

# Tips for future STOC PC chairs

## Julia Chuzhoy (chair, STOC 2020)

### (comments and additions by Anupam Gupta, 2022)

These were tips that Julia Chuzhoy wrote (passed on to Virginia Vassilevska Williams, and Anupam Gupta). There are other older documents: a [blog post](#) by Boaz Barak, and other documents by Sanjeev Arora, Omer Reingold, Salil Vadhan, and others. The conference SoCG (Symposium on Computational Geometry) also has a *Guideline for SoCG PC* document which might be useful, available [here](#).

#### Things to do before the deadline:

##### 1. Assembling the program committee:

- Start early. People are more likely to say “yes” if they need to commit to something farther in the future.
- As the PC becomes larger, recruiting people takes longer; the last few people take the longest. This is because at the end the invitations need to be sent sequentially, and one needs to carefully select a small number of people to cover all remaining areas, so finding the right people can be hard. Also, people have more commitments.
- More senior people are harder to recruit, it may be better to start from them.
- Before the first invitation emails are sent, need to decide:
  - Do you allow PC submissions? It may be a good idea to include this info in the email to avoid misunderstandings.
  - Approximate submission deadline (normally you’d want to follow standard timeline)
  - Approximate notification deadline
  - (important): dates of the PC meeting. You want to state these dates upfront in the invitation email, together with a statement that all PC members are expected to attend; coordinating meeting date later can be difficult. If the location of the meeting is known, add it to the invitation letter.
  - I have also said that I expect to reimburse up to \$500 per person for travel to the PC meeting (this was after figuring out the budget for the PC meeting; see below).
  - Decide on rough PC size, figure out how much budget for PC member you have, maybe good to add this to the invitation email. Note that PC meeting involves additional costs, such as booking the meeting room, meals etc. A host institution or some other local institution may agree to sponsor part of the cost.
- I have collected statistics from the last few conferences about the number of papers submitted/accepted in each area, and tried to recruit PC members whose

expertise matches those. It is still a good idea to run an almost-final list by several people from different areas to see if some area is not sufficiently covered. I found this very helpful, as I have discovered that I have missed a few areas and did not calibrate the coverage of several important areas correctly.

- I talked to several previous PC chairs and other senior members of the community who gave me valuable advice. Among other things, the PC should be diverse in several ways, e.g. in terms of gender and geography, representation of various areas, and various levels of seniority. I also used a list (that was sent to me by the SIGACT chair and is possibly maintained by Ryan Williams) of PC members of all previous STOC/FOCS/SODA conferences. (Anupam: I did not find this list, so created my own.) I also talked to several senior people from various areas in order to collect suggestions for potential PC members, including people who declined the invitation, and people who agreed to serve on the committee.
- One mistake that I made was letting the recruiting process continue for too long. Though I started early on (around Jan), I was not very consistent and was making breaks around conference deadlines etc. I ended up having to recruit the last 2 members in June-July, and that was very difficult. I got the largest number of declines at this latest stage, and these last few people took several months to recruit. By that time most people had commitments for the required time frame. People are much more likely to commit if asked well in advance. This is especially true for senior people who tend to have a lot of responsibilities.
- [Anupam: I increased the PC size to 45. Managing more people requires more effort, but then I have much more help. This has allowed me to keep the paper counts per PC member to below 35, which is manageable for each person.]
- *Time frame*: in the invitation email, I asked that they get back to me within 1 week, even if they need more time to think. For those who did not respond in 1 week, I sent a follow up email saying that if I don't hear from them in the next 4-5 days, I'll assume they are not interested. If no response was received till then I sent a final email saying I'm assuming they are not interested.

2. **Conference Proceedings**: around May-June (1 year before the conference) you should get in touch with outside vendor to coordinate electronic publication fees and deadlines. [Anupam's notes: SIGACT has created a "proceedings chair" position, so this is something the PC chair will not have to do, at least for now.] Last few years "Conference publishing" was used. Their contact person is Dirk, email: [info@conference-publishing.com](mailto:info@conference-publishing.com). Previous PC chair has put me in touch with him. They will send a proposed contract. You are supposed to approve it and CC Irene Frawley from ACM [frawley@hq.acm.org](mailto:frawley@hq.acm.org), who does the final approval.

