unleash a new class of applications, currently stalled behind the limited functionality of today's best effort Internet service. A new architecture could also better leverage technology trends towards incredibly high bandwidth optical networks and increasingly capable computation devices embedded in the network. Radically new networks such as ad hoc wireless networks and sensor networks would also be enabled.

In sum, if successful, this effort will directly benefit virtually every member of our society, enhancing homeland defense, ensuring that communication over the Internet is as reliable and secure as physically possible, delivering the raw performance of the network to demanding scientific and engineering applications, reducing the cost of Internet access for all users, and enabling the next generation of innovation in our shared network infrastructure.

However, achieving this vision will not be easy. Any replacement for the Internet's architecture must demonstrate its value via widespread use, or this effort will be pointless. This will require both NSF and the network research community to change business as usual, focusing on the construction of practical and usable systems, in addition to cutting edge research. This means that NSF must be willing to put forward sufficient funds, sustained over an extended period of time, for the community to build and operate an alternative architecture in live use by large numbers of users. And the research community must be willing to put the effort into moving their ideas into practice, rather than being satisfied with paper designs. Although progress can be made immediately, and funding ramped up in response to success in meeting milestones, make no mistake: the total scope of the effort needed is much larger than can be supported under NSF's current networking research budget. The corresponding benefit will be enormous—nothing less than putting the world's communication infrastructure on a secure, robust, flexible, and efficient basis for the foreseeable future.

Specifically, the workshop participants make the following recommendations to the National Science Foundation:

Recommendation 1: Immediately initiate a research program on experimental architectural research in networking.

Given the current architectural limitations, and the encouraging prospects for overcoming them, NSF should provide significant multi-year funding for architectural research. If successful, the potential benefits are enormous, easily justifying the modest initial outlay. The

amount of research funding may need to ramp in future years; as we gain experience with using these new architectures, additional opportunities will open up to leverage that experience. A dedicated research program will pull together the research community to tackle the broad scope of thorny architectural questions that we have outlined in this report, and that must be addressed for the program to be successful.

Recommendation 2: Foster experimental validation of new architectural research in networking.

Paper designs, although thought provoking, are unconvincing, both to the companies that need to adopt them, and to the research community in evaluating ideas and in gaining insight into design tradeoffs. Thus, to maximize our chance of success, NSF must foster an expectation within the experimental architectural research program that research ideas should normally be validated under real use. Not every research idea will be worth validating, but every successful research idea must be validated before it can be considered to have proven its worth. The burden this places on the research community should not be underestimated – working systems are much more expensive to build than paper designs, and useful systems are even more expensive. In many cases, this will involve deployment on an appropriate testbed and usage by the population-at-large. Although we should leave the specifics of how to accomplish experimental validation to the creativity of researchers and the intelligence of peer review, this is likely to mean that larger, multi-principal investigator and center-scale efforts will be needed to make progress along the most promising architectural directions.

Recommendation 3: Fund the development and deployment of suitable testbeds.

Since experimental validation is an important component of this research program, it is essential that researchers have access to suitable testbeds. The alternative, requiring each new effort to create their own testbed, is both impractical and inefficient. Instead, NSF should endeavor to build a meta-testbed that reduces the barrier to entry for new architectural ideas. To meet short-term needs, NSF should support an initial meta-testbed that can be deployed immediately. At the same time, NSF should initiate a deliberative process through which the community can identify long-term solutions to its meta-testbed requirements. NSF should support the deployment and on-going operation of the resulting meta-testbed, but this funding should not decrease that devoted to architectural research itself.

Recommendation 4: Start a process that will lead to substantial increases in funding for a broad multi-disciplinary effort in this area over the next few years.

To design, construct and widely deploy a new architecture for the Internet is an enormously difficult and, at the same time, an enormously important undertaking. It cannot be done on the cheap. Nor can it be done by the networking research community alone. To be successful, we will need to enlist the efforts of distributed systems researchers, e-scientists, application developers, computer architects, and network hardware technologists. Gaining consensus among those communities of the need and value of our efforts will be a difficult, but necessary step towards our eventual success. Equally important will be the ability to put sufficient funds behind that emerging consensus, to be able to demonstrate the value of our work in widespread practice.

Recommendation 5: Find ways to promote synergy and convergence among architectural visions.

Academic research focuses on novelty and, in so doing, often accentuates differences rather than identifying commonality. The past success of the Internet strongly suggests that we will be the most successful if we can coalesce around common architecture features. Architecture, by its very nature, "defines that on which we must agree." Thus, to be effective, architectural researchers should seek convergence rather than divergence. NSF should find ways to promote the necessary community practices, such as encouraging participation in working groups and funding synthetic research.

