

Reflections on SIGCOMM’s Fiftieth Anniversary

Bruce Davie
VMware
bdavie@vmware.com

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 It has NOT been peer reviewed.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

ABSTRACT

On the occasion of the fiftieth anniversary of SIGCOMM, this article contains the reflections of a past SIGCOMM chair on the field of networking over the past few decades.

KEYWORDS

Ramblings

1 INTRODUCTION

My introduction to the field of networking began in 1988, when I started work at Bellcore, as a freshly minted Ph.D. graduate from the University of Edinburgh. At that point, most of what I knew about networking was that email took a long time to get from Edinburgh to Melbourne, passing through a computer in London which connected to its peer in Melbourne about every 12 hours over a dialup line.

At Bellcore, what I initially thought we were doing was designing the next generation network for the phone companies¹. Fortunately, my horizons were broadened early on when we entered into a collaborative project with David Clark’s group at MIT and David Farber’s group from U. Penn. This subsequently grew into the Aurora Gigabit Testbed project [10], one of several projects funded by the NSF to push the boundaries of networking in the late 1980s and early 90s.

It was through these collaborations that I was introduced to the SIGCOMM research community, attending my first SIGCOMM conference in 1990. David Clark and Len Kleinrock were the winners of the SIGCOMM lifetime achievement awards that year. I still have some clear memories of that event—especially David Clark’s talk—and I became a regular attendee of the conference from then on.

So I have been around the SIGCOMM community for about 30 of the 50 years that it has existed. To many readers of this note I guess that makes me sound like an old-timer but since there were already established giants in the field when I joined it—some of whom are still active today—I always felt like a relative newcomer. I was honoured by the chance to become chair of SIGCOMM for 4 years from 2009 to 2013. This note contains a few of my thoughts on networking as I look back over my time in the field and as SIGCOMM observes its fiftieth anniversary.

¹Bellcore was the offshoot of Bell Labs that had been formed to meet the needs of the regional telephone companies after they were divested from AT&T.

2 RESEARCH AND INDUSTRY

I think the most impactful talk I gave at SIGCOMM was an Outrageous Opinion session in 2003 entitled “MPLS Considered Helpful”² [1]. It is no exaggeration to say that people still come up to me today and remind me of that talk (generally with the phrase “I’m not bitter”). Without rehashing the talk here, one very telling episode came the day after I presented it, when someone asked me if MPLS [2] was actually deployed anywhere. At this point, it was running in the networks of well over 100 service providers around the globe—something that was known to my industrial peers at Cisco, but to the SIGCOMM community, was invisible.

I mention this story because it illustrates some of the challenges in keeping the research and industrial parts of the networking community well-connected. I have spent most of my career at the intersection of research and industry, and fostering better collaboration between industry and academia was an area that the Executive Committee focused on when I was SIGCOMM chair. Activities such as the industrial demos at the annual conference and the SIGCOMM Networking Systems Award are outcomes of that focus.

I view it as a positive development that today many, if not most, of the hyperscale operators are active in the research community, thus providing a strong connection between some parts of industry and research. The area where I think more work is needed is around enterprise networking, which still seems largely invisible to the SIGCOMM community. I’m not claiming to have a solution to this but I do encourage the SIGCOMM community to keep finding ways to engage with industrial participants beyond the hyperscalers.

3 THE VALUE OF “OLD” IDEAS

If 30+ years in the field makes me a relative old-timer, the upside of this has been seeing various old ideas re-appear periodically. I’ve observed a few responses when this happens, some of which involve debates about who deserves credit. While that may be either entertaining or annoying depending on your position, a more problematic response in my view is to point out that the old idea didn’t work last time and therefore can be safely ignored now.

My favorite example of this is the idea of a logically centralised control plane that is separate from the data plane. As noted by Greenberg *et al.* [6], this idea goes back at least to the telephone network, and was proposed in various contexts for packet switching in the ’90s and 2000s (e.g. [5]). When it appeared as the basis for Software Defined Networks (SDN), I encountered plenty of people who made the case that since it hadn’t been successful in its earlier

²I also believe I have the distinction of having chaired the first Outrageous Opinion session circa 1995.

incarnations, there was no reason to expect any impact this time around. But in fact SDN has had huge impact, ranging from Google’s backbone [8] to Azure [9] to thousands of enterprise data centers [7]. There are a lot of reasons for that, and I’d argue that the key reason an old idea can take off after many decades is that other enabling technologies make it possible. In the case of SDN, the ability to build robust and scalable logically centralised controllers was enabled by, among other things, a few decades of distributed systems research.

