

On changing STOC

Oded Goldreich

June 18, 2010

Summary. I strongly object the proposed transformation of STOC into a “traditional large science meeting”; I doubt it will increase attendance and interaction, and I speculate that its only effect will be the elimination of STOC, which in turn will increase pressures on other conferences (esp., on FOCS) and will severely harm them and the entire TCS community. Indeed, a main characteristic of the reviewed proposal is that it suggests a *drastic* transformation that breaks away from the entire history of TCS. Instead, I propose two alternative changes that fit within the history of TCS and still promise to improve its state.

A methodological comment. In my opinion, finding solutions for TCS’s trouble does not reduce to mimicking formats that are used by other scientific disciplines. Instead, one should start with a clear analysis of the perceived problems, determine the desired goals, and move to possible solutions. The process should be guided by the history of TCS and by its current culture and state, while the experiences of other disciplines are less important (due to fundamental differences in scientific culture).

1 The problem with STOC and FOCS

For sake of clarity, I will use a couple of methodological dichotomies, which confront notions that do not exist in pure form in reality. In reality, the conflicting notions are mixed, and the issue is one of balance between them.

My main distinction is between competition and contents. Indeed, any competition refers to some contents, and any human activity can be viewed

through the lens of competition. Still, there are contexts where the competition aspect is more dominant and others where it is less dominant. My starting point is my subjective feeling that, in recent years, *the competition aspect became more dominant in STOC/FOCS*.

Note that I'm not saying that the competition aspect was not present in STOC/FOCS in the past. On the contrary, my own experience (which dates to the early 1980s) is that STOC/FOCS were always marked by a (flavor of a) competition. But, in the recent years, I see a greater obsession with the competition aspects and a decline of interest in anything that is not competition-oriented, including the actual contents. Concretely, I hear and overhear more discussions of which paper "got in" and which paper did not "get in", which paper got which awards, etc. And, I hear and overhear less discussions of the contents of various works, what makes them interesting, and what can be "carried home" from them.

An alternative way of demonstrating what is currently wrong with STOC/FOCS is to make a distinction between the interests of the authors/speakers and the interests of the readers/attendees. In PCs (and in outside references to the PCs), you often hear references to the legitimate (or illegitimate) interests of the authors, and the entire discourse of fairness evolves around these interests. In contrast, you rarely hear any discussion of the interests of the conference's attendees, which is indeed very odd because the conference is supposed to serve the attendees.

Still... In spite of the foregoing critique, one must bear in mind that STOC/FOCS does serve a major need of our community. For better or worse, our community views conferences (rather than journals) as the primary publication outlet, and STOC/FOCS serves as the first tier of that system. Thus, eliminating one of these outlets (as suggested by the reviewed proposal) is likely to have a devastating effect.

In general, one should realize that organizations and cultures develop hand-in-hand, and that a *drastic* change in one that is not coupled by a change in the other is likely to have bad effects. Note that the reviewed proposal does not detail complementary measures aimed to change the culture of TCS, but rather seems to assume that this culture will remain intact. In such a case, the elimination of STOC as a *first tier publication outlet* will cause a real problem, since the capacity of the first tier publication outlets will shrink dramatically. The most obvious consequence will be an increased pressure

on FOCS (and on other conferences), which will hinder their functioning.

I also speculate that attendance of STOC as restructured under the reviewed proposal will drop significantly below the current attendance. Just ask yourself who is likely to attend such a conference. Note that currently more than half the attendance is due to co-authors of accepted papers, and that these people are unlikely to submit and attend the restructured STOC (which will lose all its prestige). Attendance by relatively local graduate students and other researchers is also likely to drop. Bear in mind that the TCS culture is different from the culture of other disciplines, and the difference is largest w.r.t the attitude towards conferences. These difference include the perception of conferences as publication outlets and the existence of a large number of conferences (of which one typically attends only a small subset).

