

About velocity and dealing with “fake” scientific news

Christophe Diot
christophe.diot@gmail.com

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 It has NOT been peer reviewed.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

When I embarked on the CCR adventure 15 years ago I did not expect it to be so exciting, fruitful, and life changing. I am describing here our motivation and approach, our successes and failures, with I hope a perceptible sense of humor.

CCS CONCEPTS

• **General and reference;**

KEYWORDS

CCR and SIGCOMM history, Community services.

1 INTRODUCTION

As information gets to us faster than ever (but most of the time unverified), it is one’s duty as editor of a technical newsletter and member of the scientific community to shorten the time between submission and publication of technical contributions, while guaranteeing that the material published has scientific integrity and might be of interest to the community. This has been my driver for the 4 years I served as CCR editor, and it keeps motivating me today. Peer reviewed conferences generally do a good job at filtering flawed or already published work; some of them being a bit too zealous at rejecting slightly “imperfect” work. In such conferences, having to wait one year to get a second chance to have a paper accepted is definitely too long at the current pace of information circulation, and authors take the risk of similar work to be published in another venue before they get a chance to resubmit (very likely indeed, given how fast epidemic dissemination happens these days). Similarly, I have seen technically incorrect papers accepted after multiple submissions, that were referenced for years because it took that long to publish the technically sound results invalidating the original paper. Journals are more consistent at managing technical robustness/innovation but one probable reason for their lack of popularity in our community is the delay between first submission and eventual publication due to multiple rounds of modifications. I remember having to deal with papers that had remained under review for more than 2 years (in both CCR and ToN). I have had to ask authors to withdraw their submission as there was no way the paper could be accepted given the progress made in the area since they first submitted. Of course, such long reviewing delays should never happen in the first place, and I am proud that our editorial staff was able to meet a maximum 2 months response time during my CCR tenure.

Christophe is currently affiliated with Google LLC. This article represents his own opinion and does not represent the views of Google.

2 CCR AND VELOCITY

Between 2004 and 2008, I tried a number of things to increase the pace of scientific information dissemination while ensuring decent quality and innovation, more or less successfully. The quarterly submission system was probably the most successful. We managed to clear the buffer every three months and not carry any paper forward. This guaranteed timely publication while increasing the pressure on area editors and reviewers to deliver thorough reviews and decision in such a short time period. We also asked each editor to write a short intro to explain why the paper had been accepted. CCR readers do not see reviews and sometimes have a hard time understanding why a paper is published. Having editors write a short intro for each accepted paper had two benefits. First, editors had chance to explain what they (or the reviewers) liked in the papers, and even what they did not like in a paper. Second, the reader would understand why a paper was published. It was a lot of work to get these notes before closing the issue, but definitely worth it.

Editorials have been quite successful as well, at the price of a significant effort to convince people to invest time in papers that would not get them brownie points for their career/promotion. The success of editorial notes resulted from a combination of three main factors, namely (1) timeliness of the publication, (2) interaction (we could have people react to a paper in consecutive issues), and (3) diversity of content (from workshop reports to scientific ideas through personal opinions). It took some work and a lot of juggling, as many articles were solicited by the editorial team and last minute cancellations were commonplace. But it worked and by the end of my tenure, the reputation of editorials was good enough that 80% of editorials were submitted without solicitation. Among all editorial papers, the “10 favorite paper” series was definitely the most popular and I did not have to push hard to get contributions. Per Jim Kurose, *“all of these innovations helped make CCR a valuable alternate and more timely way (besides the best conferences and journals) to bring high-quality scientific, technical, and opinion material to the research community. In today’s “information-overload” setting, that is an incredibly valuable service.”*

Two cool editorial memories: When Jon Crowcroft sent me the first “automatically generated” paper [1] I decided to publish it immediately. Apparently, despite the unrecoverable gap between French and English humor, few people found that funny and this article received the highest volume of readers’ complaints in 4 years. I received many (sometimes unnecessarily insulting) messages for publishing this paper, with the main reason being lack of content, waste of paper, and a demonstrably poor sense of humor. We also got some thumbs up, from Keshav in particular who told me he loved

the paper. So in the end I am happy I published it and generated some scientific debate in our community.

Publishing topic series in consecutive issues was also interesting. The main (and maybe the only) success was the buffer sizing thread (where some of the articles in the series are [2–6]). Three month is far too long, and Twitter would have permitted a much better interaction, at the cost of the kind of emotional reactions we cannot get used to. However, we reached the point where multiple opinions were expressed and authors managed to interact, make each other unhappy, while giving time to the editorial team to filter some of the content.

