

---

# The First Fifty Years of ACM SIGCOMM

Network researchers know that packets are not always evenly spaced, they sometimes arrive in bursts. This burstiness is also present in history. Important events sometimes occur almost simultaneously even if there is no direct relationship between them. Fifty years ago, several historical events took place within a period of a few months. In July, after years of efforts, NASA engineers successfully launched Apollo 13 and the Eagle landed on the Moon. The first footsteps of Neil Armstrong and Buzz Aldrin have opened the era of space exploration and maybe another human will land on another planet during this century. One month later, hundreds of thousands of music fans gathered at the Woodstock festival. In parallel, a few computer scientists and engineers started to deploy prototype nodes of the ARPANET network. The first packets that they exchanged have initiated the revolution that brought the Internet and today's connected society. At that time, research in computer networks had already started and several members of this emerging community gathered in an interest group that later became ACM SIGCOMM.

SIGCOMM's 50<sup>th</sup> birthday was celebrated at SIGCOMM'20 in Beijing in August with a special panel. This fiftieth birthday was a good opportunity to look back at the evolution of both the networking field and the SIGCOMM community over half a century. Earlier this year, after a very interesting teleconference with Vint Cerf, I contacted all the former SIGCOMM chairs and CCR editors to share their reflections on the evolution of our community. Many of them wrote an invited editorial. As our community is driven by scientific innovations, I also encouraged former recipients of lifetime SIGCOMM awards, test-of-time awards and best paper awards to also share their vision with an invited editorial. The thirty-three papers published in this special issue address a wide

range of topics. Some provide a very unusual viewpoint.

## SIGCOMM turns 50

Wesley Chu chaired ACM SIGCOMM between 1973 and 1977. In *Recalling the early days (first decade) of SIGCOMM and some thoughts of the future research directions*, he discusses early networking conferences and the evolution of our field. In *It's Not About the Internet*, Lyman Chapin who served as Chair in the 1990s argues that we should collaborate with professionals from other fields such as behavioral psychology, linguistics, sociology, education, history, ethnology, and political science to solve the policy issues that arise in today's Internet.

In *Patience*, Jon Crowcroft, who chaired SIGCOMM during the second half of the 1990s recalls his thirty years of participation in our community. He first recalls three important decisions made by SIGCOMM during his term: (i) avoid pangloss, (ii) embrace globalism and (iii) have values. These decisions are still very relevant today. He then looks back and provides useful guidance for young researchers and reports several anecdotes.

Jennifer Rexford served as SIGCOMM chair between 2003 and 2007. In *Never Waste a Mid-Life Crisis: Change for the Better*, she encourages our networking community to identify a few key problems that require a solution and work together to propose solutions which can eventually be widely deployed.

Mark Crovella, who chaired SIGCOMM after Jennifer Rexford, reflects in *The Skillful Interrogation of the Internet* on the mindset that researchers should have when studying the Internet. He argues that we have only started to apply a scientific approach to study the Internet and makes a parallel with the evolution of the field of biology.

---

In *Reflections on SIGCOMM's Fiftieth Anniversary*, Bruce Davie, who chaired SIGCOMM from 2009 to 2013 discusses the relationships between networking research and the networking industry. He first argues that it is often useful to revisit old ideas. Since our environment continues to change, ideas that were considered as unrealistic one or two decades ago, could be deployed today. He also encourages us to think about the incremental deployment of new ideas. Several past examples show that incrementally deployable ideas had more impact than stronger ideas that were too difficult to deploy.

Srinivasan Keshav ends the reflections of our past Chairs. In *Reflections on being SIG Chair 2013-2017*, he recalls some issues that the SIGCOMM Executive Committee had to tackle during his term and the initiatives that they took. He ends his editorial by discussing how the current boom in robotics, AI, big data and deep learning could affect our community by reducing the number of networking students and networking researchers. He argues that, in the long term, this could be beneficial.

## The evolution of CCR

The first issue of Computer Communication Review was published in 1969. In 2013, Craig Partridge published an informal history of the first forty years of our newsletter. This issue complements this article. In *My view of Computer Communication Review, 1969-1976*, Dave Walden explains how he ensured both the regular publication of CCR issues and the quality of the articles. He also presents some key articles published during his term as CCR Editor.

Craig Partridge was CCR Editor from 1988 to 1991. He describes in *Changing ACM Computer Communication Review (1988-1991)* how Vint Cerf encouraged him to convert CCR into an “entry-level journal”. This conversion was a success and Craig consumed two-thirds of the annual page budget for the January 1989 issue. After that, pa-

pers continued to flow.