- The budget is the responsibility of Theoryfest general chair and local committee chair and should be approved by them. They may negotiate with the company about it.
- The rest is your responsibility; check that the dates are reasonable (you'll need to approve a preliminary proceedings version a few days before the conference, make sure these are dates when you are available); check that the estimated

number of pages/papers makes sense, etc. The contract mentions “Proceedings chair” and “general chair”; in context of STOC these mean PC chair.

- Important: there is an option to have the proceedings available for free online access through the ACM digital library (see more details in “call for papers” part. If you go with it, you need to make sure that the proceedings are ready 2 weeks before the conference deadline (or whenever you want to make them available).
3. **Conference website:** is maintained by Paul Beame and/or local committee. You can send an email to him with the edited file and he uploads it.
  4. **Selecting software:** I received an email from ACM early on (early Jan) saying that ACM has contracts with both easychair and hotcrp and asking me to choose one of the two systems using a link they provided. In both cases, the hosting is done by the system, you don’t need to worry about that. I chose hotcrp because this was used by previous PC chairs and based on some comparisons I found online. You still need to go to hotcrp website and create an account. At this point it may be helpful to create a new email address, like [stoc2021chair@gmail.com](mailto:stoc2021chair@gmail.com) that becomes a contact for hotcrp and all STOC-related emails.

[Anupam’s comments: each software has pluses and minuses. I chose hotcrp, even though it did not have some features I wanted, because it had a clean interface and powerful search. Then if I wanted statistics that it did not allow direct access to, I would export the data and then manipulated it using python scripts.]

5. **Preparing the Call for Papers:** I was told that this should be ready by the second half of August, Labor Day at the latest. Note that the format is different for STOC’s that are part of FCRC. I mostly copied the CfP from 2 years ago (the last year one was part of FCRC).
  - Information that you’ll need ready: conference dates, deadline and notification date (for notification date: may need to coordinate with ICALP/COLT that have deadlines close to the STOC notification date. I followed the previous year’s timeline but had to move it to accommodate ICALP). Is there a theory fest this year? Will there be a poster session?
  - Important: there is an option to make the conference proceedings available online up to 2 weeks before the conference. This is done by ACM. In order to do this, you need to give the heads up beforehand to ACM (Irene Frawley and Adrienne Griscti), and the publishing company, so that the proceedings are ready on time. Also, CfP needs to contain a specifically worded paragraph:

AUTHORS TAKE NOTE: The official publication date is the date the proceedings are made available in the ACM Digital Library. This date may be up to two weeks prior to the first day of your conference. The official publication date affects the deadline for any patent filings related to published work. (For those rare conferences whose proceedings are published in the ACM Digital Library after the conference is over, the official publication date remains the first day of the conference.)

The wording may be different for your STOC. You should check with ACM. My contact person for this issue was Adrienne Griscti, [griscti@hq.acm.org](mailto:griscti@hq.acm.org). To be on the safe side, it's better to run the CfP by her and by Irene Frawley from ACM. I also ran it by Theoryfest chair, SIGACT chair, Paul Beame, and the whole PC.

6. **Software settings:** I used hotcrp. The PC chair from the previous STOC sent me screenshots of her settings. This was very helpful because the formatting is not always self-explanatory. *[Anupam: the hotcrp config changes, naturally, but having this starting list of setting options was very useful.]* At this point you'll need to decide how you define "hard" and "soft" Col. For example, is being in the same institution a hard or a soft Col? The formatting also includes deadlines for internal/external reviews, review formats, the submission form, and so on. As an administrator you can test the submission form by submitting some papers even before the server is open. I opened submissions at the end of September. The CfP needs to be updated at this point. (For Col: one thing overlooked by current definition is postdoc advisor/advisee; I would add those to strong Col for up to 2-3 years after the end of the postdoc.)
7. **Prepare "[Info for PC](#)" document:** This is optional but is probably a good idea and has been passed down from year to year for a while now. The document is self explanatory but will require finalizing all dates, review format, etc. It includes links to a "[Guideline for Reviewers](#)" document that also needs to be updated. You'll need to decide about whether to have shell reviews, and whether to have primary/secondary reviews. I have prepared "confidentiality policy" for both the "info for PC" document and reviewers' guide. I feel it should be repeated and stressed several times during the PC deliberations both because some people are not aware of what is expected of them and some people tend to forget and need to be reminded about it.

