Internet Research: Comments on Formulating the Problem

Gathered by Sally Floyd,
with contributions from Deborah Estrin, Greg Minshall,
Vern Paxson, Lixia Zhang, and others.

January 21, 1998

1 Introduction

This note contains a discussion about formulating the research problem for Internet research.

The goal of this note is to further the discussion of implicit and explicit assumptions in network research. In particular, this note tries to articulate one such set of assumptions for Internet research. Each of these assumptions is shared by some subset of the network research community, though perhaps none of these assumptions are shared universally. The goal of this paper is not to argue the validity of the assumptions, but to articulate them, to invite discussion of conflicting or shared sets of assumptions, and to consider the implications of these assumptions in formulating problems in Internet research. This process would aim for both a greater convergence of underlying assumptions, and a more explicit and examined discussion of those assumptions.

In some sense, this note is in the tradition of Shenker et al.'s paper on "Pricing in Computer Networks: Reshaping the Research Agenda" [Shenker96], which is a discussion about formulating the research problem for the specific area of network pricing. This note does not address questions about how to design good protocols or write good papers; many of these questions are addressed elsewhere [ADVICE, Lampson84]. This note simply addresses the question of how to formulate the research problem.

The first part of the paper briefly describes some underlying assumptions. The second part discusses implications of these assumptions in formulating research problems in various areas of Internet research.

2 Outlining the underlying assumptions

A summary of the assumptions:

• *Robustness is more important than efficiency.*

• *The Internet is characterized by heterogeneity on many levels.*

• *Complex systems are not best designed from scratch and on paper.*

• *Development and deployment in the infrastructure is of necessity incremental.*

• *Explicit examined assumptions are better than implicit unexamined ones.*

• *Changes in the Internet can be unanticipated and uncontrolled.*

• *The Internet architecture and scale make requirements for global consistency problematic.*

• *Some research problems have their own natural time scales.*

3 Discussion of the assumptions

• *Robustness is more important than efficiency.*

Robustness has been one of the great strengths of the Internet, integral to its design from the very beginning [Clark88]. One of our overriding assumptions is that it is critical not to subordinate robustness to the goal of more closely approximating optimal efficiency. [This is discussed in more detail in the section on TCP.]

• *The Internet is characterized by heterogeneity on many levels.*

The Internet is characterized by heterogeneity on many levels: router scheduling algorithms and queue management mechanisms, routing protocols, levels of multiplexing, protocol versions and implementations, etc. There is heterogeneity of underlying link layers (e.g., point-to-point, multiaccess links, wireless, ATM, FDDI, etc.). Because the Internet is composed of autonomous organizations and internet service providers, each with their own separate policy concerns, there is a heterogeneity of administrative domains and pricing structures. There is a heterogeneity in the traffic mix and in the levels of congestion at different times and places. As is discussed later in the paper, this heterogeneity has to be taken into account in different ways with different areas of research.

• *Complex systems are not best designed from scratch and on paper.*

Our observation is that large complex systems are not best designed on paper as a complete system and then implemented, but are best designed incrementally, with new insights and development growing out of practical experience with existing functionally. Our assumption is that this is true for subsystems such as the Mbone, reliable multicast protocols, or integrated services as well as for larger systems such as the global Internet.

• *Development and deployment in the infrastructure is of necessity incremental.*

Unlike the development and deployment of applications, which can sometimes spread like wildfire, the development and deployment of the infrastructure is of necessity incremental. This is due to both the scale of the infrastructure and to the autonomy of individual components. This restriction to incremental deployment is problematic for proposed mechanisms (e.g., some scheduling algorithms, routing protocols, and QoS mechanisms) that will only be of use when they are ubiquitously deployed throughout the Internet.

Although there have been 'flag days' in the Internet in the past[1], it seems unlikely that this can happen again. The 1983 flag day in particular was possible because at the time there was some centralized control of the ARPAnet by a single agency (i.e., DARPA).

• *Explicit examined assumptions are better than implicit unexamined ones.*

We all make implicit assumptions in our research, and such implicit assumptions are often wrong. For example, because my intuition is partly shaped by my experiences with my simulator, which until recently had a rather simple view of the world, I tend to implicitly assume that TCP connections follow fixed symmetric paths on point-to-point links, with data flowing only from the `sender' to the `receiver', and not vice-versa. The `real world' is of course much more complicated, with dynamic and asymmetric routing, multiaccess wireless links, and two-way data traffic within a single TCP connection. Implicit assumptions about things such as the nature of traffic, the level of congestion, and the requirements of applications are in general highly problematic, and are best made as explicitly as possible.

