

Workshop Report

Network Research: Exploration of Dimensions and Scope

August 25, 2003
Karlsruhe, Germany

Karen R. Sollins
Massachusetts Institute of Technology
sollins@csail.mit.edu

Introduction

The **Network Research: Exploration of Dimensions and Scope** was intended to be a next step beyond the National Academy Press report **Looking Over the Fence: A Neighbor's View of Network Research** (National Academy Press, 2001). Further, it was organized in a context in which a number of the funding agencies in the United States were also funding their own workshops and reports intended to explore the future of networking research, as seen from the perspective of each of those agencies. Among these agencies are the DoD Advanced Research Projects Agency, the Department of Energy, and the National Science Foundation. The intention of this workshop was to separate the discussion from a particular agency. It was also understood that this would be a one-day workshop, under the auspices of SIGCOMM on August 25, 2003, at Karlsruhe, Germany. This had two implications. First, clearly that was likely to have an influence on participation, although some effort was made to include participants who do not normally attend SIGCOMM and whose research fields are not typically central to SIGCOMM kinds of topics. Second, in one day we could not expect to make significant progress with consensus or conclusions, but rather encourage discussion, trying to get ideas out onto the table.

The workshop solicited 5 page position papers initially, for two reasons. The first was to raise ideas and topics that the organizing committee might not otherwise have recognized. The second was as a start on who might be invited to the workshop. Of the papers submitted, 6 were accepted and distributed to the participants of the workshops. Attendance included authors from these 6 as well as several others of the papers,

and additional participants to broaden representation from the community.

In order to generate discussion, the committee identified five questions:

1. Do we have a shared meaning for "network research"?
2. Where is the science in network research?
3. Where is the research beyond the current tipping point?
4. How do we value and evaluate research? How does/should our field evolve?
5. Where do we go from here?

For each question except the last, we invited a speaker to raise some issues briefly (10-15 min) and a respondent who was given 5-10 min. At that point the session was opened up to general discussion, chaired by one or another of the committee members. Notes were taken by 3 student scribes. The final topic was only a discussion. The intention was to expose questions and concerns; this report is a summary of those that arose during the day. It makes no claim to completeness or conclusions.

The workshop committee was Mark Allman (ICIR), Balaji Prabhakar (Stanford), Stefan Savage (University of California San Diego), and myself, Karen Sollins (MIT) as chair. The scribes were Steve Bauer (MIT), Mayank Sharma (Stanford), and Renata Teixeira (University of California San Diego). Although the initial speakers in each session are identified below, the participants are not because many of the points were part of the larger discussion and thus attribution is impossible. Without the invaluable contributions of the other committee members, scribes and participants the workshop and this report would not have been possible.

It is our intention to create a web page from the SIGCOMM site on this workshop.

Question 1: Do we have a shared meaning for “network research”?

The initial speaker Dave Cheriton and respondent Nick McKeown presented clearly different opinions on the role and value of network research, especially academic research. Cheriton made his case based on a vision of the history of networking consisting of a series of technology developments and the impact that the “research” community did not have on that history. From this he concluded that the work that is acknowledged as research by our community is evaluated more on innovation than industrial impact. One conclusion to draw from those statements is the position that network research should have reasonably direct impact on industry. Cheriton also is a strong proponent of the position that the over-riding challenging problem for network researchers is scaling. Cheriton would use this as a driving criterion for evaluating efforts in networking research.

In contrast, McKeown proposed that network research is more axiomatic in its basis. Hence he suggested that valuable network research proceeds by leaps with radical ideas that challenge previous ideas, but also progresses linearly in the absence of such non-linear transitions. This led McKeown also to consider the production of such radical ideas. One problem is that they are difficult to predict and therefore argues for supporting wide diversity, in order to learn through experimentation with new ideas. Experience also suggests that such ideas are more likely to come from younger researchers than older ones, as in many other fields of research. As an aside, and a comment to which we will return, it was noted that often younger researchers are the most critical of both themselves and others, often making it more difficult to include a diversity of ideas in peer reviewing situations (both in terms of funding review and publications review).

These two viewpoints opened the discussion to a broad cross section of opinions and viewpoints. For example, one aspect of Cheriton’s position was a valuing of research utility on the basis of direct and long-term impact on industry. One of the issues that gets lost in such a metric is the more ephemeral but possibly quite significant impact on thinking that in turn may lead to yet

other ideas. For example, a question arose over the value of Ethernet, in particular whether the value of it was in the idea of CSMA/CD or in the long-lived preserved interfaces. An example such as this highlights some of the breadth of differences of opinion in the room, and different ways of looking at the question of what networking research is.

