

Computer Science should stay young

Boaz Barak

(June 2016, see final version of this op-ed on <https://cacm.acm.org/magazines/2016/6/202644-computer-science-should-stay-young/abstract>)

Unlike most other academic fields, in Computer Science refereed conferences are generally the most prestigious publication venues. Some people have argued that computer science should “grow up” and adopt journals as the main venue of publication, and that chairs and deans should base hiring and promotion decisions on candidate’s journal publication record, as opposed to conference publications.^{1,2}

While I share a lot of the sentiments and goals of the people critical of our publication culture, I disagree with the conclusion that we should transition to a classical journal-based model similar to that of other fields. In fact, I believe that conferences offer a number of unique advantages that have helped make Computer Science dynamic and successful, and can continue to do so in the future.

First, let us acknowledge that no peer-review publication system is perfect. Reviewers are inherently subjective and fallible, and the amount of papers being written is too large to allow as careful and thorough review of each submission as should ideally be the case. Indeed, I agree with many of the critiques leveled at Computer Science conferences, but also think these critiques could apply equally well to any other peer-reviewed publication system. That said, there are several reasons I prefer conferences to journals:

- 1) **A talk is more informative than a paper** – At least in my area (theory), I personally find that I can get the main ideas of a piece of work much better by hearing a talk about it than by reading the paper. The written form can be crucial when you really need to know all the details, but a talk is better at conveying the high-order bits that most of us care about. I think that our “conference first” culture in computer science has resulted with much better talks (on average) than those of many journal-focused disciplines.
- 2) **Deadlines make for more efficient reviewing** - As an editor in the Journal of the ACM, I spend much time for every submission chasing down potential reviewers. At this rate, it would have taken me decades to process the amount of papers I handled in 6 months as the program chair of the FOCS conference. In a conference you line up a set of highly qualified reviewers (i.e., the program committee) ahead of the deadline, which greatly reduces the administrative overhead per submission.

People often lament the quality of reviews done under time pressure, but no matter how we organize our refereeing process, if X papers are being written each year, and the community is willing to dedicate Y hours to review them in total, on average a paper will always get Y/X hours of reviewer attention. I have yet to hear a complaint from a reviewer that they would have liked

¹ Moshe Vardi, Editor’s letter, CACM May 2009, <http://cacm.acm.org/magazines/2009/5/24632-conferences-vs-journals-in-computing-research/fulltext>

² Lance Fortnow, Time for Computer Science to grow up, CACM, August 2009.

<http://cacm.acm.org/magazines/2009/8/34492-viewpoint-time-for-computer-science-to-grow-up/fulltext#R2>

to spend a larger fraction of their time refereeing papers, but have not been able to do so due to the tight conference schedule. Thus I don't expect an increase in Y if journals were to suddenly become our main avenue of publication. If this happened, then journals would have the same total refereeing resources to deal with the same mass of submissions that conferences currently do and it is unrealistic to expect that review quality would be magically higher.

- 3) **Conferences have rotating gatekeepers** – A conference program committee typically changes at every iteration, and often contains young people such as junior faculty or postdocs that have a unique perspective and are intimately familiar with cutting edge research. In contrast, editorial boards of journals are much more stable and senior. This can sometimes be a good thing but also poses the danger of keeping out great works that are not appealing to the particular board members. Of course, one could imagine a journal with a rotating board, but I think there is a reason that this configuration works better at a conference. It is much easier for program committee members to judge papers in batch, comparing them with one another, than to judge each paper in isolation as they would in a journal. This holds doubly so for junior members, who cannot rely on extensive experience when looking at individual papers, and who benefit greatly from the highly interactive nature of the conference decision process.

