

Final Thoughts on the STOC'13 "Experiment"

Joan Feigenbaum, PC Chair, ACM STOC'13

<http://www.cs.yale.edu/homes/jf/>

June 30, 2013

1. Introduction

In the Spring of 2012, I was honored to be asked by Lance Fortnow, then Chair of ACM SIGACT, to serve as Program-Committee Chair for ACM STOC 2013. Lance strongly encouraged me to "experiment" and, in particular, to try a two-tiered PC, and I agreed to do so. This break with tradition proved to be controversial and has been discussed at length by the theory community on and off for more than a year. The blog entry <http://mybiasedcoin.blogspot.com/2013/02/discussing-stoc-2013-pc-with-joan.html> and associated comments contain many highlights of the discussion (together, unfortunately, with some of the deplorable anonymous vituperation that pervades the blogosphere). The purpose of this essay is to record points that are absent from or given short shrift in the blog piece but that I think are worth remembering. The essay is *not* self-contained; if you are interested in this subject, start by reading the blog piece and the Foreword to the proceedings.

Most noteworthy is the fact that, in the event, the question of 2-tiered PCs vs. 1-tiered PCs did not seem to matter. The conference was very successful because of the excellent job done by the authors of a bumper crop of great submissions and by Dan Boneh in his capacity as local-arrangements Chair. The number of registered participants was almost identical to the number at STOC'12 in New York (364 in 2013 and 365 in 2012). Because fewer than half of the PC members (in both tiers) actually attended the conference, fewer than 10% of the participants had been involved in the PC (even if one includes indirect participants, *e.g.*, SIGACT Chair Paul Beame and other SIGACT Executive Committee members, who participated actively in some important decisions). The other 90+% of the participants showed no sign of caring about the PC procedures that had produced the program. Some people gave me unsolicited compliments for having tried something different, but these were generally just brief, friendly remarks. No one made any derogatory remarks to me, and no one complained at the business meeting. All in all, I believe that it is safe to conclude that future STOC PC Chairs should be free to convene 2-tiered PCs but should not be compelled to do so. The quantity and quality of STOC submissions are both very high, and any reasonable PC process will produce a great technical program.

Section 2 of this essay contains observations about my 2-tiered PC experience and some suggestions for future Chairs who may wish to improve on it. Section 3 contains suggestions about how the theory community could experiment more broadly than simply tinkering with committee processes.

2. Notes on the 2-tiered PC process

The primary benefits of a 2-tiered PC are the significantly lighter workload for the (second-tier) PC members and their ability to submit. Because of the lighter workload, many members of the theory community (including many distinguished, senior members) who are not willing to serve on traditional PCs are willing to serve on 2-tiered PCs. Because of the ability to submit, young-hotshot members of the theory community will serve eagerly rather than serve grudgingly in order to make sure that STOC-PC service is on their CVs before they come up for tenure. For STOC'13, I hoped that the lighter workload would lead to higher-quality reviews.

Unfortunately, the effect on reviews was minimal. We got more high-quality reviews than STOC has gotten in the past, but not hugely more; the average review quality went up a little but not a lot. That was not for lack of trying on my part. I sent explicit but not overbearing instructions about how to write good reviews, and I reminded people several times well before the notification date that they had ample time to flesh out and polish their reviews. Many PC members ignored these repeated, polite entreaties. The theory community should think about whether it really wants good reviews and, if so, how it can incent people to write them.

On the bright side, appreciation of the light workload was abundant. In my blog post (and in subsequent email to the PC), I said that I was interested in PC members' answers to the following four questions:

- i. Did you spend less time overall but more time per submission than you have in the past on "traditional" PCs?
- ii. Did you find the whole experience to be manageable and unstressful enough to be willing to do it often, by which I mean significantly more often than you'd be willing to serve on traditional PCs?
- iii. Was the opportunity to submit a factor in your willingness to serve?
- iv. Did you find it awkward to review your fellow PC members' submissions?

