

Viewpoint

Collusion Rings Threaten the Integrity of Computer Science Research

Experiences discovering attempts to subvert the peer-review process.

THE DISCIPLINE OF computer science has historically made effective use of peer-reviewed conference publications as an important mechanism for disseminating timely and impactful research results. Recent attempts to “game” the reviewing system could undermine this mechanism, damaging our ability to share research effectively.

I want to alert the community to a growing problem that attacks the fundamental assumptions that the review process has depended upon. My hope is that exposing the behavior of a community of unethical individuals will encourage others to exert social pressure that will help bring colluders into line, invite a broader set of people to engage in problem solving, and provide some encouragement for people trapped into collusion by more senior researchers to extricate themselves and make common cause with the rest of the community. My motivation for writing this Viewpoint is because I became aware of an example in the computer-architecture community where a junior researcher may have taken his own life instead of continuing to engage in a possible collusion ring.^a



Collusion rings extend far beyond the field of computer architecture. I will share another data point, from artificial intelligence and machine learning. I will keep some of the details (like the identity of the specific conference) vague because I think naming names could do more harm than good. Since my goal is to raise awareness of the issue and help

people understand how widespread it is, I do not think such details are essential.

Let me start with a reminder about several salient attributes of the review process. What I describe is not precisely what is used by any specific conference but it matches well with the three or four big conferences I have been involved in organizing.

^a See <https://bit.ly/3duc9rY>

► The peer-review process is carried out by a program committee consisting of one or two program chairs, several-hundred area chairs, and approximately 5,000 reviewers. Reviewers are asked to declare conflicts of interest so they are not assigned to review papers that would compromise their partiality.

► Authors submit papers with their names withheld for reviewing (“blind”). One notable conference received 10,000 submissions last year, up from an all-time high of 1,000 only six years earlier.

► Reviewers “bid” on specific submitted papers based on the paper titles/abstracts to indicate those they are qualified to review.

► Reviewers are assigned papers by the program chair(s), attempting to respect their bids while avoiding disclosed conflicts of interest.

► Reviewers read their assigned papers and submit reviews. They share their reviews with one another and try to reach a consensus recommendation (accept/reject) for each paper, which the area chairs and program chairs use to build the conference’s technical program.

Overall, stakes are high because acceptance rates are low (15%–25%), opportunities for publishing at any given conference are limited to once a year, and publications play a central role in building a researcher’s reputation and ultimate professional success. Academic positions are highly competitive, so each paper rejection—especially for graduate students—has a real impact on future job prospects. Some countries correlate promotion and salary decisions to the number of papers accepted at a specific set of high-profile conferences (and journals).

Given the intensity of the process, researchers push themselves very hard to do the best work that they can. The week or two leading up to a conference deadline is exceptionally stressful, with researchers neglecting other responsibilities, running their computers at capacity, and getting very little sleep. Even so, hard work does not appear to be enough to guarantee success—the review process is notoriously random. In a well-publicized case in 2014, organizers of the Neural Information Processing Systems Conference formed two independent program committees and had 10% of submissions reviewed by both. The re-

Without better investigative tools, we may never be able to hold the colluders to account.

sult was that almost 60% of papers accepted by one program committee were rejected by the other, suggesting that the fate of many papers is determined by the specifics of the reviewers selected and not just the inherent value of the work itself.

In response, some authors have adopted paper-quality-independent interventions to increase their odds of getting papers accepted. That is, they are cheating.

Here is an account of one type of cheating that I am aware of: a collusion ring. Although the details of this particular case have not been publicly disclosed, the program chairs who discovered and documented the behavior spent countless hours on their analysis. The issues are complicated, but I have no reason to doubt their conclusions. Here is how a collusion ring works:

► A group of colluding authors writes and submits papers to the conference.

► The colluders share, amongst themselves, the titles of each other’s papers, violating the tenet of blind reviewing and creating a significant undisclosed conflict of interest.

► The colluders hide conflicts of interest, then bid to review these papers, sometimes from duplicate accounts, in an attempt to be assigned to these papers as reviewers.

► The colluders write very positive reviews of these papers, perhaps even lobbying area chairs through back channels outside the view of the other reviewers.


► Colluders occasionally send threatening email messages to non-colluding reviewers if the colluders discover their names and believe the non-colluding reviewers can be influenced.

► Some colluding reviewers temporarily change their names on the online

conference management system during the discussion process, perhaps to avoid getting a reputation for supporting weak papers.

The outcome of this attack, if undetected and successful, is that some authors are rewarded with paper acceptances for very unethical behavior. Given that many conferences have to cap the number of accepted papers due to limits on the number of papers that can be presented at the conference, that means other deserving papers are being rejected to make room. The quality, and perhaps even more importantly, the overall integrity, of the conference suffers as a result.

The research community must respond forcefully to collusion rings, sending a clear message to misbehaving authors and reviewers that what they are doing is unacceptable. Beyond unambiguous messaging, however, it is not yet clear what interventions should be adopted to squelch collusion rings. Conference organizers behind the scenes are weighing dozens of proposals, all of which have potential pitfalls. Better paper-assignment technology would help close one loophole that is being exploited. But, without better investigative tools, we may never be able to hold the colluders to account.

Scientific research is a deeply cooperative endeavor. Researchers compete for attention and funding resources, but also build their ideas on top of those of their rivals. Most researchers see their work as a quest for deeper understanding, not just a way to pay the bills. At present, the peer-review process consists largely of honest participants. But, once unethical behaviors are sufficiently widespread, the incentives for continuing to engage in a community of discovery evaporate. The cheaters run the risk of destroying the very system they depend on for their professional success. It is time to take a close look at the peer-review process and to align the incentives so everyone is working toward sharing the best research work possible. 

Michael L. Littman (mlittman@cs.brown.edu) is The Royce Family Professor of Teaching Excellence in Computer Science at Brown University in Providence, RI, USA.

Copyright held by author.