• *Changes in the Internet can be unanticipated and uncontrolled.*

The history of the Internet has been one of unanticipated change. This includes changes in applications (e.g., the web, realtime traffic) and in the infrastructure (e.g., past changes in routing, current changes to incorporate multicast, integrated services, and IPv6). This can be particularly frustrating when the Internet changes, while you are in the middle of a research project, in ways that threaten to invalidate your work. As an example, work on integrated services can be difficult when ongoing changes in the traffic mix (e.g., the emergence of the web, realtime applications with congestion control and adaptive playback times, etc.) and pricing structures make it difficult to find firm ground on which to stand. We believe that this period of change is not a temporary phenomena, but will continue to characterize the future Internet.

[1] There was a `flag day' in the ARPAnet for the transition of the ARPANET host protocol from NCP to TCP/IP in January 1, 1983 [Leiner97].

• *Some research problems have their own natural time scales.*

Some fundamental research problems are remote from concerns of actual implementations, but many other research problems have natural time scales related to the problem under consideration. Theses time scales might be the time scales on which decisions will be made either in the marketplace or in the IETF, the time scales after which conditions will have changed enough to require reformulating the problem, etc. Again, this is discussed further in the section on TCP.

• *The Internet architecture and scale make requirements for global consistency problematic.*

A far-flung and complex system that requires consistency or synchronization among all the components is problematic. This is particularly true given the scale and heterogeneity of administrative control in the current Internet.

4 Some implications

These assumptions have several implications about formulating research problems. One implication is that starting from scratch and designing a new self-contained system is not always the most effective approach. In many cases it is more productive to take a piece of a larger problem, and to try to understand it and its relationship to the larger problem. These might be exactly the problems that need attention. And in many cases progress can be made in understanding a problem even if a complete answer is likely to be elusive and to require the combined effort of the larger research community.

One way to design for incremental evolution is to start with IP and the current Internet. IP lends itself to the addition of rich new functionality and to the process of evolving to meet new challenges. New transport protocols and applications are added, routers implement new queueing and scheduling mechanisms, and new underlying technologies such as ATM or wireless are used. Other incremental changes either planned or already in progress include the incremental deployment of IPv6, IP multicast, and integrated services. While it often looks attractive to design a new network architecture from scratch to accomplish a specific purpose (e.g., OSI, X.25, ATM, active networks), our inclinations and experience are with deploying new functionality in the IP network.

5 TCP

In the next two sections we consider how these assumptions about the nature of the Internet guide our own formulation of specific research problems. This section examines some of the ways that our assumptions have shaped our framing of research questions on TCP.

Implicit assumptions:

Early research on TCP dynamics (including our own) focused on behavior with a small number of competing, unlimited-demand TCP connections in one-way networks (e.g., the data traffic going in one direction, and the acks going in the other), in a simple topology with static routing over point-to-point links. This is a fine starting point for understanding traffic dynamics in the most simple scenarios, and does not necessarily reflect assumptions that this is a faithful model of real-world environments. However, it has also been necessary to acknowledge more complex environments, including short Web connections, two-way traffic, rich topologies, and dynamic routing. Understanding traffic dynamics in large complex topologies with a rich traffic mix remains a difficult and unsolved problem.

Robustness and heterogeneity:

A second area of TCP research is the addition of new functionality or modified congestion control algorithms to TCP. A central requirement is to maintain TCP's robustness in the presence of packet losses and a wide range of offered load, link speeds, packet sizes, and congestion levels. Thus, while one goal is to evolve TCP to more effectively use the available bandwidth in a range of scenarios (e.g., short web transfers, large high-bandwidth file transfers, TCP over satellite and other wireless links, TCP over ATM, TCP over integrated services, etc.) this has to be done without sacrificing effectiveness in other environments, and without affecting the robustness of the congestion control algorithms for the network as a whole. This robustness is far more important than full efficiency in using all of the available bandwidth.