I received 15 official answers from PC members, either in blog comments or by email. These answers were:

- i. 13 YES. One of the 2 who didn't answer YES was a junior person who had never served on a traditional, one-tiered STOC PC.
- ii. 12 YES
- iii. 5 YES. Of the 10 who said NO, most added that it *would* be a major factor if 2-tiered PCs became the norm. They'd expect to serve often and thus would expect to be able to submit. Indeed, PC submissions are standard practice in CS communities in which 2-tiered PCs are the norm.
- iv. 14 NO

In addition to those 15 formal responses, I heard informally from many PC members that they appreciated the light workload and would not have agreed to serve on a traditional PC because of the time commitment it would entail.

The most important lesson learned from this “experimental” PC was the tremendous value of the EC (*i.e.*, first-tier PC) members’ having devoted our time and effort to coordination and oversight of the process rather than to reviewing submissions. On a traditional PC, it is only the PC Chair who performs these functions, but coordination and oversight is simply not a one-person job or even a two- or three-person job. In STOC’13, EC oversight combined with the requirement that a first round of reviews be done early so that clear accept/reject decisions could be made early ensured that potential problems were discovered early and that we had none of the last-minute scrambling that PCs often run into. Every pair of comparable submissions was actually compared, submissions that looked too good to be true (or that required vetting by outside experts for some other reason) were sent to outside experts way before the notification date, and substitutes for delinquent reviewers were found before it was too late. I strongly recommend that any PC Chair who chooses *not* to have a separate executive subcommittee that concentrates only on this type of oversight figure out a way to devote sufficient committee time and effort to it – no Chair should try to do it all himself.

The STOC’13 committee did not have any face-to-face meetings, even of the EC. We did all of our business via HotCRP, email, phone, and Skype. This worked extremely well and actually had a few big advantages over traditional, two-day, F2F PC meetings. During traditional F2F meetings, PC members often have to sit idly through long, heated discussions about submissions that they haven't read; this is an extremely annoying experience that makes people antsy and leaves them in bad moods by the time they rejoin the deliberations. With an all-online PC process, multiple heated discussions can take place in parallel, and no one has to sit around losing patience. More interestingly, the knowledge that there will be a F2F meeting shortly before the notification date encourages the PC to put off tough decisions until the meeting; unfortunately, they don’t delay only the decisions but also some of the work that needs to be done in order to make informed decisions. Hence the (unreasonable) requests that some PC members go back to their hotel rooms to read and form opinions (literally) overnight about submissions that they’ve not read by the end of the first meeting day (so that they can discuss these submission on the second day). There should be nothing in the process that invites delay of tough decisions – those are the decisions that take the most time and that the PC should start on as early as possible.

If 2-tiered PCs are to be a regular feature of STOC, we will need to (and will most certainly be able to) get some kinks out. The three-month period between submission and notification has three major phases: (a) bidding on submissions to review, (b) first-round reviews and early decisions, and (c) second-round reviews of and decisions about survivors of the first round. For STOC’13, we budgeted too much time for (c) and not enough for (a) and (b). Between phases (a) and (b), the EC should assign both first-round and second-round reviewers for each submission but instruct second-round reviewers to hold off until after early decisions are made;

we learned from experience that it is more difficult to convince people to take on second-round reviewing right after they have finished their first round. It will be necessary to clarify the external-reviewer policy; although PC members have much lighter workloads than PC members in traditional PCs and thus should not typically need to make as heavy use of external reviewers as traditional-PC members do, there are many submissions for which extremely specialized expertise is needed and cannot be found even on a very large PC. Finally, the PC Chair should formalize the procedures that the EC will use to make decisions that the PC members cannot make from their limited vantage points, and these procedures should be spelled out for prospective PC members at the time they are invited to serve.

3. Suggestions for further experiments

Now that STOC has survived one experiment intact, I hope that our community will continue to try new things. As I've said at two STOC business meetings in a row, STOC has barely changed during the 45 years of its existence, but computer science overall has changed dramatically. Theoretical computer science is an extremely intellectually strong community, not to mention an extremely academically successful one. We need not fear change, and we might benefit from it.

Two natural arenas for experimentation are the program and the proceedings. Here, I list some suggestions that community members have made, both before and during the STOC'13 process, about how we might approach them differently.