In 2004, Christophe Diot started his four-years term as CCR Editor. In *About velocity and dealing with “fake” scientific news*, he shares his insight on some changes he brought to CCR, including the public reviews and the editorials. These are still very relevant today. He ends by suggesting, again, to adopt a continuous submission model for the SIGCOMM conferences and the importance of volunteering for our community.

Before chairing SIGCOMM, Srinivasan Keshav served as CCR Editor during a four-years term. In *Reflections on being CCR Editor 2008-2012* he explains how he replaced the email-based workflow to handle papers with an instance of `hotcrp` that he modified and why he introduced a 6 pages limit for CCR papers.

Finally, Dina Papagiannaki shares in *Looking back at being the editor of CCR (2013-2016)* the initiatives that she introduced during her term as CCR Editor. These include an interview section and regular columns.

## Fifty years of networking

In 1999, SIGCOMM celebrated its thirtieth birthday with a tutorial on *The Technical History of the Internet* organized by Vint Cerf who assembled, as indicated in the tutorial announcement, *a stellar cast of original technical contributors to the history of internetworking, including Paul Baran, Bob Braden, Dave Clark, Danny Cohen, Dave Farber, Sandy Fraser, Van Jacobson, Steve Kent, Peter Kirstein, Len Kleinrock, Larry Landweber, Dave Mills, Craig Partridge, Louis Pouzin, Larry Roberts, Dave Walden, Steve Wolff, and Hubert Zimmermann*. To prepare this tutorial, Vint Cerf and his friends collected historical documents that were gathered by the late Chris Edmondson-Yurkanan. In *‘Capture it While You Can’: Revisiting SIGCOMM 99’s Technical History of the Internet*, Frances Corry and Anna Loup discuss the importance of collecting historical information as computer

---

networking gets older. They are currently planning an exhibition of the material assembled by Chris Edmondson-Yurkanan at a future SIGCOMM conference.

Our second technical editorial, *Five Decades of the ACM Special Interest Group on Data Communications (SIGCOMM): A Bibliometric Perspective* was written by a team lead by Waleed Iqbal. They contacted us earlier this year to propose a bibliometric study of the SIGCOMM publications based on a large dataset that they had collected. As this study gives a broad view of our publications, it was interesting to also include it in this special issue. They analyze our publications by using different metrics and provide an interesting viewpoint on the output of our community measured by our publications. Furthermore, they release their entire dataset to enable other researchers to compute their own metrics.

We then explore the evolution of computer networking research through the lenses of SIGCOMM award winners and authors of best and test-of-time papers. We start with two articles co-authored by Jeffrey Mogul and Christopher Katarjiev. In *Retrospective on "Fragmentation Considered Harmful"*, they first look at the challenges of fragmenting IP packets in 1987. They analyze how things have changed since the publication of their paper with the deployment of Path MTU discovery. In *Retrospective on "Measured Capacity of an Ethernet: Myths and Reality"*, they revisit the measurements described in their 1988 paper. As this 1988 paper focused on Ethernet networks that used CSMA/CD, it is not anymore relevant to today's switched Ethernet where CSMA/CD is disabled. However, some lessons learned are still relevant.

Then, Srinivasan Keshav, in *Reflections on "Analysis and Simulation of a Fair Queueing Algorithm"* by A. Demers, S. Keshav, and S. Shenker, *Proc. ACM SIGCOMM 1989*, recalls how the Weighted Fair Queueing algorithm was designed and eventually published. This paper initiated a huge amount of work on packet schedulers and is

still very relevant today.

Our next witness is Radia Perlman. In *Network Protocol Folklore: Whys and What-ifs*, she looks at the evolution of Ethernet networks, with some anecdotes on the design of Spanning Tree Bridging. She then analyzes the evolution of the network layer protocols, starting from CLNP. Her editorial should be read by today's students who might believe that there is no alternative to IP.

In *Reflections on "A Control-Theoretic Approach to Flow Control"*, Srinivasan Keshav explains how he applied principles from control theory to improve flow control in reliable protocols and the importance of regularly noting his ideas on paper.

Walter Willinger, Murad Taqqu and Daniel Wilson describe in *Lessons from "On the Self-Similar Nature of Ethernet Traffic"* how the networking community has reacted to their SIGCOMM'93 paper. Thanks to a long and precise dataset, this paper demonstrated that the traffic in an Ethernet network could not be modeled using a Poisson process. This editorial provides lots of insight on this important discovery and the impact that it had on the field. It should definitely be on the reading list of any networking researcher who analyzes measurement data.

