

# Principles of POPL

Derek Dreyer, John Field, Roberto Giacobazzi, Michael Hicks, Suresh Jagannathan,  
Mooly Sagiv, Peter Sewell, Phil Wadler  
(on behalf of the POPL Steering Committee)

## Goal

Our goal is create a contract between POPL organizers and POPL authors that defines essential organizational and reviewing policies. We wish to establish clear expectations for authors while allowing plenty of leeway for organizers to innovate. We anticipate that the contract will change over time, but when it does, we'll inform the community and provide a justification for the change.

The remainder of this document is organized topically. Each topic has two subsections:

*Prescriptions* and *Suggestions*. *Prescriptions* are firm policies; we expect that the organizers for each incarnation of POPL will adhere to the policies. In cases where the organizers feel an exception or change is warranted, they must first consult with the SC. *Suggestions* are best (or at least very good) practices that we expect organizers to strongly consider.

## Definitions

*CFP*: Call for papers

*COI*: Conflict of interest

*DBR*: Double-blind reviewing

*EC*: SIGPLAN Executive Committee

*OC*: POPL Organizing Committee

*PC*: POPL Program Committee

*SC*: POPL Steering Committee

*SIGPLAN*: ACM Special Interest Group on Programming Languages

## Topics

# Conference Organization

## SC Composition

### *Prescriptions*

The POPL steering committee for the period from POPL  $X$  to POPL  $X+1$  consists of:

- The current and past SIGPLAN chairs
- The current and past SIGPLAN vice-chairs
- The General and PC chairs for POPLs  $X-1$ ,  $X$ , and  $X+1$

In addition to the above formal members, the general and PC chairs for POPL  $X+2$  are typically invited to participate in most SC discussions. The chair of the SC for calendar year  $X$  is the General chair for POPL  $X$ .

## Selection of Organizing Committee

### *Prescriptions*

The General and PC chairs for year  $X+2$  are selected by the SC for year  $X$ , subject to the approval of the EC. The General and PC chairs select other members of the OC.

### *Suggestions*

In recent years, POPL's "junior" OC members (Treasurer, Publicity, Student Events, Co-located Events) have held three year terms to provide continuity, and have generally chosen their own successors, in consultation with the SC. The General and PC chairs may at their discretion appoint other members of the OC.

## Conference Venue

### *Prescriptions*

The conference venue is chosen by the General Chair, in consultation with the SC. In general, we will strive to identify venues that reflect the diversity of the POPL community.

### *Suggestions*

In recent years, POPL has rotated between the East Coast of North America, the West Coast of North America (in both cases, we interpret "coast" broadly), and Europe. The SC expects to organize an instance of POPL in Asia in the near future.

## Program Committee Composition

### *Prescriptions*

The PC is selected by the PC chair in consultation with the general chair and the POPL SC. The composition of the PC is subject to approval by the EC Vice-chair.

Topical, personal, and institutional diversity is critical to the long term vitality of POPL. In general, POPL adheres to the SIGPLAN Diversity Policy. Some particulars:

- The group consisting of the PC plus the General and PC chairs should have no more than 10% of its members from any single institution<sup>1</sup>, except in cases where the PC chair makes a compelling case to the SC and EC that there is a reason to deviate (e.g., to ensure an adequate coverage of expertise). The definition of "institution" for multi-department or multi-site entities has been the topic of much debate, and it seems that no satisfactory simple definition is possible. When in doubt, we will use the broadest reasonable definition for which conflicts can reasonably be said to exist. For example, we consider Microsoft Research Redmond and Microsoft Research Bangalore to be part of the same institution because publishing success at the former boosts the reputation of the latter, and vice versa.
- PC members should not serve on the PC more often than every four years<sup>2</sup>.
- The PC should aim to achieve a roughly uniform distribution of seniority (from junior to senior).
- The PC chair should strive to avoid the appearance of favoring current and former students, postdocs, colleagues, or collaborators for membership on the PC or ERC (if any).
- The size of the PC should be such that, given the expected number of submissions (currently, around 250), each PC member will review 18-25 papers. Exceeding the upper bound is strongly discouraged.
- If there is an ERC, the same diversity criteria listed above should apply, scaled to the size of the ERC.

