



Answering Three Questions About Networking Research

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

Scott Shenker
UC Berkeley and ICSI
shenker@icsi.berkeley.edu

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

Researchers often talk about specific technical trends or research topics. But we rarely talk about how and why we do the research that we do. The process of submitting and reviewing papers puts our ideas through a particular kind of filter that may make all of the research seem like it follows some standard rubric, a SIGCOMM Normal Form if you will. During a panel at HotNets'21, five researchers—Hari Balakrishnan, Jon Crowcroft, Jennifer Rexford, Scott Shenker, and David Tennenhouse—each answered three questions about how they pick their own research topics, what areas they would like to see more research on, and how they evaluate conference papers. Due to the unexpectedly positive response to that panel, CCR will be publishing a series of answers to these three questions, starting with two participants from the panel but reaching out to others to provide answers from a broader cross-section of the SIGCOMM community.

CCS CONCEPTS

• **Networks** → **Network architectures**; • **General and reference** → **Surveys and overviews**;

KEYWORDS

Networking research, evaluating paper submissions

JENNIFER REXFORD

Q1: *How do you choose what to work on and what kind of impact are you hoping your work will have?*

I want to make the Internet infrastructure—the underlying platforms, the plumbing—worthy of the trust society increasingly places in it. That is the driving force that gives me a sense of purpose and animates much of the research in my group. The basic goal is to make future networks more efficient, performant, reliable, secure, cost-effective, and so on. A big part of achieving that long-term goal is to plant the seeds of the network's continual reinvention—to create networks capable of change. The Internet is constantly reaching new users, supporting new applications, and adopting new technologies. The key to making the Internet better is ensuring that the network can change over time. That's why I focus on how to make networks more programmable, and how to make it easier for more people to program the network.

The constant evolution of networking applications and technologies means networking researchers are always struggling to keep up. I quite enjoy the process of getting up to speed on new technologies, and identifying the new challenges and opportunities they create. The "getting up to speed" part is similar to what I enjoy about teaching—searching for a better way to organize and convey

complex material, and hopefully come to understand it more deeply myself along the way. I enjoy talking to practitioners (like network operators) who are struggling to make the new technologies work in the field. They help me learn how the technology really works (or doesn't work). Trying to make their incredibly difficult jobs easier is a big motivation for my work.

To make computer networks worthy of society's trust, we need to put the design and operation of these networks on a stronger foundation. That's why I like to collaborate with researchers in neighboring disciplines, like programming languages, software engineering, data structures and algorithms, and so on. These areas have so much to teach us about how to design and manage networks with greater care and rigor. And, researchers in other fields have a different aesthetic that can help us better understand what it is that we, as networking researchers, bring to the problems we study. Ultimately, I am opportunistic about these collaborations. So many disciplines have something important to teach us, but in the end these interdisciplinary collaborations rely on personal chemistry and mutual interest—on finding fun people who are game to engage!

Q2: *What research topics or approaches do you wish the community did more of?*

Networking is an exciting and ever-changing field full of unsolved practical problems ("nails looking for a hammer"). Often the best solutions come from leveraging and extending techniques ("hammers") from neighboring fields. I love when the community embraces ideas from other fields to make our ideas better. Sometimes that is distributed systems, or control theory, or programming languages, or algorithms, or something else. We see waves of research papers that look at networking problems through a particular disciplinary lens. Once a few researchers get the ball rolling, more researchers follow suit and a larger body of results come to fruition. I wish the community did more research that got those initial balls rolling, by identifying, and wielding, new hammers we don't yet know that we need.

Also, the constant change in the field means that researchers repeatedly need to scale the learning curve on new technologies, applications, use cases, and so on. I wish we did more as a community to help each other get up to speed, through books, survey papers, tutorials, and so on. Several members of our community devote significant attention to systematizing knowledge, through writing textbooks, surveying recent research literature, and educating us about emerging technologies. I wish more researchers did this kind of "heavy lifting", so that we can get better at teaching students and preparing researchers to tackle important challenges.

Beyond writing books or survey papers, I wish more researchers helped blaze a path into new kinds of networks and use cases. For example, my recent interest is in 5G/NextG networks. These networks bring together wireless communication, computer networking, and cloud computing. Very few researchers have mastery of all three of these areas. (I know I don't!) But that makes for a great opportunity! The researchers who are willing to dive in will help shape the area for the rest of us. A key part of exploring a new area is developing "good taste" for what research problems are interesting and important. In uncharted areas, having real data and real operational experience can be hugely valuable. Gaining experience by deploying and running a real network—eating your own dog food—can be hugely valuable. Also, one thing that is exciting about 5G networks is the huge wealth of use cases. This creates an opportunity for researchers to partner with local research collaborators (e.g., working on drones and other unmanned vehicles, augmented/virtual reality, Internet of Things, robotics, etc.) and local and regional partners (e.g., non-profits, companies, etc.). That is, researchers can take "unfair" advantage of the unique resources in their own locales to do research that few others are as well-equipped to do.