*[Anupam's comment: I would change the review form and the guideline for reviewers to ask them not to put the recommended decision into the main body of the paper; instead to have it in the comments for the PC. There is a tension between this policy and having more transparency: the problem is that authors understandably get agitated when their rejected paper has many of the subreviewer recommendations being "accept". The authors are not privy to the PC discussions subsequent to the reviews where maybe the PC felt that the subreviewers were not sufficiently calibrated, or other issues came up.]*

8. **PC Workshop:** Decide whether to hold a "PC workshop" together with the PC meeting. This makes the PC meeting trip more exciting for PC members and also gives them more options to cover the trip. *[Anupam: not relevant for me, since we did not have a physical PC meeting.]*

#### **While Submissions are coming in:**

1. *[Anupam: If PC member submissions are allowed: you may want to not let PC members see other submissions while they can themselves submit. In hotcrp, go to "Tags and*

*tracks”, and set “Who can see these submissions” to not be “Whole PC”. Change this back after the deadline.]*

### **After the deadline:**

1. Following advice of previous PC chairs I went over all submissions and checked if all pdf files look OK. There was only 1 corrupted file. Not sure if this is a helpful step. There were 4-5 papers in 2-column format (not allowed by CfP). I have decided to keep them and review as usual because they were short but it would probably have been OK to reject them out of hand.

*[Anupam: Also check for duplicate submissions. And for PC papers that their co-authors forgot to mark as PC papers. Finally, hotcrp allows you to check formatting. I am not sure what this does, but it does return a CSV file with paper titles and lengths, which I found useful to catch papers with zero pages, etc.]*

2. In the first couple of days I received some amount of emails from authors who wanted to update their papers, fix errors etc. It's up to you how you deal with those. I posted comments in the system as placeholders but did not update the papers themselves. Note that people who are assigned to review the papers can't see the comments until they submit the reviews. So in some cases when the comments are important I emailed them.

### **3. Bidding:** I made several mistakes.

- First, not all PC members understood correctly what was expected from bidding (or what I was expecting). Several people selected 30-40 papers they wanted to read, gave them high rating and didn't touch the rest of the papers. So if I did this again I would have better conveyed how I expect them to bid. The timing for bidding/assignments is very tight so a delay due to such miscommunication is unfortunate
- The system allows PC members to rate on whatever scale they prefer; the algorithm takes care of it. However, after the algorithm produced an initial assignment, I went over it manually and adjusted it. This was needed both because of the issue mentioned above, and because different people bid differently (e.g. for some “0” means they are OK with reviewing the paper, others left “0” on papers which they clearly cannot review). Since I'm not a machine, having different people with different bidding scales was not helpful. So if I did it again, I'd ask everyone to bid on the same scale (-20 to 20).
- I asked people to indicate papers that are related to each other so I can make sure their reviewers overlap. For such papers I added tags and comments to indicate the connection. (but keep in mind that comments are not visible to PC members who are assigned to the papers until they submit reviews).

*[Anupam: based on Julia's suggestions: I asked people to bid between -20 and 20, to put 80 positive bids, and to put -100 for COIs. Not everyone did 80, but enough did 60+ to make it feasible to assign mostly positive papers to everyone. I also told people how to sort papers by number of positive bids, which may have helped.]*

