Never Waste a Mid-Life Crisis: Change for the Better

Jennifer Rexford
Princeton University
jrex@cs.princeton.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 It has NOT been peer reviewed.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

Creating a better Internet—a global communications infrastructure that is more secure, reliable, performant, flexible, and so on—is one of the grand challenges of our time. Yet, making substantive change to such a large, distributed, operational network is inherently difficult. This position paper argues that the networking research community should come together and adopt a sort of "ambitious pragmatism" that tackles the big problems while identifying the practical steps to take along the way. The community can work together to (i) identify and precisely formulate the main problems we need to address, (ii) more deeply understand a diverse array of practical constraints (including business drivers, economic incentives, government policies, and more), and (iii) create new deployment platforms and institutional structures to enable good research ideas to "cross the chasm" to deployment.

KEYWORDS

Internet, network architecture, security, programmability

1 INTRODUCTION

Like SIGCOMM, I turn fifty this year. Middle age is a time to lament how good things used to be, whether or not the past ever resembled what we see in the rear-view mirror. To me, the early days of SIGCOMM—and of the Internet, as the two were close siblings growing up hand-in-hand together—seem like they must have been a time of limitless possibility, with all the wild and crazy ideas on the table, the freedom to try things out, and people working together to conceive of something bigger than any one person could build alone. Kids in a candy shop, yet with an abiding sense of purpose.

Fast forward to now, and it is hard to ignore the many signs that we are past that early period of supposed idealism and zeal. Security vulnerabilities and cyberattacks on, and across, the Internet are a constant struggle, and they are poised to get much worse with the Internet of Things and cyber-physical systems. Nation states, and even companies, use the Internet to compromise human rights, violate user privacy, and attack each other. The Internet so clearly needs fixing. And, yet, the Internet infrastructure is notoriously difficult to change, leaving so many good research ideas to grow dusty on the proverbial bookshelf (or web site!). So many researchers (and prospective students) want to work in other areas—particularly machine learning—that now give off that aura of limitless possibility, instead of computer networking. The list of symptoms goes on and on.

Are our field's best days behind us? Beyond mourning the bygone days of yore, middle age is a time to decide whether to give up or to double down. With the many challenges facing the Internet, it is more important than ever to have computer networking research that is driven by a keen sense of purpose—to leave behind a better Internet for the next generation:

- A more secure Internet, built with more robust software and hardware, and with better ways to defend itself from inevitable attacks.
- A more empowering Internet, where users can make real choices about the services they adopt and how much of their personal information they share.
- A more programmable Internet, better able to deploy new ideas over time.
- A better understood Internet, through stronger intellectual foundations and better ways to teach computer networking.
- A more affordable and usable Internet, to lower the barriers to getting more people and organizations online at reasonable cost.
- A more energy-efficient Internet, to support increasingly global computation and communication while doing our part to prevent global warming.
- And so on...

In fact, many of us in the SIGCOMM community *already* work on these challenges and others like them. But, how can we succeed in going beyond our research ideas to making the Internet better when substantive change is so difficult? How, as a community, can we best do this? How can we succeed in being "the civil engineers of the Internet" [1]? I believe the answer lies in *how* we work on these problems, and what we do with our solutions. We need to have the right combination of *ambition* (to tackle the grand challenges and outline a vision for the right long-term solutions) and *pragmatism* (to identify the right steps to take along the way).

Crystalizing the grand challenges: We spend far too little time as a community thinking precisely about our collective goals. As an opportunistic discipline—driven by constantly changing technologies and applications—we are far from having a "top ten list" of our most important goals for the Internet. Should we be striving to (finally) make the Internet routing system fundamentally more secure [2, 3], or

ensure that future cellular access networks are open to innovation [4], or something else? To the extent we have these goals, they remain informal, like the bullet list above, rather than having sharp technical formulations that account for technology constraints, economic incentives, and more—factors that affect the feasibility of our solutions and our ability to deploy them in practice. Perhaps we should talk meaningfully about these and other goals, so we reach some understanding of most important challenges for our field—driven by the larger goal of having an Internet worthy of the importance society places on it. We can do this through panels, workshops, and position papers (in CCR!), or we could go further and make a more concerted effort—much like the influential *Looking Over the Fence at Networking Research* report of an earlier era [5]. Doing a better job crystalizing the grand challenges can also help in motivating students to tackle them.

