

# Emery Blogger

## MISCELLANEOUS

# Reviewing Guidelines for Program Committee Members

MARCH 22, 2018 | EMERYBERGER

I prepared these guidelines for the PLDI Program Committee when I was program chair in 2016. I am posting this lightly-edited version in the hopes that it will be useful to other program chairs. Feel free to adapt some or all of it for your own use; if you do so, I just ask that you cite this page and point me to your page in the hopes of providing an easy go-to for future program chairs and PC members. Thanks!

## **Reviewing Guidelines**

All committee members are expected to:

- Personally read and write reviews all of their assigned papers. The reviews should be entirely of the reviewer's own devising. If you want to invite someone for a supplementary review, let me know and I will handle it.
- Write positive, detailed, and constructive reviews that will help the authors revise their papers and make them better.
- Not seek to break double-blind reviewing by Googling (or Binging, for those of you who do that) or other means.
- Turn your reviews in ON TIME.
- Actively participate in all on-line discussions of the papers, and (for EPC members), participate in a teleconference to discuss PC papers prior to the physical PC meeting.

Reviews will take the following questions into account:

- Is the paper well-motivated? What problem does it address, and is it an important problem?
- Does the paper significantly advance the state of the art or break new ground?
- What are the paper's key insights?
- What are the paper's key scientific and technical contributions?
- Does the paper credibly support its claimed contributions?
- What did you learn from the paper?
- Is the paper sufficiently clear that most PLDI attendees will be able to read and understand it?
- Does the paper clearly establish its context with respect to prior work? Does it discuss prior work accurately and completely? Are comparisons with previous work clear and explicit?
- Does the paper describe something that has actually been implemented? If so, has it been evaluated properly? Is it publicly available so that these results can be verified?
- What impact is this paper likely to have (on theory & practice)?
- Is the work of broad appeal and interest to the PLDI community?

A key part of ensuring quality reviews is making sure that papers are reviewed by experts. All reviewers will indicate specifically what the nature of their expertise is with respect to each paper, e.g., "I have written papers (X, Y, & Z) on this topic."

*Guardians:* One reviewer will be appointed as a "guardian" to lead all discussions and ensure that author responses are read and addressed. The guardian will also ensure that final reviews include a summary of the online and/or PC discussion to explain decisions for acceptance/rejection.

## **BIDDING**

*This section is adapted from the bidding instructions from POPL 2015.*

Bidding will be carried out in HotCRP (<https://pldi16.hotcrp.com/reviewprefs>). Each PC, EPC, and ERC member should enter a bid for every paper. A bid (called a *review preference* in HotCRP) is a combination of two things:

- an integer between 3 and -3, inclusive, or -100 for a conflict, which indicates how much you would like to review the paper.
  - 3: I would really like to review this paper!
  - 2: I would like to review this paper a lot, but it isn't one of my absolute top favorites.

- 1: I would like to review this paper more than average
- 0: I don't care one way or the other
- -1: I would like to review this paper less than average
- -2: I do not want to review this paper very much at all, but doing so won't kill me.
- -3: I really don't want to review this paper!
- a letter (X, Y, Z) that indicates how much expertise you *expect* to have concerning a paper
  - X: I expect to be an expert reviewer for this paper. Experts should be able to understand the technical content of the paper (unless the paper is particularly poorly written) and are acutely aware of related research in the area (i.e., you have written a paper on the topic).
  - Y: I expect to be a more knowledgeable reviewer than most for this paper, because I generally follow the literature in this area.
  - Z: I am an outsider. I do not expect to have any special knowledge of the topics discussed in this paper.

Positive numbers in your review preference mean you have greater than average desire to review the paper. Negative numbers mean you have a less than average desire to review the paper. A score of -100 means you think you have a conflict. Examples:

- A preference of 3X means you really want to review the paper a lot and expect to be an expert.
- A preference of 2Z indicates you want to review this paper somewhat less than the one you scored a 3, but you still want to review it a lot and you expect to be an outsider.
- A score of -3X means you really do not want to review this paper at all, but expect to be an expert on the topic.
- A score of 0Z means you do not care very much one way or the other whether you are assigned this paper or not, and expect to be an outsider.

There are probably a number of ways to game this preference system. Please don't try. For example, if you assign 20 papers a 3 and every other paper a -3, you won't get those 20 and I won't know which ones you really want or don't want. (We will automatically check the distribution of scores, and if you do this, you will make the program chair unhappy.) I'd like everyone to be excited about the stack of papers they receive to review but naturally, I will have to balance that against the need to ensure papers have proper expertise assigned to them.

When bidding, you won't have to read the entire paper (though, of course, you are free to look at any paper in depth when bidding), so you are only estimating your expertise. If, when you review a paper, you find you have made a mistake in your estimate, that is just fine, and is bound to happen from time to time. If you find yourself downgrading your expertise from an X, we might find a paper suddenly lacking expert reviewers. In such a case, feel free to alert the PC chair. I'll see what I can do.