Recommendation 6: Help the community learn from industry.

Disruptive architectural research should not be fettered by today's problems and practices, but it must be informed by them if we are not to simply repeat the mistakes of the past. The large gap between the research and commercial communities often prevents effective communication between the two, to the detriment of both. To bridge this gap, NSF should facilitate interactions between researchers and practitioners. This extends not only to the operational community of ISPs, but also to network application developers and those developing new physical layer technologies, as the success of any new architecture will be measured by how well it matches the requirements of its hardware and the needs of its users.

References

[AL99] C. Alaettinoglu, et. al., "Routing Policy Specification Language (RPSL)," RFC 2622, June 1999.

[AKA] Akamai, www.akamai.com

[AN01] D. Andersen, H. Balakrishnan, F. Kaashoek, R. Morris, "Resilient Overlay Networks," *Proceedings of ACM SOSP*, 2001.

[AR92] R. Arun, P. Venkataram, "Knowledge Based Trouble Shooting in Communication Networks," *Proceedings of Symposium on Intelligent Systems*, November 1992.

[BE02] S. Berson, S. Dawson and R. Braden. "Evolution of an Active Networks Testbed," *Proceedings of the DARPA Active Networks Conference and Exposition*, 5/02.

[BA02] M. Balazinska, H. Balakrishnan, D. Karger, "INS/Twine: A Scalable Peer-to-Peer Architecture for Intentional Resource Discovery," *Proceedings of the 1st International Conference on Pervasive Computing*, 2002.

[BA04] A. Bavier, M. Bowman, B. Chun, D. Culler, S. Karlin, L. Peterson, T. Roscoe, T. Spalink, and M. Wawrzoniak. *Proceedings of the 1st Symposium on Network System Design and Implementation (NSDI '04)* San Francisco, CA (March 2004).

[BE00] Y. Bernet, et. al., "A Framework for Integrated Services Operation over DiffServ Networks," RFC 2998, November 2000.

[BR97] B. Braden, et. al., "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification," RFC 2205, Sepetember 1997.

[BY98] J. Byers, M. Luby, M. Mitzenmacher, A. Rege, "A Digital Fountain Approach to Reliable Distribution of Bulk Data," *Proceedings of ACM Sigcomm*, 1998.

[CA00] Z. Cao, Z. Wang, E. Zergura, "Performance of Hashing-Based Schemes for Internet Load Balancing," *Proceedings of IEEE Infocom*, 2000.

[CH02] B. Chun, J. Lee, H. Weatherspoon, "Netbait: a Distributed Worm Detection Service," Project website, http://netbait.planet-lab.org, 2002.

[CL03] D. Clark, C. Partridge, J. Ramming, J. Wroclawski, "A Knowledge Plane for the Internet," *Proceedings of ACM Sigcomm*, 2003.

[CL05] D. Clark, et. al. Making the world (of communication) a different place. January 2005. http://www.ir.bbn.com/~craig/e2e-vision.pdf

[CR03] M. Crovella, E. Kolaczyk, "Graph Wavelets for Spatial Traffic Analysis," *Proceedings of IEEE Infocom*, 2003.

[DA01] F. Dabek, M. Kaashoek, D. Karger, R. Morris, I. Stoica, "Wide-area cooperative storage with CFS," *Proceedings of ACM SOSP*, 2001.

[DET] DETER project web site. http://www.isi.edu/deter/

[DIG] Digital Fountain, www.digitalfountain.com

[DO02] A. Doria, F. Hellstrand, K. Sundell, T. Worster, "General Switch Management Protocol (GSMP) V3," RFC 3292, June 2002.

[EMU] Emulab project web site. http://www.emulab.net

[FA05] A. Farrel, J. Vasseur, J. Ash, "Path Computation Element (PCE) Architecture," Internet-Draft, March 2005.

[FE98] P. Ferguson, D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," RFC 2267.

[FE00a] A. Feldmann, A. Greenburg, C. Lund, N. Reingold, J. Rexford, F. True, "NetScope: Traffic Engineering for IP Networks," *IEEE Network*, March/April 2000.

[FE00b] A. Feldmann, A. Greenburg, C. Lund, N. Reingold, J. Rexford, "Deriving Traffic Demands for Operational IP Networks: Methodology and Experience," *Proceedings of ACM Sigcomm*, 2000.

[FE02a] N. Feamster, J. Borkenhagen, J. Rexford, "Controlling the Impact of BGP Policy Changes on IP Traffic," *Proceedings of NANOG25*, 2002.

[FE02b] N. Feamster, J. Rexford, "Network-wide BGP Route Prediction for Traffic Engineering," *Proceedings of ITCOM*, 2002.