Grappling with reality: We should embrace the fact that the big, practical problems facing the Internet are multi-faceted and interdisciplinary [6]. Addressing these challenges effectively requires understanding technology constraints and opportunities, business drivers, economic incentives, government policies, international politics, and more. We clearly cannot do that alone, and any one person can only contribute meaningfully to these challenges in a focused way. We need better ways to educate each other, so we can come together more coherently at the complex intersections of fields, and to build fruitful interdisciplinary collaborations. Coming up to speed is hard, but we can make it easier through better survey papers and books (like the SIGCOMM eBook on *Recent Advances in Networking* [7]), better models for describing the Internet architecture [8], tutorials (like we have at the annual SIGCOMM conference), summer schools (to go deeper into a topic), and open-source software packages, so we can build on each other's expertise and understanding more readily. Participating in practitioner-oriented events (like NANOG or the IETF) can help, too. Researchers in industry, or with strong connections to industry, can play a crucial role here, in educating the community on technology enablers, business drivers, and practical constraints on what kinds of change are possible and useful. In addition, researchers in neighboring fields can help us tackle the grand challenges in networking by making their methods and tools more broadly accessible.

Crossing the chasm to deployment: Fixing the Internet requires much more than designing a clever solution, publishing a paper, and making software available on GitHub. In a field dominated by proprietary commercial equipment, we need to do more to build and support open-source components [9–11], evaluation tools [12, 13], and testbeds [14–18]. In addition to evaluating and deploying our research ideas, these testbeds enable our students to learn how to take their research to the next level by building real systems—an important skill for the next generation of researchers. Plus, we need to identify and support ways for researchers to effect substantive change, whether in doing government service, or running a consortium striving to achieve a larger shared goal for the Internet, or doing service in standards bodies. Academic researchers, in particular, have tremendous freedom to chart a course for how they focus their energy on research and teaching, and to spend sabbaticals contributing to causes larger than themselves. Plus, often academic researchers can play leadership roles in initiatives to effect change across an industry that company employees would have difficulty playing due to their perceived allegiances. Yet, people are naturally pulled in many directions, and we lack the institutional structures to enable and celebrate these kinds of work. Is this something the SIG could explore—whether through awards that acknowledge these kinds of contributions, or specific initiatives that support researchers in these roles—perhaps in collaboration with other organizations like the Internet Research Task Force (IRTF)?

These three issues are inter-related. Grappling with reality can help us crystalize the grand challenges and then prototype and evaluate our solutions as we prepare to take them across the chasm into real deployment.

In closing, we should not, as the hand-wringing title of this editorial worries, waste a mid-life crisis. We are a community that could redouble its efforts to truly make the Internet better, both in its intellectual foundations and in its operational infrastructure, for the betterment of the larger society. I believe we have plenty of good candy left at the old candy shop! However, change is not easy, especially when it comes to fixing the Internet or shifting the culture of a research community. But, as Margaret Mead once said, "Never doubt that a small group of thoughtful, committed citizens can change the world. Indeed, it's the only thing that ever has."

REFERENCES

[1] H. Schulzrinne, "Networking research — a reflection in the middle years," *Computer Communications*, vol. 131, pp. 2–7, October 2018.

[2] FCC Communications Security, Reliability and Interoperability Council, "Secure BGP deployment," March 2013. http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG6_Report_March_%202013.

[3] S. Goldberg, "Why is it taking so long to secure internet routing?," *Communications of the ACM*, vol. 57, October 2014.

[4] L. Peterson, T. Anderson, S. Katti, N. McKeown, G. Parulkar, J. Rexford, M. Satyanarayanan, O. Sunay, and A. Vahdat, "Democratizing the network edge," *ACM SIGCOMM Computer Communication Review*, vol. 49, April 2019.

[5] National Research Council, *Looking Over the Fence at Networks: A Neighbor's View of Networking Research*. The National Academies Press, 2001. https://doi.org/10.17226/10183.

[6] J. Rexford, "The networking philosopher's problem," *ACM SIGCOMM Computer Communications Review*, July 2011.

[7] H. Haddadi and O. Bonaventure, eds., *Recent Advances in Networking*, vol. 1. ACM SIGCOMM eBook, August 2013. http://www.sigcomm.org/content/ebook.

[8] P. Zave and J. Rexford, "The compositional architecture of the internet," *Communications of the ACM*, March 2019.

[9] R. Morris, E. Kohler, J. Jannotti, and M. F. Kaashoek, "The Click Modular Router," in *ACM Symposium on Operating Systems Principles*, SOSP '99, pp. 217–231, 1999.

[10] "NetFPGA." https://netfpga.org/.

[11] "XORP." http://www.xorp.org/.

[12] "ns-3." https://www.nsnam.org/.

[13] "Mininet." http://mininet.org.

[14] "PlanetLab." https://www.planet-lab.org/.

[15] "Emulab." https://www.emulab.net/.

[16] "Orbit." https://www.orbit-lab.org/.

[17] "GENI." https://www.geni.net/.

[18] "CloudLab." https://www.cloudlab.us/.