

Behind the Scenes of the HotNets 2020 Program Committee Process

Harsha V. Madhyastha

Venkat Padmanabhan

While serving as PC chairs for HotNets 2020, we encountered various situations beyond the normal course that we thought would be useful to document. Our hope is to create awareness among future authors, reviewers, and chairs, so that they are prepared, should similar situations recur. While we explain how we dealt with these situations, arising from before, during, and after the PC reviewing process, it is our hope and expectation that readers will evolve possibly better approaches to dealing with such situations if/when they are in the “hot seat.” Also, as it turns out, the ACM has policies in place to deal with at least some of these situations. While we only learnt of this after the fact, we would encourage future chairs to consult the ACM policies or, since the policies might not be straightforward to find or interpret, to consult the SIGCOMM leadership, proactively.

1. Request for anonymous authorship

A month or so before the paper submission deadline, a professor wrote to us asking if they could submit a paper where the identity of a student author would be kept hidden (by means of a pseudonym), not just in the submission but also in the final paper, if accepted for publication. The reason for this request was that the student’s apparent concern that the content of the paper (which touched on network policy matters) might jeopardize their impending employment with an entity that might frown upon it.

We deliberated on this, discussed the matter with the steering committee, and also considered the one [precedent of an anonymous publication](#) that we were aware of in the SIGCOMM community. We also considered the view that we really don’t have any procedures in place in our community today for verifying that the authors of papers are who they say they are.

Ultimately, we decided that it would not be healthy for our community to allow authors to publish papers without them being willing to reveal their identities and stand by their (scientific) work. Furthermore, in the particular case of the request we received, it wasn’t clear as to how the student would benefit (in terms of receiving credit) from being listed as a pseudonymous author versus not being on the author list at all. In fact, we suggested to the professor that they could acknowledge the student’s contributions in their paper and perhaps also include a link to a webpage that could be updated subsequently to indicate the identity of the student author at a later date, if/when the student became comfortable revealing their identity.

In the event, this paper was not submitted to HotNets 2020. And after the conclusion of the PC process, it came to our notice that the ACM’s [policy on author anonymity was spelt out in the April 2020 issue of the Blue Diamond, ACM’s publications newsletter](#).

2. Non-anonymized submission

After the submission deadline, it came to our attention that 3 submissions had ignored the anonymity requirement and had included the authors’ identities. Perhaps this was an inadvertent mistake, but since the submissions had already been available to the PC members to look through and therefore the anonymity of the authors had been compromised, we proceeded to reject the papers without review and notified the authors of our decision. One of the authors requested us to reconsider, but we were not able to since the “damage” (in terms of compromising anonymity) had already been done. To our knowledge, HotCRP automatically makes the submissions visible to the PC members, leaving the chairs with no option to hide the submissions until they have scanned through all the papers. In any case, it is only fair for the responsibility to lie with the authors to follow the clearly

posted instructions, not with the chairs (who are quite stretched as it is) or with the submission system to backstop the authors' mistakes.

3. Violations of formatting requirements

In contrast with the typical requirement that authors format their submissions for US Letter paper, we chose to mandate that the PDFs, both during the initial submission and at the time of the camera-ready preparation, be formatted for A4 paper. After all, A4 is an international standard whereas US Letter is not. In working with the ACM and Sheridan, it became clear that their existing publication processes are not well-equipped to deal with A4-formatted papers, though ultimately the proceedings were produced in A4 format. We expect that the process will become smoother in the future as more venues hopefully adopt the international standard.

Despite this change, all submissions were correctly formatted for A4 paper, though a few violated the font size and margin requirements. One submission egregiously exceeded the specified margins. Given the time and effort that the authors had presumably put into their submission, we decided against rejecting their submission outright. Instead, we wrote to the authors and asked for their paper's source. We then compiled their paper with the style file we had provided to authors, and then clipped off the paper's main text at 6 pages. Doing so proved to be worthwhile as this paper did advance to the second round of reviewing, even though it was eventually not accepted.

4. Withdrawal of paper because of author disagreement

Early in the review process, we were contacted by an author, requesting that their submission be withdrawn. Apparently, a visiting student from another institution, who was a co-author, had included certain data in the paper without the consent of the student's (non-author) advisor. When this came to light, the advisor had objected to the inclusion of this data in the submission.

We agreed to allow the authors to withdraw their submission, while impressing on the author who had written to us that this was an undesirable situation, especially given the possibility of wasted effort by reviewers on account of the lack of coordination among the authors. We requested the author to also communicate this to the visiting student and their advisor, so that such a situation doesn't recur.

