



DOI:10.1145/1506409.1506410

Moshe Y. Vardi

Conferences vs. Journals in Computing Research

An old joke tells of a driver, returning home from a party where he had one drink too many, who hears a warning over the radio about a car careening down the wrong side of the highway. “A car?” he wondered aloud,

“There are lots of cars on the wrong side of the road!”

I am afraid that driver is us, the computing-research community. What I'm referring to is the way we go about publishing our research results. As far as I know, we are the only scientific community that considers conference publication as the primary means of publishing our research results. In contrast, the prevailing academic standard of “publish” is “publish in archival journals.” Why are we the only discipline driving on the conference side of the “publication road?”

Conference publication has had a dominant presence in computing research since the early 1980s. Still, during the 1980s and 1990s, there was ambivalence in the community, partly due to pressure from promotion and tenure committees about conference vs. journal publication. Then, in 1999, the Computing Research Association published a Best Practices Memo, titled “Evaluating Computer Scientists and Engineers for Promotion and Tenure,” that legitimized conference publication as the primary means of publication in computer research. Since then, the dominance of conference publication over journals has increased, though the ambivalence has not completely disappeared. (In fact, ACM publishes 36 technical journals.)

Recently, our community has begun voicing discomfort with conference publication. A Usenix Workshop on Organizing Workshops, Conferences, and Symposia for Computer Systems

(WOWCS), held in San Francisco in April 2008, focused on the paper selection process, which is not working too well these days, according to many people. (You can find the proceedings at <http://www.usenix.net/events/wowcs08/> and a follow-up wiki at <http://wiki.usenix.org/bin/view/Main/Conference/CollectedWisdom>.)

Two presentations at the workshop evolved into thought-provoking *Communications* Viewpoint columns. In the January 2009 issue, we published “Scaling the Academic Publication Process to Internet Scale” by J. Crowcroft, S. Keshav, and N. McKeown (p. 27). In this issue, you will find “Program Committee Overload in Systems” by K. Birman and F.B. Schneider (p. 34). The former attempts to offer a technical solution to the paper-selection problem, while the latter points us to the nontechnical origins of the problem, expressing hope to “to initiate an informed debate and a community response.”

I hope the outcome from WOWCS and the Viewpoint columns published here will initiate an informed debate. But I fear these efforts have not addressed the most fundamental question: Is the conference-publication “system” serving us well *today*? Before we try to fix the conference publication system, we must determine whether it is worth fixing.

My concern is our system has compromised one of the cornerstones of scientific publication—peer review. Some call computing-research conferences “refereed conferences,” but we all know

this is just an attempt to mollify promotion and tenure committees. The reviewing process performed by program committees is done under extreme time and workload pressures, and it does not rise to the level of careful refereeing. There is some expectation that conference papers will be followed up by journal papers, where careful refereeing will ultimately take place. In truth, only a small fraction of conference papers are followed up by journal papers.

Years ago, I was told that the rationale behind conference publication is that it ensures fast dissemination, but physicists ensure fast dissemination by depositing preprints at www.arxiv.org and by having a very fast review cycle. For example, a submission to *Science*, a premier scientific journal, typically reaches an editorial decision in two months. This is faster than our conference publication cycle!

So, I want to raise the question whether “we are driving on the wrong side of the publication road.” I believe that our community must have a broad and frank conversation on this topic. This discussion began in earnest in a workshop at the 2008 Snowbird Conference on “Paper and Proposal Reviews: Is the Process Flawed?” (see <http://doi.acm.org/10.1145/1462571.1462581>).

I cannot think of a forum better than *Communications* in which to continue this conversation. I am looking forward to your opinions.

Moshe Y. Vardi, EDITOR-IN-CHIEF