Unanticipated change and time scales of research problems:

Because TCP is a changing target, research problems on TCP also have their own natural time scales. As an example, investigations of the unnecessary retransmit timeouts of Reno TCP (1990) were of significant interest a few years ago, but should become less relevant in a few years when most TCP implementations have begun to incorporate the Selective Acknowledgement option (SACK). The problem of unnecessary retransmit timeouts will be even less relevant when most routers have implemented intelligent queue management (e.g., RED) that avoids many of the bursts of packet drops characteristic of Drop-Tail gateways. Thus, tailoring the design of a future protocol (e.g., some form of ATM rate control) to address the problems of current Reno TCP implementations with unnecessary retransmit timeouts is not helpful; such problems of TCP are somewhat transient concerns, and the first-order fix, already in progress, is to fix the problem in TCP itself.

6 Integrated services

In this section we consider as an example the particular research area of providing qualities of service in the Internet. One possibility would be to start with a set of axioms, and to derive an architecture from that set of axioms. An alternate and more fruitful approach, in my judgement, would be to start instead with IP, the IP architecture, and the real world of the Internet. (Or, to rephrase, the only axiom is to start with IP and the real world, characterized as it is by heterogeneity, change, and the restrictions of incremental deployment.)

Heterogeneity and implicit assumptions:

One implication of our assumptions is that we can't make strong assumptions about the likely traffic mix for "realtime" traffic. (We use the phrase "realtime" traffic loosely to refer to traffic that needs per-flow quality of service guarantees.) The "realtime" traffic will not necessarily be mostly audio and video, and the fraction of the traffic that is video will not necessarily be mostly MPEG video. Some of the traffic will be unicast and some will be multicast. Some of the traffic will be using end-to-end congestion control and some won't. Some of the traffic will be extremely amenable to aggregation (because of a Poisson packet arrival process or a small peak-to-mean ratio for ON/OFF traffic) and some will not. Some of the realtime applications might want jitter control or strict delay guarantees from the network, but most will not. Any assumptions about the traffic should be as explicit as possible. And any such assumptions are likely to restrict the resulting design in some fashion.

One process of defining the research problems, based on these assumptions, has proceeded as follows: because the traffic cannot easily be characterized, measurement-based admissions control has strong advantages. One approach that makes minimal assumptions about traffic characteristics is to assume that flows applying for admissions will characterize themselves simply by a policed peak rate.

This realtime traffic is also not likely to become the dominant traffic in the Internet. The current Internet is dominated by best-effort TCP traffic that does not require strict bounds on per-packet delay or on packet drop rates. We would conjecture, based in part on the heterogeneity of Internet traffic, that there will continue to be a large fraction of the overall Internet traffic that does not require either per-flow QoS guarantees or aggregated differential services. In this case, the admissions control procedure does not have to take on the difficult job of trying to use 90% of the link bandwidth for admitted traffic while avoiding packet drops or large delays. When only a moderate fraction of the link bandwidth is used by realtime traffic, the admissions control procedure has much more room to maneuver.

Thus, in the case of integrated services, taking into account the assumptions of heterogeneity and the requirements for robustness does not necessarily make the research problem less tractable; in some respects it makes the research problem more tractable, by avoiding altogether the goal of high bandwidth utilization by a large body of carefully-characterized realtime traffic.

Incremental design:

The process of incremental design, taking into account lessons learned from actual deployment, can be illustrated by the ongoing design process for the Mbone tools. Audio/video/whiteboard tools were first deployed in limited environments, even though the tools are not fully-formed, with some major pieces of the functionality (e.g., the congestion control mechanisms needed for more widespread deployment for these tools to peacefully co-exist as best-effort traffic) still in the research stages. Gaining experience with the limited deployment of the Mbone tools has been one preparation for developing the tools for the next stage.

Similarly, the approach of the Internet research community is that an essential next step is to gain experience with a limited deployment of the current framework for integrated services (i.e., RSVP and the guaranteed service and controlled load service templates). This operational experience should be critical in evaluating the current framework, and in keeping track of the evolving needs and characteristics of applications.

7 Issues with incremental deployment

How should we let the realities of incremental deployment impact our research?

The deployment of multicast routing in the Mbone and of IPv6 [Hinden96] in the 6bone illustrate strategies of incremental deployment of new functionality in the infrastructure. Integrated services and RSVP are examples of projects facing the challenges of incremental deployment. This non-ubiquity and incremental deployment of "advanced features" such as multicast, active scheduling and queue management, IPv6, and so on, poses a real chicken-and-egg problem, since mainstream application developers generally steer clear of advanced features in the infrastructure until some critical mass has been achieved.