My point here is that while there is enormous value in understanding the history of our field, it’s also important to appreciate that the environment is always changing and ideas may become more feasible or relevant over time. Of course the converse is true as well: some ideas that we may accept as settled might need to be re-examined as the environment changes. Indeed, to draw on the same example, the idea that routing algorithms must be fully distributed was considered a settled fact when I was learning networking, was challenged by the 4D paper, and now we see widespread use of centralised approaches in a range of SDN systems.

4 INCREMENTAL DEPLOYMENT

One aspect of networking that has proven critical in my experience over the years is the need for incremental deployment. During the years that I sat on the End to End Research Group [3], Sally Floyd was a frequent defender of incremental deployment [4]. This is not to say that all research should be incremental, but that for research to have impact, there needs to be a path to deployment that doesn’t require a “flag day”.

One interesting example of this from my work was a failed project to build an MPLS router. The idea was that MPLS could be incrementally deployed by enabling MPLS on just a single link at a time, and so you could deploy a pair of line cards—one at each end of a link—that would only ever see MPLS packets. Then as you deployed more MPLS, you would have core MPLS routers that only ever saw MPLS packets, so they could safely have no IP lookup engines at all. Because the MPLS packet forwarding operation was somewhat simpler than IP forwarding, we hoped to be able to make a simpler and less expensive router. But it became apparent that, during the incremental deployment phase, most routers would need a mix of MPLS and IP interfaces. Since there were lots of other costs beyond packet header lookups, it really wasn’t a useful optimisation to have line cards that only understood one type of packet. In the end we had to be able to handle both types of packet on any line card, and just selectively enable MPLS on certain interfaces as the incremental deployment proceeded. So we did succeed in getting MPLS deployed, but we never built an MPLS router without IP forwarding capability³.

All the technologies I’ve worked on that had an impact also had an effective incremental deployment strategy. (I’ve worked on some failed projects that had such a strategy too, so no guarantees.) I think of the Akamai CDN—recent recipient of the networking systems award—as another great example of a technology that succeeded because of its solid incremental deployment strategy.

³In an interesting coda, MPLS forwarding got more complex as time went by, as the general purpose idea of pushing and popping labels up to arbitrary depth led to more use cases. So rather than building simpler and cheaper routers, we actually made the forwarding path more complex. Fortunately the benefits seemed to outweigh the costs.

It’s also worth noting that I don’t view incremental deployment to be synonymous with incremental research. Research can be disruptive or ground-breaking (i.e. non-incremental) and still lend itself to incremental deployment. It’s also important that some research be untethered from the constraints of deployment so that we can envision what might be possible. Ultimately, however, it is deployment that leads to impact.

5 CONCLUSIONS

One aspect of SIGCOMM that I really came to appreciate during my time as SIGCOMM chair was how much more there is to SIGCOMM than an annual conference. As well as being an organisation that is responsible for half a dozen conferences and this publication, it is also a community that continues to shape the world through its work on networking. I look forward to watching the future impact of this community as it moves into its next 50 years.

REFERENCES

- [1] Davie, B. S. *MPLS Considered Helpful*. <https://www.slideshare.net/drbruced/mpls-outrage> August 2003
- [2] Davie, B. S., and Rekhter, Y. *MPLS: Technology and Applications*. Morgan Kaufmann Publishers, 2000.
- [3] End-to-end Research Group (END2END) <https://irtf.org/concluded/end2end>
- [4] Floyd, S., et al. *Internet Research: Comments on Formulating the Problem*. January 21, 1998. <https://www.icir.org/floyd/papers/assumptions.ps>
- [5] Khosravi, H. and Anderson, T, Eds. *Requirements for Separation of IP Control and Forwarding*. RFC 3654, Nov. 2003
- [6] Greenberg, A. G., et al. *A clean slate 4D approach to network control and management*. Computer Communication Review 35(5) Oct. 2005.
- [7] Marko, K. *VMware NSX - network virtualization as VMware’s secret weapon?* Diginomica, Sep. 2018. <https://diginomica.com/vmware-nsx-network-virtualization-as-vmwares-secret-weapon>
- [8] Jain, S. et al. *B4: experience with a globally-deployed software defined WAN*. Computer Communication Review 43(4) Aug. 2013
- [9] Greenberg, A. G. *SDN for the Cloud*. SIGCOMM award keynote, August 2015. <https://conferences.sigcomm.org/sigcomm/2015/pdf/papers/keynote.pdf>
- [10] Clark, David D., et al. *An overview of the AURORA gigabit testbed*. In Proc. IEEE INFOCOM’92: The Conference on Computer Communications. IEEE, 1992.