Reviewing Papers for ACM TOCS

Remzi H. Arpaci-Dusseau
University of Wisconsin–Madison

Abstract: *We briefly discuss how to review papers for ACM TOCS, as opposed to typical conference venues. We present what is similar (high-quality and thorough reviews that provide clear recommendations) and what is different (when to accept with revisions, how to assess the journal version versus previously-published conference papers, and making clear recommendations to the editor). Finally, we apologize for the pedantic and preachy nature of this article; we know reviewing is a lot of work, and want the community to benefit as much as possible from your expertise and effort.*

Most reviewers in Computer Science review a large number of papers for conferences. In the systems world, this includes SOSP, OSDI, ASPLOS, and many similar, high-quality venues. The question then arises: how should one review a paper for ACM TOCS (or some other prestigious journal)? What is the same as reviewing a conference paper, and what is different?

Let's start with what is the same: a TOCS review should provide the detailed and carefully generated feedback you expect in any systems review. Typically, a well-done, thorough review contains the following components:

- **Summary.** Start with a one paragraph summary of the paper. A colleague of yours should be able to read this and understand what the paper is about.
- **Recommendation.** One paragraph that clearly states what you think the outcome of this paper should be (accept/reject) and some backup sentences informing the editor(s) why you think so. If it's an accept, list the most interesting contributions that the paper makes. If it's a reject, list the major reasons for rejection.
- **Strong points.** Summarize the major strong points of the paper. One paragraph per major strong point is good, e.g., "The paper presents an excellent and new idea for disk scheduling: to reorder requests so as to lower the average disk service time. etc. ..."
- **Major weak points.** Summarize the major weak points of the paper. Again, you should be thinking one paragraph per major weak point.
- **Minor points (nits).** Finally, write down your lesser points about the paper, e.g., "page 3: graph 6 is hard to read." Include here statements about formatting, readability, writing, and related topics. They can also be technical points, but lesser ones. Note that these can be positive too (e.g., "I liked graph 6 a lot – nice job!").

Just like conference reviews, journal reviews should be substantial (and perhaps even more so, as the papers are longer), typically a few pages or longer. A paragraph or two is not acceptable and is the hallmark of a poorly-done review. Unfortunately, doing such a review takes time; there seems to be no easy way around this reality.

Of course, there are also differences when reviewing a journal paper. Unlike conferences which have a yes/no binary decision to make, with a journal there is a chance for feedback and improvement. Thus, if you feel that a paper is of high quality but could use a bit more clarity in exposition or additional experimentation, you can recommend acceptance with revisions. However, take care when using these options: do you really think the paper is of the same quality as an SOSP, OSDI, or ASPLOS paper? That is, is the problem important? The approach novel? The experimentation thorough? If not, revisions are likely not enough to save the paper; a rejection is appropriate.

Be especially wary of accepting a paper with major revisions. Why do you think that the substantial changes will save the paper, and make it of the high quality expected at TOCS? If the paper is so far away from being acceptable, perhaps it should be rejected, and perhaps it is not destined to be published within a venue such as ACM TOCS. Only if the path to revision is clear (if long) should this option be utilized.

Another difference is that that journal submissions often include work from previously published conference paper(s). Thus, as a reviewer, you must judge: what value is added to the field by the publication of the journal version? Often, there is an expectation of additional experimentation or other new material not found in the published conference version (see `http://tocs.acm.org/authors.cfm` for explicit guidelines). Thus, assess how much new information has been added, and then judge whether you think it is "enough" to warrant publication; we typically expect it to be (roughly) 25% more content (not "fluff") than the conference version. If there is little added beyond what is already published, a rejection is warranted. Further, the paper should explicitly state what is new; if it does not, please notify the editor(s).

Finally, recognize your role as a reviewer is to give feedback to the editor and associate editors as to whether to accept the paper (with or without revisions) or reject it outright. The best way to do this is to make a clear and concise summary of your recommendation, as stated earlier. For example, saying "I think this paper should be accepted to TOCS" and then stating a few reasons why is tremendously helpful; similarly, putting "I believe the paper should not be published within TOCS" at the top (and again, adding reasons why) is worthwhile. A good review takes a stance; do not just write down a series of thoughts, but rather take the time to form your own position, as if you were the final decision maker. Doing so will greatly ease the task of the editorial staff and generally lead to a better review.

Overall, reviewing is arduous, demanding, and (often) thankless work. The ACM in general and editorial boards in specific greatly value your time and hope not to burden you too often. When you do partake in the process, the entire community benefits from your thorough and careful review. Thank you again for your assistance in this critical task.