

# ACM SIGCOMM Student Mentoring Column

## On Program Committees and Social Networking at Conferences

Dear students: This edition of the Student Mentoring Column focuses on program committees (their composition and how they work) and the importance of social networking at conferences. The questions below don't provide comprehensive coverage of either topic; as such, we may revisit them in future editions.

I got plenty of help in preparing this edition. In particular, many thanks to *Kyle Jamieson (Princeton)*, *Ethan Katz-Bassett (USC)*, *George Porter (UCSD)*, *Vyas Sekar (CMU)*, and *Minlan Yu (USC)* for contributing answers.

### Program committees and their inner workings:

*Q: How does one get selected to be on program committees (PCs)?*

A: The conference PC chairs select a program committee; some times they do this in consultation with the conference's "steering committee". PC chairs try to balance a range of requirements when determining the PC composition, including expertise, junior vs. senior reviewers, industry vs. academia, diversity across geographical regions, etc.

One way to increase your chances of being on a PC naturally is to be visible at the venues where you would like to contribute to the PC; e.g., writing high quality papers at these and related venues, contributing to the community by asking questions after talks, signing up for volunteer/organizational work, having engaging discussions at conferences, and socializing with the community. Later in this column, we will talk the topic of social networking at conferences in greater detail.

In addition, you should let your collaborators and others who know you well--especially those who are more visible or serve on PCs you aspire to--if you are looking to serve or are happy to help with reviews. It is likely these people are asked to serve on more PCs than they accept, and PC chairs often appreciate suggestions of people to ask when someone declines serving (and it can feel easier to decline if one has suggestions ready). Also, these people may appreciate your help as an outside/co-reviewer for your expertise on a particular topic, and the PC members are generally told who helps with reviews, helping spread your name.

*Q: How does a program committee select papers to accept? I got high scores but still got rejected.*

A: First of all, a bit of background on reviewing: as you probably know, most conferences have multiple rounds of reviewing. This helps the PC narrow down the submission pool and focus its reviewing effort on the most competitive submissions. PC chairs typically filter papers in early rounds based on scores and reviews, often based on active input from reviewers; in general, papers with low scores or weak reviews across the board end up getting filtered in early rounds. Other papers get additional reviews in subsequent rounds. Most conferences have two rounds of reviewing; some may have three. Beyond the reviews themselves, most conferences have a discussion period at the end of reviewing. For top conferences, this discussion will have both an online component and a meeting. The discussion may involve weighing a paper's strengths vs. limitations, clearing up misunderstandings, or hashing out of differences between people who liked the paper and those who did not. A paper's fate depends on the outcome of this discussion and not directly on the scores in the reviews, and the discussion may change a reviewer's opinion of a paper in either direction. Often--but not always--reviewers update their reviews if the discussion changes their opinions. For papers that made it to discussion, authors are often provided with a discussion summary. This summary may be the best signal for what determined the paper's fate. So read the summary carefully!

Now back to the question. Note that rejections happen routinely to the best researchers among us. A key reason is that most top conferences reject over 80% of their submissions as there is often limited room in the program.

Realize that not all venues use the same scoring scale, and even those that use the same scales have very different submission pools. Importantly, scores are only part of the equation on what papers ultimately get accepted into the program. High scores from one or two reviewers may indicate niche excitement about the paper, but the paper also needs to be calibrated on other fronts and there are really



multiple dimensions to this. For example, relative excitement with respect to other papers in play; whether there is someone “championing” the paper or the high scores are just lukewarm; the expertise of the reviewers; how many papers the conference is willing to accept within a reasonable program, etc.

Don’t get hung up on the scores. It is often important to try to “debug” what happened and understand why the paper was rejected. Be brutally honest with yourself: are the scores high, or are they a mixed bag, or even just middling? Crucially, does your approach have an “Achillies’ heel” that a reviewer called out? Or is there just limited excitement about the problem you’re solving? In the former case, armor yourself against the objection, considering whether to do this right up front in the introduction, later in the design section, or elsewhere in related work. Ask your advisor for help armoring yourself against technical limitations or related work if you’re unsure. In the latter case, think about other applications of your approach and consider proposing them right up front in the introduction.

Take the positive feedback as encouragement, take reviewer concerns seriously instead of dismissing them, and try and “push it over the bar” the next time!

*Q: What do you look for in a paper to accept/reject it?*

A: There’s broad consensus in our community that choice of problem is a key factor. That is to say: is the problem that you’ve chosen to work on timely, pertinent to application to the field, and intellectually deep. It is important to orient your writing to highlight these aspects.

After that, we are looking for the quantum of contribution to knowledge that your submission contributes. The goal of research is to expand our body of knowledge, and so we are judging your paper on how effectively it does just that.

Next, we’re looking at the technical approach you take: is it sound, and is your experimental methodology sound (i.e. convinces the reader that your claimed solution actually fulfills its claims) and complete (i.e., you evaluate against appropriate alternative techniques, at suitable scale, and using appropriate workloads)?

Finally, a necessary but not sufficient property of accepted papers is that the writing quality is above a certain bar. This is necessary because in order to convince reviewers

(and the community) that your technical approach is sound, the reader must understand it in its entirety. This isn’t necessarily about grammar. It is more about how you make your case and the paper’s overall flow.

It often helps to prepare your paper in advance and get external feedback. This can dramatically improve your paper’s readability, and also help you address the other issues highlighted above.

### **Social networking at conferences:**

*Q: How important is the “social” networking aspect of conferences?*

A: The social aspect of the conference is almost as important as the presentations and publications at the conference! As a student, talking to other researchers is an opportunity to make new connections, get feedback for your own work, and a great chance to sharpen your “elevator pitch” for your own research. You can learn from industry participants about the importance of your problem and the constraints you may face in practice. You can learn about what others are working on and what problems others find interesting/important. This might inspire new problems that you can work on and/or create new collaborations.

On a pragmatic note, the social aspects can also help you find internships and connect with future potential employers! It is also a natural way of getting on the “radar” of senior researchers, which might be helpful for other intangible reasons (e.g., TPC invitations, finding letter writers or potential thesis committee members). Finally, it is also a great way to get to know strong Ph.D. students from other schools. You will likely meet them again at other conferences in the future, and you may develop strong long-term friendships and/or fruitful research collaborations with them.

Of course, the social aspect is easiest when you indeed have something concrete to back it up in terms of good research results, or ideas, or even questions; e.g., it helps to be “visible” at these conferences by presenting your own original papers or posters, or asking good/deep questions after talks.

*Q: I have heard a lot about the hallway track at conferences. Should I listen to talks or talk to people instead?*

A: If you’re new to the field and/or a younger student, it would be best to focus on actually listening to the talks and participating in the Q&A right after. This will help you when you prepare your own talk, and will also help you develop a broad knowledge of current state-of-the-art



in the field. Attending sessions could also help you kick off a conversation, by bringing up a paper you particularly liked, for example. But, not all sessions may be well suited for you. Thus, in some situations, especially as you get older and/or you are on the job market, it might make

sense to spend some of the time talking with others. This may help you continue conversations that you started during a break and/or provide additional opportunities for social networking.

**Aditya Akella**

*University of Wisconsin-Madison, USA*