### *Suggestions*

- The PC chair of year X+1 can be invited to serve (with possibly a lighter reviewing load) for year X as soon as he or she is selected.

---

<sup>1</sup> Currently, the SIGPLAN policy is no more than two members from the same institution; however, the EC is contemplating a change that would allow the number to be relative (to PC size), rather than fixed.

<sup>2</sup> For this provision, we do not consider the PC chair to be a "PC member"; indeed, the PC chair is often a recent PC member.

## Timeline

- January, Year X-2: General chair selection
- May, Year X-2: PC chair selection
- August, Year X-2: Submit candidate PC (and ERC, if any) to SIGPLAN VC and POPL SC
- November, Year X-2: Finalize PC
- January, Year X-1: Publish PC and CFP
- July, Year X-1: Submission deadline
- September, Year X-1: Author Response Period
- October, Year X-1: Author notification
- October, Year X-1: Camera ready deadline

## Review Process

### Author Anonymity

#### *Prescriptions*

POPL will commit to use of a lightweight double-blind reviewing, starting in POPL '14. The data from Mike Hicks' [report from POPL '12](#) indicate that the community favors DBR. The primary goal of light DBR is to *help PC members review papers with minimal bias, not to make it hard for them to discover authorship if they try*. The process should be such that authors are able to withhold their identity, and reviewers are able to avoid learning their identity. "Lightweight" means at least the following:

1. Paper submissions should not have the author names listed and references to previous work should be in the third person.
2. Apart from the above, authors are not required to "hide" their submissions: they can put them on web pages and give talks about them, as usual.
3. Authorship will be revealed to the reviewing PC member after he/she has submitted a review (which they can subsequently update).

Given the use of DBR, the PC chair must vet any external reviewers suggested by a PC member who has not yet reviewed the paper.

#### *Suggestions*

Provision (2) above is intended to prevent DBR from inhibiting normal dissemination of scientific

ideas. Authors should not, however, take it as a license to explicitly lobby the PC or likely reviewers on behalf of their work; doing so would be contrary to the intent of DBR. While we will not attempt to codify or police such behavior, PC chairs may wish to remind authors of their obligation to live up to the spirit as well as the letter of light DBR.

## **Conflicts of Interest**

### *Prescriptions*

Authors and PC members must adhere to SIGPLAN's [conflict of interest policy](#).

### *Suggestions*

Conflicts of interest can be tricky to manage in the presence of DBR. PC chairs may find the following conflict of interest policy, from Michael Hicks' [POPL '12 DBR FAQ](#), useful.

*Using DBR does not change the principle that reviewers should not review papers with which they have a conflict of interest, even if they do not immediately know who the authors are. Quoting (with slight alteration) from the [ACM SIGPLAN review policies document](#):*

*A conflict of interest is defined as a situation in which the reviewer can be viewed as being able to benefit personally in the process of reviewing a paper. For example, if a reviewer is considering a paper written by a member of his own group, a current student, his advisor, or a group that he is seen as being in close competition with, then the outcome of the review process can have direct benefit to the reviewer's own status. If a conflict of interest exists, the potential reviewer should decline to review the paper.*

*As an author, you should list PC and ERC members (and any others, since others may be asked for outside reviewers) which you believe have a conflict with you. While particular criteria for making this determination may vary, please apply the following guidelines, identifying a potential reviewer Bob as conflicted if*

- *Bob was your co-author or collaborator at some point within the last 2 years*
- *Bob is an advisor or advisee of yours*
- *Bob is a family member*
- *Bob has a non-trivial financial stake in your work (e.g., invested in your startup company)*

*Also please identify institutions with which you are affiliated; all employees or affiliates of*

*these institutions will also be considered conflicted. If a possible reviewer does not meet the above criteria, please do not identify him/her as conflicted. Doing so could be viewed as an attempt to prevent a qualified, but possibly skeptical reviewer from reviewing your paper. If you nevertheless believe that a reviewer who does not meet the above criteria is conflicted, you may identify the person and send a note to the PC Chair.*

The PC chair may find it useful to identify another PC member to handle papers for which the chair is in conflict. In the past, the General Chair has sometimes played this role.

