

Educating The Next Generation of Reviewers

Bruce Davie
Systems Approach
bsd@systemsapproach.org

Matthew Caesar
University of Illinois at Urbana-Champaign
caesar@illinois.edu

ABSTRACT

While our focus as educators is often on the technical education of students, there is also a need to educate the next generation of reviewers. Reviewing is often treated as a “learn by experience” activity, with service on program committees, or shadow PCs, being the main way that people learn how to write reviews. This position paper advocates for a focused effort, leveraging both shadow PCs and a new “reviewer academy” process, to teach the practice of paper reviewing. Our hope is that this will help to address one of the recognized pain points of our community. This paper is intended to stimulate discussion rather than to present a fully formed solution.

KEYWORDS

Reviewing, Education

1 INTRODUCTION

There are many concerns about the process by which papers are reviewed in the SIGCOMM community and many proposals to address these issues. Recent discussion on the Networking Channel [1] presents a number of the problems, and proposes a range of solutions, as do earlier papers by Crowcroft *et al.* [2], Anderson *et al.* [5], Allman [3], and Mogul [6].

This paper does not claim to have a comprehensive solution to these wide-ranging problems, which are also spurring discussions about substantial changes to how we organize our conferences (e.g., [4]). What we propose is to focus on one aspect only: the quality of reviews. And while there are many things that might be done to improve review quality, it seems that a reasonable place to start is to educate the next generation of reviewers on how to write good reviews. If we have a bigger pool of people who know how to review well, it seems reasonable to think review quality might improve.

To the best of our knowledge, there isn’t robust data to suggest that review quality has declined. What we do know is that there are plenty of complaints about review quality and most of us will remember the one bad review more than we remember the large number of good ones. There is an opportunity to gather such data, such as by allowing authors to rate their reviews. At the same time, it seems there will always be room to improve review quality.

There are a few ways that people learn how to review papers. The *Shadow PC* concept [9] was created primarily to provide an environment in which junior researchers could experience a PC and learn how to review papers prior to their participation in “live” PCs. It could be argued that if we had enough shadow PCs and ran enough junior researchers through them, we’d end up with a large enough pool of well trained reviewers. Since that doesn’t seem to have happened over a lengthy period of running shadow PCs, we are looking at ways to scale the process.

One opportunity to improve upon the shadow PC approach would be to make an effort to collect best practices, which could then be used:

- as a resource for those running and taking part in PCs, shadow or otherwise;
- as background information for those wishing to serve on PCs;
- as a starting point for a sort of “reviewer academy”, as we elaborate below.

We also note that not everyone will agree on what constitutes a “good” review. Reviews that offer minimal feedback beyond numerical scores are one example of poor reviewing practice. Further examples (good and bad) are provided in [2].

Note that resources on reviewing best practices already exist, but may not be visible to those who would benefit from them. Keshav’s *How to Read a Paper* [7] is justifiably well known and widely assigned in CS classes (or so it seems from its widespread presence on course web sites). Another great resource that probably deserves more prominence is Roscoe’s *Writing Reviews for Systems Conferences* [8]. Developing more resources along these lines and creating a central repository for them seems like something the SIGCOMM community could undertake.

2 CHALLENGES

There are many challenges to improving review quality, many of which are highlighted in the works we have already cited. For the specific approach that we are proposing, which is to focus on educating the next generation of reviewers, the following set of problems are apparent:

- (1) People are busy. We are asking people to give their time as both mentors and mentees to activities that don’t directly advance their own interests, but contribute to

a community goal of improving overall review quality. Furthermore, lack of time for reviewing is surely one reason that quality is not higher, so we must ask if we might be just adding more work to already overloaded people.

- (2) Cultural norms. Mogul [6] highlights the prevalence of hypercritical reviewing and the likelihood that those who receive critical reviews will themselves turn around and provide critical reviews. Similarly there are norms around “what is sufficiently innovative” or “what is in scope for SIGCOMM”. Changing these norms is no small task and indeed is one reason cited for either significant changes to the SIGCOMM conference [4] or the formation of new conferences such as NINeS [10].
- (3) Scalable, sustainable deployment. If we hope to make a meaningful difference in review quality, we need to scale up the education of reviewers to reach a substantial fraction of the reviewer population. We also need to be able to sustain the process of education over multiple years.
- (4) The impact of Large Language Models (LLMs). LLMs are impacting the review process in both paper submission (papers written largely or partly by LLMs) and in the review process (reviewers outsourcing their efforts to LLMs). Any attempt to educate the next generation of reviewers has to acknowledge this situation and address it in some way.

We don’t claim to have solutions to all these issues. It may be possible to improve review quality without fully addressing them all. With this paper we are outlining the scope of the challenge to begin a discussion about plausible solutions.

