

# The Health of Research Conferences and the Dearth of Big Idea Papers

**R**esearch conferences are often the most desirable venues for presenting our research results. For academic computer scientists and engineers, preferring conferences over journals is so common that we even lobby administrators to ensure that conference papers can be viewed in the same light as journal papers in other fields [1]. Hence, the health of conferences is vital to our research mission.

One conventional indication of health is the number of submissions and the acceptance rate at the conference. The accompanying figure shows both statistics for four ACM conferences. Clearly, these conferences appear healthy from this perspective.


I am concerned, however, about the overall impact of increasing workloads on program committees and conferences and of decreasing acceptance rates on authors, especially authors of papers focusing on big ideas or new directions.

Calls for papers often include encouraging words for big idea or new direction papers. The problem is that reviewers see so many regular papers it is just too difficult to switch gears and be more understanding when evaluating bolder papers with holes in arguments or missing measurements.

Program committees typically

start with a ranked list of papers based on the average of numerical ratings in order to cope with the large number of submissions. Big idea papers are sure to get some poor evaluations, which cause them to drop down the list.

chairs and authors with a record of producing such papers. It can be small, as I wouldn't expect a flood of big idea or new directions papers. This committee could meet after the regular program committee in case the latter would



The number of paper submissions and acceptances for four ACM conferences from 2000–2004.

Hence, the increasing workload makes it exceedingly difficult for big idea or new direction papers to be accepted when selecting tens of papers out of hundreds. Occasionally, a senior member will dive in to save such a paper from its low rankings, but it's rare.

## An Experiment

I have a concrete suggestion for an experiment that I hope some conferences will consider and try. Let's set aside one session for such papers, *and* have a separate program committee to select them. This committee could consist of a few former program committee

like to pass along a few of its submissions.

Evidence for evaluating this experiment might include attendance at the session, whether it led to effective discussions at the conference, whether it led to regular papers in later conferences, and so on. My guess is we will need three to five years to evaluate the merits of this experiment before deciding whether it should continue.

Although a single session could take the place of three regular papers at a conference, I would propose instead to drop one keynote address or one panel session. Based on the conferences I've attended, I doubt they would be sorely missed.

I hope the Big Idea experiment

# President's Letter

will be discussed at the business meeting of your next conference. I look forward to hearing what happens.

## How Well Do Conferences Cope With Increasing Popularity?

My second concern is the impact the avalanche of papers might have on many aspects of a conference from three perspectives:

### • Impact on the Program

**Committee:** Either many more papers must be evaluated per program committee member or the program committee must get much larger. (As an extreme example, one conference has a 300-person program committee!) Large size makes it difficult for everyone attending the program committee meeting and it makes it difficult to have a good, single conversation about a paper. I fear either approach affects the quality of the reviews and decisions.

• **Impact on the Conference and Field:** I also can't help but wonder if given the increasing number of submissions it is wise to keep accepting the same number of papers we did 10 or 20 years ago. When there are hundreds of papers submitted and, say, 30 papers are accepted, are the ones ranked 31–60 really that bad? It is research, after all, and I'm not sure of our precision in evaluating current work without the test of time. I also wonder whether researchers will avoid bolder ideas when it's tough to publish even the more conservative ones.

• **Impact on the Authors of Rejected Papers:** Nothing is more frustrating than getting sparse or inaccurate reviews of a rejected paper. Some will wonder if the decision was arbitrary or even political. If you think the process is undependable, one reaction is to submit many papers to the conference, or to submit your unfairly evaluated paper to a related conference. Either reactions result in more papers per conference.

To illustrate this point, let's look at funding of research by NSF in the U.S. It's likely that NSF proposal acceptance rates are lower now than they were 10 years ago; today some acceptance rates are under 10%. Although the ones that win are likely quite good, I wonder if they are also more conservative. I believe that both the field and society would be better off if NSF could afford to fund more than 25% of the proposals, both in encouraging bold research and in being sure worthy ideas are funded.

By analogy, it might also be desirable to increase the percentage of authors participating at conferences. Some conferences have taken this step by accepting more papers but restricting presentations of some papers to only five minutes. For example, the 2004 Principles of Distributed Computing accepted 75 papers for a three-day conference, with half being 25-minute presentations and half being five-minute presentations.

Perhaps the most novel approach to the whole problem is

being taken by the database community under the leadership of SIGMOD. The three large database conferences are going to coordinate their reviewing so that a paper rejected by one conference will be automatically passed along to the next one with the reviews. Should the author decide to revise and resubmit the paper, the original reviewers will read the revision in light of their suggestions. The next program committee would then decide whether or not to accept the revision. Hence, database conferences will take on many of the aspects of journals in their more efficient use of reviewers' efforts in evaluating revisions of a paper.

ACM's research conferences are run by its Special Interest Group (SIGs). I've been working with the SIG Governing Board to help form a task force to study this issue, looking at why submissions are increasing and documenting approaches like those discussed here, and to evaluate their effectiveness. They plan to report back in early 2005. If you have any comments or suggestions, please contact task force chair Alexander L. Wolf (alw@cs.colorado.edu).

I'm sure we'll look forward to their observations.

## REFERENCE

1. Patterson, D., Snyder, L., and Ullman, J. Evaluating computer scientists and engineers for promotion and tenure. *Computing Research News* (Sept. 1999); [www.cra.org/reports/tenure\\_review.pdf](http://www.cra.org/reports/tenure_review.pdf).

---

DAVID A. PATTERSON (pattsrn@eecs.berkeley.edu) is president of ACM.