

Time for Computer Science to Grow Up

Lance Fortnow
Northwestern University

Abstract

Unlike every other academic field, Computer Science uses conferences rather than journals as the main publication venue. While this made sense for a young discipline, our field has matured and the conference model has fractured the discipline and skewered it towards short-term deadline-driven research.

Computer Science should refocus the conference system on its primary purpose of bringing researchers together. We should use archive sites as the main method of quick paper dissemination and the journal system as the vehicle for advancing ones reputation.

Moshe Vardi in his May 2009 CACM editor's letter [2] challenged the computer science community to rethink the major publication role that conferences play in computer science. We continue that discussion and strongly argue that the computer science field, now more than a half-century old, needs to adapt to a conference and journal model that has worked well for all other academic fields.

1 Introduction

Why do we hold conferences?

- To rate publications and researchers.
- To disseminate new research results and ideas.
- To network, gossip and recruit.
- To discuss controversial issues in the community.

The de facto main role of Computer Science conferences is the first item, that of rating papers and people. When we judge a candidate for an academic job, we first check the quality and quantity of the conferences where their work has appeared.

The current climate of conferences and program committees often lead to rather arbitrary decisions even though these choices can have a great impact particularly on researchers early in their academic careers.

But even worse, the focus on using conferences to rate papers has led to a great growth in the number of meetings. Most researchers don't have the time and/or money to travel to conferences where they do not have a paper. This greatly affects the other roles, as conferences no longer bring the community together and thus we are only disseminating, networking and discussing with a tiny subset of the community.

Other academic fields leave rating papers and researchers to academic journals, where one can have a more lengthy and detailed reviews of submissions. This leaves conferences to act as a broad forum and bring their communities together.

2 A Short History of CS Conferences

The growth of computers in the 50's led nearly every major university to develop a strong computer science discipline over the next few decades.

As a new field computer science was free to experiment with novel approaches to publication not hampered by long traditions in more established scientific and engineering communities. Computer science came of age in the jet age where the time spent traveling to a conference no longer dominated the time spent at the conference itself. The quick development of this new field required quick review and distribution of results. So the conference system quickly developed serving the multiple purposes of papers distribution through proceedings, presentations, a stamp of approval and bringing the community together.

With the possible exception of JACM, journals in computer science have not received the prestige levels that conferences do. Only a fraction of conference papers do eventually get published in polished and extended form in a journal. Some universities insist on journal papers for promotion and tenure but for the most part researchers feel they have little incentive for the effort of a journal submission.

As the field went through dramatic growth in the 1980's we started to see a shift in conferences. The major CS conferences could no longer accept most qualified research papers. Not only did these conferences raise the bar on acceptance but for the papers on the margin a preference for certain subareas emerged. Researchers from the top CS departments dominated the program committees and, not necessarily consciously, helped set the agenda with areas that helped their faculty, students and graduates. Over the years these biases became part of the system and unofficially accepted behavior in the community.

As CS grew the major conferences became even more selective and could not accept all the good papers in any specialized area. Many new specialized conferences and workshops arose and grew to capture these papers. These days we have about a dozen US-based conferences in theoretical computer science alone.

The large number of conferences have splintered our communities. Because of limitations of money and time, very few conferences draw many attendees beyond the authors of accepted papers. Conferences now serve the journal role

of other fields, leaving nothing to serve the proper role of conferences.

Other disciplines have started to recognize the basic importance of computation and we have seen strong connections between CS and Physics, Biology, Economics, Mathematics, Education, Medicine and many other fields. Having different publication procedures discourages proper collaboration between researchers in CS and other fields.

3 The Current Situation

Most CS researchers would balk at paying significant page charges for a journals but think nothing of committing well over \$1000 for travel and registration fees for a conference if their paper were accepted (not to mention the time to attend the conference). What does that monetary commitment buy the author? A not particularly fair review process.

With the tremendous almost continual growth in computer science over the past half century combined with desire of each conference to remain small and “competitive”, even with the increase in the number of conferences we simply have too many papers chasing too few conference slots. Each conference has a program committee that examines submissions and makes decisions on which papers will appear at a conference and which will not. The great papers almost always are accepted, the worst papers mostly get rejected. The problem occurs for the vast majority of solid papers landing in the middle. Conferences cannot accept all of these papers and still maintain their high reputations.

