Writing reviews for systems conferences

Timothy Roscoe
ETH Zürich

March 2007

1 Introduction

These notes arose from a conversation with Rebecca Isaacs about how to review papers. She asked me to write some notes that the 2007 SOSP Shadow Program Committee might find helpful.

It is a bit presumptuous to claim to know how to write a good review for a systems conference. In particular, I'm inherently unqualified to say whether I write good reviews or not, other than that I continue to be asked to be on the occasional PC. I think I write better reviews now than I did when I was younger and less experienced, and some of the mistakes I've tried to avoid in recent years are tackled here.

Consequently, these notes should be read as, firstly, a description of how I aspire to write reviews, and secondly, as a discussion of the kind of review I would like to receive from a PC. Although they're written in an imperative style, they're not intended to be prescriptive.

2 What's a review for?

A paper review actually serves a remarkable number of different purposes, though it's rarely mentioned. Here are some examples:

- A review serves as a record which partially documents and justifies the committee's decision to accept or reject the paper, and its reasons for doing do.
- A review provides feedback to the authors on how to improve the paper or otherwise proceed with their work, regardless of whether it's accepted or not.
- A review communicates your thoughts about the paper to other PC members. There usually isn't time to talk about all the details of a review in the PC meeting itself, and reading fellow PC members' reviews is an important time-saving technique. Typically, one reviewer will be called upon to summarize the paper, all the paper's reviews, and their personal opinion in the meeting, and others chime in with other thoughts.
- A review is a chance to get your own thoughts on the paper straight by writing them down. It's surprising how your opinion of a paper can change by being

forced to explain it. I've sometimes found a paper utterly repugnant on first and second reading, but in writing down why I hated it, found it was mostly my own prejudice and actually got to quite like it. Conversely, I've found trying to writing a positive review for a paper I thought I liked turned me off it. People who have written down a review of a paper are generally better qualified to assess it than those who have merely read the paper.

Note that a review has several audiences: the paper authors, your fellow PC members (your peers in the research community), and, finally, yourself.

3 Steps in reviewing

Everyone has their own methodology for reviewing a daunting pile of papers. I do the following, which is probably no better or worse than many, but works for me:

- As soon as you get the assignments, print them all out, and write the paper numbers on the first page of the paper. If you're a real workaholic, write the conference name as well so you don't get mixed up between the many PCs you're currently serving on.
- Preferably (if you have lots of time before the reviewing deadline), read each one in turn quickly from start to finish to get a general sense of what it's like.
- I tend to process my papers in numerical order - I've got to read them all anyway, so there's no point in cherry-picking. Since paper numbers are often assigned sequentially in submission order, often the first papers are submitted very early and can be quite random, so don't get dispirited. The middle bunch are often quite solid, and the final ones can sometimes be excellent, but can often be hurried as well. This isn't the case where pre-registration of abstracts is required (as with SOSP, for example).

  If you like, randomize the paper order now, but then retain this order. I don't do this because sooner or later I'll drop them on the floor and mix them all up.
- Starting as soon as possible, read each paper in turn carefully, scribbling notes on them in the margins with whatever comments you think of at the time. I tend to carry a bunch of papers around with me instead of a novel, and mark them up while on the train, plane, or sitting at home. You can do this in a fairly leisurely fashion (in fact, it's best to do this while you're relaxed), and if you start early there won't be much of a feeling of time pressure. If you need to look up some reference or do some other research at this stage (e.g. "I'm sure Multics did something like this..."), don't bother but just note it down. It's better to keep making progress than get derailed into reading more papers that have already been published.
- Finally, start at the beginning of the pile of papers and type up a full review of each paper (see below). By now, you will have read all the papers assigned to you at least once (perhaps twice), and have a pretty good idea how good the field is. The papers will also have steeped a bit in your subconscious for a while, and your rage at reading the ones that don't cite your work (or, worse, that scoop it) will have subsided a bit, enabling you to think a bit more clearly and write

a more reasonable and constructive review. This is the time to go off and look up references that you might need to put in your review. If you're on a plane or otherwise separated from the Internet or your trusty copy of Organick, then skip the paper and move on, and come back to it when you're better connected.