- 4. Paper Assignment:** I used the automatic algorithm from hotcrp, which I then adjusted manually.
- One thing to decide is what to do with papers that are clearly junk. I heard several different opinions, but at the end decided to assign 1 reviewer to each.
  - I spent several days going over all assignments and making sure they make sense. As you do that PC members see their assignments throughout this process so it's good to let them know to ignore them until they hear that the assignments are final.
  - As you go through the assignments pay special attention to Col's. Some PC members forget to mention some Col's and paper assignments are hard to switch later. The system allows you to mark Col's so that the PC members won't see the discussion of those papers.
  - You can also mark your own Col's. You can still override them, unless you assign paper administration to someone else. I had another PC member who administered all papers I had conflict with. Once I assigned him as paper administrator I couldn't see any info on the papers. *[Anupam: I then coordinated with the paper admins so that we could all do the same thing at the same time. They were also useful sounding boards when I needed some advice. I felt that having joint chairs is a good move in the long run.]*
  - **Important:** When doing the assignment pay attention to papers that are related to each other; try to make sure their reviewers overlap. *[Anupam: there were many groups of related papers (including some 2-part papers). I messed up on this bit, so some related papers had disjoint reviewers. (I had to email the reviewers to ask them to look over the other parts: it worked out fine, but was a bit more work for everyone.) Also, assigning overlapping reviewers is not easy with the hotcrp built-in system. This is where the gurobi-based assignment system can be useful and should be able to handle some level of overlap, say at least one or even 2 reviewers in common.]*
  - As you go over the assignments, you can mark all "-100" preferences as Col's (otherwise you'll need to take care of them later separately).
  - Repeating a point from earlier: some PC members will even end up bidding strongly on papers with COIs, so be vigilant and confirm when you have doubts. Else this will be annoying later.

*[Anupam: I used a python script (with a gurobi backend), because I wanted extra PC reviews for PC submissions, and the way to do this in hotcrp seemed less smooth. BTW, it would be good to understand the exact algorithm that hotcrp uses for the assignment, so that we can add in bells and whistles when needed. I understand that it is based on*

*work that Samir Khuller and Matt McCutchen did, which was incorporated into hotcrp by Eddie Kohler when Mike Hicks was the chair for POPL.]*

**5. Subreviewing:** keep in mind that several of the PC members probably have never served on any PC. So explicitly mention various things: e.g., confidentiality policy, how to do bidding, and what are the expectations re subreviewers (they should be used as additional expert advice, not as a way to reduce the load, e.g. when the PC member is an expert). I had several issues with the system:

- If someone declines to sub-review (for one PC member) but then wants to agree to sub-review the same paper for another PC member, or undo the decline, the system won't allow this to happen. You need to go and delete the review decline in the system.
- If someone does not agree to sub-review, PC members get an email notification. If they agree to sub-review, no notification is sent (but the system shows "review started"). I have emailed Eddie Kohler about this but I don't know if this will be fixed. *[Anupam: apparently there were issues with bounces and information leaking, etc. Issues like this arose for us too.]*
- There is a general problem when people are sent subreview requests but then can't see the paper. This usually happens when they have 2 hotcrp accounts associated with 2 different email addresses. If review requests come to 1 email address but they are logged into an account associated with another email address then clicking on the paper link won't show them the paper. Either send review request to the other address, or email them the paper directly.
- **Important:** the way hotcrp is set up is that when a PC member requests a sub-review, there are 2 mechanisms once the sub-review is completed: (1) approve the subreview, and (2) adopt the subreview. In case of (1), the sub-review won't count towards the PC member's reviews. So the system will show that the PC member still did not complete the review, the PC member won't be able to see other reviews/participate in discussions, and it is not straightforward to see whose subreview that was. I did not know about this difference and my current "info for PC" document did not reflect this. I ended up asking PC members with "approved" reviews to submit a "fake" review with just the scores and the name of sub-reviewer. This was somewhat messy and added some overhead. *[Anupam: I asked people to adopt reviews, and this worked well. As soon as someone approved a review, I would unapprove it and ask them to adopt. Happened only 2-3 times.] [Anupam: It would be great if hotcrp would automatically add in the name of the subreviewer into the "comments for PC" when you adopt a review.]*

**6. Logistics of PC meeting:** where the PC meeting will be held (the physical room), hotel for PC members, etc. Start early! I did not keep track of this, and ended up having to rush at the last moment. Get in touch with Irene Frawley from ACM to get started (I didn't realize they need to be involved until the last moment). They need to approve contract with the hotel. Try to get more funding for the PC meeting from local

Universities and such. I had an admin person at TTIC working on this but apparently ACM can also help with this. I was worried about the logistics of fitting a large group (36 PC members) into a single room where they can all talk and hear each other comfortably so I did site visits at hotels and also experimented with room configurations at TTIC. Budget-wise: I had \$20k from ACM for the PC meeting. I have committed to paying \$500 per PC member for travel expenses. This left only \$2500 for the meeting itself which is way too little. I got help from TTIC to cover the rest. Depending on whether you hold the meeting itself at a University or a hotel the cost may vary.

