



Editorial

JACM at the Start of a New Decade

It has been almost six months since I took over as editor-in-chief of *JACM*. It is also the start of a new decade. For both reasons, it seems like the right time to share some thoughts on how *JACM* is doing and where it might be headed.

Overall, I believe that *JACM* is going strong as the flagship scientific publication of the ACM. It is widely respected; according to at least one bibliometric authority (<http://www.eigenfactor.org/map>), it is the top-ranked journal in computer science. Yet *JACM* is facing nontrivial challenges. Its charter is to publish the best research across computer science, broadly construed. With the field expanding and becoming increasingly diversified, this is a tall order. Much of the focus of *JACM* has traditionally been on theory of the STOC/FOCS flavor. Past editors-in-chief, such as Joe Halpern and Prabhakar Raghavan, have worked to expand the scope of *JACM* beyond this core. Several editors have been appointed in areas not traditionally covered, such as bioinformatics, Web systems and algorithms, software engineering, and computational economics. But publications in some of the emerging or cross-boundary areas have been slow to follow. A quick look at the 93 articles published in *JACM* over the past three years shows that about 35 are in core algorithms and complexity, that is, theory in the STOC/FOCS mold. In contrast, only one bioinformatics paper and one computer architecture paper were published in the same period, and no papers were accepted in software engineering. Even some areas with very strong theoretical sides have minimal representation including cryptography (2 papers), logic in computer science (3 papers), machine learning (3 papers), and computer-aided verification (0).

There are several possible explanations for the difficulty of attracting top quality papers in some areas. Conference publications are increasingly favored over journal publications in many subfields, a trend legitimized by highly visible position statements such as CRA's memo on Evaluating Computer Scientists and Engineers For Promotion and Tenure (<http://www.cra.org/resources/bp-memo>). More specifically to *JACM*, authors accustomed to publishing in specialized journals of their own community may not be easily convinced to submit to a journal with a less focused constituency and exacting standards. There also seems to be a perception that some areas are simply not welcome to *JACM*. One way to counteract such perceptions is to ensure visible representation of these areas on the editorial board and to state explicitly *JACM*'s interest.

A proactive approach to ensuring top-quality representation from a wider spectrum of areas, initiated by Prabhakar, consists of inviting a small number of papers selected from top conferences in targeted subfields. When I took over, *JACM* had such arrangements with STOC, FOCS, and PODS. We are now in the process of exploring such arrangements with several additional conferences

including PODC, LICS, RECOMB, and CAV. This is appealing because it no longer leaves coverage of underrepresented areas entirely up to the chance of spontaneous submissions.

An orthogonal obstacle to comprehensive coverage is simply due to the proliferation of areas to be covered coupled with the physical limitations and cost of *JACM* as a print journal. Let us say, hypothetically, that we aim to publish annually the 3 best papers in 10 subfields of computer science. At an average of 40 pages per paper, this would quickly exceed the current annual page budget of about 1000 pages allocated to *JACM*. While there may be some flexibility on the part of ACM on the specific page limit, this does not scale in the long run, and makes increasingly tempting the idea of forgoing the print edition altogether and moving entirely online. While I believe that we are unavoidably headed that way, this remains for the time being a controversial idea and it will likely take some time before the ACM Publications Board makes the switch.

In the meantime, *JACM* papers are already published in electronic form in ACM's Digital Library. Besides advantages of cost and scalability, this already provides substantial added value over the print edition in the form of cross-links and searchable metadata. Another potential advantage of the Digital Library is the ability to post additional content on the home page of each article such as errata, links to online appendices, notes from the author, even slides or videos of talks. I believe that such content would be appealing to both readers and authors, and I plan to work towards providing such material on a regular basis. *JACM*'s new information director, Pierre Senellart, will play a central role in shaping *JACM*'s online presence. A new Web site for *JACM* has also been launched.

One of the often-cited drawback of journals versus conferences is the long time from submission to publication. In the case of *JACM*, the publication queue has been around 3 issues, or 6 months, for quite some time. This is reasonable, since it is considered risky to have fewer papers in the publication pipeline. In terms of time from submission to editorial decision, there is still room for improvement. One would like the expected processing time to be much shorter for papers that wind up being rejected than for those that result in acceptance because of the exacting standards of *JACM*. This is indeed borne out by the record over the past 3 years. Of the 409 papers submitted to *JACM* since 2006, 69 were accepted and 340 rejected. The average turnaround time for papers that were eventually accepted was 18 months, with a third accepted within a year. I am working with the editorial board towards keeping the processing time under a year for most papers that are eventually accepted.

Papers that wind up being rejected are of special concern because of *JACM*'s extremely high standards. If an author has to wait for a year only to be told that their paper does not rise to *JACM*'s high standards, they might hesitate to submit even a very strong paper. In order to let authors know as soon as possible if their papers do not meet *JACM* standards, we have adopted a 3-tier filter. First, the editor-in-chief can immediately reject a paper without review (historically, about 50% of all submissions have been rejected this way). If a paper makes it past the first stage, it is assigned to an associate editor who may in turn reject it without review. Finally, the associate editor may decide to engage a quick review process, by which external referees are asked to provide within a month a preliminary opinion on whether the paper meets the standards of *JACM*, even assuming all technical results are correct. The 3-tier procedure has resulted in dramatically faster processing time for

many submissions that are eventually rejected. Of the 340 papers in this category submitted since 2006, the average time to decision was 10 months and the median 1 month. One third of the papers were rejected within 2 weeks. At the other extreme, 36 papers took over a year to be rejected, and 7 took over 2 years. Clearly, there is still room for improvement, and we are working to ensure a more consistently lower processing time.

Finally, I would like to report on several changes to the editorial board. Bernard Chazelle, Lance Fortnow, Maurice Herlihy, Mark Manasse and Moni Naor have stepped down. I thank them all for their dedicated service. They are replaced respectively by Jean-Daniel Boissonnat (INRIA, covering computational geometry), Leonard Schulman (Caltech, quantum computing) Rachid Guerraoui (EPFL, distributed computing), Ravi Kumar (Yahoo!, web systems and algorithms) and Dan Boneh (Stanford, cryptography). Phokion Kolaitis (UC Santa Cruz and IBM) has replaced me as area editor for database systems and theory. Also joining the board is Thomas Henzinger (IST Austria), who will provide for the first time coverage of computer-aided verification. The board now includes several members from Europe, which I think is consistent with *JACM*'s ambition as a publication representative of computer science research worldwide. I extend to all a warm welcome.

JACM has a wonderful editorial board, consisting of leaders in their field who are willing to spend many hours serving authors and readers. Some editors have been dedicated enough to serve for lengthy periods, even several decades. Indeed, *JACM* has traditionally had no limit on length of service on the board. This has advantages and disadvantages, but it is generally considered healthy to have some built-in rotation on editorial boards. We are now transitioning towards a lower bound of 3 years and an upper bound of 10 years for service on the board.

I would like to conclude by thanking Heidi McCormick, our long-standing editorial assistant, for her discreet and effective monitoring and troubleshooting of our daily operations. It is largely due to her efforts that *JACM* runs smoothly.

VICTOR VIANU
Editor-in-Chief