Another question that arose was the present and future role of mathematics in network research. Although a position such as Cheriton’s did not encompass increased rigor in network research, there was less disagreement about the value of increasing the role of mathematics. Opinions on this topic ranged along a spectrum represented by two extremes. One extreme holds that some form of mathematical expression of the phenomena we see possibly or probably requires new mathematical theory is the central problem, leading to a call for a theory comparable to information theory or thermodynamics for network complexity. A midpoint position was the opinion that we are currently making progress on developing mathematical models using currently existing mathematical techniques. At the other extreme was the position that there is significant value, often lost on students, of non-mathematical research. This last position was reflected in the comment that often the most challenging problems we face in networking are the ones we don’t yet know how to express mathematically. This was accompanied by a concern that we often teach our students to undervalue this aspect of network research in favor of problems that can be expressed mathematically. Another position related to this was the concern that, to the extent we focus on modeling and explaining current phenomena mathematically we may be losing sight of the fact that the current approaches may reflect at best imperfect engineering solutions, rather than the intrinsic complexity of networking.

Closely related to the question of the relationship between mathematics and networking research was the question of to which fields we might compare networking research. One participant laid out three possibilities of the nature of the field:

1. A performance discipline, solving the problem of making “the network” increasingly faster or more efficient in some other way, followed by a clean-up

- activity that involves optimization, often by means of mathematics.
2. An infrastructure field in support of applications that is embarrassed to give credit to the fact that the applications' arena is where the most challenging problems are currently arising.
 3. The work that funding agencies will fund, with the clear implication of who might be driving the definition in this case.

An alternative to this was an extension of the discussion about mathematics above, in which one of the participants suggested that networking research is much more like economics than science. Science is about discovering underlying principles and rules, whereas in economics and perhaps networking research, as man-made phenomena, we can change the rules in our models and explore real, alternative possibilities. In the case of networking, this can be done by changing the actual mechanisms and the bases on which they operate, as well as doing this in our theoretical models.

One of the concluding comments in this discussion was that as with our field generally, we should allow for a variety of definitions of what we mean by network research and a rich mixture that integrates more theoretical aspects such as proofs of correctness, viability, or models with implementation and engineering. But more than that both in this abstract sense of the definition of network research, but also in our more specific thinking there should be a sense of cooperation rather than competition. One participant urged us to distinguish between styles or methodologies and actual topics. We returned to this question of topics later in the day.

Question 2: Where is the science in network research?

In this session Walter Willinger provided the initial talk, with Tony Ephrmedes as the respondent. Willinger questioned the "science" in networking research in several ways. One significant concern is with the application of modeling and evaluation as it is currently practiced in the networking research community. Willinger does not believe there is "science" in such efforts as traffic modeling, topology modeling, performance evaluation, network

simulation, protocol design, or network architecture. In particular, he pointed out that the majority of the curve fitting sorts of activities are not interesting because there is no possibility of failure; one can always fit a curve to a set of points, and without rigorous validation, such an activity is not interesting. The problem as Willinger sees it is that the application of the technique is more or less blind. He sees the same story in topology modeling, although the theory applied is graph theory in this case. With respect to the more design-oriented aspects of our field, such as protocol design, again, since we do not understand optimality, there is little scientific about protocol or more broadly architecture design. As part of this line of argument, Willinger addressed the question of the relationships among networking research, math/physics/statistics, and other related fields. It is his opinion that those fields have little to contribute to ours, but, if we can get it right, we can make contributions to theirs, at a minimum by means new interesting examples to challenge their tool sets.

Willinger then asked whether there is value in including "science" in networking research. He does not have a clear answer, but finds a contradiction in examples. The design of TCP was reasonably unscientific, but after the fact we can demonstrate that it is approximately optimal for what it was designed to do. This would suggest that in the business of protocol design perhaps "science" is not needed to do well. On the other hand, if one considers BGP, as a community we have no idea where it stands in relation to optimality. To this, one of the other participants suggested that the optimality of TCP derives from extensive study and design of TCP's responses to network dynamics, whereas nothing of that sort has been applied to BGP.

Willinger concluded with two significant points. First, as a research community we should not only be fitting models to measured data, but should provide or include an understanding of complex network systems. Without such an understanding, neither validation nor extension is readily possible. Furthermore, there is a need for a systematic and thorough model validation process. Without formalizing and systematizing this so that it can be trusted, as one of the tools of our field of research, the field lacks a form of rigor that ought to exist.