- 7. Shell reviews:** shell reviews are asked to indicate who reviews the paper and the confidence level. I am not sure if the sub-reviewer's name is helpful, because this can usually be seen in the system (unless the invite was not sent through the system). For confidence levels, there are 2 ways to enter them: in the confidence level field of the review form, or by adding some text to the "comments for PC" section. When asked I said that either one is fine. But when going over the shell reviews, the first option works much better. You can see the confidence level grades immediately when you open the paper, reading the text requires more time and effort. Also make sure that the shell reviews are not submitted but only saved as drafts. Once a review is submitted the PC member can see other reviews for same paper which we don't want to happen until they submit the actual review. At the end of the day I did not find the shell reviews helpful. The final confidence levels were often different from the ones in shell reviews. It may be the case that shell reviews help the PC members get more organized but I'm not sure overall if they are worth it.

*[Anupam: I liked the shell reviews quite a bit. It was a good check-in point for the PC, and I could prod people who had been less prompt with asking for subreviews. Also, for some papers, it became clear at this point that we had low confidence in the current set of requested reviews, and I had time to ask for an extra review or two.]*

#### **Electronic Discussion:**

- The way STOC timeline is set up is that the deadline for the reviews is just before the Christmas break (mine was Dec. 20). Keep in mind that some people will not submit reviews on time and will need reminders, and that a number of people disappear completely during the Christmas break, until the beginning of Jan. So it may be good to set the deadline somewhat early so there are at least a few days between it and the Christmas break.

*[Anupam: This is a non-trivial issue for STOC. People have stuff happening over the winter break, and the period between Dec 20 and Jan 10 was very slow. Not sure what one can do about it, we just have to build in the slack into our plans.]*

- Normally, whenever a review pointed out a mistake in a paper, or suggested that results of the paper follow from prior work, I ran this by the authors and tried to get their



response, in order to avoid making mistakes. The only exception was when the paper was clearly going to be rejected regardless of this correctness issue.

- At the PC meeting, the plan was to go over all papers that were accepted, and all undecided papers. I was told beforehand by past PC members that one cannot hope to go over more than 150 papers in a single PC meeting. I ended up with more papers than that (around 156).
- HotCRP has a very helpful system of tags. But I have made several mistakes when using it. First, I did not realize until late that there is an option to have read-only tags (these can be used to prevent PC members from accidentally deleting or modifying tags). Second, I didn't realize until late into the process that the accept/reject decisions were not visible to the PC members (this needs to be set in the settings), so they could see the tags but they could not see the final decisions. Third, I only realized at the end that one can use colors. This is very convenient: one can make sure that all accepted papers appear green, all rejected red, provisionally accepted blue etc. This makes life much easier.
- I asked each discussion lead to send me 1 slide summarizing each of their papers. The slide had to have a short summary, pros, cons and any additional comments (that should include related papers). I also asked that the slide is posted at the paper's page in the system. Before the PC meeting I figured out the order in which we will go over the papers and put all slides into files in this order. The slides were projected on a screen during the meeting.

### **Discussion Rounds (Anupam's additions)**

Here are details of the process we used, maybe more details than you wanted. They are not unique in any way, other PC chairs have used variants of the same multi-stage process. Our scores were 1-9, with 5 being the borderline score. Remember: any PC member could "revive" any paper at any time, until the very end. (And some did.)

1. We first rejected papers with no positive scores (all scores 4 or less): 100 papers, and accepted papers with all scores 7 or more: 20 papers.
2. Then we rejected papers that were not in the top 35% for any PC member: 100 more, and accepted scores of 7/7/6 (40 more). This left ~200 papers (which I referred to as "advocated"), with ~80 to accept out of these. [If I did it again, I would combine this step with the previous step. I wrote a python script to figure out this top 35% list.]
3. Then I focused on "solo" papers (with a single positive score), asked the PC member with the positive score to convince the others (or say that they were OK with the paper being rejected). This way ~60 more papers were rejected. We also accepted some more papers with significant support (e.g., 7/7/5 and 7/6/6 papers): ~25 more. So ~55 more

to accept out of ~120.

