COSC 91/191, Spring 2019
Lecture 13
April 26, 2019

Scribe: J. Peter Brady

1 Conferences and reviewing

When you write a paper, you are usually planning on submitting it to a conference, and when you send papers to conferences, they do get reviewed. Today's lecture gave an overview of what conferences are like and how the reviewing process might apply to your paper.

1.1 How conferences are organized

A conference usually has several committees and officers. There's a Steering Committee, which gives overall guidance to the conference. The Steering Committee selects a General Chair and ratifies significant decisions made by the General Chair; for example, it approves where the conference will take place, when the conference will be, and whom to select as the Program Chair. The General Chair and the Program Chair are the most important organizers in the conference.

The General Chair is like the producer of a movie — they make sure the conference happens, chooses the date and location of the conference, and selects the Program Chair. Additional officers might handle some duties of the General Chair; for example, many conferences will have a Local Arrangements Chair that handles the on-site logistics including finding and booking the venue. Tom was the Local Arrangements Chair for Symposium on Parallelism in Algorithms and Architecture (SPAA) in 2000 and started the process 18 months earlier, scouting locations in fall 1998. SPAA also had a Treasurer to handle finances, a Secretary to manage the submissions and produce the conference proceedings, and a Publicity Chair. The Program Committee, headed by the Program Chair, selects the submissions for the conference. Some large conferences have other committees to handle tutorials, workshops, exhibits, and courses.

If you are submitting a paper to a conference, you are the most interested in the Program Committee. The General Chair selects the Program Chair, and usually, the Program Chair chooses the rest of their committee. For some conferences, the Program Committee is chosen by more than one individual. For example the Program Chair of the Symposium on Discrete Algorithms (SODA) selects one-third of the committee, representatives from ACM SIGACT[1] choose one-third of the committee, and a representative from the SIAM[2] Discrete Math Group selects the final one-third of the committee.

It is hard to select a Program Committee — you want to get a balance among the constituencies of the conference in terms of areas/subareas, seniority (having junior people helps give a fresh look at topics and helps them in their careers), academia versus industry, and geographic range.

Tom related a story about being asked to be the Program Chair for a conference on optical interconnects for parallel systems. He had no background with optical interconnects and wanted to make the right decision, so he read the previous three years of proceedings from the conference mainly to find out if he knew any of the authors (he did not). Since Tom knew nothing about the area and the people in it, he decided not to

[1] Association for Computing Machinery Special Interest Group on Algorithms and Computation Theory.
[2] Society for Industrial and Applied Mathematics.

accept the Chair. Tom preferred to be on the SPAA Program Committee since that was his area and it would look good on his *vita* for tenure. He turned down the Program Chair position and was accepted to the SPAA Program Committee, so it all worked out.

Program Committee sizes vary; they could be as small as a half-dozen members or as large as several dozen. Some conferences have tracks for papers with full Program Committees for each one. The Program Committee meets and decides which papers will be accepted into the conference. The Program Chair then constructs the actual program, the work of what presentation is going when. Building a session plan is more complicated than it may seem, as the sessions should be coherent; the session should be about something that all hangs together. They don't have to be cohesive, just coherent. There may also be additional constraints such as co-located conferences where presentations happen at staggered times over the conference time frame.

1.2 How submissions are evaluated

The Program Chair receives the paper submissions and divides them up among the Program Committee members for review. Each paper is assigned to three to five program committee members, who will review the papers, outsource the papers to others (for example, grad students), or do both. Usually three to six weeks after the submission deadline, the Program Committee will meet either in person or virtually to discuss the papers and decide which to accept. Each reviewer provides numerical ratings for each submission reviewed along with written comments. The first decision made is whether the paper is appropriate for this conference; typically most submissions are, since authors understand which conferences are a match for their paper.

Tom mentioned one particular type of scoring system that had two scores: the reviewer's opinion of the paper on a scale from 0–10, and confidence in that opinion on a scale from 0–3. Some scales allow negative numbers to express negative opinions. A software scoring system normalizes each reviewer's ratings; if a reviewer is easy or tough, it normalizes scores for that reviewer.

If you were selected to be a reviewer for 16 papers, you would find 2–3 will be very good and are clear accepts, 2–5 will be very bad and are definite rejects, and the remaining papers are going to be somewhere in the middle. Most reviewers will agree on the clear accepts and clear rejects, so that the Program Committee spends the bulk of its time figuring out the strengths and weaknesses of the papers in the middle. Occasionally reviewers are deadlocked on whether to accept or reject, and so they will ask additional members to review the paper to come to a consensus. Tom has seen cases in SPAA where two papers were very good, but they were similar, and the Program Committee asked them to merge the papers into a single paper. Wojciech mentioned in his field that would be unheard of, that both could be accepted and you would have two similar papers. Tom opined that the number of papers presented between a SPAA conference (small) and a SIGGRAPH conference (very large) might explain the difference.