In responding, Ephrimedes took a different approach. As he explained, his background is in control theory. It is clear that coding, control and information theory have not been part of major contributions in networking research, but there are beginnings and significant potential there. There are now beginning to be significant contributions in network coding from Medard and others. Other areas include capacity regions and physics, such as improved understanding of network phenomena that includes a model of the physical layer. In contrast, he warned that we should be careful about including models from biocomputing and molecular biology. That may be too far afield to be usefully applicable. From his perspective, one of the key components of networking research and science is curiosity. Intellectual curiosity should be the significant driver.

These positions were representative of the nature of the discussion in this session, exploring three distinctive components of the field of networking research: measurement, formalization and validation, and the design process. Neither of the speakers discussed measurement to any significant extent, but the participants raised the issue. The issues of measurement fall into two categories, first, measurement and experimentation directed at validation or testing of particular hypotheses, and, second, time series or long-range measurement with archiving. For both styles of measurement, but especially for the second, there was a call for more effective instrumentation of the network. There was recognition that CAIDA is attempting to provide long-range archiving, but that it cannot be expected to do it alone. There was clear recognition within the group that measurements are often either impossible for legal and commercial policy reasons, or “cleansed” in order to provide privacy in such a way that relationships among the data are lost, thus lowering the long-term value of the data. As will be discussed further below, there was a call for data to be made available much more broadly, in order to allow for repetition and revalidation of published scientific conclusions.

The presentations by Willinger and Ephrimedes focused to a large extent on issues surrounding the application of formalisms to networking as part of making it more scientific. Although Willinger suggested that “curve-fitting” does not make networking research scientific, the opinion was expressed that “curve-fitting” can and

should be part of the process. There was deep concern that much of the modeling and simulation work that comes from academic researchers is written off by industry because it is poorly grounded. There is little work done to validate models and little or no work has been done to address the problem that we do not understand the effects of scaling, moving from a small simulation to simulations or conclusions about much larger scale situations. There were suggestions that there is a discontinuity or at least lack of understanding in moving from small scale to large scale. Willinger and others agreed that there is a need for rigorous “model verification process” as well as following through and including as part of a model an understanding of why the model is appropriate. It is clear that there are formalisms that can valuably be applied to particular aspects of networking, many of which are only in early stages. On this latter point, the community needs to make an effort to explain not only how a formalism can fit the data, but also how it is part of a better understanding of the phenomena being observed in the measurements.

As highlighted by Willinger, it is not clear that there is “science” in the design of networks and the particular protocols that comprise a network system. It is also not clear that there should be. One of the participants suggested that there might be “pockets” of science in networking research, but not overall. For example several people expressed the opinion that science is often driven by engineering questions. Science may provide bounds on engineering problems or possible solutions where existing models and understanding are inadequate. Another participant called for basic theory, the core of networking research. As evidenced by the breadth of different opinions, it is clear that there was no unanimity among the group about what the core of networking research is or should be. One interesting characteristic discussed is the fact that in networking, a researcher can imagine something and simply program it, while in physical sciences the researcher is limited to phenomena in the real world. (Hence, the physicists are led to arguments over whether string theory is physics or philosophy as long as it remains unobservable.) Although there was some agreement that our work needs to be based on some intrinsic principles or invariants, we are left with questions of identifying a small number of elemental ones.

Question 3: Where is the research beyond the tipping point?

The initial speaker in this section was John Wroclawski, with Tolga Uzuner responding. Wroclawski explored what he has identified as the tipping point in networking research. Network research began in a period of comfortable funding that allowed curiosity to be the driver. Over time, workable technologies were developed that provided useful functionality. With further support networking, devolved into an often critical role in achieving other goals. Networking took on a social and economic role, that has led, as with tipping points in other fields, to the point at which the economic investment in not changing outweighs the economic incentive to change or provide new services. This brings us to the point at which success has bred a resistance to change, which in turn means that newer technologies will not be accepted, despite the improvements they may bring. This leads to questions such as whether one can design to understand or select the tipping point and whether we should be teaching about the evolution of the process, to explain this tipping point to students. More specifically, we can consider possibilities for responding to the idea of a tipping point in several ways. First, we can try to explain and quantify the effects. Second, we could incorporate the concept into our design principles, by recognizing that there will be pull in several contradictory directions, and design specifically to enable and isolate some of these tussles. Third, we can intentionally design the playing fields for these tussle spaces, so that they will or will not tip at certain points.