One problem with proposals that depend on the ubiquitous deployment of a new mechanism is that there is no viable path from getting from here to there; that is, no path for transitioning from the current Internet to the world of ubiquitous deployment of the new proposed mechanisms. Fortunately, many new functionalities don't need to be implemented *everywhere* to give value. One path of incremental deployment would be for such functionality to first be deployed in intranets or other controlled environments, resulting in *pools* of feature-rich nets in the Internet. This is almost the opposite of the Mbone diffusion model, where the Mbone is layered *on top of* the Internet by Mbone-capable routers connected by tunnels.

Of course, an acknowledgement of the reality of incremental deployment in the intrastructure should not cut off research on proposals that would require flag days or entirely new networks for deployment. A good discussion of concepts, simulations, and experiments can be useful even if it does not lead to eventual implementation.

The Internet intrastructure also does not necessarily evolve by smooth increments, but also by fits and starts, or by punctuated equilibrium. The necessity for having a viable path of evolution of the global internet from the present reality to the envisioned future can be viewed not only as an unfortunate constraint preventing certain discontinuous jumps, but also as a source of power, enabling an envisioned future internet to be successful because it builds upon the present reality [Dawkins96, Kauffman95].

In our own research on adaptive web caching we have taken a two-pronged approach, thinking both about the possible state of a global data dissemination architecture of the future, and also about incremental changes that can be make in the current web-caching architecture to advance in that direction.

8 Conclusions

There are plenty of hard problems concerned with understanding and engineering a large, decentralized, complex system such as the Internet. Working on the incremental deployment of new functionality in a changing global information intrastructure does not imply that the work must be mundane or lacking in fundamental intellectual challenges. There is a wealth of fundamental intellectual challenges involved in understanding and managing large complex systems such as the Internet.

One approach to choosing a research topic is to consider potential problems that have not yet manifested in the Internet. For example, in areas such as routing and traffic dynamics there could be emergent behaviors of synchronization or self-organization that only manifest with a certain size of the Internet or a certain level of multiplexing. As the Internet grows in scale, the development is not likely to be a simple extrapolation from the current dynamics.

Although there can be considerable difficulty and frustration involved with working in a constantly-changing environment such as the Internet, there is also a great opportunity for creativity and power that comes from working in an environment of other newly-emerging applications and network services.

9 Acknowledgements

The note has resulted from conversations with a number of people, including Deborah Estrin and Greg Minshall (who has in fact contributed some of the text). This note has been heavily influenced by the ideas and architectural approaches of Dave Clark and Van Jacobson. Thanks also to Jon Crowcroft, Craig Partridge, Vern Paxson, Scott Shenker, Lixia Zhang, and other members and participants of the End-to-end Research Group for feedback. The note began in response to a proposal from K. K. Ramakrishnan. We emphasize that none of these contributors can be assumed to agree with all of the contents of the note.

This note is not a finished product, but is one stage of a conversation that hopefully will result in further development by other contributors.

References

[ADVICE] Leone, M., Web page on ”Collected Advice on Research and Writing”, http://www.cs.cmu.edu/afs/cs.cmu.edu/user/mleone/web/how-to.html.

[Clark88] Clark, D.D., The Design Philosophy of the DARPA Internet Protocols, SIGCOMM '88, pp. 106-114.

[Dawkins96] Richard Dawkins, Climbing Mount Improbable, W.W. Norton & Company, 1996.

[Hinden96] Robert Hinden, IP Next Generation Overview, CACM, June 1996, V.39 N.6.

[Kauffman95] Kauffman, S., At Home in the Universe: The Search for Laws of Self-Organization and Complexity. Oxford Univ Press, October 1995.

[Leiner97] Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, Stephen Wolff, A Brief History of the Internet, http://www.isoc.org/internet-history/.

[Lampson84] Lampson, B., Hints for Computer System Design, IEEE Software, January 1984, pp. 11-28.

[Shenker96] S. Shenker, D. Clark, D. Estrin, S. Herzog, Pricing in Computer Networks: Reshaping the Research Agenda, SIGCOMM CCR, V. 26 N. 2, April 1996.