Related to the last point, it is worthwhile to mention the NIPS 2014 experiment, where the program chairs, Corinna Cortes and Neil Lawrence, ran a duplicate refereeing process for 10% of the submissions, to measure the agreement in the accept/reject decisions. The overall agreement was roughly 74% (83% on rejected submissions and 50% on accepted ones, which were about a quarter of the total submissions) and preliminary analysis suggests standard deviations of about 5% and 13% in the agreement on rejection and acceptance decisions respectively.³ These results are not earth-shattering - prior to the experiment Cortes and Lawrence predicted an agreement of 75% and 80% (respectively) – and so one interpretation is that they simply confirm what many of us believe – that there is a significant subjective element to the peer review process. I see this as yet another reason to favor venues with rotating gatekeepers.

Are conferences perfect? Not by a long shot – for example, I have been involved in discussions⁴ on how to improve the experience for participants in one of the top theory conferences and I'll be the first to admit that some of these issues do stem from the publication-venue role of the conferences. The reviewing process itself can be improved as well, and a lot of it depends on the diligence of the particular program chair and committee members.

The boundaries between conferences and journals are not that cut and dry. A number of communities have been exploring journal-conference “hybrid” models that can be of great interest. My sense is that conferences are better at **highlighting** the works that are can be of broad interest to the community (a.k.a. “reviewing” the paper), while journals do a better job at **verifying** the correctness and completeness of the paper (a.k.a. “refereeing”), and iterating with the author to a more polished final results.

³ See blog post by Neil Lawrence, March 2015, <http://inverseprobability.com/2015/03/30/nips-experiment-analysis/>

⁴ See blog post by me, May 2015, <http://windowsontheory.org/2015/05/07/turning-stoc-2017-into-a-theory-festival/>

These are two different goals and are best achieved by different processes. For selecting which works to highlight, comparing a batch of submissions by a panel of experts relying on many short reviews (as is the typical case in a conference) seems to work quite well. But fewer deeper reviews, involving a back-and-forth between author and reviewer (as is ideally the case in a journal) are better at producing a more polished work, and one in which we have more confidence in its correctness. We can try to find ways to achieve the best of both worlds, and make the most efficient use of the community's attention span and resources for refereeing. I personally like the "integrated journal/conference" model where a journal automatically accepts papers that appeared in certain conferences, jumping straight into the revision stage, which can involve significant interaction with the author. The advantage is that by outsourcing the judgment of impact and interest to the conference, the journal review process avoids redundant work and can be focused on the roles of verifying correctness and improving presentation. Moreover, the latter properties are more objective, and hence the process can be somewhat less "adversarial" and involve more junior referees such as students. In fact, in many cases these referees could dispense with anonymity and get some credit in print for their work.

Perhaps the biggest drawback of conferences is the cost in time and resources to attend them. This is even an issue for "top tier" conferences, where this effort at least pays off for attendees who get to hear talks on exciting new works as well as connect with many others in their community. But it's a greater problem for some lower ranked conferences where many participants only come when they present a paper, and in such a case it may indeed have been better off if those papers appeared in a journal. In fact, I wish it was acceptable for researchers' work to "count" even if it appeared in neither a conference nor a journal. Some papers can be extremely useful to experts working in a specific field, but have not yet advanced to a state where they are of interest to the broader community. We should think of ways to encourage people to post such works online without spending resources on refereeing or travel. While people often lament the rise of the "least publishable unit", there is no inherent harm (and there is some benefit) in researchers posting the results of their work, no matter how minor they are. The only problem is the drain on resources when these incremental works go through the peer review process. Finally, open access is of course a crucial issue and I do believe⁵ that both conferences and journals should make all papers, most of which represent work supported by government grants or non-profit institutions, freely available to the public.

To sum up, I completely agree with many critics of our publication culture that we can and should be thinking of ways to improve it. However, while doing so we should also acknowledge and preserve the many positive aspects of our culture, and take care to use the finite resource of quality refereeing in the most efficient manner.

⁵ See my December 2012 blog post <http://windowsontheory.org/2012/12/26/occupy-acm-we-are-the-99/>