- Submit the reviews as you've done them. The chair(s) of the conference will thank you for this, because they can see progress being made. Chairs frequently agonize over whether they will have enough reviews by the PC meeting.

As you can see, I personally tend to thread paper reviewing through the gaps in the rest of my so-called life, though in practice the final writing stage is often an intensive session in front of the computer. For this style of working, the "offline review" interface of many online review packages is really useful.

Serving on a PC can be a serious time commitment. The time you spend on an individual paper varies considerably based on the paper length and quality (very bad or very good papers take less time, those in the middle take longer). In total, though, you should expect to devote several days of your time (spread out over reviewing period) to reviewing papers for a good conference.

4 Structure of a review

Reviewing forms can have varying degrees of structure, according to the taste of the program chair, or his or her motivation to wrestle with configuring the review system. However, it's a good idea to structure your review irrespective of whether the review form does this for you. The rest of this section makes no assumptions about the review form in front of you in the interests of generality.

First of all, *summarize the paper*. Give a neutral description of what you think the paper is about, where the authors are coming from, why they view the problem as important, and what they've done. This is a great way to start writing a review, particularly when you're not sure how to get started.

Second, *state what you think the contributions are*. It's rare to find a paper that completely fails to make any contributions, but it's much more common to find one that makes no useful contributions, or contributions that turn out to be flawed. In any case, state what the authors think the contributions are, or, if you think there's a contribution you think they've missed, what it is. A surprisingly common mistake among paper authors is to fail to state the contribution - if that's the case, gently point this out (see "Tone" below), but try to fill it in.

Next, give your *specific comments* on the paper by working through the scribbles you made on first (or second) reading. Often you may find your opinion has changed in the meantime, which is fine (you may even have learned something!). Some reviewers like to separate their comments into technical discussion, and then small points like typos and other mistakes, often referred to as "nits".

Specific comments which aren't nits fall into several categories, e.g.:

- Novelty: what's new about the work? Is there some related work that the authors have missed? Does the related work invalidate the contribution, or (more likely) simply change it's context or emphasis?

- How well written is the paper? Could it be made clearer? Suggestions here range from running a spell checker or improving the language, to rearranging entire sections of the paper to make it flow better.
- Are there any apparent technical flaws? If you think there are, it's better to express these concerns as questions than flat assertions (see "Tone" below).
- Are there gaps or unaddressed issues? These need not render the paper worthless, but they may be suggestions for points the authors should address next time or in the camera-ready submission.
- Was there anything you thought was really cool about the paper? It's always good to mention any moments of delight you had reading it.
- Is the paper likely to prompt interesting discussion at the conference or workshop? Are the audience likely to take away some interesting ideas they can work with, beyond simply a feeling of good, solid, but otherwise uninspiring work?
- Is the paper appropriate for the venue, given the Call for Papers? For example, papers suitable for IEEE Infocom are rarely appropriate for SOSP (and vice versa). Also, in theory at least conference papers report on complete, mature work whereas workshop papers present early work or argue a position.

Finally, provide some kind of *conclusion* at the end. If you like, summarize the good points and bad points separately, but the important thing is to give a brief recommendation for the paper and your reasons for it.

This may sound complex, but the point is not to write a lengthy essay on the paper. The review can be quite brief, but if it contains most of the above then both the PC and the authors are much more likely to find it useful (which is the object of the game).

5 Tone

Remember your review will be read by people who don't know who you are, don't know what your personality or sense of humor is like, and can't see your face or hear your voice. Consequently, tone is important.

Caution is the order of the day. Reviews are almost always anonymous, so the reviewer is in a position of considerable power and privilege, without much means of being held to account (other than by a diligent PC chair). This power and protection should not be abused - abuse in the form of bad or unfair reviews devalues the conference venue (and, consequently, one's own prestige in being on the PC). Furthermore, the protection of anonymity doesn't extend to one's fellow PC members - they know who you are, and their opinion of you will be conditioned by the reviews that you write and they read.

Of course, you should still raise any and all concerns with the paper. Most review forms have an entry of "reviewer confidence", and it's OK to put a low score here if you're not sure you're right. You can also clarify your thinking to the rest of the PC in the "PC-only comments" section which is always there: "I think there's a complete show-stopper with how they handle escaped giraffes, but I don't know enough about cookery to say for sure".