Most papers are evaluated only on the written submission, but Tom also believes that it can also be based on how good a presentation can be made for the work. Not all committee members have agreed with him on this point, that the overriding worth is in the paper. Submissions are usually only one-shot — the paper is accepted or rejected. Authors will generally get comments no matter what the outcome, but Tom mentioned that the theory community stopped sending remarks for a few years, but have gone back to sending them again. Some conferences have a rebuttal period where the authors can rebut criticisms before the Program Committee makes a final decision. Journals are different, as authors usually get multiple times to fix issues with a paper. Journals typically have four ratings: accept, accept with minor revisions, accept with major revisions, and reject.

1.3 How to review a paper

Zobel discusses contribution as "the main criterion for judging a paper" if it contains "*originality* and *validity*" [1, page 27]. He defines originality as the extent to which the ideas presented in a paper are "significant, new and interesting" and validity as the extent to which the ideas are correct. Zobel goes on to say, "Most papers are to some degree extensions or variations of previously published work; really groundbreaking ideas are rare. Nonetheless, interesting or important ideas are more valuable than trivial increments to existing work. ... Only a truly excellent presentation, thorough and written well, can save a paper with marginal new ideas, while a revolutionary paper must be appalling in some respect to be rejected."

Tom says he's seen revolutionary papers get rejected because they are too "wacko." There has to be some tempering of crazy ideas, some excellent reasons, or great writing to get the reviewer over to the author's side. Zobel continues, "Moreover, the ideas in a well-presented paper often seem less sophisticated than those in a poorly presented paper, simply because the authors of the former have a better knack for explanation." (Or as Tom said, why we are taking this course.)

To demonstrate validity, the reviewer can ask whether the ideas of the paper are sound and are shown in a way that other scientists can verify. Zobel mentions that the quality of a paper can be determined by its bibliography [1, page 28]; listing only a few papers may be "bad scholarship," whereas citing papers without citing related literature may be hiding "a core bibliography that is far too short."

Reviewers should look for errors that don't affect the quality of work but should be fixed before publication such as spelling, grammar, written expression, bibliography errors, and errors in math. Some errors, such as consistent math errors may warrant rejecting the paper. Zobel writes that poorly written papers should be accepted if the content is significant, but incompetent papers should be rejected. Good writing alone is not enough to accept a paper [1, page 29]. For borderline papers, don't use "revise and resubmit" as a way of making a soft rejection. If you believe with some revision the paper can be a success go ahead, but if the paper is a complete mess, reject it.

If you are preparing a referee's report, you have two audiences: one is the author and the other is a journal editor or the Program Committee. Your report to accept the paper needs to persuade the committee that it is good enough, give sufficient detail, and to convince them that you've read it. If you are rejecting the paper, the report needs to provide a clear explanation about the faults.

Tom gave an example of how not to reject a paper; when he first came to Dartmouth, he wrote an article on BMMC permutations along with his first Ph.D. student and a Dartmouth mathematics professor. Tom thought that even though it was algorithmic work it had applicability to the systems realm, and so they submitted it to IEEE Transactions on Parallel and Distributed Systems. The paper sat there for two years before being rejected. The referee's report said that it had appeared elsewhere. Tom was sure that this work had not been done anywhere else and wrote a letter to the editor asking them to contact the reviewer and asking for where this type of paper had been published before. Tom never heard anything back from them; that and some other reasons led him to cancel his membership in IEEE. The article was finally published in the SIAM Journal on Computing [2].

The comments that you make in a referee's report to the author when you accept should include any errors found and ways to improve the paper. Comments for when you reject might consist of ways to get the paper up to an acceptable level or explain to the authors why their work is not good.

Review comments need to be constructive. Zobel writes that "Some referees construct flaws in papers where none exist. For example, an assessment may include generic statements that could be made almost regardless of relevance to the paper's topic ... If there is a general problem, then describe it, preferably with examples; otherwise say nothing." [1, pages 31–32] You should be reasonably polite; as Zobel says impolite

comments are “not acceptable.”

References

[1] Justin Zobel. Writing for Computer Science, Third Edition. London: Springer London. 2014.

[2] Thomas H. Cormen, Thomas Sundquist, and Leonard F. Wisniewski. Asymptotically Tight Bounds for Performing BMMC Permutations on Parallel Disk Systems. SIAM Journal on Computing, 28:1, pp. 105–136. 1998.