Uzuner responded with the position that research is driven by economics. Innovation may occur in either the process itself or specific product innovation. He believes we are at or past the tipping point with networking technology, so that further interesting research and development will be somewhat limited and bounded by economics. The areas in which research can continue to have impact are theory (understanding and perhaps bounding complexity, noting that minimizing complexity may not be optimal), product strategies and yield management, and finance. In considering the different sorts of commercial players in the field,

smaller companies will often benefit the most from research, to which one of the other participants responded that that is often by necessity. Uzuner also proposed that late comers to a technology are often the least likely to succeed; Uzuner was not contradicting Wroclawski, but rather agreeing and suggesting the directions in which research may still have an impact.

One of the participants brought up the term “network externalities”. One of these is economics. One participant suggested that innovation is driven by need rather than economics. Once the solution is “good enough” then innovation stops. Another suggested that economics drives innovation in order to allow for “lock-in”, although this participant also suggested that one needs to include the network architect in this analysis and for this person economics may not be the driver.

Wroclawski pointed out that he was calling for something more significant than network researchers becoming economists, in reaction to the view that economists are generally analysts, modeling existing phenomena. Rather, one of the roles the network researcher can play is as the shaper or molder of evolution, and as such we should do research on ecosystems in order to understand the interconnections. The network is something that is architected, designed and built. Understanding at many levels of abstraction how the network researcher can influence these is important. Wroclawski used an S-shaped curve; a number of the participants found interesting points with respect to these curves. One of the researchers pointed out that the transition points in such a curve are important, especially, the point at which innovation stops and product development becomes dominant. Another way of saying this is that at different points on such a curve different kinds of research may be done. There is some research that explains, other research that expands the possibilities and their benefits and costs. Then there is research that goes the next step beyond where we know how to go at present. For example, there was a point in time when it was understood that routers needed to be speeded up by orders of magnitude. The research on this topic was focused on engineering that speedup. Another participant pointed out that such curves can be seen in many other disciplines as well. It was also suggested that families of such curves allow for an exploration of the “evolution of evolvability”. In

response, one of the other participants suggested that often it will be companies that stay on the existing path, while researchers are more likely to lead the way to a paradigm shift. These paradigm shifts are what move one from one curve to another. Although there was not unanimity on this subject there was further discussion about whether or not “research” should be limited to the curiosity, or early stages of such an S-curve, considering the rest to be something else. There was also a strong point made that there should be funding support for radical ideas.

Question 4: How do we value and evaluate research? How does/should our field evolve?

The initial speaker in this session was Craig Partridge, followed by Steve Wolff responding. Partridge considered the influences on research, especially environmental. One issue is the physical environment, which affects both the sorts of people involved and the roles they play. He considered a set of somewhat different kinds of facilities including: universities, not-for-profit research labs, for-profit labs, government labs (although he had little to say on this topic because it is outside his experience), and subsidized labs. The average cost per person in the academic environment (faculty and students) is about one third that of not-for-profit labs, in part because the faculty member is typically raising only two to three months of salary and graduate students are much less expensive. One of the clear distinctions is that in the not-for-profit the typical researcher is working full time. On the other hand the faculty member is much more of a small entrepreneur, raising money, producing output, mostly through students, and leveraging that to raise more. The difference in cost for researchers in the other sorts of labs is less different from the not-for-profit, although typically, the researcher in a not-for-profit and often in government labs is also raising money as the academic is, although in these cases for full salary.

One of the other clear distinctions in environment that is reflected in the nature of the research is the presence or absence of students

and other kinds of staff. Typically, in a research lab, there are senior researchers with many years experience and a large number of recent PhDs doing the bulk of the research. There are few people in between. In labs, which are increasingly commercial, there are increasing numbers of support and administrative staff. The faculty member does the bulk of the management of funding and projects alone. In a laboratory there is likely to be fiscal, technical, and administrative staff as well as other support for the research operation.

Another difference arises from the sources of funding. There tend to be larger amounts of money for research that is expected to have more direct product results. In addition, there are increasing amounts of money for increasingly classified work. That said, corporate research labs are in deep trouble to the extent they still exist at all. There are three problems. The first is that they often were not doing things useful to the company. Second, often the company does not understand how to take advantage of possibly useful results. Third, these labs often do not know how to stop projects when their usefulness is past.

Wolff addressed questions of how we evaluate research. One distinction is between basic and applied research, which distinguishes based on whether “we” care about ownership of the intellectual property involved; if “we” care, then we can categorize the work as applied. Wolff also noted three distinct scales that may form an evaluation: prestige, funding, and academic peer.

The discussion fell into several major topics: the effects of funding raising requirements, peer and other reviewing, motivations for research, in what ways are we training our graduate students.