2 Desired goals

While the reviewed proposal seems to set increasing the attendance as its primary goal, my view is that this attendance is not a goal at all but rather a means for some other goals. The goals are serving our community.

As argued above, in my opinion, the reviewed proposal will definitely fail to provide a *central* meeting place for our community, and its net effect will be merely the elimination of STOC.

As stated above, TCS conferences serve both as primary publication outlets and as meeting place (not necessarily central) for exchange of ideas. My concern is that STOC and FOCS seem to do a less good job in providing the second service (i.e., fostering exchange of ideas), and my goal would be to *regain this service without seriously disrupting the other service* (i.e., publication outlet). It seems that the reviewed proposal also has the second service in mind (as a secondary goal...), but it neglects the publication outlet service aspect.

If one wishes to drastically change the culture of TCS such that conferences are no longer viewed as a publication outlet, then one should say so explicitly, justify the need (and benefit) in such a drastic transformation, and detail how it can be effected and what will be its cost (e.g., in terms of damage during the transitional period).

3 Two possible solutions

I wish to commend the EC of SIGACT for deciding to investigate this matter (i.e., the nature of STOC), and I share its desire to see STOC regain its character as a meeting place for TCS and a forum for learning and interaction. I suggest two alternative changes that aim to restructure STOC but in a less drastic manner than in the reviewed proposal. The two suggestions go in opposite directions, but I consider each of these directions better than the current state of affairs.

3.1 A single-session program that is aimed to serve learning

This proposal is aimed to make STOC a central meeting of TCS with emphasis on the unity of the community. The plan consists of two measures.

1. Mandating a single-session program. This will establish a feeling of a shared experience both at the declarative and actual levels. I consider declarative actions as important, and note that (if candid then) they tend to effect reality.

The “cost” of this measure is cutting the number of presentations to around 60, which seems to be non-dramatic and arguably even good. Note that I assume that invited talks are avoided, which in my opinion is not a big loss because most of these talks tend to be quite useless (especially when compared to good items of the regular program).

2. Instructing the PC to compose the program while using the perceived benefit to the attendees as the main selection criteria; that is, typically, acceptance should not be viewed as an award and/or a token of recognition for deserving achievements, but rather as a service to the non-authors who may benefit (i.e., learn and/or get inspired) by the work.

I do realize that such an instruction is quite vague and open to misunderstandings and/or hard to implement. Still, I think that a clear statement to that effect will not be ignored and will have a positive effect.

Indeed, this proposal will reduce the number of papers accepted to STOC, and thus increase the pressure on STOC and FOCS, but these effects are

relatively small (especially when compared to the corresponding effect of the reviewed proposal).

I believe that this proposal will best serve my declared goal, and will thus best serve the TCS community.

3.2 Moving to 3–4 parallel sessions

Here no attempt is made to unify the community, but rather to regain real interest in the contents of the conference (i.e., serve as meeting place for ideas). Concretely, here I propose a significant increase in the size of the program, by using 3–4 parallel sessions. Indeed, this will further increase the breaking apart of TCS, unless one also introduces plenary sessions that will more than compensate for the “segregation” caused by more parallelism. In any case, the significant benefit is in reducing the pressure on STOC and FOCS (and on the authors and attendees), which is likely to result in more attention to the actual contents and less obsession with the competition aspects.

I am actually very much in favor of having a plenary session with 10–20 presentations, but realize that this proposal may raise much objection I wish to decouple it from the other proposal (of using 3–4 parallel sessions), although the two fit together very nicely.

Let me end with some back-of-the-envelope calculations of the various options. Assuming that 50–60 presentations can fit nicely into a single non-parallel session, whereas parallelism introduces an overhead of 25%, consider the following possibilities.

- 3–4 parallel sessions allow 120–200 presentations.
- Allocating 25% of the program to plenary sessions means having 15 presentation in a plenary session and another 90–150 presentations in 3–4 parallel sessions.

Let me stress that the current proposal differs from the reviewed one in maintaining a non-trivial notion of program selection.