## **Expert and External Reviews**

POPL is committed to identifying expert reviewers for every submission, insofar as possible. By "expert", we mean a reviewer who is very well versed and current in related work in the field. Authors gain confidence in the outcome of review decisions when expert reviewers are involved. That said, we believe that well-informed, but non-expert reviews also play a significant role in acceptance decisions: they represent the majority of the future readership of a paper, and involving them mitigates against topical balkanization.

### *Prescriptions*

The PC chair should strive to identify at least two (PC or external) expert reviewers for each paper that the PC chair deems to be a serious contender for acceptance, but reserve the right not to do so in rare cases where a sufficient number of willing external reviewers can't be identified. The chair should also strive to identify at least one (PC or external) informed non-expert reviewer.

The PC chair should strive to ensure that all reviews are available prior to the author response period, unless (in rare cases) the response itself raises critical new issues that justify an additional review.

External reviewers will be encouraged to participate in online discussion of the papers they've been assigned; however, they should not see unrelated online PC discussions.

### *Suggestions*

Several successful formulas for obtaining expert reviews have been used in past incarnations of POPL:

- The PC chair designates and announces an External Review Committee (ERC) in the CFP.
- The chair (sometimes in consultation with the PC) identifies a roster of external reviewers after submissions have closed.

- The chair appoints a "guardian" PC member for each paper, whose job is to identify (in consultation with the chair) external reviewers

Each of these processes has different strengths and weaknesses, and we leave it to the discretion of the PC chair to choose the process he/she is most comfortable with.

The PC chair may encourage (or even require) PC members to identify candidate external reviewers for each paper they bid on. Authors may also be allowed to suggest candidate expert reviewers. This can be helpful for niche areas where experts may not be well known to the PC. The PC is under no obligation to use suggested reviewers.

## **Submission of Supplementary Material**

### *Prescriptions*

Authors will be allowed to submit supplementary material (proofs, software, datasets, etc.) at the time of submission. The PC is allowed, but not required, to consult this material.

The PC chair must allow two forms of supplementary material to be submitted: anonymized material, which may be made available to reviewers along with the submission, and non-anonymized material, which may only be made available to reviewers after they have submitted their initial reviews. Both are useful in different circumstances: the former is appropriate for materials (e.g. proofs, technical appendices) that are easy to anonymize and may aid expert reviewers in assessing the technical correctness of a paper, whereas the latter is appropriate for materials (e.g. software, datasets) that are difficult to anonymize.

## **Evaluation Criteria and Acceptance Ratio**

### *Prescriptions*

There will be no numerical limit on number of acceptable papers. In 2010, there was a community discussion which concluded that a desirable acceptance rate is approximately 25% (at the time, this corresponded to about 50 papers). The consensus was that this rate could be achieved without lowering the quality of accepted papers. For the record, the acceptance rates for POPLs since 2010 are as follows: 2010: 18%, (39/207); 2011: 23% (49/209); 2012: 21% (44/205); 2013: 18% (43/238). More historical acceptance data can be found [here](#).

### *Suggestions*

- The primary job of the PC is not to attempt to assess the long-term "importance" of work, but to vet the papers for exposure to the wider PL community, which will ultimately render

judgment on importance.

- The PC should focus its deliberations on:
  - whether there is a genuine research contribution
  - whether the approach is fundamentally sound
  - whether the community will benefit from reading the paper
  - whether the paper (and supporting material, if any) contain sufficient information for others to reproduce and build on the results
- The committee should lean toward accepting papers that are controversial, that is, ones that, after discussion, still have both a strong advocate and a strong detractor.
- The committee should lean toward accepting papers that explicate their results clearly.

## PC Submissions

PC submissions can be problematic to manage, but we believe that on balance, the benefits of allowing PC submissions exceed the costs.

### *Prescriptions*

POPL will allow PC submissions, but disallow submissions by the General and PC chairs.