Program: As has been pointed out many times, a STOC PC has conflicting goals when it puts together the program. On the one hand, it wants to confer *prestige* upon the results presented at STOC; it is thus led to be *exclusive, e.g.*, by rejecting many good submissions in order to keep the acceptance rate down and by shying away from inviting talks (or scheduling any events) that are "high-risk / high-reward" in that they can't be counted on to deliver anything mathematically impressive but may nonetheless be interesting and provocative and thus lead the theory community in fruitful directions. On the other hand, the PC wants to promote *community formation* by attracting as many conference participants as possible and keeping them engaged and interested for three or four days; it is thus led to be *inclusive, e.g.*, by not pushing the acceptance rate down as far as some other CS communities have done and by using one of the four days for workshops, tutorials, and invited talks.

For many years, STOC PCs have traded off these conflicting goals by devoting three days almost entirely to two parallel streams of 20-minute talks. Plenary sessions have been used very sparingly and mainly for talks by award winners. In my opinion, 20-minute talks are not, in general, valuable enough to warrant three whole days. This is not the fault of the speakers; people in our community are honored to have their STOC submissions accepted, and they take very seriously their responsibility to give good talks. It just doesn't seem to be feasible for many of them to do so any more, given the immense breadth and depth of modern theoretical computer science. Talks aimed at specialists would be utterly incomprehensible to most of the people at the conference, but talks aimed at non-

specialists often contain 15 or more minutes of definitions and related work, and the remaining five or fewer minutes don't suffice to explain the main contribution of the paper.

This was not the state of affairs when I first started going to STOC in the 1980's; it is a result of the stupendous success that our field has had in attracting mathematically strong researchers and in using a big and ever-expanding technical toolkit. We should celebrate the fact that our community produces a vast array of technically intricate results every year, but we must grapple with the fact that the 20-minute PowerPoint presentation is not an effective medium in which to present these results.

The following potential changes to the program structure that have been suggested over the years might provide opportunities for more effective talks. These are just a start, and more creative thinking about how to put together a great STOC program is in order.

1. Devote part of each day to plenary talks that are long enough to reach a "general theory audience." Accepted submissions should be considered for plenary talks based both on the quality of the results and on the speaking ability of the authors. PCs could also invite exceptional speakers to give plenary talks rather than restricting consideration to submitted results.
2. Devote part of each day to 20-minute talks aimed at specialists. Increase the number of parallel streams of 20-minute talks so that, in combination, suggestions 1 and 2 maintain the current acceptance rate of between 25% and 30%. Note that American Mathematical Society conferences feature both plenary talks aimed at "general mathematical audiences" and a large number of parallel mini-conferences aimed at specialists.
3. Eliminate 20-minute talks altogether; accepted submissions that are not chosen for plenary talks could be presented in poster sessions. Note that this is what's done at the NIPS (neural information-processing systems) conferences, and machine-learning researchers (including those who go to STOC and well as NIPS) love it; a young machine-learning researcher told me this year that NIPS poster sessions are his favorite venue in which to present papers, "because [he gets] tons of immediate, high-quality feedback."
4. Reconsider the question of whether acceptance rates should be between 25% and 30%. Does this range achieve the right tradeoff of the PC's conflicting goals?

Proceedings: Although they were not supposed to be part of "the experiment," proceedings proved to be a source of lively discussion and considerable controversy for STOC'13. The two major issues were our requiring 10-page, 2-column submissions and our forbidding one-page proceedings contributions.

Recently, STOC has allowed 20-page proceedings contributions. The EC for STOC'13 decided to return to the traditional 10-page limit on proceedings

contributions. Our reason for doing so was quite straightforward: If the theory community wants to continue to publish most major results first in unrefereed conference proceedings and, later, in refereed journals, then it should maintain a substantial and obvious distinction between a proceedings contribution and a journal paper. Twenty-page proceedings contributions make this distinction harder to maintain. Many authors will be able to fit an entire paper into 20 proceedings pages; when asked why they're not bothering with refereed journal papers, they'll be tempted to say "the STOC paper has all of the results, and it's essentially refereed." STOC submissions are *not* refereed! No one who has our community's best interests at heart should say that they are. Although the same misrepresentation can happen with 10-page proceedings contributions, a 20-page upper bound just exacerbates the problem.