[HA02] M. Handley, O. Hodson, E. Kohler, "XORP: An Open Platform for Network Research," *Proceedings of ACM HotNets-I*, October 2002.

[HI03] R. Hinden, S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture," RFC 3513, April 2003.

[HJ00] G. Hjalmtysson, "The Pronto Platform – A Flexible Toolkit for Programming Networks using a Commodity Operating System," *Proceedings of OpenArch*, 2000.

[HO93] J. Hochberg, et. al., "Nadir: An automated system for detecting network intrusion and misuse," *Computers & Security*, 12(3), 1993.

[HS03] H. Hsieh, et. al., "A Receiver-Centric Transport Protocol for Mobile Hosts with Heterogeneous Wireless Interfaces," *Proceedings of Mobicom*, 2003.

[I2] Internet2 project web site. http://www.internet2.edu

[I3] Internet Indirection Infrastructure project web site. http://i3.cs.berkeley.edu

[KU02] J. Kurose (editor). Report of the NSF Worksop on Network Research Testbeds. www.gaia.cs.umass.edu/testbed_workshop, 11/02.

[KA04] J. Kannan, A. Kubota, K. Lakshminarayanan, I. Stoica and K. Wehrle, "Supporting Legacy Applications over *i3*," UCB Technical Report No. UCB/CSD-04-1342, May 2004.

[KO00] E. Kohler, "The Click Modular Router," Ph.D. thesis, MIT, November 2000.

[KU00] J. Kubiatowicz, et. al., "Oceanstore: An Architecture for Global-Scale Persistent Storage," *Proceedings of ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2000.

[LA04] A. Lakhina, M. Crovella, C. Diot, "Diagnosing Network-Wide Traffic Anomalies," *Proceedings of ACM Sigcomm*, 2004.

[LBO] L-bone. www.loci.cs.utk.edu/.

[MO05] R. Moskowitz, P. Nikander, P. Jokela, T. Henderson, "Host Identity Protocol," Internet Draft, Feb. 2005.

[NLR] National Light Rail Project web site. www.nlr.net.

[PER02] C. Perkins, Ed. "IP Mobility Support for IPv4," RFC 3344.

[PET02] L. Peterson, T. Anderson, D. Culler and T. Roscoe. "A Blueprint for Introducing Disruptive Technology into the Internet," *Proceedings of ACM HotNets-I Workshop* (October 2002).

[PET04] L. Peterson, S. Shenker, and J. Turner. "Overcoming the Impasse Through Virtualization," *Proceedings of ACM Hotnets-III* (November 2004).

[PA04] K. Park, V. Pai, L. Peterson, Z. Wang, "CoDNS: Improving DNS Performance and Reliability via Cooperative Lookups," *Proceedings of Symposium on Operating Systems Design and Implementation (OSDI)*, 2004.

[PO81] J. Postel, ed., "Internet Protocol," RFC 791, September 1981.

[PO97] P. Porras, P. Neumann, "EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances," *Proceedings of NIST-NCSC National Information Systems Security Conference*, 1997.

[PRE05] President's Information Technology Advisory Committee. Cyber Security: A Crisis of Prioritization, February 2005. [QU01] B. Quinn, K. Almeroth, "IP Multicast Applications: Challenges and Solutions," RFC 3170, September 2001.

[RA02] S. Ratnasamy, M. Handley, R. Karp, S. Shenker, "Topologically-Aware Overlay Construction and Server Selection," *Proceedings of IEEE Infocom*, 2002.

[RE99] R. Reddy, D. Estrin, R. Govindan, "Large-Scale Fault Isolation," *IEEE Journal on Selected Areas in Communications*, March 1999.

[RO01a] A. Rowstron, P. Druschel, "Storage Management and Caching in PAST, A Large-Scale Persistent Peer-to-Peer Storage Utility," *Proceedings of ACM SOSP*, 2001.

[RO01b] A. Rowstron, P. Druschel, "Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems," *Proceedings of IFIP/ACM International Conference on Distributed Systems Platforms*, 2001.

[SA99] S. Savage, A. Collins, E. Hoffman, J. Snell, T. Anderson, "The End-to-End Effects of Internet Path Selection," *Proceedings of ACM Sigcomm*, 1999.

[SH97] S. Shenker, C. Partridge, R. Guerin, "Specification of Guaranteed Quality of Service," RFC 2212, September 1997.

[SH99] A. Shaikh, J. Rexford, K. Shin, "Load-Sensitive Routing of Long-Lived IP Flows," *Proceedings of ACM Sigcomm*, 1999.

