



Computer Science Research Community: The Death Spiral?

Michael Rabinovich • Case Western Reserve University

Much has been written about the deficiencies of peer review in academic research (the article, “We Are Sorry to Inform You...” in particular will provide good therapeutic reading after a paper rejection).¹ Peer review is a perpetual source of consternation and good humor, an inexhaustible topic for conversations over beer, and a subject of serious research in its own right. The computer science research community in particular is constantly seeking to improve this process through various tweaks, including double-blind submissions, author rebuttals, public reviews, and online repositories with reader annotations and ratings.

Although peer review attracts much attention, I believe other issues around research publishing present far more direct and serious challenges to the community. These include unscrupulous authors trying to “game” – and in the process undermine – the peer review system, and irresponsible professional service volunteers who want to amass a service record without contributing an adequate amount of effort to perform the service properly.

My predecessor at the IC helm, Fred Douglass, discussed these issues in a two-part column five years ago.² Unfortunately, the situation hasn’t improved since then, and has in my view gotten worse. Thus, revisiting these issues seems to be in order. Here, I examine these topics based on years of my own professional service. Whereas Douglass focused on detecting duplicate submissions, I concentrate on deterrence and discuss procedures that would occur *after* detection. Moreover, Douglass argued for “a new independent agency, jointly sponsored by the professional organizations” that would act as a neutral reputation agency for authors and volunteers, akin to a credit bureau. I believe that the neutrality

of such an agency does not necessitate organizational independence, and propose expanding the scope of an existing IEEE mechanism to assume this role.

Duplicitous Authors

The constant pressure to publish causes a well-known behavior in which authors publish multiple papers on the same topic with very little new insight. Academics habitually joke about the changing definition of the “least publishable unit,” and how it grows smaller in the modern age of nano-everything. However, a more duplicitous and harmful scenario involves double submissions, in which an author submits essentially the same paper to multiple venues simultaneously in the hope of increasing the paper’s likelihood of being accepted at one of them.

Unlike papers with small novelty deltas, which are in print for everyone to see and judge, double submissions are by intent hidden. Authors bend the rules to “double-dip” into the reviewing system. This practice is extremely harmful to the community for two reasons: First, it clogs the reviewing system, leading to publication delays and a general decline in review quality. Second, if a paper is accepted at more than one conference, its authors typically choose which venue to publish with and withdraw from the others. This leaves the other conferences with holes in their schedules and shortchanges other worthy submissions that could have been accommodated instead.

When caught, authors often profess ignorance of the rules against duplicate submissions. I personally believe these are often untrue excuses. Many years ago I chaired a

small workshop that could only accommodate around 20 papers. Shortly after decisions went out (and many worthy submissions had to be rejected), two authors withdrew their papers, leaving 10 percent of the schedule empty. One submission was withdrawn for “personal reasons” and the other because the author lacked the funding to attend. In the latter case, because the author was from a developing country, we worked hard to put together a travel grant, only to be told that the author still wouldn’t attend. I then searched the Internet for the two papers, and sure enough – both were accepted elsewhere. That the authors tried to hide the true reasons for withdrawal clearly shows they knew their behavior was wrong.

Another incident occurred with a recent WWW 2012 submission. This time, shortly after decisions went out, one author of an accepted paper informed us that due to additional results obtained since submission, they wanted to change the paper’s title in the camera-ready version to reflect its broader nature. With some hesitation, we granted this request (and in the hindsight, we shouldn’t have done so because the additional results weren’t peer-reviewed). However, just a couple of days before the proceedings were to be sent for typesetting, we were alerted that the same paper was being presented at another conference. We verified that the other paper was indeed essentially the same as our submitted version. In the aftermath of double acceptance, these authors had gone a step further: instead of withdrawing from one of the conferences, they tried to slip in an extra publication by replacing a peer-reviewed version with a different paper that would then appear at our conference without peer review.

Again, the contact author claimed ignorance, but small telltale signs reveal a more deliberate planning.

First, the submission was blinded, although this wasn’t a requirement at WWW – I suspect to reduce the likelihood that the paper would be matched with the duplicate submission. Second, additional authors weren’t even entered into the submission website, so if the paper were rejected, their identities would never be revealed.

Why does this type of behavior propagate? An easy guess is that authors suffer few consequences. At worst, the paper is simultaneously rejected at both venues affected, but because the revision is always confidential, detection is only accidental – for example, when the same reviewer happens to handle both submissions – and thus rare. With virtually no long-term consequences, and ethics aside, the “rational” game in this environment is to pursue double submissions. Furthermore, other authors observe such behavior’s “benefits” and can feel pressure to follow suit to even the playing field with their less-scrupulous colleagues. Trying to avoid “being a sucker” can be a powerful motivator for social behavioral change.

Professional Service Delinquency

Besides unscrupulous authors, the other side of the worrying state of our community is the attitude toward professional service. Most of us have had experience chasing delinquent reviewers. Admittedly, reviewing individual papers is a thankless job (unless you get lucky and the paper turns out to be an enjoyable read), usually done as a favor to the requesting editor or program committee member. Still, when you accept a commitment to write a review by a certain date and then are late, you’re obviously abdicating responsibility.

