Moshe Y. Vardi

DOI:10.1145/2544169

Scalable Conferences

Adapting computing-research conferences to the growth of the field.

IN A RECENT social-media posting I quoted a blog entry by Michael Mitzenmacher, titled “Easy Now,” which opened with the sentence “In the past year or so, I’ve gotten reviews back on multiple papers with the complaint that the result was too simple.” He went on to assert: “From my standpoint, easy is a plus, not a minus.” Both the original blog entry and my own posting were heavily commented on, with the general sentiment strongly sympathizing with Mitzenmacher. This unhappiness with the current state of computing-research conferences seems to reflect the general mood in the community, as has been discussed on these pages over the past few years.

A three-day Perspective Workshop on the subject of “Publication Culture in Computing Research” was held at Schloss Dagstuhl in November 2012 (for details, see <http://bit.ly/1c9jxAS>). A key motivation for the workshop was the observation that in spite of the pervasive dissatisfaction with the status quo, “the community seems no closer to an agreement whether a change has to take place and how to effect such a change.” I would have liked to report that we reached agreement that change must take place and we figured out how to effect such a change. Unfortunately, we did not. We did, however, reach agreement on many issues.

One of the main insights developed at the workshop was the computing-research publishing ecosystem—both conferences and journals—has simply failed to scale up with the growth of the field. Consider the following numbers. Between 2002 and 2012, Ph.D. production in computer science and engineering in North America doubled, roughly from 800 to 1,600 (numbers for other

parts of the world are not available, regrettably). The number of conference papers published by ACM also roughly doubled, from 6,000 to 12,000. How did we respond to this growth in research production? Simple; instead of doubling the *size* of our conferences we doubled the *number* of conferences. The number of ACM conferences during this period grew from about 80 to almost 160!

We are all aware of the adverse effects of “conference inflation.” Instead of serving as community-building events, many conferences have become paper-publishing events, the infamous “journals that meet in hotels.” Matching papers and conferences has become more difficult, as reviewers struggle to find reasons to reject papers, such as “the result is too simple.” Papers bounce from conference to conference, creating an ever-increasing review workload. It is not uncommon to hear of a paper being rejected summarily from one conference only to receive a best-paper award from another conference.

I find this failure to scale extremely ironic considering how much our discipline is about scaling: higher complexity, larger volumes of data, and larger problems. We have built the Internet, which is about to go interplanetary, but we have failed to scale our own institutions. Considered from that perspective, one path forward in the publication-culture debate is to note the growth of the field and resolve to grow our conferences rather than to continue proliferating them. Imagine SIGPLAN, for example, having, say, two large biannual meetings, rather than the 14 conferences SIGPLAN sponsors now.

A bold proposal along these lines is expressed in the Viewpoint

“Publish Now, Judge Later” by Doug Terry on page 44 of this issue. Terry starts with the observation that computing-research conferences today face a reviewing crisis with too many submissions and not enough time for reviewers to carefully evaluate each one. The result is the process, meant to identify the papers of the “highest quality,” is itself of questionable quality. In fact, there is evidence that while reviewers may reach consensus on the small fractions of the strongest submissions and the weakest submissions, there is no consensus on the main bulk of the submissions, and the final accept/reject decisions are essentially random.

Terry, therefore, proposes an approach where conferences accept any paper that extends the current body of knowledge, as it is extremely difficult to judge the true significance of any new research result. In this approach, a conference publication is not the *final* publication of a research result, but its *first* publication. Through discussions and follow-on journal publication, the community will eventually reach judgment on the significance of the result.

The change from “reject as default” to “accept as default” would be a significant change to our publication culture. I do not expect to see such a change be adopted quickly or widely. It would be nice, however, to see one computing-research subcommunity be brave enough to experiment with it. To quote a Chinese proverb, “A journey of a thousand miles begins with a single step.”

Follow me on Facebook, Google+, and Twitter.

Moshe Y. Vardi, EDITOR-IN-CHIEF

Copyright held by Author.