Make the review constructive. With a bit of thought, it's very easy to transform every negative comment into a constructive suggestion, and you shouldn't need to have to do with is by actually inventing new ideas, just suggesting that the authors should invent some. For example: "This system doesn't deal with unexpected vegetables" can be turned into the more positive "The paper would be much stronger if it discussed how the system deals with unexpected vegetables."

Criticize the paper, not the work itself. The paper is a description of the work, but in theory, the only thing you know about the work is what is written in the paper. You should generally assume this regardless of whether or not you know who the authors are, or whether or not you're already familiar with the work, since you're deciding whether the paper should be presented to an audience who are (mostly) unfamiliar with it.

Rather than saying "this paper doesn't cite Multics, which did everything you do and more", it's better to say something like "This paper reminded me of Multics, which seems quite similar. I would find the paper more persuasive if it stated what the authors do over and above Multics."

Also bear in mind that, in *all* cases, there is a still a remote possibility that you've misunderstood the paper. Hence, a flat assertion like "The algorithm given in the paper breaks in the presence of Byzantine faults" is risky, and may be unfair to the authors. It's better to write "The description in the paper left me worried that algorithm breaks in the presence of Byzantine faults.", preferably followed by a sketch of why you think this is the case.

Finally, it's OK to be humorous in a review, but try to stay emotionally mature, and never ridicule the paper. Bear in mind that witty comments lose much of their comedic impact when read by a disappointed author (often a fresh-faced graduate student, and most of us remember what it was like to be one of those). I recently received the following comment in a review of a paper:

> "...Also, cut Figure 1. Holy, moly, Worst Figure Ever. It's so bad it transcends bad, it even transcends being so-bad-on-purpose-it's-good, and ends up a sort of badness vortex."

While out of context of a review for a respected publishing venue this has a certain amusement value, I don't think on balance it reflects well on the anonymous reviewer. As it happens, the paper was accepted, and this comment convinced me that Figure 1 had to remain in the camera-ready version of the paper.

6 The PC meeting

The day of the PC meeting is usually long and tiring, often spent in a crowded room with the door closed. It is not uncommon for PC meetings to overrun by several hours. On the positive side, systems PC meetings are mostly good-natured and typically end with a meal at a fine restaurant. After many hours of intense discussion, a relaxing conversation with your colleagues about something other than papers is very welcome.

Remember that the purpose of the PC meeting is to select the best possible program, not necessarily to pass judgement on each paper.

As a PC member, you may feel that your work is largely done once you submit the last review, but this is not the case. Helping the committee come to a consensus on a program is as much about your verbal contributions at the PC meeting as about your written reviews. It pays to be prepared.

Before the PC meeting, you may find it helpful to do the following:

- Reread, or skim, the papers you reviewed.
- Reread your reviews. Ensure you can articulate why you are for or against the paper. Work out which papers you feel strongly about, and why.
- Read the other reviews of your papers, if available. If there is disagreement among the reviews, try to understand why.
- Read or skim related work that may have come to light since you first wrote your review.

Most importantly, be prepared to speak about each paper that you have reviewed. The PC chair will usually ask one of the reviewers to summarize each paper for the rest of the PC (many of whom have not read the paper). If this is you, first say what the paper is about, then summarize the reviews, then give your personal opinion of the paper.

Note that this is not the same as paraphrasing your own review – it's important to be familiar with the other reviews of the paper. Time is always tight at PC meetings, so brevity is important. Try to strike a balance between providing enough information while keeping the summary brief.

If you are in the enviable position of being well prepared in good time, then you may be able to look over a few papers that you didn't review as well. An informed contribution to the discussion on a paper is always appreciated, however small.

7 Conclusion

Writing a good review is important: it helps the authors do better work, it helps you to learn more about the subject being reviewed, and it makes you look good in front of your peers. Surprisingly, it can also be fun – even writing a review for a really bad paper can be rewarding if it forces you to explain concepts afresh in a new context.

8 Acknowledgements

Many thanks to Rebecca Isaacs for making me do this and for contributing the PC meeting section, and also to Tim Harris for helpful suggestions.