[SP03] N. Spring, D. Wetherall, T. Anderson, "Scriptroute: A facility for distrubted Internet measurement," *Proceedings of USENIX Symposium on Internet Technologies*, 2003.

[SXB] 6 Bone. http://www.6bone.net/.

[ST93] R. Steenstrup, "Inter-Domain Policy Routing Protocol Specification: Version 1," RFC 1479, July 1993.

[ST01] I. Stoica, R. Morris, D. Karger, F. Kaashoek, H. Balakrishnan, "Chord: A Peer-to-Peer Lookup Service for Internet Applications," *Proceedings of ACM Sigcomm*, 2001.

[ST02] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, S. Surana, "Internet Indirection Infrastructure," *Proceedings of ACM SIGCOMM*, August, 2002.

[SU02] L. Subramanian, I. Stoica, H. Balakrishnan, R. Katz, "OverQoS: Offering Internet QoS Using Overlays," *Proceedings of ACM HotNets*, 2002.

[TO01] J. Touch. "Dynamic Internet Overlay Deployment and Management Using the X-Bone," *Computer Networks*, July 2001, pp. 117-135.

[TO03] J. Touch, Y. Wang, L. Eggert, G. Finn. "Virtual Internet Architecture," ISI Technical Report ISI-TR-2003-570, March 2003.

[TU01] P. Tullmann, M. Hibler, J. Lepreau, "Janos: A Java-oriented OS for Active Networks," *IEEE Journal on Selected Areas in Communications*. Volume 19, Number 3, March 2001.

[WH02] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb and A. Joglekar. "An Integrated Experimental Environment for Distributed Systems and Networks," *Proceedings of OSDI*, 12/02.

[YA04] L. Yang, R. Dantu, T. Anderson, R. Gopal, "Forwarding and Control Element Separation (ForCES) Framework," RFC 3746, April 2004.

[ZH04] M. Zhang, C. Zhang, V. Pai, L. Peterson, R. Wang, "PlanetSeer: Internet Path Failure Monitoring and Characterization in Wide-Area Services," *Proceedings of Symposium on Operating Systems Design and Implementation (OSDI)*, 2004.

List of Workshop Participants

David Andersen	Massachusetts Institute of Technology	dga@csail.mit.edu
Tom Anderson	University of Washington	tom@cs.washington.edu
Andy Bavier	Princeton University	acb@cs.princeton.edu
Jeff Chase	Duke University	chase@cs.duke.edu
K. C. Claffee	Coop. Assoc. for Internet Data Analysis	kc@caida.org
Patrick Crowley	Washington University	pcrowley@cse.wustl.edu
Mike Dahlin	University of Texas	dahlin@cs.texas.edu
Dave Clark	Massachusetts Institute of Technology	ddc@csail.mit.edu
Constantine Dovrolis	Georgia Institute of Technology	dovrolis@cc.gatech.edu
Joe Evans	National Science Foundation	jbevans@nsf.gov
Darleen Fisher	National Science Foundation	dlfisher@nsf.gov
Paul Francis	Intl. Computer Science Institute	francis@aciri.org
Sergey Gorinsky	Washington University	gorinsky@cse.wustl.edu
Roch Guerin	University of Pennsylvania	guerin@ee.upenn.edu
T. V. Lakshman	Bell Laboratories	lakshman@dnrc.bell-labs.com
John Lockwood	Washington University	lockwood@cse.wustl.edu
Guru Parulkar	National Science Foundation	gparulka@nsf.gov
Adrian Perrig	Carnegie-Mellon University	perrig@cmu.edu
Larry Peterson	Princeton University	llp@cs.princeton.edu
Mothy Roscoe	Intel Research	troscoe@intel-research.net
Srini Seshan	Carnegie-Mellon University	srini@cmu.edu
Scott Shenker	Intl. Computer Science Institute	shenker@icsi.berkeley.edu
Alex Snoeren	University of California, San Diego	snoeren@cs.ucsd.edu
David Taylor	Washington University	det3@arl.wustl.edu
Jonathan Turner	Washington University	jon.turner@wustl.edu
Joe Touch	Information Sciences Institute	touch@isi.edu
Arun Venkataramani	University of Texas	arun@cs.utexas.edu
Xiaowei Yang	Massachusetts Institute of Technology	yxw@mit.edu
Raj Yavatkar	Intel Architecture Labs	raj.yavatkar@intel.com
Hui Zhang	Carnegie-Mellon University	hzhang@cs.cmu.edu

The above participants attended the workshop on January 13-14, 2005, and contributed to the writing of this report. The workshop was organized by Tom Anderson (University of Washington), Larry Peterson (Princeton University), Scott Shenker (University of California, Berkeley), and Jonathan Turner (Washington University).