In short, I wish the community did more research that leads us into the unknown, by introducing us to new disciplinary "hammers" or uncovering new practical "nails."

Q3: *How do you approach evaluating conference papers?*

I take a long, hard look at the problem formulation. Is the problem well-motivated, and is it clearly explained? I look at how the problem is framed and whether that framing passes the smell test, such that a good solution to the problem would be of broad interest. Then I look at the solution. Does it have a spark of creativity? Does the approach, or the outcome itself, teach me something that could be used to solve a wider range of problems, beyond the current one? The output of the paper should be the knowledge, wisdom, and insight, so (unless the specific problem is super important in its own right) I look for papers to shed light on how to solve a broader class of problems. I confess that I don't care all that deeply about the performance evaluation, except perhaps to understand the evaluation *framework*—whether the parameters and performance metrics are well-considered. I'm much more interested in understanding the problem formulation and the novel ideas underlying the solution.

Saying that I value the "problem formulation" does not necessarily mean that the problem statement needs to be mathematical, with a set of specific parameters. Many networking problems are more qualitative, and that's okay. But, I do think it is valuable to make the problem formulation precise and well-motivated, so we know what problem the paper is tackling, and why. Separating the problem formulation from the problem solution allows the reader (and the reviewers!) to interrogate them separately, as the formulation and the solution should be judged rather differently. The problem formulation should be judged for its "good taste," and the problem solution should be judged for its creativity and effectiveness. Separating the problem formulation from the problem solution also helps other researchers build on the work, by solving the same problem a different way, or adjusting the problem formulation to consider new constraints or opportunities.

In reviewing papers, I wish we did a better job clarifying which criticisms of a paper are "important" and which are not. Reviewers

often point out many weaknesses of the papers they review. We're all far too good at "debugging" each other's ideas! When reviewer provide a long list of criticisms, the authors cannot tell which items were the death blow for the paper, and which were smaller concerns or something readily addressed in revising the paper. More granular review forms, with separate sections asking for the main reasons to accept or reject the paper, can help with this. Still, I wish reviewers were more thoughtful in trying to distinguish one kind of concern from another, to make better paper-reviewing decisions and to provide more useful feedback to the authors.

SCOTT SHENKER

Q1: *How do you choose what to work on and what kind of impact are you hoping your work will have?*

My choice of problems is driven by many factors, including the interests of students and collaborators, my own limited capacity for understanding, and who I talked to that morning. The process is more serendipitous than systematic, so instead of focusing on my problem choices I'll focus my remarks on the impact I hope to have (which, implicitly, strongly shapes what problems I work on).

For that question, the answer is straightforward: I want to change the way people think about a problem. When there is a strong but unfounded conventional wisdom, I want to overturn it. When a topic has no satisfying intellectual framework, I want to bring structure to the discussion. Ultimately, I am seeking conceptual clarity. To be clear, very little of my work actually meets this bar of conceptual clarity, but this is what I aspire to. There are others in the field, such as my colleague Sylvia Ratnasamy, who are masters at this.

In some cases, the work that my collaborators and I do might change what gets deployed, but that is a secondary consideration to me. Why? Because there are a whole set of business and organizational issues that filter what gets deployed, and those issues do not particularly interest me. I primarily want to change the nature of the conversation, how people think about the problem.

I will end this answer by noting that I am not recommending this for others. In fact, I cannot defend my chosen style of impact on any grounds other than this is what gives me joy. This joy can come in grappling with some practically important design issue, or straightening out the reasoning in a paragraph that doesn't quite make sense, or toying with an intriguing but unrealistic mathematical model. Regardless of the size or import of the problem, the search for conceptual clarity is what feels rewarding to me and is what I hope my work brings to the field.

Q2: *What research topics or approaches do you wish the community did more of?*

These remarks apply more to the SIGCOMM conference than to the HotNets Workshop, whose content I find more stimulating. Fundamentally, I wish the networking community, as represented by SIGCOMM, made more room for addressing "bigger" problems.