As if this isn’t bad enough, I’m seeing an increasing trend in which the loose interpretation of commitment

percolates up the professional service chain. Within IC, we regularly interact with guest editors (GEs) for special issues. Although most GEs take their responsibility seriously, a few don’t; the staff and editors spend countless hours and nerve cells trying to hold the feet of these latter GEs to the fire and ensure that they’re moving their special issues along to meet publication schedule constraints. What makes this especially peculiar is that GEs aren’t necessarily solicited to perform these services, but rather submit their special issue proposals to the magazine.

However, a particularly egregious kind of delinquency occurs with PC members. We’re used to PC members being late with their reviews. But increasingly, some members *never* do the work, basically leaving it to the PC chairs and the more responsible fellow members to scramble and put out the fires thus created. At the last WWW conference, several track chairs reported that some PC members never submitted any reviews! We ended up having to downgrade some members to external reviewers (when they submitted at least some reviews) and exclude others altogether.

Again, this behavior has a natural tendency to snowball due to built-in positive feedback. People think that simply accumulating the professional record on their resumes will help their careers (although I believe this is a nearsighted view: real stature can ultimately be earned only through consistent, high-quality research and service contributions, not through numbers on your CV). Yet because the community has no institutional memory, a poor job at one service assignment doesn’t necessarily affect future service invitations. This stimulates seeking more commitments but discourages expending the appropriate effort. Moreover, as long as everyone muddles

through with the immediate task at hand (reviews are obtained, the special issue is salvaged, and so on), all aggravation is forgotten as everyone moves on, which can lead some to infer that their irresponsible stance is really an expected, or at least easily forgiven, behavior. For authors, this behavior most often manifests in delayed decisions and superficial reviews; but they also receive a signal that such inadequate quality of professional service is acceptable, which influences them when it's their turn at the service helm.

What Can We Do?

Clearly, we need to change the system to break the cycle I've described. I believe the crucial missing component is "backpressure" from the community – that is, a deterrent to undesirable behavior. One rudimentary form of such backpressure already exists in isolated islands. Some conference series have steering committees that actively collect feedback from conference organizers, maintain this feedback across conferences within the series, and vet future organizers accordingly. Journal reviewing sites typically let board members rate their reviewers. Colleagues exchange past experiences with regard to candidates for service. However, these mechanisms are isolated and insufficient.

It might appear that the embarrassment of public disclosure could provide the needed backpressure, but this isn't the case. An individual in a position to make the disclosure (for instance, the conference chair) might be understandably reluctant to assume the role of sole "judge, jury, and executioner" when it comes to hurting someone's career. Myriad questions would arise: How did the individual decide the case? Did the affected person receive enough opportunity to explain his or her

circumstances and perhaps offer a rebuttal? How do we avoid personal bias affecting the decision? On the balance, I believe such routine disclosures would badly fracture the community and bring far more harm than benefits, and I deliberately didn't include any names in the examples I cited previously.

A reasonable alternative might be to create an umbrella committee to handle these tasks – an approach Dougkis advocated in part 2 of his column.² Yet getting such an agency off the ground would require a coordinated action from a critical mass of professional organizations that would act as cosponsors. I believe a more practical approach would be for one influential organization to take the lead. Individual professional societies strive to serve broad communities, not just their members. For example, IEEE's mission statement is "to foster technological innovation and excellence for the benefit of humanity" (www.ieee.org/about/vision_mission.html). Improving the health of the research community would certainly fall within this scope.

In fact, Dougkis alluded to the IEEE policy of dealing with plagiarism as an inspiration for the umbrella reputation agency. This policy has since matured into a well-established mechanism: formal procedures are in place in which dedicated individuals handle plagiarism and self-plagiarism complaints and maintain a database of violators, including a "prohibited authors" list, which is supposed to be checked for every submission to any IEEE-sponsored venue. This mechanism, or a similar one from another society, could expand its scope to both make its services available to the entire community and cover professional service lapses by volunteers. Those responsible for running the mechanism could be viewed as members of a special committee, and

their appointment could itself follow some formal procedure.

Individual conferences and journals could choose to associate themselves with this committee (that is, assume obligation to provide feedback, abide by its "black lists," take into account any ratings information the committee maintains, and so on) and would advertise this association in their calls for papers. Over time, such an association would be expected from any respectable venue. As more conferences associate themselves with the committee, it would provide an increasing deterrent against misbehavior. In fact, the committee *could* make its findings public when appropriate because they would be the result of a proper and uniformly applied process.


Such a committee could help reverse the negative trends discussed here. In the meantime, it falls to senior community members to be diligent in addressing these ills in whatever small ways we can and not taking the easy path of simply getting the task at hand over the hump and moving on. □

Acknowledgments

I thank Fred Dougkis for comments that contributed to this column. The opinions expressed here are my personal opinions.

References

1. S. Santini, "We Are Sorry to Inform You..." *Computer*, vol. 38, no. 12, 2005, pp. 127–128.
2. F. Dougkis, "Collective Wisdom: A Modest Proposal to Improve Peer Review," parts 1 and 2, *IEEE Internet Computing*, vol. 11, nos. 5–6, 2007.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.