**Entering bids in HotCRP:** There is more than one way you can enter bids in to HotCRP. One way to begin is to go the reviewer preferences page (<https://pldi16.hotcrp.com/reviewprefs>). There, you will see a list that shows all submitted papers. You may enter your preferences in the text boxes here. Alternatively, you may flip through the paper pages (use keys k and j to flip forwards and backwards through the paper pages efficiently). If you go through the papers in numeric order, flip a coin first to decide whether you will go through them back to front or front to back. You may also upload preferences from a text file; see the "Download" and "Upload" links below the paper list on your review preferences page.

## SHEPHERDING

PLDI 2016 will be shepherding *all* accepted papers. Please write your reviews taking this into account. We are doing this for a variety of reasons, including improving paper quality and letting us accept papers with flaws that can easily be fixed. Shepherding will enforce that all *\*minor\** changes requested by reviewers are incorporated in the final paper. Making this apply to all papers means that there will be no stigma attached to having a paper shepherded.

This approach will let us accept papers, for example, that do not cite certain related papers that reviewers feel should be discussed. We can also require that authors address minor stylistic issues to enhance readability, so those kind of things should not be deal-breakers for acceptance.

However, there is a limit to how much we can expect of the shepherding process. For example, if a paper's evaluation is unacceptable, that is probably not something that can be salvaged during shepherding. The same is true for cases when the technical core of the paper is impenetrable.

In your reviews (either in the comments to authors, the PC, or both), please feel free to point out things that will need to be addressed during shepherding.

## FAQ

1. Who is reviewing PC member submissions?
  - The new *External Program Committee* will be responsible for reviewing all PC submissions.
2. What's the role of the External Program Committee?
  - The External Program Committee is an innovation that was recently approved by the PLDI Steering Committee. Its aim is to guarantee an extremely high quality reviewing process by forming a committee composed of the leaders of our field. This approach is inspired by the standard practice in the systems community, where senior members of the community are invited to serve on a "light" Program


Committee that reviews fewer papers (e.g., 10) than the “heavy” Program Committee and does not attend or participate in the physical meeting. In addition to providing a group of distinguished experts who can be counted upon to provide expert reviews across the areas of PLDI, the External Program Committee will review and make the decisions for all Program Committee submissions, making its job incredibly important.

3. What is the role of the ERC?
  - In a departure from recent PLDIs, this year the ERC will serve primarily as a stable for obtaining expert reviews as needed, and not as a load-shedding mechanism or as a means of handling PC submissions (which now are handled by the EPC). The ERC is also going to be wider than usual, including experienced senior graduate students.
4. How long should my reviews be?
  - You should aim for your reviews to be approximately 500 words long. HotCRP has a feature that enables searching for reviews by the number of words: you can see all of your reviews with fewer than 500 words by entering this in the search bar: “re:me:words<500”.
5. What does “expertise” mean for the purposes of reviewing?
  - You should enter a sentence or two explaining what your expertise level is for each paper you review. The working definition of “expert” is that you have written one or more papers on the topic – you should indicate the titles and dates. Knowledgeable means that you follow the literature on this topic but may have missed recent developments.
6. Are we doing two rounds of reviewing?
  - Yes. There will be two rounds, immediately followed by an author response period.
7. When will authors be unblinded?
  - Only *accepted* papers will be unblinded to preserve the integrity of double-blind reviewing for future submissions. The author responses will also be anonymous.
8. Can I enlist a student (or trusted colleague) to work with me on reviewing a paper?
  - Briefly, no. In more detail: every committee member is expected to write their own, independent reviews for every paper. If you believe that having a student do an expert review of a paper would be helpful, please let me know and if it is appropriate (e.g., there are no conflicts), they can always be invited as an expert reviewers. Also, please do not distribute your assignments to anyone without vetting by the chair; because of double-blind reviewing, there may be authorship conflicts you are not aware of. Just send the chair a note. Because having students write reviews is an important part of the training process, this year we are specifically opening up the ERC to senior graduate students for this exact reason. If you do not want a student to actually submit a review but just want them to write a “test” review (one that will not actually be sent to the authors), that can also be arranged, but again, please vet this with the chair.
9. How should I avoid breaking double-blind reviewing when searching for related work? For instance, reading the cited previous paper that sounds like it is most technically similar, and finding a figure that’s identical to the paper I’m reviewing is a pretty strong indication about authorship.
  - Inadvertent discovery of authorship is sometimes unavoidable. Here are some ways of reducing the risk of stumbling across something and thus breaking double-blind.
    - Initially read the paper off-line and write a preliminary review assuming that the authors have done their homework properly (in terms of scholarship, citing previous work, etc.).
    - Look for related work as a matter of due diligence after writing that review, and revise if needed.
    - To avoid accidentally unblinding the authors, don’t type the title into Google.
10. How long will the author response be? Will there be a hard limit on the number of words? (I have lots of questions I’d like the authors to answer.)
  - The author response will only have a soft limit, but reviewers are **required** to only read the first 600 words (roughly one page of text). Here’s the message that will be sent out to the authors.

*The authors’ response should address reviewer concerns and correct misunderstandings. In particular, respond to explicit questions by reviewers. Make it short and to the point; the conference deadline has passed. Try to stay within 600 words. You may write more words but reviewers are only required to read the first 600 words, so address your key points first.*

Advertisements

New



[Report this ad](#)



[Report this ad](#)