Even if the best decisions get made, several good papers will not make the cut. But several factors make the process imperfect at best.

- Being on a PC requires a large time commitment because of the number of papers involved. With the increase in conferences, many researchers, particularly senior scientists, cannot serve on many of these committees, leaving these important decisions mostly to those with less experience.
- As our research areas get continually more specialized a few to none of the PC members can properly judge the importance of most results.
- These specialized areas have a small number of researchers, meaning the appropriate PC members know well the authors involved and personal feelings come to play.
- PC members tend to favor papers in their own areas.
- The hardest decisions are made by consensus. This leads to an emphasis on safe papers (incremental and technical) versus those that explore new models and research directions outside the established core areas of the conference.
- No or limited discussions between authors and the PC means papers often get rejected for simple misunderstandings.

Various conferences have implemented a number of innovative and sometimes controversial ideas to try to make the process more fair (author information removed from papers, author responses to initial reviews, multi-level program committees, separate tracks for areas and quality, higher/lower acceptance ratios) but none can truly avoid most of the problems above.

In the extreme many of the best scientific papers slip through the cracks. For example, nearly half of the Gödel Prize winners (given to the best CS theory papers after they've appeared in journals) were initially rejected or didn't appear at all in the top theoretical computer science conferences.

We end up living in a deadline-driven world, submitting a paper when we reach an appropriate conference deadline instead of when the research has been properly fleshed out. Many also just publish "least-publishable units," doing just enough to get accepted into a conference.

4 The Road Ahead

How do we move a field mired in a long tradition of conference publications to more journal-based system? Computer science lacks a single strong central organization that can by itself break the inertia in our system.

The Computing Research Association, in its 1999 tenure policy memo [1], specifically puts conference publications above journals.

The reason conference publication is preferred to journal publication, at least for experimentalists, is the shorter time to print (7 months vs 1-2 years), the opportunity to describe the work before one's peers at a public presentation, and the more complete level of review (4-5 evaluations per paper compared to 2-3 for an archival journal). Publication in the prestige conferences is inferior to the prestige journals only in having significant page limitations and little time to polish the paper. In those dimensions that count most, conferences are superior.

A decade later the CRA should acknowledge that the growth in computer science and advances in technology changes the calculus of this argument. Quick dissemination via the web makes time to print less relevant and 2-3 careful journal referee reports give a much more detailed level of review than 4-5 rushed evaluations of conference reviewers. The CRA needs to make a new statement that the current conference system no longer fully meets the needs of the computer science community and support the growth of a strong journal publication system. This will also encourage chairs and deans to base hiring and promotion more on journal publications as it should be.

The ACM and the IEEE Computer Society sponsor many of the strongest computer science conferences in the US. These organizations need to allow SIGs and TCs to restructure or perhaps eliminate some conferences even if it hurts their publication portfolio and finances in the short run.

But most importantly, leaders of a major conferences have to make the first move, holding their conferences less frequently and accepting every reasonable paper for presentation without proceedings. By deemphasizing their publication role, conferences can once again play their most important role: Bringing the community together.

5 Conclusion

Our conference systems forces researchers to focus too heavily on quick, technical and safe papers instead of considering broader and newer ideas. Meanwhile we have focused much of our time and money on conferences where we can present our research that we can rarely attend conferences and workshops to work and socialize with our colleagues.

Computer science has grown to become a mature field where no major university can survive without a strong CS department. Time for computer science to grow up and publish in a way that represents the major discipline that computer science has become.

Acknowledgments

I would like to thank Ming-Yang Kao, Richard Ladner, Rakesh Vohra and the reviewers for helpful comments and the many people whom I've shared discussions on this topic. The views presented remain solely my own.

References

- [1] Evaluating computer scientists and engineers for promotion and tenure. http://www.cra.org/reports/tenure_review.html.
- [2] M. Vardi. Conferences vs. journals in computing research. *Communications of the ACM*, 52(5):5-5, 2009.