Not all of the innovations we undertook were successes. CCR online was definitely a failure, but worth trying. We built CCR online as a community that allowed contributors to exchange ideas, comment on papers, provide unsolicited reviews, interact in real time without any moderation. It was designed as the highest velocity component of the CCR ecosystem. Despite Ernst Biersack’s and Moritz Steiner’s efforts to make it a success, we failed to establish CCR online in these 4 years. First the technology was not ready for the level of interaction needed. Second, we were all too busy with the other innovations to dedicate enough energy to CCR online. We were definitely too ambitious and should have introduced CCR online after having established the new features described earlier. On the other hand, we should be glad it did not succeed, as we were not prepared for the level of editorial oversight such a platform requires.

3 ON CONTINUOUS SUBMISSION AT SIGCOMM

While CCR editor and later as a SIGCOMM Executive Committee member, I have lobbied for SIGCOMM to adopt a continuous submission model similar to VLDB¹.

With one deadline per year, and super deep validation of the technical content, the SIGCOMM conference acts as a bottleneck in (1) the diffusion of new research results, and (2) young graduates’ career. I proposed to build SIGCOMM’s program as a best-of CCR articles published in the past 12 months of CCR (or all papers ; there are different ways to build the SIGCOMM program and this is not the point of this discussion). Moving to this model would have many advantages, and no obvious drawbacks (to me at least) – it would (1) increase CCR submission quantity and quality, (2) give authors a chance to improve papers beyond the reviewer feedback, (3) reduce the pressure on the SIGCOMM TPC, (4) offer multiple chances to authors to publish a SIGCOMM paper, (5) reduce the risk of accepting borderline papers, (6) provide more freedom to the TPC chairs and committee to build an exciting program as they would pick from a large set of thoroughly reviewed papers. I still believe this is the right model! The CCR session that has been introduced 8 years ago at SIGCOMM was a good first step, but apparently it did not manage to convince people to make the leap; and the organization of a parallel session last year (which seems to be institutionalized in 2019) appeared as a big step backward to me. Since their pioneer move, VLDB has been followed by many other

¹<http://vlldb.org/pvldb/pvldb-faq.html> mentions the motivation for "continuous" submission, and <http://vlldb.org/pvldb/submission.html> provides process details. More information is available at vlldb.org.

conferences, SIGMETRICS and NSDI being noticeable examples in our community. To date, I have not seen any of these top tier conferences taking the leap backward. So it might be time to wake-up, SIGCOMM!

4 ABOUT VOLUNTEERING

I want to conclude with a few thoughts on the CCR editor’s job in particular, and community volunteering in general. I tried many things as CCR editor - much more than I had time to do. As noted above, many were successful, and we learned a lot from the others. We always had a unique objective: to make scientific information accessible to *all*, and to reduce the barriers so that everyone can benefit from the shared knowledge. With changing times, and practices, it is always important to innovate (something our community should definitely appreciate as much as any). And it takes volunteers who are willing to freely contribute their time and expertise to undertake these efforts, and the risks inherent in innovating. Our community needs such volunteers, and needs to appreciate the efforts (both successes and not) of these volunteers. My volunteer experience at SIGCOMM was very formative and I encourage all of you to contribute as the dynamic and the impact of our community are a consequence of its volunteers’ engagement. Long Life to CCR and SIGCOMM.

ACKNOWLEDGMENTS

Thanks to Jim Kurose, Dave Wetherall and Renata Cruz Teixeira for helping me make this paper better and sound less grumpy; and to Keshav, Jon Crowcroft, Jehangir Amjad, Olivier Bonaventure, Bruce Davie, Divesh Srivastava and Augustin Soule as well for additional thoughtful comments.

REFERENCES

- [1] Jon Crowcroft and Christian Kreibich. 2006. Only 365 Days Left Until the Sigcomm Deadline. *SIGCOMM Comput. Commun. Rev.* 36, 5 (Oct. 2006), 57–62.
- [2] Amogh Dhamdhere and Constantine Dovrolis. 2006. Open Issues in Router Buffer Sizing. *SIGCOMM Comput. Commun. Rev.* 36, 1 (Jan. 2006), 87–92.
- [3] Mihaela Enachescu, Yashar Ganjali, Ashish Goel, Nick McKeown, and Tim Roughgarden. 2005. Part III: Routers with Very Small Buffers. *SIGCOMM Comput. Commun. Rev.* 35, 3 (July 2005), 83–90.
- [4] Nick McKeown and Damon Wischik. 2005. Making Router Buffers Much Smaller. *SIGCOMM Comput. Commun. Rev.* 35, 3 (July 2005), 73–74.
- [5] Gaurav Raina, Don Towsley, and Damon Wischik. 2005. Part II: Control Theory for Buffer Sizing. *SIGCOMM Comput. Commun. Rev.* 35, 3 (July 2005), 79–82.
- [6] Damon Wischik and Nick McKeown. 2005. Part I: Buffer Sizes for Core Routers. *SIGCOMM Comput. Commun. Rev.* 35, 3 (July 2005), 75–78.