PC members will not be allowed to review or discuss other PC papers; hence *all* reviews of PC submissions must be external. The PC chair will designate the roster of external reviewers for PC papers, manage the review process, and make the final acceptance decision, in consultation with the reviewers. In cases where the PC chair has a conflict with the author, the chair must designate a non-conflicted senior PC member to manage the review process and make the final acceptance determination.

SIGPLAN requires that PC papers be held to a [higher standard](#) than other papers. For POPL, the criterion for acceptability of a PC paper is *clear accept*. We will not prescribe a specific floor on review scores, since such numbers tend to be poorly calibrated. However, the PC chair must ultimately be able to convince the SC and EC that those PC papers that were accepted were comfortably within the envelope of accepted, non-PC papers.

Acceptance decisions for PC papers will be announced at the same time as other author decisions are announced; i.e., the PC will not be aware of any PC paper decisions until the PC meeting is complete.

The policy on PC submission must be clearly explained to candidate PC members when their



participation on the PC is solicited.

## **Author Response**

### *Prescriptions*

POPL will allow author review responses. Authors will be allowed to read all reviews and scores. Authors will be strongly encouraged to be brief, but we will not prescribe a maximum response length. Authors must be allowed at least four days to respond.

### *Suggestions*

- PC chairs may choose to state that the PC is not obliged to read or respond to rebuttals beyond a certain designated length.
- The PC chair may choose to allow authors to provide feedback on review quality, e.g., via a simple rating scale, in addition to rebutting the content of the review.

## **PC Member Responsibilities**

### *Prescriptions*

PC members must commit to reading all of their assigned papers and writing their own reviews. PC members may also suggest additional reviewers, but they should not subcontract reading or review writing duties to others. The PC chair must always be consulted before additional reviewers are contacted to avoid conflicts and to ensure topical balance. PC members may share and discuss papers with students (subject to the usual confidentiality and COI provisions), and incorporate information from such discussions into their review, but the PC member is still responsible for writing the review. When a student is involved in this capacity, the PC chair should be informed so that the student receives appropriate credit.

## **PC Management**

### *Prescriptions*

POPL will continue to have an in-person PC meeting. PC members will be allowed to participate in the discussion (both online and in-person) of papers that they didn't review, and with which they have no conflicts.

The PC meeting should be preceded by an extensive online discussion period. The online discussion period should be a minimum of one week, but two (or more) is strongly preferred. During the discussion period, the PC will be allowed to see all papers for which they have no conflicts, and participate in the discussion. Online discussion must be actively facilitated by the PC chair, and (if

used) per-paper "guardians" (see above).

The PC chair will strive to ensure, prior to the PC meeting, that every *competitive* paper has at least

- three PC reviews
- at least one, and preferably two expert reviews

### *Suggestions*

- Nierstrasz's Identify the Champion (ABCD/XYZ) scoring process has stood the test of time, and we strongly encourage PC chairs to use it. However, the chair should feel free to add other reviewing criteria, with the goal of streamlining the discussion process and better calibrating reviewer baselines.
- PC members should be strongly encouraged to submit reviews as they are completed; this makes it easier for the PC chair to monitor progress and identify problems (e.g., the need for additional reviewers) early.
- The PC chair may wish to identify major reviewer disagreements prior to the PC meeting and seek second opinions where necessary to help resolve the disagreements.
- Papers that are clearly below the bar for acceptability may be identified during the online discussion period, and excluded from discussion at the PC meeting.
- Experience has shown that a two day PC meeting provides adequate time for deliberation while avoiding PC burnout.
- Instead of considering the papers in order from highest-ranked to lowest, consider the papers in a quasi-random order (see article by Kathleen Fisher in SIGPLAN Notices, 46(4):17, April 2011).
- If the most positive reviewer for a paper is external, the chair may wish to ensure that the paper is discussed on day one of the PC meeting and the outcome summarized by the PC chair so that the external reviewer can provide additional feedback prior to a final decision.
- The PC chair is encouraged to use a multi-round discussion process for papers where an initial consensus is not apparent.
- Reviewers are encouraged to use a proxy when visiting the author's web site, to preserve their own anonymity.
- We recommend that the PC chair allow all reviewers (PC *or* external) to see decisions for the papers they have reviewed before decisions are publicly announced.