

POPL'11 Program Chair Report

Mooly Sagiv*
Tel Aviv University

February 15, 2011

Abstract

This document summarizes my experience serving as a POPL'11 program chair. It is meant to assist future program chairs. See [1] for POPL'11 general chair report.

I describe the process and estimate the time devoted for each phase following a brief conclusion.

1 Goals

I set two main goals for the conference: (i) to improve the quality of the reviews and (ii) to increase the number of accepted papers.

Quality of the reviews I was hoping that all papers would be carefully reviewed by at least two expert reviewers in the field. This task is complicated by the following factors:

1. POPL submissions cover a wide range of topics.
2. Certain papers are highly technical, being built on a great deal of prior knowledge that is known only to few

*Parts of this work was done while visiting Stanford University supported in part by grants NSF CNS-050955 and NSF CCF-0430378 with additional support from DARPA.

researchers, in addition to which, a conflict of interest precludes many of these researchers from reading the relevant papers.

3. Some PC members interpret the ACM conflict of interest rules in a strict manner especially when it comes to visiting research laboratories.
4. Several papers combine a few topics, making it difficult to find a single person who is an expert in all the presented material.
5. Many reviewers tend to underestimate their skills.

Accepting more papers I was hoping to accept more papers without affecting quality. This was suggested to me by my colleagues from the FOCS/STOCS community and later at the POPL'10 business meeting in Madrid.

2 Selecting the Program Committee

It took me 10 days to form a program committee of 25 members including the selection of PC

members and getting positive responses. The criteria for selecting PC members were:

1. Covering as many areas as possible and including selecting reviewers in my own research areas as I wanted to avoid the need to review papers.
2. Actively publishing in PLDI and in POPL.
3. Capable of producing excellent reviews. Here I depended on my and steering committee members previous experience. There were many excellent researchers with very high publication profile which I did not include on my PC out of concerns on their reviewing skills.
4. Making sure that the PC members were active in their field and had sufficient time to review the papers themselves rather than assign them to sub-reviewers.
5. Allowing researchers to serve on the POPL PC for the first time.
6. Balancing between academia and industry and various places on the globe.

Four researchers declined my invitations: 3 because of the requirement to physically attend the PC meeting and one because he was too busy and I had to select 4 other PC members. The result was a committee consisting of 4 women and 21 men; 9 seniors and 16 juniors; 10 Europeans, 10 Americans, 2 Asians, 2 Canadians, and one from the Middle East. I was extremely pleased from the response I got and believe we put together one of the best POPL program committees. This was later reflected from the quality of the reviews. I would add that this procedure took much less time than I anticipated.

3 Bidding for papers

PC members were asked to bid for papers. I suggested that they start only after all the full papers were uploaded and bid earlier if they were out of town. Of the 25 PC members, 24 bid for papers and I filled in for the last member.

Assigning papers to PC members

I used the HOTCRP system to iteratively assign papers in the following way: The system automatically assigned all papers to three PC members and then I looked for a paper with a bad assignment in terms of expertise. I manually assigned this paper to PC members using my knowledge of their strengths and then erased all other automatic assignments and repeated the process. In certain cases I ignored the bidding in order to ensure better reviews. Eventually, the automatic assignments became more suitable and I could accept certain automatic assignments, which made the procedure easier. The whole procedure took nearly two full days. The result was an assignment which in most cases matched the research areas of PC members, with a load of 22 to 27 papers per PC member (I decided to review 6 papers myself in order to reduce the load on some PC members). Notice that I decided to allow different loads on the PC in order to better cover the research areas. I emailed the assignment to the PC members and asked them to mark the papers for which they could provide expert reviews. About half of them responded. I used the response to slightly adjust the assignments and to decide on the number of external expert reviews required.

4 Assigning the Papers to External Expert Reviewers

I skimmed all papers and I asked for 0–3 external expert reviewers per paper¹. I used the PC members to help in certain cases. I also allowed PC members to solicit extra expert reviews although the system was set up in such a way that they could not see these reviews until they submitted their own. The number of external reviewers was determined mainly by coverage of papers by PC members. I also asked for additional external reviewers for resubmissions from other conferences in order to assure unbiased opinions. This procedure took about 10 partial days which in my opinion is a reasonable amount of time and ensures expert reviews for every paper unlike for example relying entirely on an extended external reviewing committee which POPL’11 did not have.

5 Monitoring the Reviews

I read the submitted reviews and in certain cases decided to ask for yet another external expert review. For example, when a paper contained both theory and application and an expert in the application was negative I asked for more external expert review on theory side to assure that all sides were represented. I also discussed the reviews with reviewers in certain cases. I also periodically checked the number of expert reviewers per paper, via SQL queries.

¹I did not ask for external reviews for papers which were reviewed by 4 expert PC members.

6 Electronic Discussion

During the two weeks of electronic discussion, the PC and the external reviewers decided on a list of submissions to be discussed at the PC meeting. For each of the discussed papers, a champion from the PC was assigned. I encouraged the PC to discuss as many papers as possible. We also used this time to correct our reviews and to complete a form responding to the author response to our original reviews. In certain rare cases, the PC decided to ignore reviews by external experts which were off target and I had to ask for last minute external reviews. We also corresponded with the authors in order to better understand their papers. In certain cases, I forwarded revised papers from the authors to the PC and the external reviewers. The procedure was completely anonymous².

Another issue that was discussed was the scope of papers. We decided to consider all papers that meet the standard of quality criteria and were of interest to some of the POPL attendees. We considered some practical papers which are typically submitted to PLDI and some theoretical papers that are typically submitted to LICS.