4. I also spent some time trying to get parity across areas, and across PC members. It's OK to have different acceptance rates for areas/PC members---this can happen even with randomness---but I did not want this to happen unnoticed. For some areas, I asked the area PC members to discuss and figure out the top area papers still in play. After some more discussions (and ~10 more accepts and ~20 rejects), ~90 papers remained.
5. Then I asked for champions (I should have called them "supporters", since the thought of championing can dissuade some people). But it was still good, we got about 35 championed papers. Non-championed papers were regretfully rejected; these were all pretty good papers, having passed through several rounds of decisions. [Championing was done by approval voting, and people could champion any number of papers.]
6. Then I asked people to look at the championed list and at all other papers still in play, and to mark the papers that they wanted to see discussed. These papers tagged #discussatmeeting were discussed at the PC meeting (see later), the other championed papers were accepted at this point (~25).
7. There were a few other decisions (based on late feedback, and last-minute changes of heart). That ended the pre-meeting discussion process. We were left with only ~15 papers to discuss at the end of this process.

#### **PC meeting:**

- **Logistics:** we have booked a meeting venue in Chicago, with A/V and tech person on site, and had breakfast/lunch/snacks catered. TTI helped a lot with organizing that. This was paid for by the SIAM PC meeting budget, some budget donated by TTI, and additional small budget that I asked from STOC local committee. With about 35 people present, we needed a big room and had to use microphones. Because I wanted people to sit close to each other, we put the tables into a U-shape, and at the bottom part of U there were 3 rows.
- I was told that it is very important to have a plan for the PC meeting, so I had it carefully planned. One should add a lot of extra time margins because no matter how much one tries to follow time limits, sometimes a paper does need more discussion, and discussions do take more time than planned. With 156 papers to go over, one can only spend a couple of minutes per paper. We ended up working from 9am to 8pm on day 1 (with 2 coffee breaks, 1 lunch break, but we continued discussing papers over lunch). On day 2 we worked from 9am to 6pm.
- One thing that I did and turned out to be a very good idea is ask one responsible PC member to keep track of all decisions made during the meeting. While I was doing it myself too, as I was also running the meeting, it was very helpful to have this backup.
- Also a good idea to have someone who can step in and run the meeting if you need a break.

### **Anupam's notes from the 2022 virtual meeting:**

- We only had ~15 papers to discuss at this point, so we could take time over them. We had the discussion lead spend 3-4 mins on a slide, and then there was discussion, and people voted. This took ~15 mins per paper, and sometimes more. We did not have to rush, which I really appreciate.
- For the online meeting, create a zoom waiting room so that people with COIs can hang out there: then you can let them in when the discussion is over. Do not ask them to just leave and come back --- else they don't know when the discussion is over.
- Make the other admins co-hosts for the zoom meeting. Create the zoom voting forms in advance: the co-hosts can then use them.
- Announce/discuss the voting rule up-front: simple majority?  $\frac{2}{3}$  majority? Something more exotic? Abstentions or not? We allowed abstentions, which I felt was good because it did not force people to vote on things without enough info. But what if only a handful of people had opinions, and many abstained? I felt this was OK, because the papers we were discussing were the best of the remaining papers, and even a weak signal from the experts was useful.
- A hotcrp detail: the PC chair does **not** see papers on the discussion list in hotcrp that they have a COI with, so your list may differ from others'. Tell the PC to remind you when it's such a paper, and you can go to the waiting room. [I am not sure how to advance the hotcrp meeting ticker to the missing page. Maybe this can be done by the paper admins?]

### **Things to do after PC meeting:**

- Finalize and send accept/reject decisions. If there are merges, ask the authors first.
- Clean up and send reviews to authors. (Because several conferences have close deadlines, e.g. ICALP, COLT, these may need to be sent early).
- Select best paper, best student paper, special issue invitees; appoint special issue editors (they can take care of selecting special issue papers if you prefer), and get in touch with SICOMP editor-in-chief for the special issue.
- Post list of accepted papers online.
- Around the time of the camera-ready version, the publisher needs the partition of papers into sessions with numbers and names for each session.
- About 2 weeks later they need front matter (ask previous PC chair)
- Prepare the program, assign session chairs
- Slides for business meeting (ask previous PC chair)