Here, bigger might mean the Internet at large, by which I mean the public Internet, rather than private datacenters and WANs. Bigger could also mean broader implications, such as network neutrality, digital divide, privacy, climate change, and democratization.

Or bigger could mean more foundational; I don't mean more theoretical, but something that addresses the conceptual foundations of networking, even if it doesn't have immediate practical implications. I would put Internet architecture in this category, my own personal obsession.

I want to add a comment about why these bigger problems are being squeezed out. For a long time SIGCOMM struggled with the fact that our work was often in vain. There is an old adage: "Humans plan, God laughs". Well, for several decades it was: "Researchers publish, Cisco laughs, and the carriers ignore."

This stopped when the hyperscalers came along, first with their datacenters, and then with their private WANs. They could do things without impatiently waiting for vendors to implement designs and without vainly hoping that carriers would deploy them. They had use cases, they had data, and they had deployment! This was such a sweet answer to our years of bitter yearning for impact.

And now datacenters and private WANs play a very important role in the work of our community. There is no doubt that this work has been valuable, but this emphasis has also had an unfortunate side effect.

We have now gotten used to having proven use cases, extensive data, and actual deployment for papers we accept. The need for these characteristics has led us to focus on problems of interest to the hyperscalers where those requirements can be met. However, I would claim that providing incremental improvements to these private networks are not the most important problems facing the research community.

We all know the story about the person looking for their keys under the lamppost and when asked if that is where she left them, she answers "no, but this is where the light is." That is our situation: our most important problems may lie elsewhere, but the problems of hyperscalers are where the light is, and by light I mean "path to publication".

So I would urge that we recognize that we focus too much on where we can meet these requirements and not enough on problems where work is more urgently needed but where we can't meet the requirements for a SIGCOMM normal-form paper, and thus have trouble publishing.

Q3: How do you approach evaluating conference papers?

In reviewing papers, my approach is simple. I first ask whether the results, if correct, are important. That importance can arise from changing the way we think about the problem; for example exploring new possibilities that we had not yet thought about, or overturning the conventional wisdom. Or the importance could arise from the results potentially having significant and generalizable practical impact. That latter requirement of generalizability often gets less attention than it should, in that the work should teach us something we can apply more broadly, not just solve a very specific use-case.

Only if the paper passes that test do I care whether it is correct, and then my standard is not "are there any flaws?" but rather "of course there are flaws, but are the flaws fatal?". If the answer is no, I want to accept.

Thus, I only ask two questions about the paper, and the questions are considered in a specific order. Very little of my evaluation of

the paper involves mastering the details. However, there are three questions I never ask:

"Will this design get deployed?" Asking that question is again only looking under the lamppost in the proverbial story. We must not limit ourselves to designs that the hyperscalers or carriers are going to deploy. If we do, then we are just an advanced engineering team, not a research community.

"Did this paper present a novel mechanism?" Don Norman once said "Academics get paid to be clever, not to be right." We should fight this tendency to focus on cleverness. In particular, we should not be evaluating papers based on whether there is some cleverness or mechanistic novelty in a paper; instead, we should focus on whether the results would, in a way that has not been previously proposed, help us improve or better understand the world. As long as the paper has important results that contribute in one of those ways, I don't care whether the paper has developed a novel mechanism. If people can use existing mechanisms to achieve important results, why is that not worthy of publication?

"Is this science?" Often when a paper does not have enough graphs or equations, a reviewer will criticize it for not being scientific enough. I think my reaction to this complaint gets to the heart of a debate about what our conferences are for. Papers should, of course, present some evidence that their proposed solution works as intended. However, demanding our current degree of rigor in the evaluation is extremely limiting, in two ways. First, it limits us to problems where the relevant metric is easily measurable, which isn't true for many of our most pressing problems. Second, there are many topics where only the hyperscalers or carriers have the ability to characterize the relevant workloads or measure the performance of widely-deployed designs.

While some may think that demanding rigorous evaluations, despite these unfortunate side-effects, is necessary for work to be considered research. I think this is completely wrong. The goal of a research conference is not to hold ideas to such a high bar that only those lying under the publishable lamppost should succeed, it is to provide a venue for publicizing the most interesting ideas which then lets the community pursue them in more depth if warranted. Publishing a paper whose interesting idea ultimately doesn't pan out is far more valuable than publishing an incremental development that does work, but provides little in the way of general insight.

To sum up my rant, the presence of numerous graphs or equations does not mean that the results are important in the sense described above, so we should start by evaluating the importance of the results, assuming they are correct, or else we run the risk of publishing rigorous evaluations of designs that we have little reason to care about.