Having decided on 10-page proceedings contributions, we wanted submissions to contain the same technical material that authors intended to publish in the proceedings. After a long and inconclusive discussion about how many one-column, 11-point pages are equivalent to 10 two-column, ACM-proceedings pages, we decided simply to require that submissions be 10 pages in ACM-proceedings format, *i.e.*, that authors submit precisely (with respect to length and format) what they propose to publish (supplemented, optionally, by a full paper or appendices that PC members could read at their discretion). In so doing, we adopted the submission-formatting policy used by most ACM conferences and workshops.

I received no complaints from authors about the 10-page limit on proceedings contributions but numerous complaints about two-column submissions. Fortunately, the two-column proceedings format that has plagued both authors and readers for 45 years can now be discarded: ACM has developed a one-column proceedings format intended to be read on laptops and tablets. The length of proceedings contributions remains an issue for future PC Chairs.

The one-page proceedings-contribution controversy is newer and more interesting. As theoretical computer science has grown more interdisciplinary, the fact that computer science differs from the rest of the scientific world in its use of "extended abstracts" as precursors to refereed journal papers has created problems for some of our authors. In particular, members of our community who work in economics and computation, in quantum computing, and in computational science want to submit papers to highly prestigious journals that regard "extended abstracts" in conference proceedings as "prior publication" that precludes subsequent submission to the journal; for example, *Econometrica*, *Science*, and *Nature* all adhere strictly to such rules. Some of our conferences, including Innovations in Theoretical Computer Science (ITCS) and the ACM Conference on Electronic Commerce (ACM EC) have solved this problem by allowing one-page proceedings contributions (and requiring that each one-page abstract contain the long-lived URL of a preprint of the full paper). The authors of an accepted STOC'13 submission wanted to publish a one-page abstract in the proceedings so that they could submit to a prestigious journal that precludes prior publication of extended abstracts. We did not allow them to do so, because STOC has never had one-page proceedings contributions before, and the ability to do so had not been advertised in the Call for Papers; we did not think that one set of authors should be allowed to

take such a step unless and until the entire community had been given the same option at the time of submission. We do acknowledge the problem, however, and believe that STOC should adopt a policy that enables our authors to submit to the premier journals in interdisciplinary areas.

Here are some proceedings-related suggestions for future STOC PC Chairs; once again, it is a partial list, and my main goal is to encourage creative thinking and experimentation.

5. Adopt ACM's new one-column proceedings format.
6. If the theory community continues to publish both extended abstracts and journal versions, then maintain a sharp distinction between the two by, *e.g.*, limiting extended abstracts to 10 pages.
7. Consider radically different publishing procedures, *e.g.*, those of VLDB, which are designed "to bring the way [they] publish in the DB community into line with general scientific publishing" (<http://vldb.org/pvldb/pvldb-faq.html>).
8. Allow one-page proceedings contributions, accompanied by preprints of full papers posted online at stable URLs, in order to facilitate journal publication in interdisciplinary areas. See the Calls for Papers of ITCS and ACM EC for examples of how this can be done.
9. Consider abandoning published proceedings altogether. Conference proceedings and extended abstracts as we know them are vestiges of the world before the WWW, in which print was the only way to disseminate results, and traditional refereed journals took too long for a fast-moving field like computer science. This vestigial system is confusing and ambiguous (the latter because it blurs the distinction between formal refereeing and the type of reviewing that conference PCs do) and puts us out of step with the rest of the scientific world. Theoretical computer scientists should simply write complete papers, submit them to refereed journals, and post them in preprint form in the arXiv; they should also submit to conferences. Accepted conference submissions can be presented in whatever form the conference desires (short talk, long talk, poster, *etc.*), and conferences can be as prestigious and hard to get into as the PCs wish them to be. Rather than publishing proceedings, conferences should just post programs that point to arXiv'd preprints of accepted submissions.

Once again, I hope that our community will continue to experiment with various aspects of STOC. Let's empower and trust PC Chairs and encourage them to try interesting new things.