5. Rejection of paper because of undisclosed concurrent submission

Late in the review process, it came to light that there was a concurrent full-length submission to a conference in a different community, with striking similarity in content to one of the submissions to HotNets. We communicated with the chairs of the other conference and verified that the two papers had common authorship in part. Furthermore, on examining the two papers, it was clear that much of what had been claimed as novel in the HotNets submission was present in the concurrent submission. The HotNets submission centered on elaborating on a subset of what was in the concurrent submission (which, right around that time, had also been accepted for publication in the other venue).

With the consent of the chairs of the other conference, we communicated with the authors in common and asked them to explain why their HotNets submission did not make any mention of the concurrent submission. In subsequent communication, we also pointed out several specific claims of novelty made in the HotNets submission, which weren't so in light of the concurrent submission. The authors did not provide a satisfactory explanation.

In view of the above, we concluded that this was a clear case of violation of the norms against duplicate submission, which are the bedrock of the publication process for any scientific community, and proceeded to reject the paper (which, by then, had already received multiple reviews at HotNets,

so there was no doubt wasted effort on account of this paper). We also communicated our decision to the chairs of the other conference, upon their request.

Furthermore, since the authors brought it up, we clarified that the policy on concurrent submission in the HotNets CFP refers to a very different situation, one where the authors submit a position paper to HotNets introducing a novel idea or broad direction, where they might receive credit for such an idea, and concurrently choose to submit another full-length paper that digs deeper into and works out the details of various aspects of the proposal, where they might get credit for such details but NOT for the novelty of the broad idea itself. It goes without saying that even if authors were to choose to make such concurrent submissions, they would have to reveal the existence of the concurrent submission and be careful in claiming credit in each for only the “delta.” Otherwise, there would be no way for the reviewers to know what constitutes prior work, and what is new and deserves credit.

6. Authors’ complaint regarding inappropriate language in review

A few days after the decisions and the reviews had been communicated, the authors of a rejected submission wrote to us to call out certain inappropriate and seemingly biased language in one of the reviews. Upon examining this, we acknowledged to the authors that some of the language was inappropriate and apologized to them for not catching it. We also clarified that we had recruited multiple external experts as reviewers for this submission and that the ultimate rejection of the paper was based on technical concerns.

While none of the reviewers had caught the inappropriate language even during the vigorous online debate on this paper (which had centered on technical issues), it was ultimately our responsibility, as the chairs, to have caught it. Ironically, during our scan through the hundreds of reviews of the nearly 120 papers in play, we had caught a few other issues, though minor in comparison, and had leaned on the reviewers to fix those (which they did).

With the consent of the authors, we forwarded their (anonymized) complaint to the reviewer in question, adding our concurrence with the authors that certain language was inappropriate and possibly offensive. We offered to the reviewer our willingness to communicate any response that they might have back to the authors, again while maintaining anonymity. The reviewer wrote back focusing on the technical reasons for advocating rejection. We wrote back to the reviewer, explaining that while the authors might have a disagreement with technical points in the review too, the reason for their complaint was the inappropriate language. The reviewer was not persuaded. With the consent of the reviewer, we forwarded the reviewer’s (anonymized) email response to the authors. We added that we were inclined to close the matter and one of the authors indicated their concurrence.

Thereafter, we were surprised to learn that snippets from the private email exchange between the authors and the reviewer, mediated by us, had been posted on Twitter by one of the authors. We thought this was inappropriate since we had been careful to forward communication from each party to the other only with their consent. We requested the author in question to clarify that the posting of snippets from the (private) email exchange had been done without our (the chairs’) knowledge. In the same email, we also suggested to the authors that they consider publishing a rebuttal of the whole review — the technical points and the inappropriate comments. In other words, we encouraged the author to publish and rebut the review; our objection was only to snippets of the subsequent private email exchange between the authors and the reviewer, mediated by us, being posted online without our knowledge. (We learnt after the fact from the SIGCOMM leadership that [according to ACM policy, “reviewers own the copyright for their reviews.”](#) so future authors and chairs should keep this in mind should such a situation recur.)

Thereafter, the HotNets steering committee and separately the SIGCOMM CARES committee reached out to us and we had the opportunity to provide our perspective. Subsequently, it also came to light that the Twitter posts by one of the authors had made claims regarding our role as the chairs that were contrary to the facts. We had no interest or inclination in joining the “Twitter fray,” so instead we communicated the facts to the chairs of the CARES committee, for inclusion in their report. We will update this message with a link to their report once it becomes available.

If a similar situation were to recur, we would advise future chairs to bring it to the notice of the SIGCOMM CARES committee (which, somehow, wasn’t on our radar) and let them handle it from the get-go. Since CARES has experience in dealing with such situations across the various SIGCOMM conferences and workshops, leaning on them would be better than any effort by the chairs themselves, even if it is well-intentioned.