Matt Mathis and Jamshid Mahdavi take a step back at the modeling of TCP congestion control algorithms in *Deprecating The TCP Macroscopic Model*. Their original paper proposed mathematical models to describe the evolution of TCP's congestion control scheme. They argue for a deprecation of their mathematical model and encourage researchers to analyze and model the recently proposed BBR congestion control scheme that takes a very different approach to solve the congestion problem.

In *Sizing Router Buffers (Redux)*, Nick McKeown, Guido Appenzeller and Isaac Keslassy also discuss Internet congestion, but from the viewpoint of the router's buffers. They discuss how the recommendations for the sizing of router buffers have

---

changed during the last fifteen years.

In *Retrospective on “A Delay-Tolerant Network Architecture for Challenged Inter-nets”*, Kevin Fall explains the evolution of Delay Tolerant Networking during the last fifteen years and key results in this field.

Retransmissions remain a standard technique to recover from transmission errors in many protocols. Two of the editorials in this special issue describe the evolution of alternatives to retransmissions. In *XORs in the past and future*, Muriel Médard, Sachin Katti, Dina Katabi, Wenjun Hu, Hariharan Rahul and Jon Crowcroft describe the evolution of network coding. They discuss both the theoretical aspects and the deployment of protocols that include network coding techniques. In *A Digital Fountain Retrospective*, John Byers, Michael Luby and Michael Mitzenmacher explain how fast and practical erasure codes have enabled scalable reliable multicast services.

In the late nineties, researchers proposed *active networks* as a solution to bring innovation back in networks that were already considered as being ossified. In *Retrospective on “Towards an Active Network Architecture”*, David Wetherall and David Tennenhouse take a step back at network programmability and compare today’s approaches to earlier ones. Other researchers looked at the control plane protocols. In *Reflections on a clean slate 4D approach to network control and management*, David Maltz Geoffrey Xie, Albert Greenberg, Jennifer Rexford, Hui Zhang and Jibin Zhan discuss how their idea of centralizing network control has contributed to the development of Software Defined Networking (SDN). Another viewpoint of the evolution of SDN is provided by Martin Casado, Nick McKeown and Scott Shenker in *From Ethane to SDN and Beyond*. They explain how their design of a simple solution to manage their campus network demonstrated the feasibility of SDN and highlight important lessons.

In *Lessons from “A first-principles approach to understanding the Internet’s router-level topology”*, Walter Willinger,

David Alderson and John Doyle remind us that the Internet has a highly engineered infrastructure and that it is important to understand its component technologies before analyzing topology measurements. They highlight the limitations of scale-free models that some researchers have tried to apply to the Internet.

In *Don’t Mind the Gap: Bridging Network-wide Objectives and Device-level Configurations*, Ryan Beckett, Ratul Mahajan, Todd Millstein, Jitendra Padhye and David Walker discuss the evolution of Propane, a high-level language and compiler which can be used by network operators to generate low-level router configurations that meet network-wide routing objectives. They highlight the importance of working with small models when working with complex systems.

In *Perspective: White Space Networking with Wi-Fi like Connectivity*, Ranveer Chandra and Thomas Moscibroda explore the utilization of the radio spectrum that was reserved for television signal to provide Internet access, notably in rural areas. This TV spectrum has excellent propagation characteristics. They proposed to use it to provide Internet access in their SIGCOMM’09 paper. They discuss how the technology has evolved and the steps required to convert a research prototype in a widely deployed technology.

In *Perspective: Eliminating Channel Feedback in Next-Generation Cellular Networks*, Deepak Vasisht, Swarun Kumar, Hariharan Rahul and Dina Katabi discuss the impact of their SIGCOMM’16 paper that enabled cellular base stations to estimate the downlink channels without any user feedback.

Finally, Costin Raiciu and Gianni Antichi discuss in *NDP: rethinking datacenter networks and stacks, two years after* their attempts to convince industry to deploy the ideas proposed in their SIGCOMM’17 paper in datacenters. It reminds us that there are many hurdles in the journey to convert a research prototype into a widely deployed solution. This is the difference between re-

---

search and innovation. Convincing industry of the validity of research results often requires more time and effort than convincing one of our very selective TPCs...

I hope that you will enjoy reading this new issue and welcome comments and

suggestions by email at `ccr-editor at sigcomm.org`.

**Olivier Bonaventure**  
*CCR Editor*