Beyond the questions raised by the speakers, several of the participants discussed the influence of needing to raise money in the academic and not-for-profit lab environment. The problem, especially for non-faculty is the need to raise funds continuously, to cover salaries. A faculty member can simply take a break now and then, and teach. In addition, the faculty member gets a sabbatical on a regular basis. The researcher gets no such break, but in exchange can work on projects full time. One of the effects on the researcher is that the need for continuous funding leads to incremental proposals, in order to increase the probability of

success in receiving funding. The only suggested path out of this dilemma is for the researcher to run several projects simultaneously so that they are at different points of advancement, allowing for some degree of exploration at any given time.

A second concern raised by a number of participants was that decision-making, especially with respect to paper selection for the most prestigious conferences (e.g. SIGCOMM), has had a stifling effect at least on research reporting, and possibly on research output more broadly. Two dimensions of this were discussed, the limitation to certain kinds of topics and the limitation to certain styles of papers. One of the younger members of the community expressed a degree of self-denial with respect to research reported in order not to violate “sacred cows”. Another commented there is a certain amount of pressure to publish, which tends to drive at least some of the choice of subject matter for research and size of efforts into publishable units. It is important to notice that comments earlier in the day pointed out that it is often the younger members of the community who are the most critical of others and least tolerant of breadth of ideas and risk. There were some questions about whether this has any relationship to the fact that young faculty cannot afford to take risks themselves in their research, at least until they achieve tenure. This was followed by a discussion of the nature of the political structure that makes value judgments about research, questioning whether or not a democracy can be more effective. Those involved in NSF reviewing pointed out that more reviews do not generally reflect more distinct opinions. With reasonably broad reviewing representation, beyond three or four opinions on a proposal or paper, additional comments generally do not increase the number of distinct opinions.

At several points in the discussion participants brought up questions of what does and what ought to motivate researchers. Clearly, as Partridge and Wolff pointed out, for some researchers the motivators fall into such categories as promotion within one’s organization, peer acceptance (often through publication), success in acquiring funding. Several participants suggested that their research was motivated at least in part by education. Another suggested that the intellectual exercise of the research itself was the motivator. One participant pointed out that even within the

academic community this is dependent on the nature of the university. Researchers at top tier universities are more likely to have significantly more freedom in directing their own research. Those at lower tier schools find many limitations including less funding, heavier teaching loads, students who require less challenging projects and so forth.

Questions about students appeared in a number of the topics above, but one concern was discussed more fully, the question of what students are being taught broadly about the quality of research to which they should aspire. There was a sense that the bulk of the research done by students, at least in the USA, particularly because the larger numbers are not in top tier schools, is weak at best. The group was not clear about cause and effect with respect to this problem, but there was deep concern that by not setting the research standards high enough, students are not taught to set the standards high enough for themselves, a lesson that they will need later in life if they are to become researchers themselves.

Question 5: Where do we go from here?

In this last brief session, there was less discussion and more of simply throwing out ideas. They are reported here with no value judgment or particular ordering.

- As a community we should identify key foundational questions (as the mathematics community does). These may require “tools” (e.g. mathematics and other theory) that do not now exist. One response to this was that at least some of the basic understanding may not be expressible mathematically.
- We should change the model of evaluation, especially for program committees such as SIGCOMM, to make them either more democratic or more populist. One suggestion was to make reviewing not anonymous, but rather signed, allowing for better evaluation of the reviewer. Another was to post submissions publicly, allowing anyone who wanted to comment on them. There was some

- discussion about finding a conference on which to experiment with a very different model of evaluation. Note that earlier discussions explored questions of who might be either conservative or overly critical of others' work.
- There needs to be significantly more participation in the process of evaluating our field. Two additional and related issues were raised. First, small groups are significantly more effective for discussion. Second, more than one day is important, in order to do more than lay out problems as was done to some extent in this workshop. One suggestion was to run several parallel small workshops of a couple of days. One might do some coalescing of results and conclusions.
 - An alternative, less radical, suggestion was to encourage much more breadth and churn in program and other reviewing committees. Committees should include people from both traditional and newer (perhaps more radical) research directions. They should regularly include many more junior faculty. Perhaps there should be instituted maximum terms or number of terms within a longer period of participation on an individual committee.
 - One participant was quite worried about the suggestion that some directions of research are "good" or worthwhile and others "bad" or less worthwhile.
 - As a community we should make a much more significant commitment to cross-disciplinary, high risk, and disruptive ideas.

There seemed to be agreement that the discussions were only preliminary and need to be broadened to include more people and a broader set of people.