

Publish Now, Judge Later

By Douglas B. Terry

Microsoft Research Silicon Valley

Abstract

Conferences these days face a reviewing crisis with too many submissions and not enough time for reviewers to carefully evaluate each submission. Numerous good papers get rejected. One possible solution is for conferences to accept any paper that extends our body of knowledge and let the community judge the paper's significance.

The Problem

Conferences in the computing field have large numbers of submissions, overworked and overly critical reviewers, and low acceptance rates. Conferences boast about their low acceptance rates as if this were the main metric for evaluating the conference's quality. With strict limits placed on the number of accepted papers, conference program committees face a daunting task in selecting the top papers, and even the best committees reject papers from which the community could benefit. Rejected papers get resubmitted many times over to different conferences before these papers are eventually accepted or the authors give up in frustration. Good ideas go unpublished or have their publication delayed, to the detriment of the research community. Poor papers receive little attention and do not get the constructive feedback necessary to improve the paper or the work.

Because reviewers approach their job knowing that they must eventually reject 4 out of 5 submissions (or more), they often focus on finding reasons to reject a paper. Once they formulate such a reason, correctly or incorrectly, they pay less thought to the rest of the paper. They do not adequately consider whether the flaws could be corrected through modest revisions or whether the good points outweigh the bad. Papers with the potential for long-term impact get rejected in favor of papers with easily evaluated, hard to refute results. Program committees spend considerable time trying to agree on the best 20% of the papers that were submitted rather than providing comments to improve the papers for the good of all. Even if committees were able to perfectly order submissions according to quality, which they aren't, papers that are close in quality may receive different outcomes since the line needs to be drawn somewhere. People do not always get the credit they deserve for inventing a new technique when their submission is rejected and some later work is published first.

A Proposal

My proposed solution is simple. Conferences should accept and publish all reasonable submissions. Some fields, such as Physics, I'm told, hold large annual conferences where anyone can talk about

almost anything. I am not suggesting that our conferences accept every submission. I believe that computing conferences should enforce some standards for publication quality, but our current standards are far too stringent. We might argue about what constitutes a reasonable publication. Keeping in mind that the main purpose of publication is to teach others, here's what I suggest.

A submission is "reasonable", and hence publishable, if it contains something new (a novel idea, new experimental result, validation of previous results, new way of explaining something, and so on), is based on sound methodology, explains the novelty in a clear enough manner for others to learn from it, and puts the new results in a proper context, that is, compares the results fairly to previous work. Rather than looking for reasons to reject a paper or spending time comparing papers, the role of conference reviewers is (a) to assess whether each submission is reasonable according to this criteria, and, perhaps more importantly, (b) to offer concrete suggestions for improvement. Any paper meeting this criteria should be accepted for publication, perhaps with shepherding to ensure that the reviewers' suggestions are properly followed.

Ultimately, papers will be judged in the fairness of time by accepted bibliometrics, such as citation counts, and, more importantly, by their impact on the field and on the industry. The importance of a published paper is often not known for many years. The "10 years after" or "hall of fame" awards should be used as the way to honor the best papers. These awards should be noted in the ACM Digital Library. Search engines, along with collaborative filtering and public recommendations, could direct researchers to high-quality, relevant work.

Practical Issues

What if a conference accepts more papers than can be presented during the length of the conference? In the steady state, this may not be a serious problem since there are lots of conferences and not that many new papers. If papers stop being submitted (and rejected) from a half dozen conferences, we will end up with far fewer submissions to conferences. To deal with large numbers of papers, conferences may need to have parallel sessions or shorter presentations or both. Personally, I am a fan of shorter presentations. An author should be able to present the key idea behind his work in 10-15 minutes and let people read the paper for more detail. Some papers could be presented as posters only, but I am not a fan of this approach. I would prefer to see all accepted papers treated equally. Let the community judge the papers.

How do authors decide where to submit their papers? Conferences will still have topics of focus. For example, we'll still have conferences on databases, algorithms, systems, networks, etc. One additional criterion for acceptance is that the paper fits the topical scope of the conference. Some papers may fit into multiple conferences. For example, a paper on distributed storage systems could be a database paper and a systems paper, i.e. be suitable for presentation at SIGMOD or SOSP. In this case, since the criteria for accepting papers is the same for all conferences, it doesn't matter much to which conference the paper is submitted. In either case, assuming they are ACM conferences, the paper will end up in the Digital Library. Most likely, an author will submit his paper to the conference that attracts the community to which he mostly closely aligns, such as a conference that is sponsored by a Special

Interest Group (SIG) to which he belongs. Low quality conferences will likely go away, leaving one top conference in each technical area or for each technical community. To me, having fewer conferences would be a good thing.

What prevents people from submitting papers containing the “least publishable unit”? Authors can decide for themselves when they have a significant result that they want to share with the community. Getting ideas and results published quickly is a good thing. There’s no reason that someone should wait until they have a full paper’s worth of results before submitting their work. The length of the paper can be commensurate with its contributions. People who submit lots of short papers with very marginal contributions risk harming their reputations and will likely receive fewer “test of time” awards than those that submit more major results. That may be sufficient incentive to discourage overly incremental submissions.