7 PC meeting

All PC members attended the in person PC meeting which lasted two days. 91 papers were discussed.

We discussed papers in the following order: those with no conflicts, those submitted by non-PC members, and finally papers co-authored by

²In particular, the PC did not access material from the authors home page.

PC members. There was no enforced order between the groups.

The idea was not to impose an order from highest to lowest and consider each of the papers on its own. The intention is to increase the possibility of accepting more high quality papers³. Also, I wanted to allow the PC members to obtain shared experience and therefore ordered the papers without conflict first. Finally, I tried to combine papers with similar conflicted PC members to decrease the time needed for PC members to leave and join the discussion.

For each paper, we decided whether to accept, reject, or table it for further discussion. During and after the meeting, PC members recorded feedback to authors about the discussion including constructive comments (This was done in a special review form field). We felt that it is important to also provide feedback to authors of rejected papers.

The meeting ran very smoothly. We accepted 47 papers and conditionally accepted three more papers: two because of serious issues of correctness identified by the PC members and one because of presentation. Eventually, the program included 49 papers (the authors of one of the conditionally accepted paper could not complete the proof).

8 Creating the program

My initial idea was to allow only some parallel sessions following POPL'93. However, the PC rightfully convinced me that this would be unfair to those authors whose papers appear in parallel sessions and I decided to have a full parallel sessions. Talks were kept rigorously to thirty

³See notes by Kathleen Fisher.

minutes in order to enable the audience to easily switch rooms between sessions.

9 Statistics

The PC Chair report available from [2] contains statistical data on the review. In this document, I focus on the study of expert reviews which are provided by experts in the field. The system allowed 3 values for expertise: “low”, “medium”, and ‘expert’. 95% of the submissions received 3 or more reviews with medium or high expertise. 97% of the submissions were reviewed by at least one expert reviewer and 77% of the submissions were reviewed by at least two expert reviewers.

I was originally concerned that the high number of expert reviews will lead to more rejections. Table 1 shows that this is not the case. As expected none of the 3% of the papers without any expert reviews were accepted. We can see that papers with 1 to 3 reviewers have similar acceptance rates and the rates slightly increase for papers with 4 and 5 expert reviewers.

10 Mistakes

I probably made several mistakes in the procedure, but most of them were corrected by the general chair Tom Ball, and the program and steering committee.

The reviews includes 5 scores (“Reject”, “Weak reject”, “Weak accept”, “Accept”, and “Strong accept”) in contrast to the normal 4 scores used (“Reject”, “Weak Reject”, “Weak Accept”, and “Accept”). This is rather unusual and may have confused some authors and referees. I was hoping that reviewers with “Strong Accept” will fight strongly for a paper. I also used the scores to consider papers with “Ac-




Table 1: Acceptance rates vs. Number of expert reviewers per submission. The general acceptance rate is 23%.

cept”, and “Strong Accept” to be discussed. It may be better to use a different way to allow referees to support a paper.

In the future, it may make sense to determine criteria for papers to be presented in a plenary session. I now understand that the idea to conditionally accept papers was wrong as it took too much time of the PC’s time and prevented us from publicizing the program at the required time.

My biggest mistake by far was the way PC papers were handled which led to unnecessary rejections and caused some really bad feelings among all of us on the PC. For example, we had to compare a strong paper written by PC member to a rejected paper on the same topic and decided to reject both in order to be fair. Also, PC members had to lead the discussion

right after they found out that their paper was rejected. In retrospect, it would have been better and fairer if I had assigned the papers to external experts outside the PC and discussed them prior to the PC meeting with the external experts not involved the PC members at all. This procedure can be implemented with and without an extended reviewing committee. The result can still be announced after the meeting.

11 Conclusion

I was surprised by the reaction of many of the authors, especially authors of rejected papers which told me that the reviews which they received this year are the best they ever had.

The positive feedback I got from authors, including authors of rejected papers, leads me to recommend that future PC chairs maintain the principle of having at least two expert reviewers per paper.

Acknowledgments

I thank Tom Ball for inviting me to serve as a POPL PC chair and assisting me throughout the process.

The POPL steering community, Peter O’Hearn and Greg Morrisett for their help in the PC selection process.

Philip Wadler for comments on this manuscript.

Greta Yorsh for her excellent suggestions during the first (long) night papers were assigned, and for her help in drawing up the program.

Jeff Foster, who hosted the PC meeting and the following workshop.

Jan Vitek, who chaired the discussion of papers about which I had a conflict of interest.

Byron Cook for volunteering to conduct a great student presentation session at the conference.

Eddie Aharonovich and Shachar Itzhaky for the help with the HOTCRP system and with statistics.

Tom Reps reminded me to make sure that PC members respond to the author responses.

Alex Aiken and Kathleen Fisher for a wealth of good advice throughout the process.

Gilit Zohar-Oren for assistance throughout the process.

My wife Tamar Sagiv for supporting me and taking care of our lovely daughters.

I am deeply grateful to the 379 external reviewers and the PC members for providing excellent reviews which exceeded the expectations of authors of accepted and rejected papers. I am also grateful to the POPL'11 program committee for making POPL'11 a success and helping in every possible way.

References

- [1] Thomas Ball. POPL'11 general chair report.
<http://www.cse.psu.edu/pop1/11/Slides/generalchairreport.pdf>,
January 2011.
- [2] Mooly Sagiv. POPL'11 PC chair slides.
<http://www.cse.psu.edu/pop1/11/Slides/pop11chairreport.pdf>,
January 2011.