The Relevance of Journals to Computer Science

Every scientific discipline builds on the past: new ideas invariably appear from the analysis, synthesis, and repudiation of prior work. Even an innovator like Sir Isaac Newton wrote to Robert Hooke on 15 February 1676: “If I have seen further it is only by standing on the shoulders of giants.” A necessary prerequisite for building on the past is for the body of archival work to be of the highest possible quality. Work that enters the communal memory should have no errors that either that the authors are aware of, or that can be rectified by careful peer review. Of course, no process can hope to eliminate errors altogether, but archival work should be free from errors that can be avoided with reasonable care.

Conference publications, by their very nature, are susceptible to errors. The process is driven by strict deadlines, preventing authors from having a back-and-forth exchange with the reviewers in an attempt to fix problems. Program committee members, faced with a stack of 15 to 25 papers to review, naturally limit the depth of their reviews. Moreover, the selection of a paper for publication means only that a paper ranked amongst the best of those submitted for consideration by the program committee, rather than a guarantee of absolute quality. Although shepherding does improve the quality of an accepted paper, even shepherding is only mildly effective when faced with the natural reluctance of authors to do additional work for a paper that has already been accepted for publication. For these reasons, a field that treats conferences as archival publications is building on a foundation of sand.

The Computer Research Association (CRA), however, arguing on behalf of the entire field of computer science, states that: “The reason conference publication is preferred to journal publication, at least for experimentalists, is the shorter time to print (7 months vs 1-2 years), the opportunity to describe the work before one's peers at a public presentation, and the more complete level of review (4-5 evaluations per paper compared to 2-3 for an archival journal) [Academic Careers, 94]. Publication in the prestige conferences is inferior to the prestige journals only in having significant page limitations and little time to polish the paper. In those dimensions that count most, conferences are superior.” [1]

The two negatives for conferences identified by the CRA, page limits and ‘lack of polish’ are worth examining. Today, the IEEE/ACM Transactions on Networking (ToN) limits papers to ten free pages and a maximum of 14 pages [2]. This is scarcely longer than many conference papers. Thus, the situation for journal papers is even worse than what the CRA states. On the other hand, what the CRA dismissively calls a ‘lack of polish’ sweeps many issues under the metaphorical carpet: issues like inadequate experimental design, lack of rigour in statistical analysis, and incorrect proofs. It seems unwise to permit these severe problems in papers that we admit to archival status. Unfortunately, given the conference publication process, these errors are unavoidable. Perhaps it would be better to think of ways of improving the journal publication process instead.

Let's start by considering the reasons why the CRA thinks conference publications are superior to journal publications. Two of the three reasons—-number of reviews and time to publication—-are easily remedied. There is no reason why journal editors could not ask for more reviews. Few conference papers receive more than three reviews and this number could be easily matched by journal editors. Second, the two-to-three year publication delay for a journal paper, according to Henning Schulzrinne, who has had a long history of dealing with this process at ToN, arises primarily from the delay in assigning papers to reviewers and the delay in the authors' responses to the first round of reviewer comments. The equivalent processes at conferences take only a few weeks. Why can't journals match that? As a contrasting data point, journals in civil engineering have review times of 90 days and publication delays of only three to five months [3], which is shorter than even conference publication delays in computer science.

This leaves conferences with just one advantage over journals, that of permitting face-to-face meetings. Specifically, in his recent article in CACM [3], Lance Fortnow argues that conferences allow the community:

" * To rate publications and researchers.
* To disseminate new research results and ideas.
* To network, gossip, and recruit.
* To discuss controversial issues in the community."

These are tangible and valuable benefits. However, as Fortnow and others have argued, we could organize conferences where not all accepted papers are presented on stage, leaving some to be presented in the form of posters. These would result in better-attended, more inclusive conferences, which would meet the needs Fortnow identifies, while not detracting from the archival value of journals. The informal poster format would also allow the presentation of early-stage ideas, which is valuable both to authors and to the research community. If posters are clearly marked, this would not detract from the prestige of full papers already published in the conference.

I believe that we can begin to restore the integrity of archival publications by taking the following steps. First, we should increase the number and perceived prestige of posters at SIGCOMM-sponsored conferences, with more time set aside in the technical program for attendees to view posters. This would boost conference participation and better disseminate early stage ideas. Second, we should re-engineer the journal publication process to cap publication delay to six months. Third, journal editors should allow papers to be as lengthy as they need to be, instead of imposing an artificial page limit. Fourth, a greater emphasis on journal publications will be possible only if journals themselves are economically viable. If it turns out that print journals are unviable (a debatable point), we should consider moving to electronic-only journals or subsidize the production cost from conference fees.

As these changes are made, other synergies may also present themselves. For example, reducing the conference review load could free up resources for journal reviews. Similarly, increased conference attendance from a more generous poster acceptance policy could increase journal subsidies, and moving to electronic journals would not only reduce costs, but would also cut publication delay.

The net result of these changes will be to restore the integrity of our archival work. We cannot afford to let this slip much longer: the time to act is now!

S. Keshav
Editor

[1] D. Patterson, J. Snyder, J. Ullman, Evaluating computer scientists and engineers for promotion and tenure; http://www.cra.org/reports/tenure_review.html, August 1999.

[2]http://www.ton.seas.upenn.edu/submissions.html#format

[3] http://pubs.asce.org/editors/journal/resourceeditor/editorresponsibilities.htm

[4] Lance Fortnow, Viewpoint: Time for computer science to grow up, Communications of the ACM. Vol. 52 No. 8, Pages 33-35