How would this affect journals? I suspect that journal submissions would go up and more emphasis would be placed on journal publications. Journals would continue to have distinguished review boards that accept and reject papers based on quality. Thus, a journal publication will be viewed as more prestigious than a conference paper. Papers with early results that are presented at conferences may later become journal articles with more substantial results, refined ideas, or practical experiences. Results from multiple conference papers may be combined into more comprehensive journal papers. This could make the publication practices for computing research more similar to those of other scientific disciplines.

Alternative Proposals

I am certainly not the first to observe flaws in our current publication practices or to suggest changes [7,8]. Attendees at a recent Dagstuhl Perspectives Workshop on the “Publication Culture in Computing Research” spent days debating alternatives. That workshop prompted this position statement. Others have suggested modifications to our publication processes, such as open access [3] and post-publication peer reviews [5], and a number of these viewpoints have already appeared in CACM [4,6,9]. New services have been deployed for some communities, such as PubZone [2] which fosters public discussion of published papers in the database field. These practices and systems merit consideration, but are mostly orthogonal to what I propose.

Public web sites, like the Computing Research Repository (CoRR) [1], have been established to encourage the rapid dissemination of new ideas. Authors may choose to make their papers immediately available by depositing them in such a repository. This approach addresses some of the problems that I raise, but differs in three fundamental ways. First, the authors do not get the thrill or experience of presenting their work in front of a live conference audience. Second, the deposited papers generally are later submitted for publication in a more established conference or journal. Therefore, concerns remain about repeated submissions and its load on reviewers. Third, and most importantly, the papers are not peer-reviewed. My proposal retains pre-publication peer review. Thus, authors benefit from receiving constructive feedback that should be considered when revising their papers in advance of publication,

and readers benefit from the knowledge that the work was vetted by a distinguished program committee.

How to Get There

Adopting new publication policies is not simple. I do not expect established conferences to change their practices overnight. Conferences have a vested interest in protecting their hard-earned reputations by maintaining low acceptance rates. University computer science departments have succeeded at getting promotion committees to value conference publications, and are reluctant to make changes that might damage that position. Nevertheless, I believe that gradual steps are possible. As an encouraging trend, I know of a couple of recent systems conferences that accepted more papers than usual while continuing as single-track conferences. Serving as a program committee member for one of those conferences (MobiSys 2012), I observed first-hand the difficulty of getting reviewers to alter their mindsets and accept even marginally more submissions.

One way to move forward is to establish new “high acceptance” conferences in addition to the existing “low acceptance” conferences. Adding more conferences is not a good long-term solution, but could nudge the community in the right direction, provide experimental data, and spark discussion. For example, this year SIGOPS is running a new conference, the Conference on Timely Results in Operating Systems (TRIOS), in conjunction with its highly regarded Symposium on Operating Systems Principles (SOSP). This experimental conference will accept papers that are rejected from SOSP but still make a significant contribution. Lessons learned from this experiment will feed into a broader discussion of publication practices at the SIGOPS business meeting. We expect it to provide insights into whether the community values conferences with less-constrained acceptance rates and whether authors will choose to present their work at such a conference or wait for publication opportunities that might look better on their resumes.

Concluding Remarks

My main proposal is that conferences accept and publish any submission that contributes something new to our body of knowledge and that conveys its contribution in a clear and fair manner. The benefits of accepting any reasonable conference submission and abandoning low acceptance rates are clear:

- Research results get published in a timelier manner.
- Reviewers focus on providing constructive feedback.
- Program committees do not waste time reviewing the same submissions over and over again.
- Credit goes to those who first conceive of an idea and to groups that develop similar ideas in parallel.
- The community judges work by its long-term impact.

However, it does require a fundamental shift in how the research community, as well as tenure committees and other review boards, evaluates conference publications. I believe that some kind of shift is needed.

References

1. "CoRR: Computing Research Repository", <http://arxiv.org/corr/home>.
2. "PubZone Scientific Publication Discussion Forum", <http://pubzone.org/>.
3. Michel Beaudouin-Lafon. "Open Access to Scientific Publications", *Communications of the ACM*, February 2012.
4. Bertrand Meyer, Christine Choppy, Jorgen Staunstrup, and Jan van Leeuwen. "Research Evaluation for Computer Science", *Communications of the ACM*, April 2009.
5. Cameron Neylon. "Reforming Peer Review. What are the Practical Steps?", March 8, 2011, <http://cameronneylon.net/blog/reforming-peer-review-what-are-the-practical-steps/>.
6. David Roman. "Scholarly Publishing Model Needs an Update", *Communications of the ACM*, January 2011.
7. Jack Rosenberger. "Should Computer Scientists Change How They Publish?", *BLOG@CACM*, July 29, 2012.
8. Moshe Y. Vardi. "Revisiting the Publication Culture in Computing Research", *Communications of the ACM*, March 2010.
9. Dan S. Wallach. "Rebooting the CS Publication Process", *Communications of the ACM*, October 2011.