3 PROPOSAL: A REVIEWER ACADEMY

One observation made in the Networking Channel discussion on reviewing [1] is that the current approach to PC meetings, which are mostly now virtual, has eroded one of the venues in which junior researchers used to learn about good reviewing practices. When everyone sat in a room for a day or two and all PC members were present for all paper discussions, there was at least the possibility of establishing norms of behavior in PC meetings and for high quality reviews to be made visible as good examples to the whole PC. While there are certainly advantages to the default online model for PC meetings, there is now a lack of opportunity to develop shared norms around reviewing. This leaves a bit of a vacuum to be filled in educating the next generation of reviewers.

We are proposing a “reviewer academy” program, where early-career researchers are mentored by more senior members of the community on how to write high-quality reviews

and conduct review meetings. We envisage activities including tutorials and hands-on activities such as reviewing past papers, comparing reviews of the same paper, and reviews of reviews. This is an opportunity to teach norms around review tone, fairness, recognizing different research styles and contexts, and avoiding common biases (institutional, regional, or stylistic). Education on the subject of bias might draw on material such as that provided by Project Implicit [11].

We also envisage providing guidance on the appropriate and ethical use of AI tools in reviewing. Given the known limitations of LLMs (such as hallucinations and bias) it is important that reviewers understand the peril of using LLMs without sufficient oversight. The ethical implication of sharing unpublished material with third party tool providers also should be discussed.

There could be clear benefits in operating these sessions in person, possibly alongside existing conferences, workshops and Pods. At the same time, we will better scale up the reviewing community and provide a more inclusive environment if we also provide online options. Running sessions as a set of video conferences with active participation from all attendees would seem a feasible approach. Some parts of the program might be offered as recorded online training in order to lessen the workload for mentors and make the system more scalable, but we recognize that online training alone is unlikely to produce the cultural shift towards higher review quality that we are looking for.

To make the system sustainable over the long term we are in discussions with the SIGCOMM EC to support this program. We envision some of the sessions as events attached to SIGCOMM conferences to leverage the presence of mentors and mentees in common locations.

Given the increasing importance of artifact evaluation, we can imagine expanding this program to include reviewing of artifacts, which requires a somewhat different (partially overlapping) set of skills from paper reviewing. We probably should consider this as a phase two rather than taking on too broad a scope initially.

We propose that there would be some sort of certificate of completion for participants in the academy sessions. Conference steering committees and PC chairs could prefer to recruit PC members who have completed a reviewer academy or shadow PC.

4 CONCLUSIONS

We propose the creation and collection of materials documenting reviewing best practices to educate the next generation of reviewers. We recommend leveraging the existing shadow PC process as well as a new set of “reviewer academy” sessions, both in-person and online. Our overall aim is

to improve reviewing culture broadly in our community. If the academy is successful, a lightweight recognition could be provided to those who have taken part. We hope to spark a discussion that leads to the development of a larger pool of expert reviewers and improvement in overall review quality.

ACKNOWLEDGMENTS

Thanks to the reviewers for their helpful feedback on the paper.

REFERENCES

- [1] The Networking Channel *The Future of Paper Reviewing: Building a Fair, Sustainable System for the Next Decade*. <https://www.youtube.com/watch?v=41MXsbmotlc> April 2026
- [2] Crowcroft, J., Keshav, S. and McKeown, N. *Scaling the academic publication process to internet scale*. Communications of the ACM, January 2009.
- [3] Allman, M. *Thoughts on reviewing*. ACM SIGCOMM Computer Communication Review, March 2008.
- [4] Shenker, S. *Rethinking SIGCOMM's Conferences: Making Form Follow Function*. ACM SIGCOMM Computer Communication Review, October 2022
- [5] Anderson, T., Kurose, J. and Shenker, S. *Improving Sigcomm: A Few Straw Proposals*. <https://sigcomm.edu.cs.illinois.edu/advice/improvingSigcommReviewingJuly2001.pdf>
- [6] Mogul, J. *Towards more constructive reviewing of SIGCOMM papers*. ACM SIGCOMM Computer Communication Review, July 2013.
- [7] Keshav, S. *How To Read a Paper*. <https://web.stanford.edu/class/cs114/reading-keshav.pdf>
- [8] Roscoe, T. *Writing Reviews for Systems Conferences*. <https://people.inf.ethz.ch/troscoe/pubs/review-writing.pdf> March 2007
- [9] Bonaventure, O. *et al.* *The CoNEXT shadow TPC*. ACM SIGCOMM Computer Communication Review, March 2008
- [10] NINeS 2027: Statement of Purpose <https://nines-conference.org/sop>
- [11] Project Implicit. <https://www.projectimplicit.net>